IBM Storage Virtualize for SAN Volume Controller and FlashSystem Family

Concept Guide



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 9, release 1, modification x, and to all subsequent modifications until otherwise indicated in new editions.

© Copyright International Business Machines Corporation 2025.
US Government Users Restricted Rights – Use, duplication or disclosure restricted by GSA ADP Schedule Contract with IBM Corp.

Contents

Figures ..... **ix**

Tables ..... **xiii**

Chapter 1. Notices **1**
Trademarks 2

Chapter 2. System monitoring **5**
Dashboard 5
Monitoring capacity 7
Performance statistics 11
Statistics collection 13
IBM Storage Insights 27
Preventive monitoring 28

Chapter 3. Pools **31**
Arrays 34
Array configurations 35
Expanding an MDisk in a distributed array 46
RAID write response time 47
MDisks 48
Troubleshooting medium errors and bad blocks 50
External storage 51
Planning an external storage system 51
Attaching an external storage system 53
Discovering an external storage system 55
Importing an external storage system 56
Managing an external storage system 57
Removing an external storage system 59
External storage system zones 60
Data reduction pools 60
Managing pools 62
Throttles for pools 62
Linked pools 63
Easy Tier 63
Easy Tier modes of operation 63
Easy Tier automatic data placement requirements, recommendations, and limitations 66
Easy Tier reports 67
Easy Tier settings for storage pools and volumes 73

Chapter 4. Volumes **75**
Basic volumes 78
Mirrored volumes 78
Custom volumes 80
Standard-provisioned volumes 81
Thin-provisioned volumes 81
Compressed volumes 83
Deduplicated volumes 84
Volume protection 84

Expand or shrink volumes ... 85
Dependent volumes ... 86
Managing throttles for volumes ... 87
Volume groups ... 87
Transparent cloud tiering ... 88
Supported cloud service providers ... 89
Security considerations for cloud accounts ... 90
Configuring transparent cloud tiering ... 90
Cloud volumes and snapshots ... 95
Transparent cloud tiering restrictions ... 101
Snapshots ... 101
Snapshot properties ... 102
Cloning snapshots ... 102
Restore or refresh from a snapshot ... 103
Safeguarded snapshots ... 104

Chapter 5. Hosts ... 107
Mapping volumes to a host ... 107
Determining the volume name from the device identifier on the host ... 109
Volume reservations ... 110
Host portsets ... 110
Host connectivity ... 111
Host clusters ... 112
Managing host clusters ... 113
Managing WWPN changes of replaced host HBAs ... 114
Managing throttles for hosts ... 115

Chapter 6. Host attachment ... 117
Using VMware ... 117
Using OpenStack ... 118
Using Kubernetes and OpenShift ... 118
Copy offload by using ODX on Microsoft Windows ... 118
SCSI Fibre Channel host attachments ... 119
IBM Power Systems AIX ... 119
IBM Power Systems with Virtual I/O Server ... 121
IBM z Systems that run the Linux operating system ... 122
IBM z Systems running the IBM z/VSE operating system ... 124
Hosts that run the Linux operating system ... 124
Hosts that run the Microsoft Windows Server operating system ... 128
Hosts that run the VMware ESXi operating system ... 132
Hosts that run the Microsoft Hyper-V operating system ... 134
Citrix XenServer hosts ... 136
Zoning and portsets ... 137
Fibre Channel port-name identification ... 140
Queue depth in Fibre Channel hosts ... 142
iSCSI Ethernet host attachment ... 143
iSCSI overview ... 143
Example iSCSI configurations ... 145
iSCSI configuration details for host connections ... 146
Setting up a host server as an iSCSI initiator ... 150
Linux iSCSI host attachment ... 152
Microsoft Windows iSCSI host attachment ... 158
IBM AIX iSCSI host attachment ... 163
VMware iSCSI host attachment ... 166
Queue depth in iSCSI hosts ... 168
Multisession ... 170
NVMe over Fibre Channel host attachments ... 172

NVMe hosts that run the Linux operating system..... 172
VMware ESXi installation and configuration for NVMe over Fibre Channel hosts..... 177
NVMe hosts that run the Microsoft Windows operating system..... 178
IBM Power Systems AIX..... 179
NVMe over RDMA and NVMe over TCP host attachments 181
Linux..... 181
VMware ESXi..... 182

Chapter 7. Policies..... 185
Provisioning policies..... 185
Snapshot policies..... 186
Replication policies..... 187

Chapter 8. Replication..... 191
High availability..... 191
Planning high availability..... 192
Implementing high availability..... 193
Host operating systems for high availability..... 196
Highly available snapshots..... 197
Disaster recovery..... 198
Planning replication for disaster recovery..... 199
Implementing replication for disaster recovery..... 200
Managing replication for disaster recovery..... 200
Migrating replication for disaster recovery to use partitions..... 202
High availability and disaster recovery (3-site replication)..... 203
Partnerships..... 206
Pool links..... 207
Portset links..... 208
Replication policies..... 208
Replication statistics..... 210

Chapter 9. Access..... 211
Ownership groups..... 211
Configuring ownership groups..... 214
Managing ownership groups..... 216
Users..... 221
Locking user accounts..... 222
Managing users and user groups..... 224
Locking and unlocking user account automatically..... 228
Locking and unlocking user accounts manually..... 229
Locking and unlocking superuser accounts..... 230
Working with local and remote users..... 233
Audit log..... 233
Login message..... 234

Chapter 10. FlashSystem grid..... 235
Implementing a FlashSystem grid..... 236
Set up a system through FlashSystem grid onboarding..... 238
Changing coordinator of a FlashSystem grid..... 239
Multi-tenancy..... 240
Configuring multi-tenancy in FlashSystem grid..... 241

Chapter 11. Storage partitions..... 243

Chapter 12. Data migration..... 247
Migrating storage partitions between systems..... 247

Moving volumes to a data reduction pool ... 251
Migrating extents using the CLI ... 252
Migrating volumes to an encrypted pool ... 253
Migration methods for image mode volumes ... 255
Migrating volumes between pools by using the CLI ... 256
Moving a volume between I/O groups ... 258

Chapter 13. Stretched systems ... 259
Stretched system configuration details ... 259
Planning for a stretched topology system ... 260
Requirements for stretched systems with Fibre Channel connections ... 261
Stretched system configuration by using interswitch links ... 263
Stretched system configuration without using interswitch links ... 265
Configuring an enhanced stretched system ... 266
Configuring an external iSCSI storage system ... 268
Manual quorum override in a stretched system ... 268
Converting a volume to a stretched volume ... 270
Changing the system topology ... 270

Chapter 14. Legacy copy services ... 271
FlashCopy function ... 271
FlashCopy mapping guidelines for storage systems ... 272
FlashCopy mapping events ... 272
Grains and the FlashCopy bitmap ... 274
Thin-provisioned FlashCopy ... 274
Legacy Safeguarded Copy based on FlashCopy ... 275
IBM Copy Services Manager ... 276
IBM Storage Copy Data Management ... 276
Configuring Safeguarded Copy function ... 280
Managing the Safeguarded Copy function ... 287
FlashCopy statistics collection ... 291

Chapter 15. Notifications ... 293
Syslog notifications ... 293
SNMP notifications ... 295
SNMP over TLS ... 296
Examples of configuration using net-snmp ... 297

Chapter 16. Network ... 299
Ethernet connectivity ... 299
Routing considerations ... 299
Configuring VLAN ... 299
System IP addresses ... 301
Priority flow control ... 302
Modifying VLAN for priority flow control ... 304
Planning for a direct-attached iSCSI configuration ... 306
Fibre channel connectivity ... 307
Auto zoning ... 307
Fabric Device Management Interface ... 310
Fibre Channel port numbers and worldwide port names ... 310
N-Port ID Virtualization ... 310
Planning for connectivity ... 314
Zoning examples ... 320
Fibre Channel port masking ... 324
Fibre Channel SAN configuration details ... 326
System configuration by using SAN fabrics with long-distance fiber connections ... 331
Determining and changing the Fibre Channel network speed ... 332

HTTP proxy server ... 332
Portsets ... 333
Configuring Ethernet portsets ... 334
Configuring Fibre Channel portsets ... 335
Partnerships ... 336
Partnerships using Fibre Channel connectivity ... 337
Partnerships using IP Connectivity ... 339
Intersystem link for host traffic ... 354
Long-distance requirements for partnerships ... 354

Chapter 17. Cyber resilience ... 355

Chapter 18. Security ... 357
Security overview ... 357
Managing security ... 360
Security protocol levels ... 361
SSH rules ... 391
System Certificates ... 391
Configuring system certificates ... 393
Updating or creating an internally signed certificate ... 394
Requesting and installing an externally signed certificate ... 397
Export system certificate ... 399
Removing system certificate ... 400
Partition Certificates ... 400
Configuring partition certificates ... 401
Updating or creating an internally signed partition certificate ... 401
Requesting and installing an externally signed partition certificate ... 402
Export partition certificate ... 404
Removing partition certificate ... 404
Multifactor authentication ... 405
Configuring multifactor authentication ... 406
Changing MFA credentials of system ... 425
Changing the failmode parameter for multifactor authentication ... 426
User authentication ... 426
Remote Authentication ... 428
Single Sign-on ... 431
Configuring single sign-on ... 432
Managing proxy server for SSO login ... 454
Changing SSO credentials of system ... 454
Encryption ... 455
Encryption overview ... 455
Licensing encryption ... 457
Encryption with internal key management ... 459
Encryption with key servers ... 461
Encryption with USB flash drives ... 475
Encryption recovery key ... 479
Adding nodes or enclosures to an encryption-enabled system ... 483
Decommissioning encryption ... 484
Secure data deletion ... 486
Password policy ... 487
Two person integrity ... 488
Host authentication ... 490
Ransomware threat detection ... 491

Chapter 19. System ... 493
Active - active capability ... 493
I/O group ... 497

Nodes....499
Quorum....502
IP quorum application....504
System topology....506
System layers....507
Licensed functions....508
Bitmap space configuration....509
FlashCopy and volume group snapshot capacity configuration....512
Throttles....516

Chapter 20. Support....519
Call Home....519
Support assistance....520
Support packages....522

Chapter 21. System update....525
Patches....525
Auto update....525
Manual update....526
Service assistant manual update....527
Drive update....528
Obtaining packages....529

Chapter 22. Management GUI....531

Chapter 23. FlashSystem.ai....533
Upgrading FlashSystem.ai....536
Planning for the FlashSystem.ai container....536
Certificate configuration for FlashSystem.ai....538
Set up FlashSystem.ai container....539
Set up Kubernetes with wildcard certificate....540
Set up Kubernetes with custom certificate....541
Set up Kubernetes with cert-manager....544
Set up OpenShift with wildcard certificate....545
Set up Open Shift with custom certificate....547
Set up OpenShift with cert-manager certificate....549
Configure FlashSystem systems for FlashSystem.ai....551

Chapter 24. Ansible collection....553

Chapter 25. Storage Virtualize RESTful API....555
Storage Virtualize Performance Statistics....559
Performance data files....559
Getting performance statistics....562

Chapter 26. Conventions used in this IBM Documentation....577
Notifications....577
Emphasis....577
Numbering conventions....577

Figures

1. Storage pool..... 31

2. Two member drive storage array with no rebuild area..... 38

3. Two member drive storage array with a failed drive..... 39

4. Rebuild-in-place process on the two member drive.....40

5. Distributed RAID array (RAID 6 level) 41

6. Distributed RAID array (RAID 6 level) with a failed drive..... 42

7. Storage systems and MDisks..... 48

8. Storage pools and volumes.....76

9. Hosts, WWPNs, IQNs or EUIs, and volumes..... 108

10. Hosts, WWPNs, IQNs or EUIs, volumes, and SCSI mappings..... 109

11. Host cluster that contains private and shared volumes 113

12. Transmitting SCSI over TCP/IP.....144

13. iSCSI configuration with a single subnet.....145

14. iSCSI configuration that uses multiple subnets and provides alternative configuration interfaces.... 145

15. iSCSI configuration with host multipathing..... 146

16. Four two-session iSCSI multisession configurations for a single subnet..... 171

17. Two four-session iSCSI multisession configurations for a single subnet.....171

18. Four two-session iSCSI multisession configurations for multiple subnets and dual configuration.....171

19. Two four-session iSCSI multisession configurations for a multiple subnets and dual configuration.. 172

20. Ownership group inheritance..... 212

21. Example: Child pools are configured on the system.....217

22. Example: Child pools are not configured on the system..... 219

23. Storage partitions page of the FlashSystem grid..... 249

24. A stretched system with a quorum disk at a third site..... 262

25. Stretched system nodes with ISLs, DS3000, DS4000, and DS5000 connected to both fabrics..... 264

26. Recovering Safeguarded backups..... 279

27. A configuration that is not valid for use with NPIV..... 311

28. A configuration that is valid for use with NPIV and (WWPN-based) zoning.....312

29. A single switch configuration that is not valid for use with NPIV and hard (port-based) zoning..... 313

30. A single switch configuration that is valid for use with NPIV and hard (port-based) zoning 314

31. Fabric port configuration..... 317

32. An example of a host zone.....321

33. An example of a storage system zone.....322

34. An example of a system zone..... 322

35. Fabric with ISL between nodes in a system.....327

36. Fabric with ISL in a redundant configuration..... 327

37. Fibre Channel forwarder linked to existing Fibre Channel SAN.....328

38. Fibre Channel forwarder linked to hosts and storage systems without an existing Fibre Channel SAN.....328

39. Fibre Channel host connects into the Fibre Channel ports on the Fibre Channel forwarder..... 329

40. Fibre Channel host connects into the Fibre Channel ports on the Fibre Channel forwarder without an existing Fibre Channel SAN.....329

41. Simple SAN configuration.....330

42. SAN configuration with a medium-sized fabric.....330

43. SAN configuration with a large fabric..... 331

44. SAN configuration across two sites.....331

45. Redundant fabrics.....338

46. One intersite link, one I/O group per system..... 346

47. Two intersite links, one I/O group per system..... 347

48. Scaling of host I/O 347

49. Product transfer rate comparison for systems built with IBM Storage Virtualize.....348

50. Active-active with all paths optimized..... 493

51. Active-active with default optimized and non-optimized paths..... 494

52. Active-active Reads through Node 1.....495

53. Active-active read through Node 2.....495

54. Active-active - Write to Node 1..... 496

55. Active-active - Write to Node 2..... 497

56. Example of a write operation in an I/O group..... 498

57. Example of a standard system topology 506

58. Example of a stretched system topology 506

59. Example configuration with systems in different layers This figure shows an overview of system layers 507

Tables

1. Low space warning percentages for compressed arrays..... 9

2. Statistic collection per node per port of iSCSI, NVMeTCP and NVMeRDMA..... 13

3. Statistic collection for Host for Individual Nodes.....14

4. Statistics collection for MDisks for individual nodes.....15

5. Statistic collection for volumes for individual nodes..... 17

6. Statistic collection for node ports.....18

7. Statistic collection for nodes..... 20

8. Performance statistics for volume.....21

9. Statistic collection for volume cache per individual nodes.....21

10. XML statistics for an IP partnership port.....23

11. Statistics collection for an individual IP partnership..... 23

12. ODX VDisk and node level statistics..... 24

13. Statistics collection for cloud per cloud account ID..... 25

14. Statistics collection for cloud per VDisk.....26

15. Pool states.....33

16. MDisk status.....49

17. Bad block errors..... 50

18. Storage-system port-selection algorithm..... 52

19. Minimum overhead capacity requirements for data reduction pools..... 61

20. Data move statistic descriptions..... 68

21. Workload categorization descriptions.....70

22. Easy Tier settings for storage pools and volumes.....73

23. Volume states.....77

24. Volume cache modes....78

25. Parameter and recommended settings for linux hosts.... 125

26. Registry key parameters for QLogic models.... 130

27. Configuration file parameters for the Emulex HBA....130

28. Registry key parameters for QLogic models.... 135

29. Comparison of iSCSI and Fibre Channel components....143

30. System registry.... 160

31. Path status information.... 168

32. Common queue depth configurations....170

33. NVMe over RDMA minimum compatible Ethernet cards and firmware versions for VMware ESXi 7.0 UP2....182

34. Example error scenarios....204

35. Statistics collection for volume group....210

36. Statistics collection for node statistics.... 210

37. Statistics collection for volume statistics....210

38. Supported access control for management interfaces and functions with user groups....224

39. FlashCopy mapping events....273

40. Statistics collection for volume aggregate....291

41. Statistics collection for volume count....292

42. Security levels and required credentials for SNMP version 3 servers using the user-based security model....295

43. Examples of WWPN values for the system 310

44. Port allocation scheme.... 315

45. Two hosts and their ports.... 320

46. Two storage systems and their ports.... 320

47. Six hosts and their ports.... 323

48. Three storage systems and their ports....323

49. Port masking configuration....325
50. Maximum supported round-trip latency between sites....354
51. Intersystem heartbeat traffic in Mbps....354
52. Supported SSL/TLS security levels....362
53. SSH algorithms supported at each security level....373
54. Protocols supported at level 7....377
55. Java SSL ciphers supported at security level 7....377
56. OpenSSL ciphers supported at level 7 (chsecurity -sslprotocol 7)....378
57. Protocols supported at level 6....378
58. Java SSL ciphers supported at security level 6....378
59. OpenSSL ciphers supported at level 6 (chsecurity -sslprotocol 6)....378
60. Protocols supported at level 5....378
61. Java SSL ciphers supported at security level 5....379
62. OpenSSL ciphers supported at level 5 for TLS 1.3 (chsecurity -sslprotocol 5)....379
63. OpenSSL ciphers supported at level 5 for TLS 1.2 (chsecurity -sslprotocol 5)....380
64. Protocols supported at level 4....381
65. Java SSL ciphers supported at security level 4....381
66. OpenSSL ciphers supported at level 4 (chsecurity -sslprotocol 4)....381
67. Protocols supported at level 3....381
68. Java SSL ciphers supported at security level 3....381
69. OpenSSL ciphers supported at level 3 (chsecurity -sslprotocol 3)....383
70. Protocols supported at level 2....384
71. Java SSL ciphers supported at level 2....384
72. OpenSSL ciphers supported at level 2 (chsecurity -sslprotocol 2)....386
73. TCP and UDP ports that are supported....387

74. SSH algorithms supported at each SSH security level (chsecurity -sshprotocol)..... 389

75. System certificate scope values..... 393

76. Recommended delimiters..... 396

77. General tab..... 408

78. Sign-on tab..... 408

79. Add API Client action..... 409

80. Authentication methods comparison 427

81. Authentication methods summary 432

82. Register an application page..... 434

83. General tab..... 438

84. Sign-on tab..... 438

85. Welcome..... 441

86. Server application..... 441

87. Configure Application Credentials..... 443

88. Configure Web API..... 443

89. Choose Access Control Policy..... 444

90. Configure Application Permissions..... 444

91. Register an application page..... 446

92. General settings on the app integration page..... 448

93. Assignments settings on the app integration page..... 448

94. Metadata section..... 450

95. Relying Party section..... 450

96. OIDC Response section..... 451

97. Settings section..... 451

98. Comparison of methods to securely delete data from drives..... 486

99. Node status.....500
100. System topology and volume summary.....506
101. Comparison of storage layer and replication layer systems.....507
102. Supported SCU licenses, drive classes, and ratios.....508
103. Example of SCU usage.....509
104. Product comparison of required I/O group memory for policy-based replication.....510
105. Bitmap space default configuration.....511
106. Examples of memory required 511
107. Distributed RAID requirements.....512
108. Volume group snapshot capacity limits.....512
109. FlashCopy capacity limits.....514
110. FlashCopy capacity limits.....515
111. Supported network configurations for Call Home with cloud services.....519
112. Firewall rule configuration.....522
113. Firewall rule configuration for remote access.....522
114. Support packages.....523
115. Upgrade time for a four-node cluster.....527
116. Packages published to FixCentral.....529
117. Connectivity requirements.....537
118. 539
119. 556
120. Calculations.....563
121. Calculations.....564
122. Calculations.....567
123. Calculations.....570

124. Calculations..... 572

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

© Copyright IBM Corp. 2025

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.

Chapter 2. System monitoring

The following topics describe how to monitor the system.

Dashboard

The **Dashboard** on the management GUI provides a way to quickly assess the overall condition of the system and view notifications of any critical issues that require immediate action.

The **Dashboard** is the home page that displays after you log in to the system. It contains high-level information about the system and is not a replacement for the **Performance** page or the system view. The **Dashboard** contains sections for performance, capacity, and system health that provide an overall understanding of what is happening on the system.

Performance

The **Performance** section on the **Dashboard** displays a graph with up to 5 minutes of data from the performance of a specific system or a comparison of the performance of the nodes. Select either **System** or **Node Comparison** to switch the view between system-level performance and node-level performance.

The performance graph is updated with new data every 5 seconds. Select any of the available metrics to change what is displayed by the graph. To view a metric not available on the **Dashboard**, go to the **Performance** page in the navigation.

Capacity

The **Capacity** section on the **Dashboard** provides an overall view of system capacity. There are two types of view for the Capacity section, one view for FCM based systems and another for non FCM based systems. The following section displays for non FCM based systems:

Usable Capacity
Usable capacity indicates the total capacity in all storage on the system. Usable capacity includes all the storage the system can be virtualized and assigned to pools. Usable capacity is displayed in a donut graph and is divided into three categories: **Used Capacity**, **Available Capacity**, and **Total Capacity**. If the system supports self-compressing drives, certain system configurations make determining accurate usable capacity on the system difficult. For example, if the system contains self-compressed drives and data reduction pools without compression enabled, the system cannot determine the accurate amount of usable capacity that is used on the system. In this case, overprovisioning and losing access to write operations is possible. If this condition is detected by the system, the **Usable Capacity** section of the **Dashboard** page displays a message instead of capacity information. To recover from this condition, you need to ensure that all thin-provisioned volumes and thin-provisioned volumes that are deduplicated are migrated to volumes with compression enabled in the data reduction pools. Alternatively, you can migrate the volumes to standard-provisioned volumes and use the drive compression to save capacity.

Used Capacity
Used Capacity indicates the amount of capacity that is used on the system after capacity savings. The system calculates the stored capacity by subtracting the available capacity and any reclaimable capacity from the total capacity that is allocated to MDisks. To calculate the percentage, the stored capacity is divided by the total capacity that is allocated to MDisks.

Available Capacity
Available capacity is calculated by adding the available capacity and the total reclaimable capacity. To calculate the percentage of available capacity on the system, the available capacity is divided by the total amount of capacity that is allocated to MDisks.

Total Capacity
The **Total Capacity** shows all the capacity available on the system. Select **MDisks** to view more information about the usable capacity of the system on the **MDisks by Pools** page.

© Copyright IBM Corp. 2025

If the system supports self-compressing drives, certain system configurations make determining accurate usable capacity on the system difficult.

Capacity Savings

Capacity Savings indicates the amount of capacity that is saved on the system by using data reduction, thin-provisioning, and total savings. The value for each of these capacity savings methods compares the stored capacity before capacity savings is applied to the stored capacity after capacity savings is applied. Data reduction shows the physical capacity not consumed by written capacity due to savings from compression and deduplication. Thin-provisioning displays the total capacity savings for all thin-provisioned volumes on the system. Total savings indicates the total capacity savings that the system has saved from thin-provisioning and data reduction. You can view all the volumes that use each of these technologies. These requirements can include additional hardware, licenses, or both to use these functions. Verify all system requirements before using these functions. The system supports several compression technologies, such as data reduction pool and self-compressing drives, where data is compressed in the hardware as it is written to a drive. A single capacity savings value for systems that use multiple compression technologies cannot be determined.

The following section displays for FCM based systems:

Physical Capacity

Physical capacity indicates the total capacity in all storage on the system. Physical capacity includes all the storage the system manages and can be assigned to pools. Physical capacity is displayed in a meter chart and available in percentage and number (for example GiB, TiB). The capacity number for physical used and physical available displays in this section.

Logical capacity

Logical capacity indicates the total amount of storage that the system perceives as free or used for storing data, including all optimizations such as deduplication, compression, and thin provisioning. Logical capacity is displayed in a meter chart and available in percentage and number (for example GiB, TiB). The capacity number for total logical capacity and logical free capacity are displayed in this section. The Logical capacity section also displays:

- Capacity savings, which indicates the amount of capacity that is saved on the system by using data reduction, thin-provisioning, and total savings. Data reduction shows the physical capacity not consumed by written capacity due to savings from compression and deduplication. The percentage value for data reduction method compares the stored capacity before capacity savings is applied to the stored capacity after capacity savings is applied. Thin-provisioning displays the total capacity savings for all thin-provisioned volumes on the system. Total savings indicates the total capacity savings that the system has saved from thin-provisioning and data reduction.
- Breakdown of how storage is used in a system displayed in the form of a meter chart and available in number (for example GiB, TiB). The capacity number for written by hosts, written by snapshots, and logical reserved displayed in this section.
- Total capacity of all virtualized storage on the system that has been dedicated for use by volumes.

System Health

The **System Health** section on the **Dashboard** provides a holistic view of the system through tiles of data that represent different components that make up the system. A tile contains one type of component, but it can contain multiple items of the same type.

The current software version number and the system ID are displayed above the **Connectivity components** category.

System Health tiles

The tiles of data in the **System Health** section of the Dashboard represent different components that make up the system, providing a holistic view of the system.

Each tile contains one type of component, but it can contain multiple items of the same type. For example, a Fibre Channel port is a component, so it is contained in a tile, but the Fibre Channel tile can contain multiple Fibre Channel ports. Tiles with errors and warnings are displayed first so that

components that require attention have higher visibility. Healthy tiles are sorted in order of importance in day-to-day use.

The following list defines the categories of tiles:

- **Hardware components** display the health of all components that are specific to the physical hardware.
- **Logical components** display the health of all logical and virtual components in the management GUI.
- **Connectivity components** display the health of all components that are related to the system's connectivity and the relationship between other components or systems.

  Note: The *Configured* status of **IBM Storage Insights** in the **Connectivity components** section indicates that the storage system is registered with IBM Storage Insights and the **Call Home** field shows the last time when your storage system was synced with IBM Storage Insights.

Monitoring capacity

The system supports multiple ways of monitoring and adjusting capacity for your data. Use these methods to determine whether new storage resources are needed or if existing resources need to be reallocated.

Monitoring usable capacity

The system supports several ways to monitor usable capacity to ensure that storage is sufficient for host workloads.

Monitoring system-level capacity

The **Capacity** section on the **Dashboard** provides an overall view of system capacity. There are two types of view for the Capacity section, one view for FCM based systems and another for non FCM based systems. The following section displays for non FCM based systems:

If you are using the command-line interface to determine usable capacity on your system, several parameter values are used from the **`lssystem`** command to calculate stored, available, and total capacities. Stored capacity is calculated with the values in the **`total_mdisk_capacity`**, **`total_free_space`**, **`total_reclaimable_capacity`** by using the following formula:

- Total stored capacity = **`total_mdisk_capacity`** - **`total_free_space`** - **`total_reclaimable_capacity`**

To calculate the available capacity, use the values in **`total_free_space`** and **`total_reclaimable_capacity`** in the following formula:

- Total available capacity = **`total_free_space`** + **`total_reclaimable_capacity`**

The value in the **`total_mdisk_capacity`** in the **`lssystem`** command indicates the total usable capacity is on the system.

Monitoring external storage systems, pool, and MDisk level capacity

The system supports virtualizing capacity that is on external storage systems that are attached to the system. For external storage systems, administrators must configure out-of-space alerts that notify the system when usable capacity on the external storage system reaches a defined threshold. Without these thresholds defined on the external storage systems, external storage on these systems can become overprovisioned and risk running out of usable capacity that is used for host operations. Overprovisioned external storage occurs when the sum of the provisioned capacity of all the volumes in a system or pool is greater than the total usable capacity of the system or pool that is allocated from the external storage system. The system detects potential overprovisioned external storage and presents a warning on the **Dashboard** page. If this warning is displayed, select **Overprovisioned External Storage** to display pools that contain storage that is nearing these overprovisioned scenarios. Depending on how external storage is configured in pools, different strategies and best practices can be used to ensure that usable capacity is sufficient and reported accurately to the system.

Configuration Scenario 1: One pool contains storage from one external storage system

In this configuration, storage administrators have the same concern of running out of capacity as when data is written directly to external storage systems. Storage administrators must set up out-of-space alerts on external storage and monitor usage to ensure that the I/O operations do not exceed the usable capacity. Consult documentation for your external storage system for specific guidelines on these thresholds. If an out-of-space condition occurs, then usable capacity must be freed by deleting data or volumes.

Configuration Scenario 2: Multiple pools contain a single tier of storage from multiple external storage systems

In this configuration, multiple pools use the same type of storage, or tier, across several external storage systems. If capacity from different external storage systems is shared across multiple pools, then provisioning groups are created. Provisioning groups are objects that identify whether storage is shared across multiple pools. The **MDisks by Pools** page in the management GUI, displays all pools and their assigned MDisks. If you are not sure whether external storage systems are shared across several pools, right-click the pool and select **View Resources** to display the provisioning group that is associated with the pool. In this configuration, the system spreads extent allocations across all the external storage within the provisioning group to ensure that space is used evenly across all the external storage systems. However, usable capacity on external storage systems within the provisioning groups can still become overprovisioned, so storage administrators need to configure out-of-space alerts and monitor usable capacity to determine how capacity is being used on the system. If low-space warning occurs on virtualized MDisks on external storage and usable capacity is available on other external storage system in the pool, you can remove some of these MDisks from the pool until the usable capacity is within reasonable limits. This process migrates data to other MDisks in the pool. Ensure that enough usable capacity is available on the other external storage systems in the pool so that they do not run out of space during this operation.

Configuration Scenario 3: Multiple pools contain different tiers of storage from multiple external storage systems

In this configuration, different tiers of storage are present on the external storage systems within the pools. As with the previous configuration, provisioning groups are created to identify the external storage that is shared within the pools. However, in this case the system attempts to allocate the entire usable capacity in the top tier rather than spreading extents across provisioning groups. Storage administrators need to monitor the usable capacity in these top tiers of storage to ensure that space is sufficient for the workloads. If low-space warnings occur on virtualized MDisks on external storage and usable capacity is available on other external storage system in the pool, you can remove some of these MDisks from the pool until the usable capacity is within reasonable limits. This process moves data to other MDisks in the pool. Ensure that enough usable capacity is available on the other external storage systems in the pool so that they do not run out of space during this operation.

Configuration Scenario 4: Data reduction pools

Systems that support data reduction pools, which use data reduction technologies, like compression, can configure these types of pools to minimize overallocation on the external storage system. Data reduction pools also support the ability to reclaim unused capacity from host unmap operations and volume deletions. In the management GUI, select **Pools** > **MDisks by Pools** to view all the pools configured on the system. Data reduction pools are shown with **Data Reduction** set to **Yes**. When you use data reduction pools that virtualize external storage that can run out of space, always enable compression when you create volumes in these pools. The external storage system must be configured to present only usable capacity to the system, as the size of the written data cannot be reduced any further.

To assist with storage planning, Table 1 on page 9 lists the low space warnings that can be generated by the system for compressed arrays. Error 1246 is not raised on Easy Tier® pools or data reduction pools.

Table 1. Low space warning percentages for compressed arrays			
Condition	**Event ID**	**Error code**	**Percentage of usable capacity used**
Compressed array is running low on usable capacity	020009	1246	90% used
Critical level of compressed array is running low on usable capacity.	020010	1246	96%
1% usable capacity that is left for compressed array.	020011	1242	99%
Compressed array out of usable capacity.	020012	1241	100%

If your system configuration dedicates specific storage systems to individual pools, analyzing the usable capacity for the pool can determine whether more usable capacity is needed. On the **Pools** page, check each MDisk in the pool for capacity and use the value for **Storage System - LUN** to determine the external storage system that provides the MDisk with space.

Monitoring capacity savings

You can use the management GUI to determine the capacity savings that are created by thin-provisioning, compression, and deduplication.

The system supports several methods to save capacity. These methods can be used separately, but when used together they can provide greater capacity savings and efficiencies for the system. The system supports thin-provisioning, compression, and deduplication. With thin-provisioned volumes, the amount of capacity that can be provisioned to hosts is not limited by the usable capacity because the system eliminates the unused allocated capacity. A *thin-provisioned* volume presents a different capacity to mapped hosts than the capacity that the volume consumes in the storage pool. With thin provisioning, storage administrators can configure thin-provisioned volumes that can grow into the usable capacity they need based on the storage needs of the host. This flexibility saves on the costs to manage the storage capacity for host applications. Apart from standard-provisioned volumes, all volume types provide capacity savings and benefits from thin-provisioning. It can be used with compression and deduplication, or it can be used by itself. Thin-provisioning is supported in both standard pools and data reduction pools. However, data reduction pools support reclaiming capacity when it is no longer needed by hosts and then can redistribute it automatically for other uses.

Unlike thin-provisioning, which removes unused capacity, compression and deduplication increase the usable capacity by reducing the size of the data itself. The system supports several types of compression technologies that can be used on the system at the same time. Software-based compression is implemented when you create volumes, where compression can be enabled on the volume. With compressed volumes, data is compressed in the software then it is written to disk, saving more capacity. When data is read to hosts, the software decompresses the data. 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. Compressed volumes can also be created in standard pools, but data reduction pools support more capacity savings functions, such as automatic redistribution of reclaimed storage. The system also supports hardware-based compression technologies, such as self-compressing drives. This type of compression happens automatically within the hardware itself as data is written to the drive, and decompressed on the hardware as it is read to hosts. Depending on what technology handles the compression operations, a single capacity savings percentage for systems that use multiple compression technologies cannot be determined. Select **Compression Details** to view savings details for each of the compression technologies that are currently

used on the system. If compressed volumes are configured in standard pools, you cannot create compressed volumes in data reduction pools within the same I/O group. Potential compression savings can be calculated for existing volumes in the management GUI through a built-in compression estimation tool. This tool can be used to determine the capacity savings that are possible for existing data on the system by using compression.

The system also supports deduplication. Deduplication can be configured with volumes that use different capacity saving methods, such as thin-provisioning. 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.

In addition to capacity savings from these methods, data reduction pools add another benefit by reclaiming capacity and redistributing it to other uses. Reclaimable capacity is calculated as part of the usable capacity on the system. The amount of reclaimable capacity varies based on the number of write operations and how much reclaimable capacity currently exists on the system. The reclaimable capacity represents usable capacity that will eventually become available for overwriting with new data when it is "reclaimed" and made available to the system. This reclaimable capacity is calculated and displayed with the usable capacity of the system. Since reclaimable capacity is not used by production data and is eventually reclaimed for use, the amount of reclaimable capacity is subtracted from stored capacity and added to available capacity in the **Capacity** tile on the **Dashboard** in the management GUI. This calculation gives a more accurate view of capacity usage on the system. In systems that contain self-compressing drives, the actual usable capacity that reclaimable capacity consumes cannot be calculated unless it is already compressed by the data reduction pool. When the data is compressed on the drive, it is not possible to determine how much usable capacity that becomes available when the reclamation process completes. Each of these capacity savings methods requires capacity monitoring to ensure efficient capacity provisioning on the system. Capacity savings result from using these methods and establishing best practices for capacity management in your organization and actual savings is contingent on these practices. However, compression and deduplication might provide data reduction two times greater than without these methods and thin-provisioning might provide two times the data reduction. When these data reduction methods are used together, overall capacity savings are combined.

Note: On systems that support only standard pools built on top of FlashCore module based arrays, it is best practice to set the volumes to be fully allocated.

To determine capacity savings in the management GUI, complete these steps:

1. In the management GUI, select **Dashboard**.
2. Under **Capacity Savings**, you can view the total capacity savings that the system currently gains from using capacity savings methods. The total capacity savings is divided into thin-provisioning, compression, and deduplication savings. Click the link to view volumes that use each of these methods to save capacity. The management GUI calculates capacity savings for each of these methods. For thin-provisioning savings, all written capacity is subtracted from all the provisioned capacity on the system. Written capacity is the capacity that hosts write to storage before any data reduction methods are applied. Thin-provisioning savings indicate how much capacity is provisioned to hosts but is not currently used to store data from hosts. For compression savings, the system calculates savings from compressed volumes in data reduction pools. If your system supports compression in standard pools, the total compressed savings are calculated by combining the savings from both data reduction pool and from standard pools. Deduplication savings are calculated for deduplicated volumes in data reduction pools. Deduplication is not supported in standard pools.

Recovering from configuration that prevents accurate usable capacity reporting

If the system supports self-compressing drives, certain system configurations make determining accurate usable capacity on the system difficult.

If the system contains self-compressed drives and data reduction pools without compressed volumes, the system cannot determine the accurate amount of usable capacity that is used on the system. In this case, overprovisioning and losing access to write operations is possible. If this condition is detected by

the system, the **Usable Capacity** section of the **Dashboard** page displays a message instead of capacity information. To recover from this condition, you need to ensure that all thin-provisioned volumes and thin-provisioned volumes that are deduplicated are migrated to volumes with compression enabled in the data reduction pools. Alternatively, you can migrate the volumes to standard-provisioned volumes and use the drive compression to save capacity.

You can use volume mirroring to enable compression on existing volumes in data reduction pools. Because data reduction pools also support deduplication, you can add that function to compressed volumes to gain extra capacity savings. After the original volumes are deleted from the data reduction pool, the capacity from those volumes is added to the reclaimable capacity for the system, and the error is removed from the **Dashboard**. Depending on the amount of reclaimable capacity on the system, the system can take some time to remove the reclaimable capacity from the system. During that process, the usable capacity tends to shrink in size. If you want to use self-compressing drives, convert these volumes to standard-provisioned volumes that can be used in a standard pool that uses self-compressing drives for storage. In the management GUI, you can use the **Modify Capacity Savings** action to quickly convert the capacity savings method for your volumes. This action creates a mirrored volume of the original, with the new capacity savings method enabled, then automatically deletes the original version from the pool.

If you have data reduction pools that contain volumes without compression on systems that contain self-compressing drives, you need to convert all these volumes to enable compression. Accurate reporting for usable capacity is unavailable for these configurations. To change the capacity savings on volumes in data reduction pools, complete these steps:

1. In the management GUI, select **Volumes** > **Volumes by Pools**. The left navigation displays all the configured pools on the system and expands the data reduction pool that contains the volumes that you want to change the capacity savings method for.
2. Right-click the volume that you want to convert and select **Modify Capacity Savings**. Since this process creates volume copies to manage the conversion, ensure that the pool has enough storage to accommodate these volume copies. For example, if you are converting a volume that has 2 GB allocated capacity, ensure that the pool has another 2 GB to accommodate both the original version and the copy during the conversion. After the conversion completes, the original copy is automatically deleted from the pool. You can also convert volumes in standard pools to data reduction pools to use the data reduction features that are available in these pools. For that method, you need to create a new data reduction pool and volume mirroring.
3. On the **Modify Capacity Savings** page, select **Compression** to change the currently configured capacity setting to enable compression. For additional capacity savings, you can also select **Deduplicated** to remove duplicate data from the pool.
4. Click **Modify**. After the original volumes are deleted from the data reduction pool, the capacity from those volumes is added to the reclaimable capacity for the system, and the error is removed from the **Dashboard**. Depending on the amount of reclaimable capacity on the system, the system can take some time to remove the reclaimable capacity from the system. During that process, the usable capacity tends to shrink in size.

Performance statistics

Real-time performance statistics provide short-term status information for the system.

To access these performance statistics, click **Monitoring** > **Performance** in the management GUI. In addition, the management GUI displays an overview of system performance, in the **Performance** section on the **Dashboard**.

You can use system statistics to monitor the bandwidth of all the volumes, protocols, and MDisks that are being used on your system. You can also monitor the overall System CPU (total) for the system. These statistics summarize the overall performance health of the system and can be used to monitor trends in bandwidth and System CPU (total). You can monitor changes to stable values or differences between related statistics, such as the latency between volumes and MDisks. These differences can then be further evaluated by performance diagnostic tools.

Additionally, with system-level statistics, you can quickly view bandwidth of volumes, protocols, and MDisks. Each of these charts displays the current bandwidth in megabytes per second (Mbps) and a view of bandwidth over time. Each data point can be accessed to determine its individual bandwidth use and to evaluate whether a specific data point might represent performance impacts. For example, you can monitor the protocols to determine whether the host data-transfer rate is different from the expected rate.

You can also select node-level statistics, which can help you determine the performance impact of a specific node. As with system statistics, node statistics help you to evaluate whether the node is operating within normal performance expectations.

CPU Utilization

The System CPU (total) chart shows the current percentage of CPU usage and peaks in utilization.

A single spike often does not indicate a performance impact on the system; however, if a data point is consistently more than 95% of utilization and I/O input is high, the system might be overloaded, which might indicate a need for more back-end storage. However, for compression utilization alone, it can be normal to see rates at 100%, especially if compression is used frequently on the system. If both compression and I/O input are high, the system might need more storage to accommodate the utilization.

Interface

MDisks and volumes

The MDisk charts display six metrics while the Volumes charts display another six metrics on the **Performance** section. These metrics are: Read IOPS, Write IOPS, Read bandwidth, Write bandwidth, Read latency, and Write latency. You can use these metrics to help determine the overall performance health of the volumes and MDisks on your system. Consistent unexpected results can indicate errors in configuration, system faults, or connectivity issues. Both the volumes and MDisk charts contain the same metrics to compare and use to evaluate performance; however, the data points for these metrics can be different due to the impact of system cache, RAID overhead, and Copy Services functions. You can view data points in megabytes per second (MBps) and I/O per second (IOPS).

If you select the read metric, data points on the chart indicate the average amount in MBps or IOPS for read operations that have been processed over the sample period. Read metrics represent how much data the system is processing (the bandwidth). The read latency metric for volumes measures the average amount of time in milliseconds that the system takes to respond to read requests over the sample period. The MDisks latency metrics measure the response time for back-end storage. Increased workload or error recovery can cause spikes in read latency. Values that are consistently higher than expected can indicate a fault situation or an overloaded system. Volume write latency tends to be lesser than the MDisk write latency because of volume caching. If the volume write latency is equal to or greater than MDisk write latency, it suggests that the drives might be overloaded and more drives might be necessary to accommodate the increased workload. Volume latency may be lesser than MDisk latency because of the volume cache. Alternatively volume latency may be higher than MDisk latency due to hosts that use large I/O sizes. Neither of these conditions indicates a problem with the system.

Note: The management GUI supports latency metrics that are displayed in both microseconds and milliseconds. Measurements for metrics are rendered based on whether current values are less than 1.000 milliseconds. Values over 1.000 milliseconds are displayed as milliseconds and metrics under that value display as microseconds. Whereas in command-line interface (CLI) the values are always displayed in microseconds.

The difference between read and write IOPS shows the mixture of the workload that the system is executing. You can determine the average transfer sizes of data that the system is experiencing by dividing the reads and write operations in MBps by the read and write operations in IOPS. This information can be used for validating and predicting disk configuration for the system or input to a disk provisioning application. Write latency is the average time (in milliseconds) that the system writes data to volumes or MDisks but does not include the time for write operations that are used to keep volumes in Replication

relationships synchronized. As with read latency, MDisk write latency tends to be higher than volumes because of write caching and RAID overheads. For example, a write operation to a volume can result in more read and write operations on the MDisk depending on the RAID type for the array.

Related reference

"IP partnership requirements" on page 339
When you create and manage IP partnerships, consider the following requirements and constraints:

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 six 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**; one for drives, which is named **Nd_stats**; and one for hosts, which is named **Nh_stats**.

The files are written to the `/dumps/iostats/` directory on the node. 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.

Tables

The **lsportstats** command displays statistics for each node and physical port. When the **type** field is **Eth_iSCSI, Eth_NVMe_TCP** and **Eth_NVMe_RDMA**, the statistics collection is displayed per physical port for iSCSI , NVMeTCP and NVMeRDMA.

Table 2 on page 13 describe the information that is reported for Ethernet iSCSI individual nodes.

Table 2. Statistic collection per node per port of iSCSI, NVMeTCP and NVMeRDMA

Statistic name	Description
hbt	Indicates the bytes transmitted to hosts.
hbr	Indicates the bytes received from hosts.
het	Indicates the commands that are initiated to hosts.
her	Indicates the commands that are received from hosts.
hsr	Indicates the count of data transfers that were considered slow for host read operations.
hsw	Indicates the count of data transfers that were considered slow for host write operations.
har	Indicates the cumulative number of blocks of data written (since the node started).
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.

Table 3 on page 14 describes the statistics collection for hosts, for individual nodes.

Example XML showing the statistic collection for Host for Individual Nodes

```
<?xml version="1.0" encoding="utf-8" ?>
<diskStatsColl
xmlns="http://ibm.com/storage/management/performance/api/2024/10/hostIoStats"
xmlns:xsi="http://www.w3.org/2001/XMLSchema-instance"
xsi:schemaLocation="http://ibm.com/storage/management/performance/api/2024/10/hostIoStats
schema/SVCPerfStatsH.xsd" scope="node" id="node1" cluster="KK" node_id="0x0000000000000001"
cluster_id="0x000002042cc076a2" sizeUnits="512B" timeUnits="msec" contains="hostIoStats"
timestamp="2026-02-03 05:31:04" timezone="GMT+0:00" timestamp_utc="2026-02-03 05:31:04">
<host idx="0"
```

```
id="KK"
ro="34754" wo="14395" wou="0" rb="318928" wb="115160"
rl="2643" wl="1987" rlw="0" wlw="0" xl="894"
wxl="436" rxl="458"
uo="0" ub="0" uou="0" ul="0" ulw="0"
tiocnt="48760" tbw="52168" tqueuediocnt="0"
tfailediocnt="0" tmaxdelay="0" tcumdelay="0">

</host>
</diskStatsColl>
```

Table 3. Statistic collection for Host for Individual Nodes

Statistic name	Description
idx	Indicates the identifier of the Host for which the statistics apply.
id	Indicates the name of the Host for which the statistics apply.
ro	Indicates the cumulative number of Host read operations that were processed (since the node started).
wo	Indicates the cumulative number of Host write operations that were processed (since the node started).
wou	Indicates the cumulative number of Host write operations that are not aligned on a 4 K boundary.
rb	Indicates the cumulative number of blocks of data that was read (since the node started).
wb	Indicates the cumulative number of blocks of data written (since the node started).
rl	Indicates the cumulative read response time in milliseconds for each Host. The cumulative response time for Host 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.
wl	Indicates the cumulative write response time in milliseconds for each Host. The cumulative response time for Host 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.
rlw	Indicates the worst read response time in microseconds for each Host since the last time statistics were collected. This value is reset to zero after each statistics collection sample.
wlw	Indicates the worst write response time in microseconds for each Host since the last time statistics were collected. This value is reset to zero after each statistics collection sample.
xl	Indicates the cumulative read and write data transfer response time in milliseconds for each Host since the last time the node was reset. When this statistic is viewed for multiple Host and with other statistics, it can indicate whether the latency is caused by the host, or fabric.
wxl	Indicates the cumulative write data transfer response times in milliseconds for each Host 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 Host.
rxl	Indicates the cumulative read data transfer response times in milliseconds for each Host 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.
uo	Indicates the cumulative number of Host unmap operations that were processed (since the node started).

Table 3. Statistic collection for Host for Individual Nodes (continued)	
Statistic name	**Description**
ub	Indicates the cumulative number of blocks of data unmapped (since the node started).
uou	Indicates the cumulative number of Host unmap operations that are not aligned on an 8K boundary (according to the alignment/granularity setting in Block Limits VPD Page (0xb0).
ul	Indicates the cumulative unmap response time in milliseconds for each Host. The cumulative response time for Host 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 Host. 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.
tiocnt	Indicates the cumulative number of host I/O operations that were examined for I/O throttle (since the node started).
tbw	Indicates the cumulative number of blocks of data transfer that was delayed due to I/O throttle (since the node started).
tqueuediocnt	Indicates the cumulative number of host I/O operations that were delayed due I/O throttle (since the node started).
tfailediocnt	Indicates the cumulative number of host I/O operations that were failed due to excessive delay introduced by I/O throttle (since the node started).
tmaxdelay	Indicates the maximum delay introduced for any host I/O by throttle in microseconds (during stats collection interval).
tcumdelay	Indicates the cumulative delay introduced for host I/O operations by throttle in milliseconds (since the node started).

Table 4 on page 15 describes the statistics collection for MDisks, for individual nodes.

Table 4. 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).

Table 4. Statistics collection for MDisks for individual nodes (continued)	
Statistic name	**Description**
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).
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 5 on page 17 describes the volume information that is reported for individual nodes.

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

Example XML showing the statistic collection for vdisk for Individual Nodes

```
<vdsk idx="0"
...
id="CD"
ro="134" wo="2050268" wou="0" rb="862" wb="16407272"
rl="34" wl="850709" rlw="0" wlw="0" xl="62769"
wxl="62768" rxl="1" oro="0" owo="0"
orl="0" owl="0" oiowp="0"
uo="0" ub="0" uou="0" ul="0" ulw="0"
tiocnt="6123" tbw="6123" tqueuediocnt="5986"
tfailediocnt="0" tmaxdelay="0" tcumdelay="470283">
...
</vdsk>
```

Table 5. 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.
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.

Table 5. Statistic collection for volumes for individual nodes (continued)	
Statistic name	**Description**
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 or fabric.
tiocnt	Indicates the cumulative number of volume I/O operations that were examined for I/O throttle (since the node started).
tbw	Indicates the cumulative number of blocks of data transfer that was delayed due to I/O throttle (since the node started).
tqueuediocnt	Indicates the cumulative number of volume I/O operations were delayed due to I/O throttle (since the node started).
tfailediocnt	Indicates the cumulative number of volume I/O operations that were failed due to excessive delay introduced by I/O throttle (since the node started).
tmaxdelay	Indicates the maximum delay introduced for any volume I/O by throttle in microseconds (during stats collection interval).
tcumdelay	Indicates the cumulative delay introduced for volume I/O operations by throttle in microseconds.

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

Table 6 on page 18 describes the port information that is reported for individual nodes.

Table 6. Statistic collection for node ports	
Statistic name	**Description**
bbcz	Indicates the total time in microseconds for which the buffer credit counter was at zero. 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. **Note:** The cer metric is always 0.

Table 6. Statistic collection for node ports (continued)

Statistic name	Description
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.
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. **Note:** The het metric is always 0.
hsr	Indicates the count of data transfers that were considered slow for host read operations.
hsw	Indicates the count of data transfers that 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.

Table 6. Statistic collection for node ports (continued)	
Statistic name	**Description**
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 7 on page 20 describes the node information that is reported for each node.

Table 7. 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 that 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.

Table 7. Statistic collection for nodes (continued)	
Statistic name	**Description**
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.
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 8 on page 21 provides details about the statistic collection for volume performance.

Table 8. 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 9. 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.

Table 9. 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
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.
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.

Table 9. 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**
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 10 on page 23 describes the XML statistics specific to an IP Partnership port.

Table 10. 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.
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 11 on page 23 describes an IP partnership information that is reported for each partnership. It does not include the I/O statistics for high-speed Ethernet partnerships.

Table 11. Statistics collection for an individual IP partnership	
Statistic name	**Description**
index	Indicates the index for which the IP partnership is configured. **Note:** To view **index** parameter, use **lspartnership -gui** command.
cluster	Indicates the name of the remote cluster.
remote_cluster_id	Indicates the ID of the remote cluster.
port_id	Indicates the selected **port_id** for IP partnership connection on local node. **Note:** To view **port_id** parameter, use **lsportethernet** command.
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.

Table 11. Statistics collection for an individual IP partnership (continued)	
Statistic name	**Description**
ipre	Indicates the bytes retransmitted to other nodes in other clusters by the IP partnership driver.
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 12 on page 24 describes the offload data transfer (ODX) Vdisk and node level I/O statistics.

Table 12. 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 13 on page 25 describes the statistics collection for cloud per cloud account ID.

Table 13. 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 successfully transferred to the cloud
Bytes Down	bdown	Total number of bytes successfully 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 13. 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 14 on page 26 describes the statistics collection for cloud per VDisk.

Table 14. Statistics collection for cloud per VDisk			
SNo	**Statistic name**	**Acronym**	**Description**
1	blocks up	bup	Number of blocks that are uploaded in the cloud.
2	blocks down	bdn	Number of blocks that are downloaded from cloud.

Note: A block is 512 bytes.

Retrieving statistics files

To view the statistics files, you need to retrieve them from the `/dumps/iostats/` directory on the configuration node.

Use the Secure Copy (SCP) or SSH File Transfer Protocol (SFTP) to copy the statistics files from the configuration node to a local workstation. The system copies all the files from the `/dumps/iostats/` directory on the configuration node to the `/tmp/` directory on a local machine. To verify a particulate file in the individual partnership statistic details, check **`Nn_stats`**.

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.

```
scp superuser@cluster_name:/dumps/iostats/* /tmp
```

XML formatting example for IP partnership statistics

The following example shows the summed up view for an IP partnership statistics:

```
<port id="21"
type="IPREP"
type_id="1"
wwpn="0x0000000000000000"
fc_wwpn=""
fcoe_wwpn=""
sas_wwn=""
iqn=""
hbt="266" hbr="0" het="0" her="4"
cbt="0" cbr="306" cet="6" cer="0"
lnbt="0" lnbr="0" lnet="0" lner="0"
rmbt="54792236296" rmbr="281758748" rmet="27812092" rmer="27530254"
iptx="47383884117" iprx="7213917137" ipre="30408" ipsz="611"
ipbz="254" iprt="265" iptc="40307617091" iprc="83362750" />
```

The following example shows an individual IP partnership view:

```
<partnership_ipperf index="1" cluster="fab3_recovery0" remote_cluster_id="0x000002043da02218"
port_id="3"
iptx="16824170146" iprx="3571399065" ipre="0" ipsz="999"
ipbz="434" iprt="86" iptc="13298305502" iprc="23648286" />
<partnership_ipperf index="2" cluster="fab3_recovery1" remote_cluster_id="0x000002043c2061e0"
port_id="4"
iptx="30559713971" iprx="3642518072" ipre="30408" ipsz="1990"
ipbz="916" iprt="77" iptc="27009311589" iprc="59714464" />
```

IBM Storage Insights

IBM Storage Insights is a cloud-based service that allows you to analyze, protect, and optimize your storage infrastructure.

You can integrate IBM Storage Virtualize systems with IBM Storage Insights by using Call Home with cloud services or by installing lightweight applications that are called *data collectors* into your storage environment. The ability to register and connect to IBM Storage Insights is integrated with the management GUI or command-line interface on your IBM Storage Virtualize system so that you can configure it at any time.

To help you monitor a storage system, IBM Storage Insights collects status, configuration, capacity, and performance metadata about its managed disks (MDisks), volumes and volume copies, nodes, drives, and other resources. This metadata includes statistics and other information that's used for analysis, troubleshooting, and viewing in dashboards, charts, and reports.

A key part of monitoring your system includes the detection of potential ransomware attacks. To ensure that, you have the latest storage metadata for detecting those types of attacks, compression and cyber resiliency statistics for volumes are collected every 5 minutes. With these statistics, IBM Storage Insights builds a historical model of a storage system and uses its built-in intelligence and formulas to identify when and where ransomware attacks might be occurring. For more information on statistics, refer Chapter 17, “Cyber resilience,” on page 355.

There are two versions of the IBM Storage Insights service: IBM Storage Insights and IBM Storage Insights Pro. IBM Storage Insights is available at no cost if you own IBM Block Storage. With this version, you can monitor the basic health, status, and performance of various storage resources.

IBM Storage Insights Pro is a subscription-based product that provides a more comprehensive view of the performance, capacity, and health of your storage resources. In addition to the features offered by IBM Storage Insights, IBM Storage Insights Pro provides tools for intelligent capacity planning, storage reclamation, storage tiering, and performance troubleshooting services. Together, these features can help you reduce storage costs and optimize your data center.

To use Call Home with cloud services, ensure you have the following set up:

1. System memory requirement: Ensure that your storage system has at least 128 GB of system memory to support integrated IBM Storage Insights.
2. IBMid: Ensure that you have an IBMid. You can create an IBM account here Create an IBMid.
3. IBM Storage Insights registration: Register for the IBM Storage Insights service at https://ibm.biz/register.
4. Service name: Create a unique name for your IBM Storage Insights service. To ensure that the name is unique, you can use your company name and another identifying feature, such as a location or department. For example, *Bank Texas* or *Insurance Company Accounting*.

For more information about collecting metadata in IBM Storage Insights, see the following topics:

- Monitoring resources with Call Home, see Monitoring resources with Call Home
- Downloading and installing data collectors, see Downloading and installing data collectors

Related information

"Call Home" on page 519
Call Home connects the system to support personnel who can monitor and respond to system events to ensure that your system remains up and running.

Preventive monitoring

To prevent problems before they occur, you can monitor and verify the system hardware and software periodically.

Some of the following activities can be completed on a daily or weekly basis. Typically, you complete the following activities at least every six months.

System environment

Verify that the prerequisite conditions for the system environment are still met. In particular, verify that temperature and moisture limits are not exceeded, and that airflow and power requirements are still met.

Physically verify the system cabling.

- Inspect the system for any loose cables, open doors, and bent covers.
- Inspect the system for any open slots, which can affect airflow.
- Ensure that wrap plugs are installed on the unused optical ports.
- Ensure that optical cables do not have any bend that is less than 3 inches in radius.
- Check for any abnormal noise that is coming from the system.

Check the error and information LEDs on the system. A system attention light, an enclosure fault light, or a system information light indicates that a serviceable event occurred on the system.

Management GUI

In the management GUI, check the status indicators on the Dashboard. In particular, examine the **System Health** tiles section of the Dashboard. Each tile represents a type of component. Tiles with errors and warnings are displayed first.

In addition to the **System Health** tiles, you can monitor the performance statistics for processor (CPU) utilization, host and internal interfaces, volumes, and storage pools.

Check for updates to the system software, and update the software as needed.

Check that sufficient free physical space (`physical_free_capacity`) is available to meet demand until you plan to add more physical capacity.

Check that each storage pool does not share a resource provisioning group (`shared_resources`) with another storage pool. Check that each storage pool has enough free physical space (`physical_free_capacity`) to meet demand until you plan to add more physical capacity.

Check that each over-provisioned managed disk (`over_provisioned`) has enough free physical space (`physical_free_capacity`) to meet demand until you plan to add more physical capacity.

For more information on the event notifications, see Chapter 15, “Notifications,” on page 293.

Chapter 3. 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.

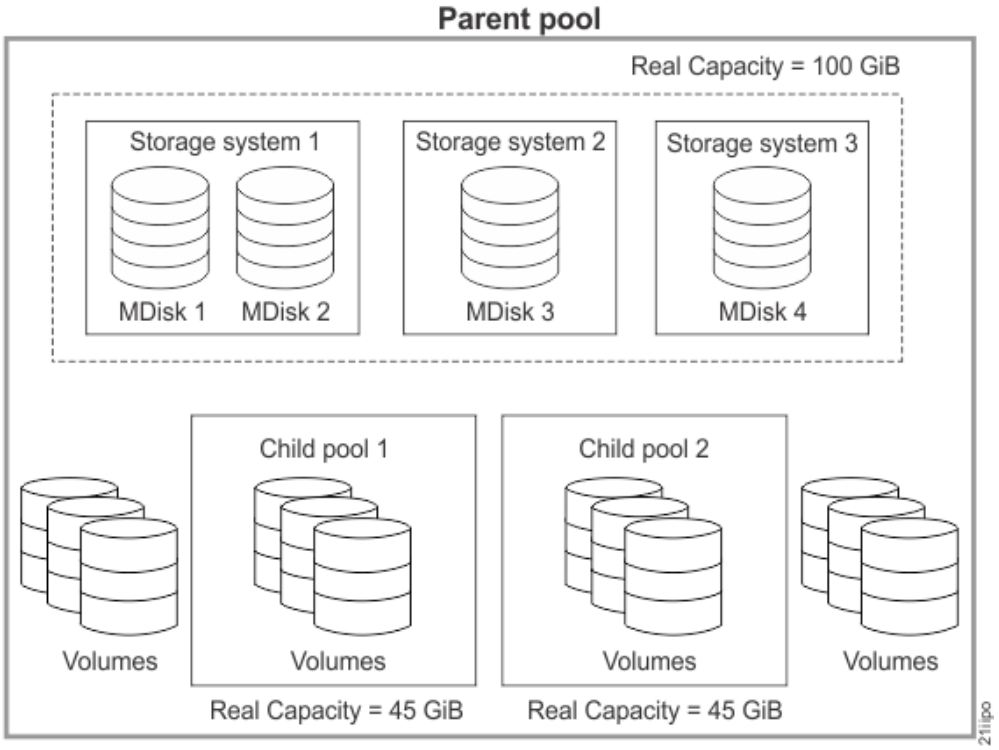


Figure 1. Storage pool

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.

You specify the extent size when you create a new parent pool. You cannot change the extent size later. The choice of extent size affects the total amount of storage that is managed by the system. Consider the following when determining the extent size for your pool:

- A system can manage 2^22 extents. For example, with a 2048 MB extent size, the system can manage up to 2048 MB x 4,194,304 = 8 PB of storage.
- When a volume is created, the storage capacity for the volume is rounded to a whole number of extents.

  Important: You can specify different extent sizes for different systems; however, you cannot migrate (volumes) between systems with different extent sizes. If possible, create all your systems with the same extent size.

You can add MDisks to a pool at any time either to increase the number of extents that are available for new volume copies or to expand existing volume copies. 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.

© Copyright IBM Corp. 2025

Attention: If you add an MDisk to a system as an MDisk, any data on the MDisk is lost. If you want to keep the data on an MDisk (for example, because you want to import storage that was previously not managed), you must create image mode volumes instead.

The **Pools** page in the management GUI displays the **Physical Capacity** and **Logical Capacity** for FCM based systems and non FCM based systems. *Physical capacity* indicates the amount of capacity that is available for volumes before any capacity savings methods are applied. *Logical capacity* indicates the amount of capacity that is available for volumes after capacity savings methods (such as deduplication, compression, and thin-provisioning) are applied.

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 remove MDisks from a parent pool if there are enough free extents are available elsewhere in the pool to move any extents that are in use from this MDisk. The system migrates all extents that are used by volumes in the pool to other MDisks in the parent pool to ensure that data is not lost.

If the parent pool is deleted, you cannot recover the mapping that existed between extents that are in the pool or the extents that the volumes use. If the parent pool has associated child pools, then you must delete the child pools first and return its extents to the parent pool. After the child pools are deleted, you can delete the parent pool. The MDisks that were in the parent pool are returned to unmanaged mode and can be added to other parent pools. Because the deletion of a parent pool can cause a loss of data, you must force the deletion if volumes are associated with it.

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

Child pools are created within parent pools. When a child pool is created without a quota, the child pool can access the entire usable capacity of the parent pool without limit. The usable capacity of the child pool and the usable capacity of the parent pool are reported as the same. Volumes created in the child pool reduce the usable capacity of both the parent pool and the child pool. Both standard child pools and data reduction child pools can be created without a quota

A standard child pool can be created with a quota by specifying its 'size'. When a child pool is created with a quota, the usable capacity for the child pool is reserved from the usable capacity of the parent pool. The usable capacity that is reserved for the child pool is no longer reported as usable capacity of the parent pool. Volumes created in the child pool reduce the usable capacity that is available to the child pool only. Data reduction child pools cannot be created with a quota.

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.

- 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.

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.

With the Safeguarded Copy function, child pools provide a Safeguarded backup location for a group of volumes that are associated with the parent pool. The Safeguarded backup location can contain many snapshots of volume data, each created at a specific interval and with a defined retention period to satisfy your recovery point objective. After the Safeguarded backup location is created, you need to create a volume group and assign a Safeguarded backup policy to the volume group. For more information, see “Creating Safeguarded backup locations” on page 285.

Pool properties

The properties of the pool are:

- Data reduction pools enable data reduction technologies on all drive types. For more information, see “Data reduction pools” on page 60.
- The system supports IBM Easy Tier, a function that responds to the presence of any combination of the drive types within the same pool. For more information, see “Easy Tier” on page 63.
- *Throttling* is a mechanism to control the amount of resources that are used when the system is processing I/Os on a specific pool. For more information, see “Throttles for pools” on page 62.
- A provisioning policy is an entity that defines a set of rules for allocating capacity from a pool. For more information, see “Provisioning policies” on page 185.

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.

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 15. Pool states

State	Description
Online	The pool is online and available. All the MDisks in the pool are available.

Table 15. Pool states (continued)	
State	**Description**
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: 1. Verify that the fabric configuration rules for storage systems are correct. 2. Ensure that you configured the storage system properly. 3. Correct any errors in the event log.
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. **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.

Arrays

The system can manage a combination of internal and external storage and represents all the detected storage as drives.

You can select to have the system automatically assign all detected internal drives to pools, or select to customize the configuration. When the system automatically assigns drives, pools are created based on

the available drive classes and recommended array configurations. In this case, the system might not use all available drives if they do not meet configuration recommendations. If you choose to customize, you can adjust the configuration based on the number of drives and drive criteria. These arrays are presented as MDisks, which can be added to pools.

Array configurations

An array is an ordered group of physical devices (drives) that is used to define logical volumes or devices. An array is a type of MDisk that is made up of drives; these drives are *members* of the array. A Redundant Array of Independent Disks (RAID) is a method of configuring member drives to create high availability and high-performance systems.

The system supports distributed RAID array configurations. *Distributed RAID array* that contains more than two drives have a minimum of one rebuild area. All the drives in an array can process I/O data and provide faster rebuild times when a drive fails. The RAID level provides different degrees of redundancy and performance; it also determines the minimum and maximum number of member drives in the array.

In order to restore redundancy for degraded distributed RAID arrays, the *rebuild-in-place* process is used for reconstructing the data (or parity) directly back into the replaced member drive. Degraded distributed RAID arrays with only two member drives use the *rebuild-in-place* process to restore redundancy, copying the data directly back into the replaced member drive.

If encryption is enabled on your system, you can create RAID arrays that are encrypted. For more information, see “Encryption” on page 455.

Distributed RAID arrays

Distributed RAID array configurations typically contain rebuild areas that are used to maintain redundancy after a drive failure. RAID array configurations that do not contain rebuild areas still maintain redundancy after a drive failure, by using *rebuild-in-place*. Data is rebuilt and redundancy is restored much more rapidly.

For more information and details about distributed RAID arrays, see “Distributed RAID array properties” on page 36.

You can expand distributed RAID arrays to increase the available capacity of the array or add usable capacity. As part of the expansion, the system automatically migrates data for optimal performance for the new expanded configuration. Expansion of distributed RAID arrays supports incremental growth of available capacity for arrays and is compatible with other functions, such as IBM Easy Tier and data migrations.

The type of RAID array configuration that the system uses is based on the type of drives that are detected through the serial-attached SCSI (SAS) or Non-Volatile Memory express (NVMe) protocols. If your system contains SAS-attached drives only, the system can configure distributed RAID arrays based on the number of drives available and the characteristics of those drives. Distributed RAID arrays are preferred because distributed RAID arrays provide better performance and rebuild times. For systems that support only NVMe drives, only distributed RAID arrays are created for array configuration since these drives have improved performance for high-demand drives. Distributed RAID arrays are also supported for compressing drives. NVMe drives are only supported within certain models of control enclosures. For systems that contain both NVMe and SAS drives, the system restricts mixing these different types of drives in the same RAID array configuration. Drive members within the array must be as similar as possible to avoid performance problems and issues replacing drives.

The system also supports NVMe drives that are self-compressing, where the drive itself compresses data as it is written to the drive. If the array includes self-compressing drives, the system uses the following rules when the array is created:

- Non-compressing drives and self-compressing drives cannot be in the same RAID array configuration.
- All drives in the array have the same usable capacity and use the same compression algorithm.
- If the array contains at least 7 drives, the system recommends a distributed RAID 6 array, which is preferred to ensure capacity for self-compressing drives during rebuild operations.

- Only one distributed RAID array of self-compressing drives is allowed per storage pool.

Distributed RAID array configuration guidelines

The management GUI automatically defaults to the recommended number of rebuild areas that is based on the current width and other settings for the distributed RAID array. The width of the array includes the number of physical drives and rebuild areas that are needed to ensure redundancy if drives fail. The system can recommend a maximum of four rebuild areas for distributed RAID arrays for SAS-attached drives and non-compressing NVMe distributed RAID arrays. Compressing NVMe distributed RAID arrays support one rebuild area. Other settings, such as drive class, are also used to determine the best configuration for arrays. For example, if you are adding storage or expanding a distributed RAID array, the management GUI disables any drive classes that are incompatible with drives already in the pool. It then analyzes the current array configuration in the pool to recommend stripe widths and drive counts.

After system setup, you must configure storage by creating pools and assigning storage to specific pools. Ensure that a pool or pools are created before you assign storage. In the management GUI, select **Pools** > **Actions** > **Add Storage**. The **Add Storage** selection automatically configures existing drives into arrays based on array characteristics.

Use the **lsarrayrecommendation** command to display the system recommendations for configuring an array. To configure a distributed RAID array, you can use the **mkdistributedarray** command. For more information on arrays commands, refer to the **Array command**.

The management GUI does not provide array recommendation if the following conditions are present:

- The selected pool already contains a distributed RAID array with compressing drives.
- The selected drive class might not be compatible with other drives that are already in the pool.
- The number of selected drives is not enough to create a valid RAID array.
- The extent size is smaller than the minimum extent size.

If these conditions are present, you can take one of the following actions:

- Select a different drive class.
- Select a different storage pool.
- Select a different number of candidate drives.
- Expand an existing array in the pool instead of adding a new array.

For more information on supported array types and RAID levels, see Configuration limits.

Distributed RAID array properties

Distributed RAID arrays contain rebuild areas that are used to maintain redundancy after a drive fails.

If enough drives are not available on the system, you cannot configure a distributed RAID array. Distributed RAID arrays solve rebuild bottlenecks in array configurations because rebuild areas are distributed across all the drives in the array.

The rebuild write workload is spread across all the drives rather than just the single spare drive, which results in faster rebuilds on an array. Distributed RAID arrays remove the need for separate drives that are idle until a failure occurs. Instead of allocating one or more drives as spares, the spare capacity is distributed over specific rebuild areas across all the member drives. Data can be copied faster to the rebuild area and redundancy is restored much more rapidly. Additionally, as the rebuild progresses, the performance of the pool is more uniform because all of the available drives are used for every volume extent. After the failed drive is replaced, data is copied back to the drive from the distributed rebuild areas. Unlike hot-spare drives, read/write requests are processed on other parts of the drive that are not being used as rebuild areas. The number of rebuild areas is based on the width of the distributed RAID array and on the technology of the member drives.

For example, a distributed RAID array that uses RAID 6 drives can handle two concurrent drive failures. After the two failed drives are rebuilt, the array can, once again, tolerate another two drive failures. If all of the rebuild areas are used, the array becomes degraded on the next drive failure.

The effective capacity of a distributed RAID 1 array that is made from the same number of underlying drives as a RAID 5 or RAID 6 array is in general less than the effective capacity of the latter. However, while capacity-wise, distributed RAID 5 and 6 are better, performance-wise, RAID 1 may be better for certain workloads.

Supported RAID levels

The system supports the following RAID levels for distributed RAID arrays:

Distributed RAID 1
Distributed RAID 1 arrays consist of two mirrored strips, which are distributed across all member drives. Distributed RAID 1 arrays can tolerate a single failed member drive when a rebuild area is in place. If a member drive fails in a distributed RAID 1 array contains two member drives, or the rebuild area is no longer available, the array becomes degraded. Unlike distributed RAID 5 and 6, distributed RAID 1 does not contain any parity strips. In order to restore redundancy for degraded distributed storage arrays, the *rebuild-in-place* process is used, reconstructing the data directly back into the replaced member drive. Degraded distributed RAID arrays with only two member drives use the *rebuild-in-place* process to restore redundancy, copying the data directly back into the replaced member drive.

Distributed RAID 5
Distributed RAID 5 arrays stripe data over the member drives with one parity strip on every stripe. These distributed RAID arrays can support 4 - 128 drives. RAID 5 distributed arrays can tolerate only one failed member drive. In order to restore redundancy for degraded distributed RAID arrays, the *rebuild-in-place* process is used for reconstructing the data (or parity) directly back into the replaced member drive.

Distributed RAID 6
Distributed RAID 6 arrays stripe data over the member drives with two parity strips on every stripe. Distributed RAID 6 arrays can tolerate any two concurrent member drive failures. In order to restore redundancy for degraded distributed RAID arrays, the *rebuild-in-place* process is used for reconstructing the data (or parity) directly back into the replaced member drive.

Example of a distributed RAID array with no rebuild areas

Note: While this example is showing a 2 member drive storage array (with distributed RAID 1), the same rebuild-in-place process occurs in any distributed storage arrays if the member drives have failed to an extent that no rebuild area is available. In these cases, the data is reconstructed, not copied, onto the replaced drive.

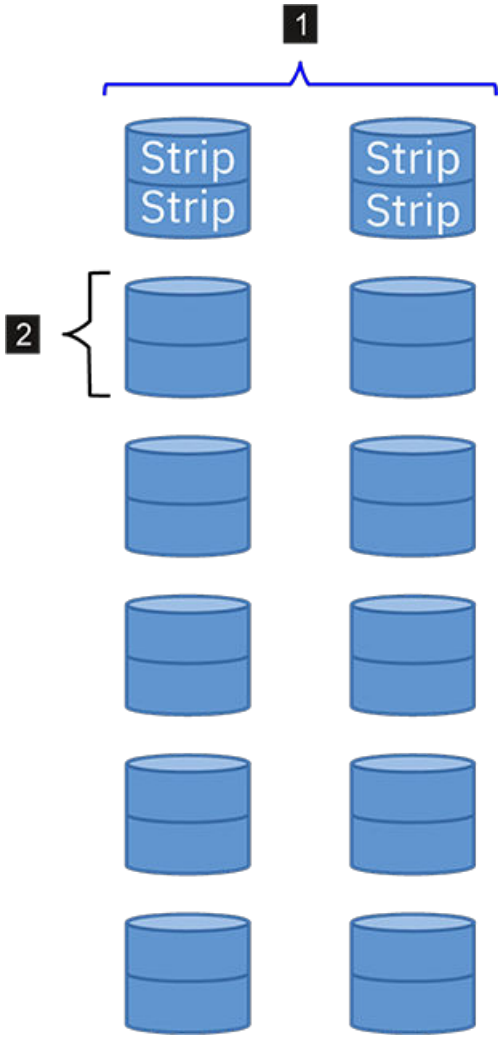


Figure 2. Two member drive storage array with no rebuild area

Figure 2 on page 38 shows an example of a distributed RAID array with two member drives and no rebuild area; both of the drives in the array are active.

1 (Minimum) two active drives and stripe width
2 Pack, with depth of two strips

Figure 3 on page 39 shows a distributed RAID array with two member drives that contains a failed drive. To recover data, the data copies the strip from the active drive to the new or previously failed drive.

1 Active drive
2 Failed drive

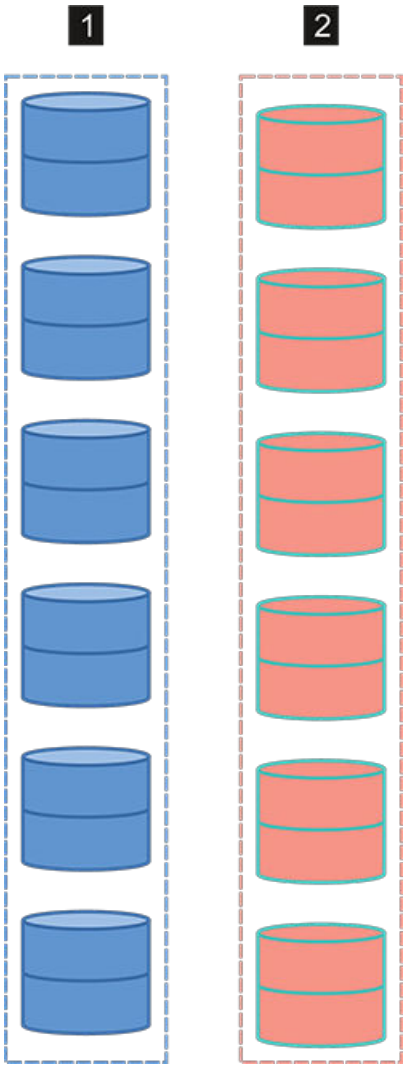


Figure 3. Two member drive storage array with a failed drive

Figure 4 on page 40 shows the rebuild-in-place process, with the new data being copied directly into the replaced drive.

1 Active drive
2 Data being copied directly into the replaced drive

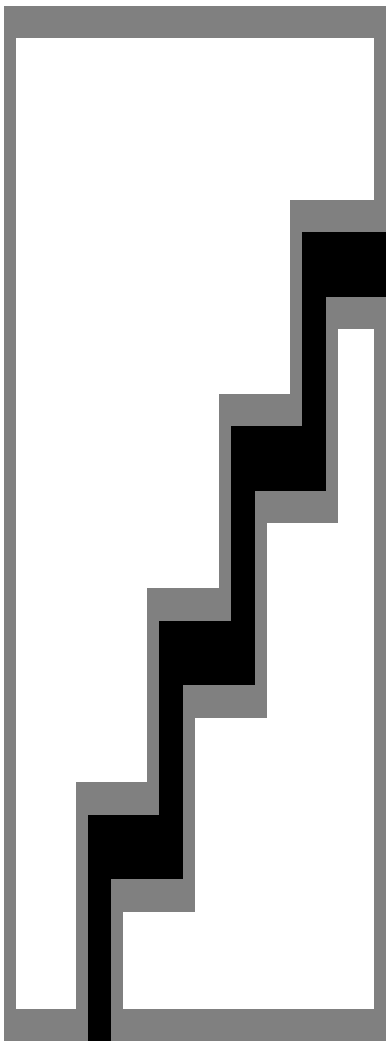


Figure 4. Rebuild-in-place process on the two member drive

Example of a distributed RAID array with rebuild areas

Figure 5 on page 41 shows an example of a distributed RAID array that is configured with RAID 6; all of the drives in the array are active. The rebuild areas are distributed across all of the drives and the drive count includes all of the drives.

1. An active drive
2. Rebuild areas, which are distributed across all drives.
3. Drive count, which includes all drives.
4. Stripes of data (two stripes are shown)
5. Stripe width
6. Pack, which equals the drive count that is multiplied by stripe width.
7. Additional packs in the array (not shown)

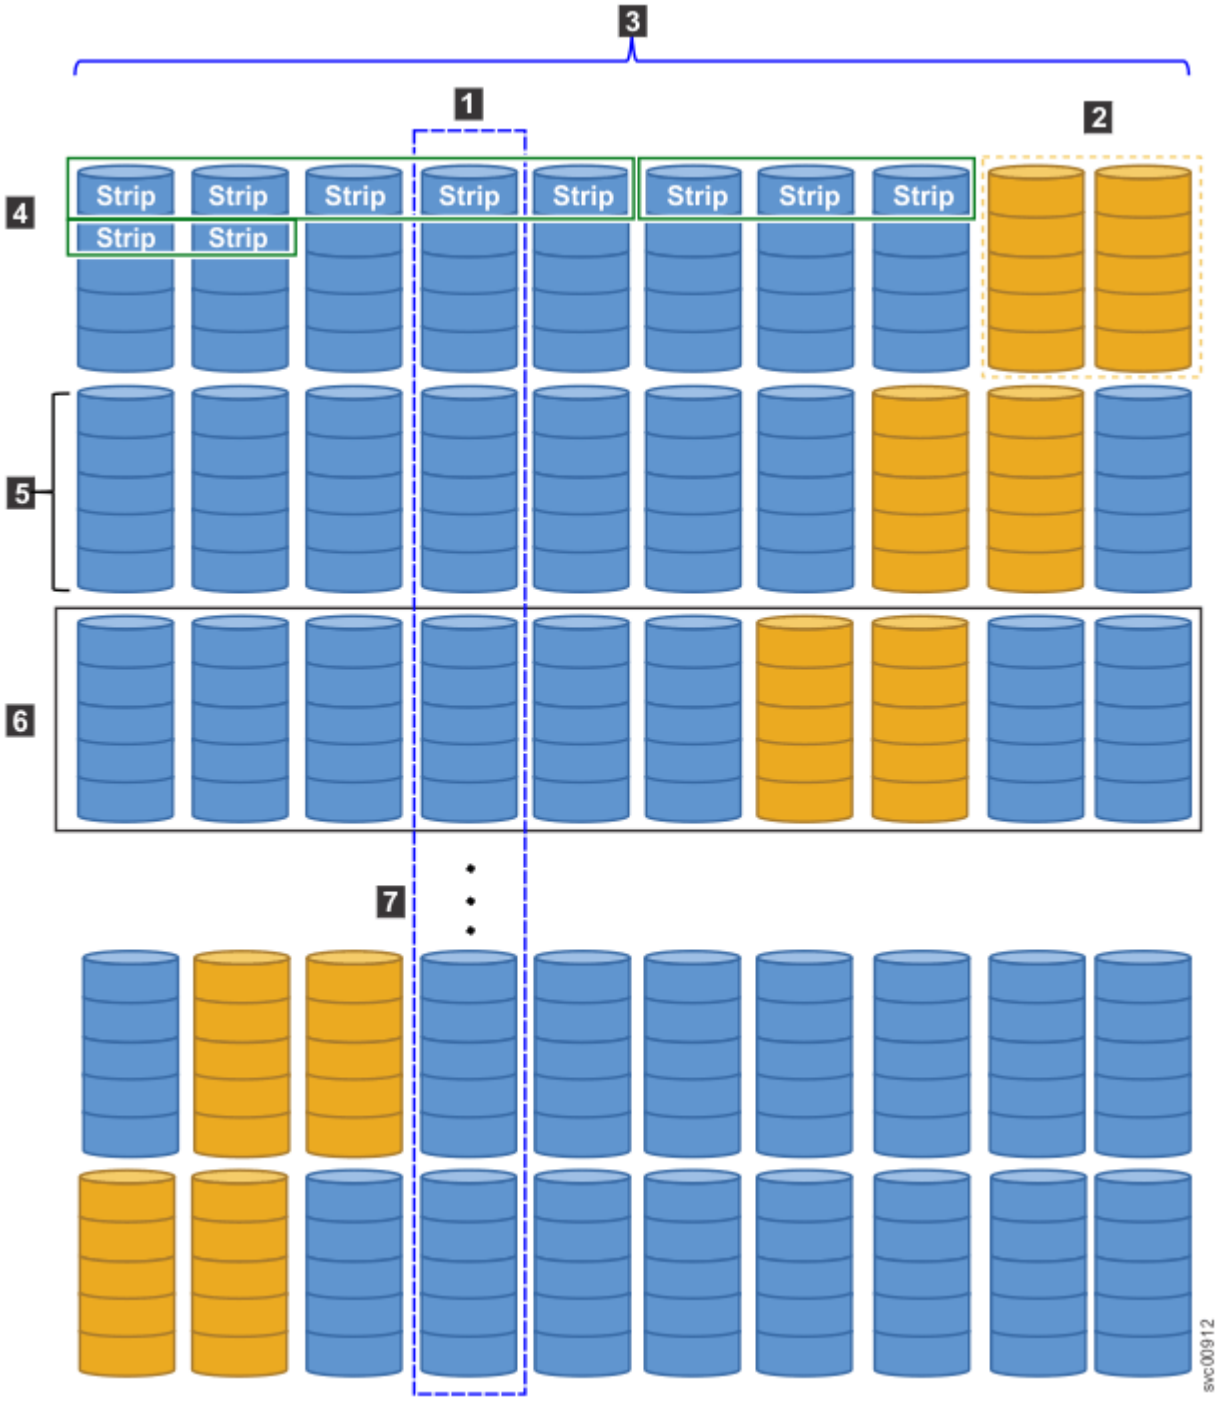


Figure 5. Distributed RAID array (RAID 6 level)

Figure 6 on page 42 shows a distributed RAID array that contains a failed drive. To recover data, the data is read from multiple drives. The recovered data is then written to the rebuild areas, which are distributed across all drives in the array. The remaining rebuild areas are distributed across all drives.

1 Failed drive
2 Rebuild areas, which are distributed across all drives.
3 Remaining rebuild areas rotate across each remaining drive.
4 Additional packs in the array (not shown)

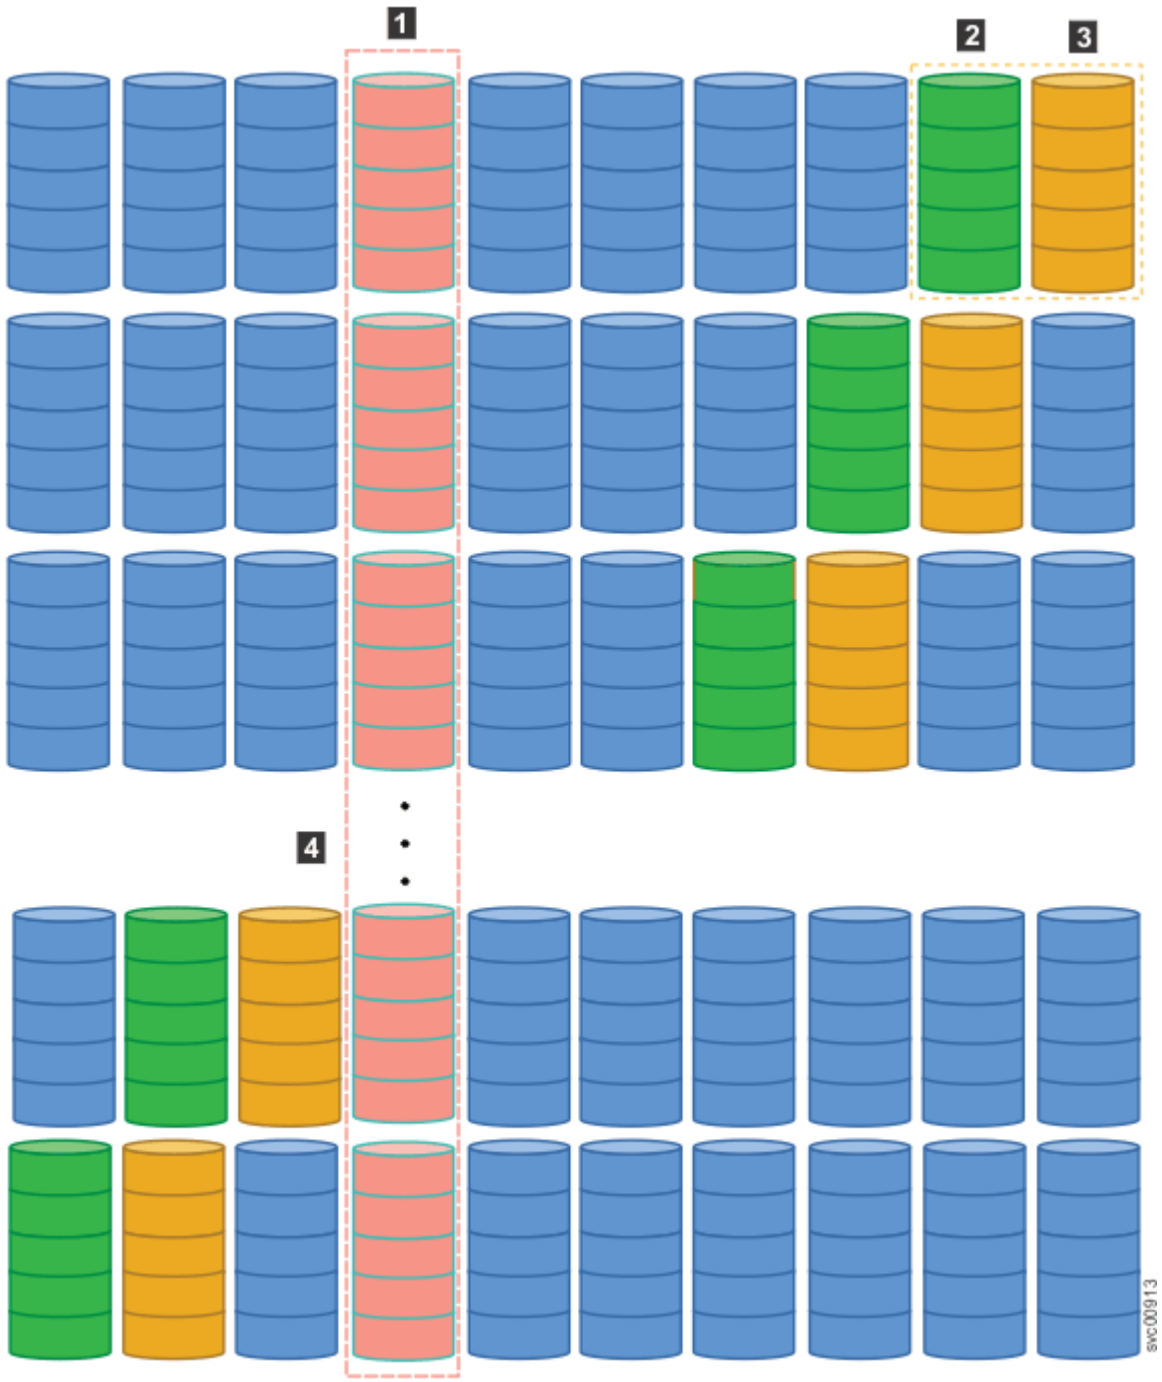


Figure 6. Distributed RAID array (RAID 6 level) with a failed drive

Array width

The *array width*, which is also referred to as the *drive count*, indicates the total number of drives in a distributed RAID array. This total includes the number of drives that are used for data capacity and parity (if there are parity strips, like in distributed RAID 6) and the *rebuild area* (if existing) that is used to restore redundancy of the data.

Rebuild area

The *rebuild area* is the drive capacity that is reserved within a distributed RAID array to reconstruct data after a drive failure; it provides no usable capacity. The rebuild area is distributed across all of the

drives in the array. As data is rebuilt during the *copyback* process, the rebuild area contributes to the performance of the distributed RAID array because all of the member drives perform I/O requests.

Note: Distributed RAID 1 systems with only two member drives do not contain a rebuild area.

Rebuild-in-place

A distributed RAID array process that takes place when no rebuild area is available, which causes the array to become degraded. The *rebuild-in-place* process restores data redundancy by copying or reconstructing the data directly back into the replaced member drive by using the original data distribution.

Note: Not all distributed RAID array storage systems might support the *rebuild-in-place* process. For more information, see "Supported RAID levels" on page 37.

Stripe and stripe width

A *stripe*, which can also be referred to as a *redundancy unit*, is the smallest amount of data that can be addressed. Distributed RAID array strip size is 256 KiB.

The *stripe width* indicates the number of strips of data that can be written at one time when data is rebuilt after a drive fails. This value is also referred to as the *redundancy unit width*. In Figure 5 on page 41, the stripe width of the array is 5.

Drive class

To enhance performance of a distributed RAID array, all of the drives must come from the same, or superior, *drive class*. Each drive class is identified by its *drive_class_id*. The system uses the following information to determine the drive class of each drive:

Block size
Indicates the block size of the drive class. Valid block size is either 512 or 4096 bytes.

Capacity
Indicates the capacity of the drive class.

I/O group
Indicates the I/O group name that is associated with the drive class.

RPM speed
Indicates the speed of the drive class. Valid RPM speed can be 7.2 K, 10 K, or 15 K. For SSDs, this value is blank.

Technology
Indicates the technology for the drive class. The following technology types are supported:

Storage Class Memory
Storage Class Memory drives use persistent memory technologies that improve endurance and reduce latency of current flash storage device technologies. All SCM drives use NVMe architecture.

FlashCore® Module
IBM FlashCore® Module drives are high-performance flash drives that process read and write operations and provide faster access to data than Enterprise or nearline drives. For most FlashCore Module drives, the system monitors their wear level and issues warnings when the drive is nearing replacement.

Tier 1 flash
Tier 1 flash drives are lower-cost flash drives, typically with larger capacities, but slightly lower performance and write endurance characteristics. The system monitors their wear level and issues warnings when the drive is nearing replacement.

Enterprise disks
Enterprise disks are disk drives that are optimized for performance.

Nearline disks
Nearline disks are disk drives that are optimized for capacity.

To replace a failed member drive in the distributed RAID array, the system can use another drive that has the same drive class as the failed drive. The system can also select a drive from a superior drive class. For example, two drive classes can contain drives of the same technology type but different data capacities. In this case, the superior drive class is the drive class that contains the higher capacity drives.

To display information about all of the drive classes that are available on the system, use the **lsdriveclass** command. Example output from the **lsdriveclass** command shows drive classes on the system. In example 1, the deduplication capable drive class 0 contains drives with a capacity of 27.7 TB. In example 2, drive class 0 contains drives with a capacity of 837.9 GB and drive class 1 contains drives with a capacity of 1.1 TB. As drives have the same technology type and block size, drive class 1 is considered to be superior to drive class 0 as it has greater capacity.

Example 1 output from the `lsdriveclass` command

```
id RPM capacity IO_group_id IO_group_name tech_type   block_size candidate_count superior_count
total_count transport_protocol compressed fips deduplicated
0      27.7TB   0           io_grp0       tier0_flash 512        1               1
11          nvme              yes        no   yes
```

A detailed output example

```
lsdriveclass 0
```

```
id 0
RPM
capacity 27.7TB
IO_group_id 0
IO_group_name io_grp0
tech_type tier0_flash
block_size 512
candidate_count 1
superior_count 1
total_count 11
transport_protocol nvme
compressed yes
physical_capacity 5.1TB
fips no
deduplicated yes
```

Example 2 output from the `lsdriveclass` command

```
id RPM   capacity IO_group_id IO_group_name tech_type        block_size  candidate_count
superior_count total_count  transport_protocol compressed  fips deduplicated
0 10000  837.9GB  0           io_grp0       tier0_enterprise 512         11
22             11          sas               no          no   no
1 10000  1.1TB    0           io_grp0       tier0_enterprise 512         11
11             11          sas               no          no   no
```

Slow write priority settings

When a redundant array is doing read/write I/O operations, the performance of the array is bound by the performance of the slowest member drive. If the SAS network is unstable or if too much work is being driven to the array when drives do internal ERP processes, performance to member drives can be far worse than usual. In this situation, arrays that offer redundancy can accept a short interruption to redundancy to avoid writing to, or reading from, the slow component. Write operations that are mapped to a poorly performing drive are committed to the other copy or parity, and are then completed with good status (assuming no other failures). When the member drive recovers, the redundancy is restored by a background process of writing the strips that were marked out of sync while the member was slow.

This technique is governed by the setting of the **slow_write_priority** attribute of the distributed RAID array, which defaults to **latency** when the array is created. When set to **latency**, the array is allowed to become out of sync in an attempt to smooth poor member performance. You can use the

command to change the **slow_write_priority** attribute to **redundancy**. When set to **redundancy**, the array is not allowed to become out of sync. However, the array can avoid suffering read performance loss by returning reads to the slow component from redundant paths.

When the array uses **latency** mode or attempts to avoid reading a component that is in **redundancy** mode, the system evaluates the drive regularly to assess when it becomes a reliable part of the system again. If the drive never offers good performance or causes too many performance failures in the array, the system fails the hardware to prevent ongoing exposure to the poor-performing drive. The system fails the hardware only if it cannot detect another explanation for the bad performance from the drive.

Distributed RAID array member drive replacement

If the fault LED light on a drive is lit, the drive is marked as failed and is no longer used in the distributed RAID array. When the system detects that a failed drive was replaced, it automatically removes the failed hardware from the array configuration. If the new drive is suitable (for example, in the same drive class), the system begins a copyback operation to make a rebuild area available again in the distributed RAID array.

When a drive fails, the deduplication process stops. The physical capacity of the failed drive without deduplication might exceed the available rebuild space, which can result in an incomplete distributed rebuild. In such case, the system notifies you to replace the drive. If the drive is not replaced, the rebuild might pause after some time because of insufficient space.

Planning array configurations

When you plan your network, consideration must be given to the RAID configuration that you use. The system supports distributed RAID 1, 5, and 6 array configurations.

Distributed array

The concept of distributed RAID arrays is to distribute an array with width W across a set of X drives. For example, you might have a 9+P+Q+R RAID 6 array that is distributed across a set of 40 drives. The array type and width define the level of redundancy. In this example, a 25% capacity overhead for parity and rebuild areas exists.

If an array stride needs to be rebuilt, two component strips must be read to rebuild the data for the third component. The set size defines how many drives are used by the distributed array. Performance and usable capacity scale according to the number of drives in the set. The other key feature of a distributed array is that instead having a hot spare, the set includes distributed rebuild area strips that are also distributed across the set of drives. The data and distributed rebuild areas are distributed such that if one drive in the set fails redundancy can be restored by rebuilding data on to the spare distributed rebuild areas at a rate much greater than the rate of a single component.

Note: When you determine the array configuration for your system, plan to create arrays of 256 KiB strip size only.

Distributed arrays of NVMe drives support are used to create large-scale internal managed disks. As a result, rebuild times are dramatically reduced, which decreases the volumes' exposure to the extra load of recovering redundancy. Because the capacity of these managed disks is potentially so great, when they are configured in the system the overall limits change to allow them to be virtualized. For every distributed array, the space for 16 MDisk extent allocations is reserved. Therefore, 15 other MDisk identities are removed from the overall pool of 4096. Distributed arrays also aim to provide a uniform performance level. A distributed array can contain multiple drive classes if the drives are similar (for example, the drives have the same attributes, but the capacities are larger) to achieve this performance. All the drives in a distributed array must come from the same I/O group to maintain a simple configuration model.

A distributed array has the following key benefits:

- Quicker rebuild times with less impact to host I/O.
- More user flexibility in defining how many drives are used by an array (for example, a user can create 9+P arrays with 24 drives without having four drives left unused).

- Distributed rebuild areas means that no idle drives exist in the system, and thus performance improves slightly.
- Degraded distributed RAID arrays can use the *rebuild-in-place* process (if supported) to restore redundancy, copying or reconstructing the data directly back into the replaced member drive.

One disadvantage of a distributed array is that the array redundancy is covering a greater number of components. Therefore, mean time between failure (MTBF) is reduced. Quicker rebuild times improve MTBF; however, limits to how widely distributed an array can be before the MTBF becomes unacceptable remain.

Distributed array expansion

Distributed array expansion allows the conversion of a small, not-very-distributed array into a larger distributed array while it preserves volume configuration and restriping for optimal performance. Expansion offers the option of getting better rebuild performance with an existing configuration without the migration steps that might require excess capacity. Expanding a distributed array is preferable over creating a new small array.

Expansion can increase the capacity of an array, but it cannot change the basic parameter of stripe width. When you plan a distributed array configuration, it is necessary to plan for future array requirements. Additionally, a distributed array that might fit within the extent limit (16*128 K extents) at a particular extent size might not fit if you expand it over time. Planning your extent size for the future is also important. The minimum (and recommended) storage pool extent size is 4096 MiB.

Expansion also benefits NVMe arrays for the same reasons. However, thin provisioned (compressing) NVMe drives add an extra layer of complexity when you calculate the available capacity in the array during an expansion. When you plan for the possible expansion of thin provisioned NVMe arrays, the drives must be the same physical and logical size. When you expand a thin provisioned NVMe array, the usable capacity is not immediately available, and the availability of new usable capacity does not track with logical expansion progress. The expansion process monitors usable capacity usage and analyzes the changes that are caused by the actions it takes during data restriping. This information is used to release the correct amount of usable capacity as it becomes available.

Note: DRAID array expansion is not supported for systems that use FCM5 drives in this release. Support for this configuration is planned for a future release.

Expanding an MDisk in a distributed array

Distributed RAID expansion

You can expand distributed RAID arrays to increase the available capacity of the array or add usable capacity. As part of the expansion, the system automatically migrates data for optimal performance for the new expanded configuration. Expansion of distributed RAID arrays supports incremental growth of available capacity for arrays and is compatible with other functions, such as IBM Easy Tier and data migrations.

You can dynamically add one or more drives and increase the capacity of an existing distributed RAID array that is already in use. The expansion process is dynamic and nondisruptive, which means that you can continue to access the data on the distributed RAID array. New drives are integrated and data is re-striped to maintain the placement of strips across the existing and new components. *Stripe width* does not change during an expansion. This means that when an array is expanded, it would allocate parity in the same ratio. Therefore, not as much of the new drives would be additional capacity compared to an equivalent array with the same number of drives with a higher stripe width, as the ratio of data strips to parity strips would be lower. Expansion does not allow you to change the stripe width or geometry. For more information about *stripe width* and *strips*, see “Distributed RAID array properties” on page 36.

Note: DRAID array expansion is not supported for systems that use FCM5 drives in this release. Support for this configuration is planned for a future release.

Note: When you expand a thin-provisioned NVMe array, the physical capacity is not immediately available, and the availability of new physical capacity does not track with logical expansion progress.

Using the management GUI

To expand a distributed RAID array, complete these steps:

1. In the management GUI, select **Pools** > **MDisks by Pools**.
2. Expand the pool and right-click an MDisk within a distributed RAID array and select **Expand**.
3. Select the drive class that is compatible with the drives in the original MDisk and enter the number of drives to expand the array.
4. Expand **Advanced Settings** to view the number of additional drives that is recommended for the new configuration.
5. Click **Expand**. After starting the expand operation, you can monitor the progress of the expansion by selecting **Monitoring** > **Background Tasks**.

   Note: The **Progress Remaining** value represents the expected time of completion for the data migrations that are required during the expansion of the array. The **Progress Remaining** value can be empty during processing when data migrations are not required or have not started. In these cases, use the percentage value in the **Progress** column or the value displayed in the **Additional Capacity Remaining** field to estimate completion of the expansion.

Using the command-line interface

RAID write response time

This function means that the RAID software layer, where redundancy exists to do so, can prevent drive bad behavior from having an unlimited impact on I/O performance. In addition, the system tries to avoid immediately committing to an array rebuild due to a brief offline event from a single drive, while there is full redundancy.

Enabling RAID latency mode

The system temporarily reduces the array redundancy when drive latency impacts a write task. There is a task that is associated with the repair operation to catch up the out of sync component. Other background tasks stop running when drives are slow to lower the workload and expedite recovery.

If the preference is that write I/O does not cause loss of sync, then the CLI supports this function that is disabled on an array-by-array basis through the **slowwritepriority** parameter to the **charray** and **mkdistributedarray** commands. This parameter does not affect the RAID layer efforts to return reads promptly when drives misbehave. The **slowwritepriority** parameter also does not prevent the array from attempting to avoid a full component rebuild by using an *incremental rebuild* when drives go offline (also displayed as a *resync* task).

If there is excessive use of the function due to I/O, then the availability of the data can be at risk. To warn of an excessive loss of sync, the array logs an event to help identify if a particular member is the cause. The system monitors drives and fails a drive that is consistently causing slow events.

The response time target is not active in the following situations:

- On a RAID 1 mirror where one mirror is rebuilding
- On a RAID 6 array with a rebuilding component, and it is that component that is slow
- During system events

MDisks

A *managed disk (MDisk)* is a logical unit of physical storage. MDisks are not visible to host systems.

An MDisk might consist of multiple physical disks that are presented as a single logical disk to a network. An MDisk always provides usable blocks of physical storage to the system even if it does not have a one-to-one correspondence with a physical disk.

Each MDisk is divided into a number of extents, which are numbered, from 0, sequentially from the start to the end of the MDisk. The extent size is a property of pools. When an MDisk is added to a pool, the size of the extents that the MDisk is divided into depends on the attribute of the pool to which it was added.

The access mode determines how the clustered system uses the MDisk.

Attention: If you observed intermittent breaks in links or if you replaced cables or connections in the SAN fabric or LAN configuration, you might have one or more MDisks in degraded status. If an I/O operation is attempted when a link is broken and the I/O operation fails several times, the system partially excludes the MDisk and changes the status of the MDisk to excluded. You must include the MDisk to resolve the problem.

Access modes

The following list describes the possible access modes for MDisks:

Unmanaged
The MDisk is not used by the system.

Managed
The MDisk is assigned to a storage pool and provides extents that volumes can use.

Image
The MDisk is assigned directly to a volume with a one-to-one mapping of extents between the MDisk and the volume.

Array
The MDisk represents a set of drives in a RAID from internal storage.

Attention: If you add an MDisk that contains existing data to a storage pool while the MDisk is in unmanaged or managed mode, you lose the data that it contains. The *image mode* is the only mode that preserves this data.

Figure 7 on page 48 shows examples of physical disks and MDisks.

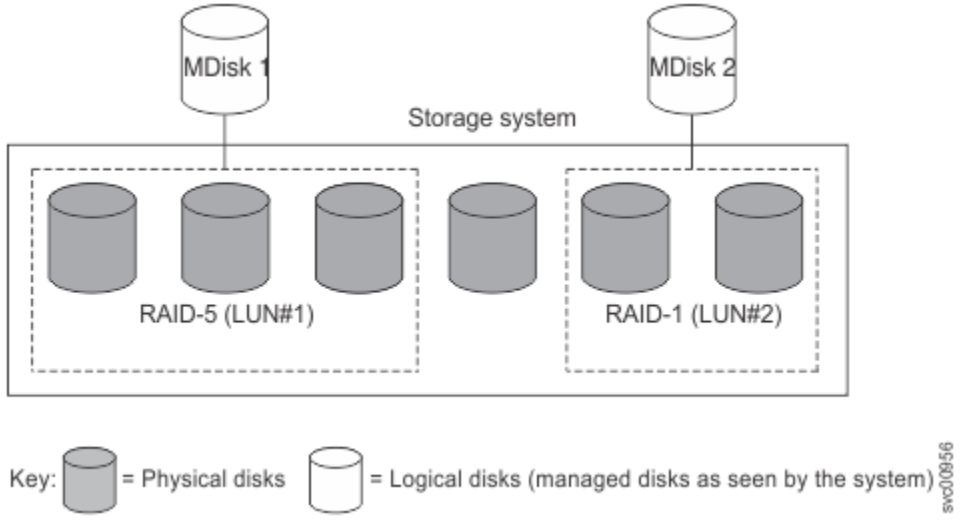


Figure 7. Storage systems and MDisks

Table 16 on page 49 describes the operational states of an MDisk.

Table 16. MDisk status	
Status	**Description**
Online	The MDisk can be accessed by all online nodes. All the nodes that are currently working members of the system can access this MDisk. The MDisk is online when the following conditions are met: • All timeout error-recovery procedures complete and report the disk as online. • Logical unit number (LUN) inventory of the target ports correctly reported the MDisk. • Discovery of this LUN completed successfully. • All of the MDisk target ports report this LUN as available with no fault conditions. • For iSCSI connections, both nodes of at least one I/O group in the system can access the MDisk.
Degraded paths	The MDisk is not accessible to one or more nodes in the system. Degraded path status is most likely the result of incorrect configuration of either the storage system or the SAN fabric. However, hardware failures in the storage system, SAN fabric, or node might also be a contributing factor to this state. To recover from this state, follow these steps: 1. Verify that the fabric configuration rules for storage systems are correct. 2. Ensure that you configured the storage system properly. 3. Correct any errors in the event log.
Degraded ports	The MDisk has one or more 1220 errors in the event log. 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.
Excluded	The MDisk was excluded from use by the system after repeated access errors. Run the directed maintenance procedures to determine the problem.
Offline	The MDisk cannot be accessed by any of the online nodes. None of the nodes that are currently working members of the system can access this MDisk. This state can be caused by a failure in the SAN, storage system, or one or more physical disks that are connected to the storage system. The MDisk is reported as offline if all paths to the disk fail.

Extents

Each MDisk is divided into chunks of equal size called *extents*. Extents are a unit of mapping that provides the logical connection between MDisks and volume copies.

MDisk path

Each MDisk from external storage has an online path count, which is the number of nodes that have access to that MDisk. The path count represents a summary of the I/O path status between the system nodes and the storage device. The maximum path count is the maximum number of paths that were detected by the system at any point in the past. If the current path count is not equal to the maximum path count, the MDisk might be degraded. That is, one or more nodes might not see the MDisk on the fabric.

Troubleshooting medium errors and bad blocks

A storage system returns a medium error response to a host when it is unable to successfully read a block. The system response to a host read follows this behavior.

The system allocates volumes from the extents that are on the managed disks (MDisks). The MDisk can be a volume on an external storage controller or a RAID array that is created from internal drives. In either case, depending on the RAID level that is used, there is normally protection against a read error on a single drive. However, it is still possible to get a medium error on a read request if multiple drives have errors or if the drives are rebuilding or are offline due to other issues.

The system provides migration facilities to move a volume from one underlying set of physical storage to another. In all these cases, the migrated volume or the replicated volume returns a medium error to the host when the logical block address on the original volume is read. The system maintains tables of bad blocks to record where the logical block addresses that cannot be read are. These tables are associated with the MDisks that are providing storage for the volumes.

The **dumpmdiskbadblocks** command and the **dumpallmdiskbadblocks** command are available to query the location of bad blocks.

Important: The **dumpmdiskbadblocks** outputs the virtual medium errors that are created, and not a list of the actual medium errors on MDisks or drives.

It is possible that the tables that are used to record bad block locations can fill up. The table can fill either on an MDisk or on the system as a whole. If a table does fill up, the migration or replication that was creating the bad block fails because it was not possible to create an exact image of the source volume.

The system creates alerts in the event log for the following situations:

- When it detects medium errors and creates a bad block
- When the bad block tables fill up

This table lists the bad block error codes. The recommended actions for these alerts guide you in correcting the situation.

Table 17. Bad block errors

Error code	Description
1840	The managed disk has bad blocks. On an external controller, this error must be a copied medium error.
1226	The system fails to create a bad block because the MDisk already has the maximum number of allowed bad blocks.
1225	The system fails to create a bad block because the system already has the maximum number of allowed bad blocks.

Clear bad blocks by deallocating the volume disk extent, by deleting the volume or by issuing write I/O to the block. It is good practice to correct bad blocks as soon as they are detected. This action prevents the bad block from being propagated when the volume is replicated or migrated. However, it is possible for the bad block to be on part of the volume that is not used by the application. For example, it can be in part of a database that is not initialized. These bad blocks are corrected when the application writes data

to these areas. Before the correction happens, the bad block records continue to use up the available bad block space.

External storage

An external storage system, or *storage controller*, is a device that coordinates and controls the operation of one or more disk drives.

A storage system synchronizes the operation of the drives with the operation of the system as a whole. IBM Storage Virtualize for Public Cloud software only supports cloud storage as external storage. Cloud storage is logical disks at the storage system level. Cloud-based storage is purchased and configured from a supported cloud-based provider.

The system can manage a combination of internal and external storage and represents all the detected storage as MDisks. The management GUI can be used to configure both internally and externally attached storage. Data that is stored on external storage can also be migrated to internal storage, simplifying data and storage management. The system supports Fibre Channel and iSCSI-attached storage systems. Fibre Channel attached storage is automatically detected through the SAN fabric, but iSCSI storage must be manually added to the system for virtualization.

External storage systems reside on the SAN fabric and are addressable by one or more worldwide port names (WWPNs). An external storage system might contain one or more logical units (LUs), each identified by a different logical unit number (LUN). External storage systems that are managed by the system typically contain multiple LUs.

Select **Pools** > **External Storage** in the management GUI to view all MDisks detected by the system. The MDisks are organized by the external storage system that presents them.

For the latest support information about external storage systems, see the IBM support website.

Planning an external storage system

To use Logical Units (LUs) of the external storage systems as Managed Disks (MDisks), several configuration requirements need to be followed to avoid performance and data integrity issues.

Use the following configuration guidelines that describe the system protocol for access to external storage systems LUs through multiple target ports.

External storage-system configuration-requirements

Follow the guidelines to configure the external storage systems arrays.

1. Avoid splitting arrays into multiple logical disks at the external storage systems level. Where possible, create a single logical disk LU from the entire capacity of the external storage systems array.
2. Depending on the redundancy requirements, create RAID-5 (RAID 5) arrays by using 5 - 8 data bits plus parity components that are, 5 + P, 6 + P, 7 + P, or 8 + P.
3. Complete the appropriate calculations to ensure that your external storage systems are configured correctly.
4. If any external storage system that is associated with an MDisk has the **allowquorum** parameter set to no, the **chquorum** command fails for that MDisk. Before you set the **allowquorum** parameter to yes on any external storage system, see IBM support website for external storage systems configuration-requirements.

Exporting external storage systems LUs to the system

Ensure that you are familiar and follow the following guidelines for exporting external storage systems LUs to the system.

- When you define the system as a host object to the external storage systems, you must include all ports on all system nodes and system candidate nodes.

- When you first create an external storage-system LU, you must wait until it is initialized before you export it to the system.
  - Do not present new LUs to the system until the array initialization format is complete. If you add a Logical Unit Number (LUN) to a storage pool before the array initialization format is complete, the storage pool goes offline. While the storage pool is offline, you cannot access the volumes that are in the storage pool.

    Attention: Failure to wait for the LUs to initialize can result in excessive system discovery time and an unstable view of the SAN.
- When you export an external storage-system LU to the system, the LU must be accessible through all ports on the external storage system that are visible to the system.
- The external storage-system LU must be mapped to all the system ports by using the same LUN for all external storage-system ports.

  Attention:

  The system uses the same LUN as the addressing method to identify an external storage-system LU through any access connection. Ensure that an external storage-system LU is mapped to all the system ports by using the same LUN for all external storage-system ports. Failure to do this mapping can cause loss of access to the associated MDisk.

After exporting the LU to the system, a system discovery might need to be requested to discover the new external storage-system MDisks and add them to the MDisks view.

External storage systems LU access through multiple ports

If an MDisk LU is accessible through multiple external storage-system ports, the system ensures that all nodes that access the LU coordinate their activity. The system also ensures that nodes access the LU through the same external storage-system port or group of ports depending on the specific access capabilities of the external storage system.

When the system can access an LU through multiple external storage-system ports, it uses the following criteria to determine the accessibility of these ports:

- The system node is a member of a system.
- The system node has Fibre Channel or iSCSI connections to the external storage-system port.
- The system node successfully discovers the external storage-system LU.
- Slandering does not cause exclusion of system node access to the MDisk through the external storage-system port.

An MDisk path is presented for all system nodes that meet these criteria.

MDisk port selection for external storage systems

When an MDisk is created, the system selects a single port or group of ports of the external storage-system ports to access the MDisk.

Table 18 on page 52 describes the algorithm that the system uses to select the external storage-system port or group of ports.

Table 18. Storage-system port-selection algorithm

Criteria	Description
Accessibility	Creates an initial set of candidate storage-system ports that include the ports that are accessible by the highest number of nodes
Slandering	Reduces the set of candidate storage-system ports to the ports that have lowest number of slanders

Table 18. Storage-system port-selection algorithm (continued)	
Criteria	**Description**
Preference	Reduces the set of candidate storage-system ports to the preferred ports used by the external storage system
Load balance	Selects the single port or group of ports from the set of candidate storage-system ports that has the lowest MDisk access count

After the initial external storage-system port-selection is made for an MDisk, the following events can cause the port selection algorithm to rerun:

- A new node joins the system.
- Requests for discovery of MDisks.
- The external storage system changes its preferred port.
- New external storage-system ports are discovered.
- The currently selected external storage-system port becomes inaccessible.
- Slandering causes exclusion of system node access to the MDisk through the external storage-system port.

Related concepts

"Attaching an external storage system" on page 53
Attach your external storage systems configurations through Fibre Channel or iSCSI connections with system nodes.

"Discovering an external storage system" on page 55
System initialization includes discovery process that identifies external storage systems and the Logical Units (LUs) that they export.

Attaching an external storage system

Attach your external storage systems configurations through Fibre Channel or iSCSI connections with system nodes.

External storage systems attachment through Fibre Channel connections

All nodes in a system must be able to connect to the same set of ports on an external storage system. A system is considered degraded if it contains any two nodes that cannot connect to the same set of external storage-system ports. In such a situation, a system error is logged that requires a repair action.

Use one of the following techniques to control the access to the external storage systems:

- Switch zoning
- LUN-masking capability of the external storage systems

You can configure certain external storage systems to safely share resources between the system nodes and direct-attached hosts. This type of configuration is described as a split external storage system. In all cases, it is critical to configure the external storage system and SAN so that the system cannot access logical units (LUs) that a host or another system can also access. You can arrange the split external storage systems configuration by external storage-system logical unit number (LUN)-mapping and masking.

Configurations where an external storage system is split between two sets of system nodes are also supported. In all cases, it is critical to configure the external storage system and SAN so that the system nodes cannot access LUs that a host or another system can also access. You can use external storage systems LUN mapping and masking to arrange the configuration, failure to which can result in data corruption.

Attention: An external storage-system LU must not be shared between the system and a host or between more than one system. This configuration is not supported and is likely to cause undetected data loss or corruption. Ensure that only a single system can access each external storage-system LU

When an external storage system is detected on the Fibre Channel SAN, the system attempts to recognize it using its Inquiry data. If the device is not supported, the system configures the device as a generic device. A generic device might not function correctly when it is addressed by a node, especially under failure scenarios. However, the system does not regard accessing a generic device as an error condition and does not log an error. Managed disks (MDisks) that are presented by generic devices are not eligible to be used as quorum disks.

External storage systems attachment through iSCSI connections

Use the following guidelines to configure iSCSI-attached storage to the system:

- Unlike Fibre Channel connections, manually configure the connections between the source system and the target external storage systems.
- Discovery of the target IPs and session establishment is manually configured:
  - Either I/O group-wide - that is nodes from a specified I/O group discover and establish sessions, or
  - System-wide - that is all nodes discover and establish sessions.
- After successful discovery, manually establish the initiator sessions from the source system to the target external storage system through:
  - Nodes in a specific I/O group, or
  - All nodes across the system.
- A one-to-one mapping of source ports to target ports is required. The same numbered ports on each source system or node connect to only one port on each target system or node. Multiple initiator ports on the same system node cannot connect to the same target port (many-to-one connectivity). Similarly, one initiator port cannot be connected to multiple target ports on same target external storage system (one-to-many connectivity).
- Zoning is not required.
- Direct attachment between the system and external storage systems is not supported and requires Ethernet switches between the system and the external storage system.
- To avoid a single point of failure, the guidelines suggest use of a dual-switch configuration. For full redundancy, a minimum of two paths between each initiator node and target node must be configured with each path going through a separate switch.
- Extra paths can be configured to increase throughput if both initiator and target nodes support more ports.
- iSCSI storage is not supported for use as a quorum device. Therefore, a SAN Volume Controller system requires at least one Fibre Channel storage controller to be configured in addition to iSCSI storage. For more information, see “Quorum” on page 502.

External storage system direct attachment

If an external storage system with multiple World Wide Node Names (WWNN) is directly attached with the system, unexpected connectivity issues might occur for associated system controller objects in the system event log. Reason for the event logging is redundancy expectation of both initiator and target port connections between the system node and all the controller objects that are associated with the external storage system. In a direct attached configuration, the connection redundancy might not exist. To disable the unexpected event logs related to connection redundancy, use **chcontroller -directattachenhanced yes** command. For more information, see **chcontroller** command.

Discovering an external storage system

System initialization includes discovery process that identifies external storage systems and the Logical Units (LUs) that they export.

The discovery process systematically recognizes all visible ports on the SAN for devices that identify themselves as external storage systems. Each LU that the external storage system exports to the system is also identified. The LUs can contain new storage or a new path for previously discovered storage. The set of discovered LUs forms the managed disk (MDisk) view associated with external storage systems.

Requesting system discovery

The discovery process runs automatically when ports are added to or deleted from the SAN and when some error conditions occur.

To request the discovery process manually, use the **Discover Storage** function from the **External Storage** view in the management GUI. You can also use the **`detectmdisk`** command-line interface (CLI) command to request the discovery process manually. The request to perform discovery causes the clustered system to rescan the SAN. The rescan discovers any new MDisks that were added to the external storage systems and rebalances MDisk access across the available external storage-system ports.

The system discovery process completes asynchronously and might take a few minutes. When using the management GUI to perform system discovery, the GUI checks whether the discovery process is complete. When using the CLI to perform system discovery, use the **`lsdiscoverystatus`** command to check whether the discovery process is still in progress.

Discover new external storage system LUs

You can add newly exported external storage-system LUs to the system by using the system management GUI or the CLI.

When LUs are exported by creating LUN mappings on an external storage system, the system managed disk (MDisk) view is not always automatically updated. To update the view, you must request the system discovery process that is performed.

Note: If the data on the associated external storage-system LU needs to be preserved, then follow the Importing an external storage system procedure.

If newly added external storage systems do not appear in the MDisks display, ensure to cable the system correctly to the external storage systems.

Note: To achieve the best system performance, follow the guidelines for adding MDisks to pools as described in Pool.

Discover external storage-system quorum-disks

During quorum disk discovery, the system assesses each external storage-system LU to determine its potential use as a quorum disk. From the set of eligible LUs, the system nominates three quorum candidate disks.

The MDisk associated with an external storage-system LU must meet the following criteria to be considered as a candidate for a quorum disk:

- It must be in managed mode.
- It must be visible to all nodes in the system.
- It must be presented by an external storage system that is approved for quorum disks.
- It must have sufficient free extents to hold the system state and the configuration metadata.

Note: Quorum disks cannot be placed on the iSCSI-attached external storage systems.

Quorum disk candidates are not found after the discovery because of one of the following situations:

- LU MDisks of the external storage system does not exist in managed space mode. An error is logged when this situation occurs.
- LU MDisks of the external storage system exists in managed space mode, but they do not meet the eligibility criteria. An error is logged when this situation occurs.

Importing an external storage system

You can use image mode volumes to import and migrate existing data that is managed by an external storage system into the system.

You can either use the command-line interface (CLI) or the Storage Migration wizard in the management GUI to migrate data from external storage systems. The Storage Migration wizard simplifies the process by creating the image mode volume of the migrated data, host objects, host mappings, and moving that volume to a storage pool to be managed by the system.

The Storage Migration wizard guides through the appropriate changes to the system. After the migration is complete, remove the external storage system from the system by using the appropriate method for each supported connection type.

If the existing external storage systems do not follow the configuration guidelines, consider completing the data migration away from the image mode volume before you resume I/O operations on the host systems. If I/O operations are continued and the external storage system does not follow the guidelines, I/O operations might fail at the hosts, which might cause loss of access to the data.

The procedure for importing managed disks (MDisks) that contain existing data, depends on the amount of available capacity of the system. You must have the same amount of available capacity in the system as the size of the data that you want to migrate into the system. If you do not have the amount of available capacity, the migration causes the storage pool to have an uneven distribution of data because some MDisks are more heavily loaded than others. Further migration operations are required to ensure an even distribution of data and subsequent I/O loading.

Image mode volumes can be used to migrate existing data from an external storage system. However, they cannot be used to export the replicated volumes to an external storage system. To configure replication on image mode volumes, add volumes to a volume group and then assign a replication policy. Image mode volumes cannot be added to volume groups in a storage partition. A volume group that contains volume with a cache mode other than **`readwrite`**, the replication direction cannot be changed.

For more information on image mode volumes, see Volumes.

Importing image mode volumes with an equivalent amount of available capacity

When you import an image mode volume that has a certain gigabytes capacity, and your system has that amount of available capacity in a single storage pool, then:

- Follow the Start New Migration wizard in the management GUI.
- Select **Pools** > **Import External Storage** to import the image mode volumes and to provide an even distribution of data.

Importing image mode volumes with a smaller amount of available capacity

When you import an image mode volume that has a certain gigabytes capacity, and your system does not have that amount of available capacity in a single storage pool, then:

- Follow the Start New Migration wizard in the management GUI.
- Select **Pools** > **Import External Storage** to import the image mode volumes.

Do not select the destination pool at the end of the wizard. If the destination pool is not selected, the system creates the image mode volumes but does not migrate the data away from the image mode volumes. Use volume mirroring or migration to move the data around as you want.

Managing an external storage system

You can manage the external storage-system resources that are exported to the system by using the system GUI or the system CLI. Managing the external storage system includes how to identify and change the system configuration of the external storage system resources.

Identifying an external storage system

To view the external storage-system serial-number in the management GUI, click **Pools** > **External Storage**. To identify the serial number, run the **lscontroller** command on CLI.

For Fibre Channel external storage systems, the serial numbers, the Fibre Channel worldwide node name (WWNN), or worldwide port name (WWPN) is displayed on the external storage system. The WWNN or WWPN can be used to identify the different external storage systems.

For iSCSI external storage systems, refer to the **iscsi_port_id** and ip fields in the output of **lscontroller**. Since the controller can have multiple Internet Protocol addresses (IP addresses), an iSCSI controller has multiple **iscsi_port_id** and ip listings. The ip field shows the IPv4 or IPv6 address for the iSCSI controller and is blank for other controllers. The **iscsi_port_id** refers to a row ID in the output of **lsiscsistorageport** command. The **lsiscsistorageport iscsi_port_id** lists the IQN of the target iSCSI controller in the **target_iscsiname** output field. Use the IP addresses and IQN gathered to determine the actual external storage system.

To verify the actual external storage system that has the IP addresses and IQN found:

- Start the local user interface, or
- Use the command-line tools that the external storage system provides.

Expanding the capacity of an external storage-system logical unit (LU)

On some external storage systems, you can expand the size of an LU by using vendor-specific disk-configuration software that is provided. The following steps are required for the system to use extra capacity that is provided in this way.

To ensure that the additional capacity is available to the system, follow these steps:

- Remove the managed disk (MDisk) from the storage pool.
- Wait for the migration to complete.
- Verify the specified MDisk is listed as unmanaged by selecting **Pools** > **MDisks by Pool** and expand Unassigned MDisk to display all unmanaged Mdisks.
- Use the vendor-specific, disk-configuration software to expand the size of the logical unit on the external storage system.
- Instigate a system discovery to rescan the Fibre Channel network by following the Discovering an External Storage System procedure.
- Wait for the discovery to complete.
- Check the specified MDisk displays the additional capacity that was expanded.
- Add the specified MDisk back to the pool.

The MDisk extra capacity becomes available for use by the system.

Changing the configuration of an existing external storage system

You can change the configuration of an existing external storage system when you want to delete and replace logical units (LUs).

To delete an existing external storage-system LU and replace it with a new LU:

- Remove the specific existing MDisk from the storage pool by following the procedure to remove an external storage-system LU outlined in the topic Removing an external storage system.

- Check the specified MDisk is no longer displayed in the **Pools** > **External Storage** view.
- Use the vendor-specific, disk-configuration software to create a new LU and add the new LU mapping on the external storage system so that the new LU is now visible to the system.
- Discover the newly exported external storage-system LU by following the procedure in Discovering an External Storage System.
- Check the newly created MDisk is displayed in the **Pools** > **External Storage** view.
- Add the newly created MDisk into a system pool.

The newly created MDisk capacity is now available for the system to use.

Important: When creating new external storage-system LU-mappings, always follow the system requirements as described in Planning an external storage system.

Modifying an external storage-system LU mapping

You can modify the LU mapping of an external storage-system LU. The LU mapping logical unit number (LUN) ID and LU unique identifier (UID) are linked and cannot be changed dynamically when the system is managing the MDisk.

To modify an LU-mapping LUN, complete the following steps:

- Before removing an MDisk from a storage pool, make sure to preserve the existing LU array data on the external storage system.
- Remove the specific existing MDisk from the storage pool by following the Removing an external storage system procedure to remove an external storage-system LU.
- Check the specified MDisk is no longer displayed in the **Pools** > **External Storage** view.
- Use the vendor-specific, disk-configuration software to add the new LU mapping on the external storage system so that the existing LU is now visible to the system through the new LU-mapping LUN.
- Discover the newly exported external storage-system LU by following the procedure in Discovering an external storage system.
- Check the newly re-created MDisk is displayed in the **Pools** > **External Storage** view correctly with the new LU-mapping LUN.
- Add the newly re-created MDisk into a system pool.

The newly re-created MDisk with the changed LUN is now available for use by the system.

Important:

When creating a new external storage-system LU-mapping, always follow the system requirements as described in Planning an external storage system.

Servicing external storage systems

External storage systems that are supported for attachment to the system are designed with redundant components and access paths to enable concurrent maintenance. Hosts have continuous access to their data during component failure and replacement.

Related concepts

“Planning an external storage system” on page 51
To use Logical Units (LUs) of the external storage systems as Managed Disks (MDisks), several configuration requirements need to be followed to avoid performance and data integrity issues.

“Discovering an external storage system” on page 55
System initialization includes discovery process that identifies external storage systems and the Logical Units (LUs) that they export.

Related tasks

“Removing an external storage system” on page 59

You can replace or decommission an external storage system or individual MDisks.

Removing an external storage system

You can replace or decommission an external storage system or individual MDisks.

About this task

During this procedure, you add a device, migrate off the data from the external storage system, and remove the old MDisks.

An alternative procedure is to migrate all volumes that are using storage in this pool to another pool. This procedure consolidates the volumes in a single or new pool. However, you can move one volume at a time. The following procedure describes migration of all the data through a single command.

You can also use this procedure to remove or replace a single MDisk in a pool. If an MDisk experiences a partial failure, such as a degraded array, and you can still read the data from the disk but cannot write to it, then replace that MDisk. Steps “1” on page 59 and “2” on page 59 detail how you can add or remove a single MDisk rather than a list of MDisks.

Complete the following steps to remove an external storage system:

Procedure

1. Add the new MDisks to the pool by completing the following steps:
   a) Click **Pools** > **MDisks by Pool**.
   b) Select the pool that you want to add the new MDisks to and select **Add Storage**.
   c) In the Assign Storage wizard, select **External**.
      Select the external storage system and the MDisks that you want to add. Ensure that you add the number of MDisks so that the capacity of the new MDisks is the same or exceeds the capacity of the old MDisks. Click **Assign**.
2. Force the deletion of the old MDisks from the pool to migrate all the data from the old MDisks to the new MDisks.
   a) Click **Pools** > **MDisks by Pool**.
   b) Expand the pool that you added the new MDisks to and right-click the MDisks to remove and select **Remove**.
      Click **Yes** to migrated data from the removed MDisks to the MDisks that you added to the pool.

      Note: If the MDisk is in image mode and you want to preserve the data on the external storage-system logical unit (LU) instead of converting the volume to a striped volume, then for step “2” on page 59:
      - Go to **Volumes** > **Volumes** view in the management GUI and delete the image mode volume, or
      - If using command-line interface (CLI), issue the **`rmvolume`** command to delete the volume.
3. When all the migration tasks are complete, verify that the MDisks are unmanaged by selecting **Pools** > **MDisks by Pool** and expand **Unassigned MDisk** to display all unmanaged MDisks.
   Ensure that all the MDisks that were removed are listed.
4. Access the external storage system and unmap the LUNs from the system ports.

   Note: If you no longer want to preserve the data that is on the LUNs, you can delete the LUNs.
5. Complete the following steps to have the system rescan the SAN:
   a) Click **Pools** > **MDisks by Pool**.
   b) Select **Discover MDisks** from the **Actions** menu.

   Note: If the unmanaged MDisk still exists in the view and is degraded or excluded, then include the MDisk and instigate discovery again to remove the MDisk.
6. Verify that the external storage system that you want to decommission has no MDisks.

7. Remove the external storage system from the SAN so that the ports can no longer access the external storage system.

External storage system zones

Ensure that you are familiar with the external storage system zoning details. Switch zones that contain storage system ports must not have more than 40 ports. A configuration that exceeds 40 ports is not supported.

Zones

You can configure zones with both host and external storage ports, However, it is recommend configuring them separately. For details on fabric zoning for hosts, see “Zoning and portsets” on page 137. The exception to this is where split host and split storage system configuration is in use.

All nodes in a system must be able to detect the same ports on each back-end storage system. Operation in a mode where two nodes detect a different set of ports on the same storage system is degraded, and the system logs errors that request a repair action. This can occur if inappropriate zoning is applied to the fabric or if inappropriate LUN masking is used. This rule has important implications for back-end storage, such as IBM DS4000 storage systems, which impose exclusive rules for mappings between HBA worldwide node names (WWNNs) and storage partitions.

Each port must be zoned so that it can be used for internode communications. When configuring switch zoning, you can zone some node ports to a host or to back-end storage systems.

When configuring zones for communication between nodes in the same system, the minimum configuration requires that all Fibre Channel ports on a node detect at least one Fibre Channel port on each other node in the same system. You cannot reduce the configuration in this environment.

It is critical that you configure storage systems and the SAN so that a system cannot access logical units (LUs) that a host or another system can also access. You can achieve this configuration with storage system logical unit number (LUN) mapping and masking.

If a node can detect a storage system through multiple paths, use zoning to restrict communication to those paths that do not travel over ISLs.

However, to accommodate the limitations of some switch vendors on the number of ports or worldwide node names (WWNNs) that are allowed in a zone, you can further reduce the number of ports or WWNNs in a zone. Such a reduction can result in reduced redundancy and additional workload being placed on other system nodes and the Fibre Channel links between the nodes of a system.

If the round-trip latency between systems is greater than 80 milliseconds, 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.

The minimum configuration requirement is to zone both nodes in one I/O group to both nodes in one I/O group at the secondary site. The I/O group maintains fault tolerance of a node or port failure at either the local or remote site location. It does not matter which I/O groups at either site are zoned because I/O traffic can be routed through other nodes to get to the destination. However, if an I/O group that is doing the routing contains the nodes that are servicing the host I/O, there is no additional burden or latency for those I/O groups because the I/O group nodes are directly connected to the remote system.

Data reduction pools

Data reduction pools reduce the amount of data that is stored on internal drives and external storage systems through data reduction techniques including compression, deduplication, and thin provisioning. A volume in a data reduction pool can be configured to use different capacity savings methods

simultaneously. Data reduction pools reclaim previously used capacity that is no longer needed by host systems. Support for the host SCSI unmap command is enabled by default.

Note:

- Review the Configuration limits and Software Feature limits for your product to determine whether data reduction pools are supported.
- Data reduction pools are supported on internal arrays and on MDisks presented by external storage systems.
- Data reduction pools are not supported on systems that support FCM5 drives. These drives provide hardware-based data reduction capabilities, including thin provisioning, compression, and deduplication.

To estimate the 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 Comprestimator, see https://www.ibm.com/support/pages/node/6209688.

If you use external storage systems that support data reduction technologies, you can also configure data reduction on the storage systems. The storage system can reclaim that freed storage and reorganize the data on other volumes to more efficiently use the capacity. For standard-provisioned volumes, the system fully controls storage on these storage systems. When a volume is deleted, capacity is freed on the system and can be reallocated; the storage system is not aware of this freed space. However, if the storage system uses compression, thin-provisioning, or deduplication, the storage system controls the use of the usable capacity. In this configuration, when capacity is freed, the system notifies the storage system that capacity is no longer needed. The storage system can then reuse that capacity or free it as reclaimable capacity. The system also supports reclaimable capacity from certain internal drives, such as the 15 TB tier 1 flash drives, which can improve performance on these types of drives.

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 19 on page 61 includes the minimum data reduction pool capacity that is required to be able to create a volume within the pool.

Table 19. 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.

For more information on creating data reduction pools using command-line interface, see mkmdiskgrp command.

Managing pools

MDisks can be added to pools when the pool is created or later.

MDisks can be removed from pools if there are sufficient free extents in the pool. Before physically removing a MDisk from a system, the MDisk must be removed from the pool by using the command-line interface (CLI) or the management GUI.

When MDisks are removed from a pool, any data on the MDisk is migrated to other MDisks in the pool. Wait for this process to complete before physically removing the MDisk. After the MDisk has been physically removed from the system, you may need to perform a rescan so that it disappears from views.

You can determine the relationship between MDisks and Volumes by using the CLI commands lsvdiskmember, lsmdiskmember, lsvdiskextent, and lsmdiskextent.

If using external storage, you can determine the relationship between MDisks and controller LUNs using the CLI commands **`lscontroller`** and **`lsmdisk`**.

Throttles for pools

You can create, modify, and remove throttles for pools by using the management GUI or the command-line interface.

Throttling is a mechanism to control the amount of resources that are used when the system is processing I/Os on a specific pool. If a throttle is defined, the system either processes the I/O, or delays the processing of the I/O to free resources for more critical I/O. You can set throttles for both parent and child pools. Throttles work in a hierarchy when they are applied to both the parent and child pool. Throttles that are defined in the parent pool applies to volumes in both parent and child pools. However, throttles that are defined on the child pool only apply to the volumes in that child pool. For mirrored volumes that are in different pools, only the throttling that applies to the primary copy is considered, regardless of which copy the I/O is directed to. In this case, any throttles defined on the secondary copy of the volume are ignored. Only throttles on the primary copy are used, whether the pool is parent or child pool. When throttles are defined in both parent and child pools, the more restrictive throttle applies to the volume in the child pool, but the I/O is counted against both pools. Consider these examples:

- A parent pool is configured with a throttling limit of 5K IOPS and its single child pool had a limit at 3K IOPS. When I/O operations to a volume in the child pool reaches the throttling limit, the volume cannot receive any more I/O during that time. However, volumes in the parent pool can still use 2K of IOPS to process during that time.
- A parent pool is configured with a throttling limit of 4K IOPS . The parent pool also has two child pool and each child pool has a throttling limit of 3K of IOPS.
- A parent pool is configured with a throttling limit of 5K IOPS . The parent pool has also has two child pool without a throttle defined. All in all, the volumes in each pool cannot handle more than 4K IOPS but each of these pools can could serve its IOPS limit if the volumes in the other pools consume at a lower rate of IOPS In this example, the volumes in the parent and both child pools can reach the 5K IOPS together, regardless of which pool the volumes are located.

Using the management GUI

You can create, edit, and remove throttle for a specific pool by using the **Pools** > **Pools** panel. On the **Edit Throttle** page, enter values for the following throttles:

Bandwidth Limit
Enter the maximum amount of bandwidth the pool can process before the system delays I/O processing for this pool.

IOPS Limit
Enter the maximum I/O operations per second the pool can process before the system delays I/O processing for this pool.

Using the command-line interface

You can use **mkthrottle** command to create a new throttle for bandwidth and IOPS for a pool.

You can use **chthrottle** command to change an existing throttle for the volume and **rmthrottle** command to remove a throttle for pool.

Linked pools

High availability and disaster recovery replication require the storage pools that are used for creating replicated volumes to be linked between systems.

When a volume is created on the active management system for HA, or on the production system for DR, pool links define the pool that is used on the remote system to create the replicated volume. For information about linking pools for replication, see Pool links.

Easy Tier

The system includes IBM Easy Tier, which is a function that responds to the presence of high performance drives in a storage pool that also contains hard disk drives (HDDs). The system automatically and nondisruptively moves frequently accessed data from slower to faster tiers of storage.

Use the following guidelines when you group similar disks:

- Group MDisks with equal performance in a single tier of a pool.
- Group similar arrays in a single tier. For example, configure all 6 + P RAID-5 arrays in one tier of a pool.
- Group MDisks from the same storage systems types in a single tier of a pool.
- Group MDisks that use the same type of underlying physical disk in a single tier of a pool. For example, group MDisks by Fibre Channel or SATA.
- When possible, include similar sized MDisks in a storage pool tier to make it easier to balance the MDisks in that storage pool 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.

If a pool contains one type of MDisk, Easy Tier goes into `balancing` mode. When the pool contains multiple types of MDisks, Easy Tier is automatically turned on. The local MDisks are classified as storage class memory, tier 0 flash, tier 1 flash, enterprise, or nearline drives. The system supports storage class memory, tier 0, and tier 1 flash drives. Storage class memory drives use Non-Volatile Memory Express (NVMe) technology that provides faster performance than the traditional SSDs but are slower than dynamic random-access memory (DRAM)-based drives. Tier 0 flash drives are higher-cost flash drives that provide high performance for read and write operations. Tier 1 flash drives are lower-cost flash drives that have large capacity but lower performance and write endurance characteristics.

Easy Tier modes of operation

Easy Tier offers several modes of operation to optimize the placement of data on flash drives and on HDDs in a storage pool. The system supports flash drives, which offer a number of potential benefits over magnetic hard disk drives (HDDs). The benefits can include faster data access and throughput, better performance, and less power consumption.

Determining the amount of data activity in an extent and when to move the extent to the proper storage tier is usually too complex a task to manage manually.

Easy Tier evaluation mode collects usage statistics for each storage extent for a storage pool where the capability of moving data from one tier to the other tier is not possible or is disabled. An example of such a storage pool is a pool of homogeneous MDisks, where all MDisks are typically HDDs. A summary file is created in the `/dumps` directory on the configuration node (`dpa_heat.`*`node_name.date.time`*`.data`), which can be offloaded and viewed by using the IBM Storage Tier Advisor Tool.

Easy Tier automatic data placement also measures the amount of data access. Easy Tier then acts on the measurements to automatically place the data into the appropriate tier of a storage pool that contains both MDisk tiers.

Dynamic data movement is not apparent to the host server and application users of the data, other than providing improved performance.

For a storage pool and volume to be automatically managed by Easy Tier, ensure that the following conditions are met:

- The volume must be striped.

Remember: Easy Tier is supported on compressed volumes, but it migrates data based only on read activity.

Volumes that are added to storage pools use extents from every type of MDisk initially. By default, the extents that are used are from the Enterprise tier, but Nearline and Flash tiers can also be used. Easy Tier then collects usage statistics to determine which extents to move to the MDisks.

Easy Tier evaluation mode

When IBM Easy Tier evaluation mode is enabled for a storage pool with a single tier of storage, Easy Tier collects usage statistics for all the volumes in the pool.

The system monitors the storage use at the volume extent level. Easy Tier constantly gathers and analyzes monitoring statistics to derive moving averages for the past 24 hours.

Volumes are not monitored when the `easytier` attribute of a storage pool is set to *off*. Volumes are monitored when the `easytier` attribute of a storage pool is set to *measure*.

You can control or view data placement settings by using the following command-line interface (CLI) commands:

`chmdiskgrp`
Modifies the properties of the storage pool. Use this command to turn on evaluation mode on a storage pool with a single tier of storage and to turn off Easy Tier functions on a storage pool with more than one tier of storage.

`lsmdiskgrp`
Lists storage pool information.

`lsvdisk`
Lists volume information.

`lsvdiskcopy`
Lists volume copy information.

`mkmdiskgrp`
Creates a storage pool.

Other MDisk commands such as **`addmdisk`**, **`chmdisk`**, and **`lsmdisk`** can be used to view or set the tier an MDisk belongs to.

Automatic data placement

When IBM Easy Tier is active, the system completes several functions to automatically place data in the appropriate storage tier. Easy Tier measures the host access activity to the data on each storage extent. It also provides a mapping that identifies high activity extents, and then moves the high-activity data according to its relocation plan algorithms.

To automatically relocate the data, Easy Tier initiates the following processes:

1. Monitors volumes for host access to collect average usage statistics for each extent over a rolling 24-hour period of I/O activity.
2. Analyzes the amount of I/O activity for each extent to determine whether the extent is a candidate for migrating to or from the higher performing flash drive tier.

3. Develops an extent relocation plan for each storage pool to determine exact data relocations within that storage pool. Easy Tier then automatically relocates the data according to the plan.

While it relocates volume extents, Easy Tier follows these actions:

- Attempts to migrate the most active volume extents first.
- Refreshes the task list as the plan changes. The previous plan and any queued extents that are not yet relocated are abandoned.

Automatic data placement is enabled by default for storage pools with more than one tier of storage. When you enable automatic data placement, by default all striped volumes are candidates for automatic data placement. Image mode and sequential volumes are never candidates for automatic data placement. When automatic data placement is enabled, I/O monitoring is undertaken for all volumes whether or not the volume is a candidate for automatic data placement. When automatic data placement is enabled and sufficient activity warrants relocation, extents will be relocated within a day after enablement. You can control whether Easy Tier automatic data placement and I/O activity monitoring is enabled or disabled by using the settings for each storage pool and each volume. Each command in the following list that can create or change the settings for storage pools can enable or disable both Easy Tier functions. Any command that can create or change the settings for volumes can enable or disable automatic data replacement, if automatic data replacement is enabled for the storage pool. You can control or view automatic data placement by using the following command-line interface (CLI) commands:

addvdiskcopy
Adds a copy to an existing volume by changing a nonmirrored volume into a mirrored volume.

chmdiskgrp
Modifies the properties of the storage pool. Use this command to turn on Easy Tier evaluation mode or I/O monitoring and to turn off Easy Tier functions on a storage pool with more than one tier of storage.

Note: When automatic data placement is active on a storage pool, set a warning threshold for the storage pool. Automatic data placement cannot function if the storage pool is used 100%.

chvdisk
Modifies the properties of a volume.

lsmdiskgrp
Lists storage pool information.

lsvdisk
Lists volume information.

lsvdiskcopy
Lists volume copy information.

mkmdiskgrp
Creates a storage pool.

mkvdisk
Creates sequential, striped, or image mode volumes.

If you want to disable automatic data placement for a volume or storage pool, set the `easytier` attribute to *off*.

Easy Tier acceleration mode

You can change the mode rates of extent migrations through the IBM Easy Tier function and pool balancing functions.

Normally, your system runs with the Easy Tier accelerated mode set to **off**. This system default rate ensures that migrations do not affect system performance.

To view the currently configured value, enter the **lssystem** command. The default view for `easy_tier_acceleration` is **off**. The maximum rate of extent migrations in normal mode (the system default) is 12 GB per 5 minutes. If `easy_tier_acceleration` is set to **on**, the maximum migration rate

is 48 GB per 5 minutes. The migration rate of 48 GB per 5 minutes in accelerated mode cannot always be guaranteed.

You can use the **chsystem** command to turn on accelerated mode and extend the limit of the migration rate from 12 GB per 5 minutes to 48 GB per 5 minutes:

```
chsystem -easytieracceleration on
```

You can return to the rate of 12 GB per 5 minutes by turning off the accelerated mode:

```
chsystem -easytieracceleration off
```

You can also temporarily enable the accelerated mode to rapidly use new capacity, such as in the following examples:

- You added more capacity to an existing storage pool either by adding to an existing tier or by adding a tier to the pool. When you enable the accelerated mode, the system can quickly spread the existing volumes into the new capacity.
- You are migrating multiple volumes between pools, and the target pool has more tiers than the source pool. Volume migration between storage pools preserves the tier of the volume's extents. When you enable the accelerated mode, the system can quickly take advantage of the additional tiers in the target pool.

To avoid the possibility of performance issues that are caused by overloading the managed disks, use the accelerated mode during periods of reduced system activity.

Easy Tier automatic data placement requirements, recommendations, and limitations

Some limitations exist if you use the IBM Easy Tier function on your system.

- To ensure optimal performance, all MDisks in a storage pool tier must have the same technology and performance characteristics.
- Easy Tier automatic data placement is not supported on volume copies that are image mode or sequential. I/O monitoring for such volumes is supported, but you cannot migrate extents unless you convert the image or sequential volume copies to striped volumes.
- Automatic data placement and extent I/O activity monitors are supported on each copy of a mirrored volume. The Easy Tier function works with each copy independently of the other copy. For example, you can enable or disable Easy Tier automatic data placement for each copy independently of the other copy.
- When a volume is migrated out of a storage pool that is managed with the Easy Tier function, Easy Tier automatic data placement mode is no longer active on that volume. Automatic data placement is also turned off while a volume is being migrated between two pools that have Easy Tier automatic data placement enabled. Automatic data placement for the volume is reenabled when the migration completes.

Limitations when removing an MDisk by using the force parameter

When an MDisk is deleted from a storage pool with the **force** parameter, extents in use are migrated to MDisks in the same tier as the MDisk being removed, if possible. If insufficient extents exist in that tier, extents from the other tier are used.

Limitations when migrating extents

When Easy Tier automatic data placement is enabled for a volume, the **migrateexts** command-line interface (CLI) command cannot be used on that volume.

Limitations when migrating a volume to another storage pool

When the system migrates a volume to a new storage pool, Easy Tier automatic data placement between the generic flash drive tier and the Flash, Enterprise, or Nearline tier is temporarily suspended. After the volume is migrated to its new storage pool, Easy Tier automatic data placement between the generic flash drive tier and the Flash, Enterprise, or Nearline tier resumes for the newly moved volume, if appropriate.

When the system migrates a volume from one storage pool to another, it attempts to migrate each extent to an extent in the new storage pool from the same tier as the original extent. In some cases, for example when the target tier is unavailable, the other tier is used. For example, the generic flash drive tier might be unavailable in the new storage pool.

If the automatic data placement is enabled in the new storage pool, pending Easy Tier status changes are assigned after the volume completes its move to the new storage pool. Although the status changes are based on volume use in the old storage pool, the new status is honored in the new storage pool.

Limitations when migrating a volume to image mode

Easy Tier automatic data placement does not support image mode. No automatic data placement occurs in this situation. When a volume with Easy Tier automatic data placement mode active is migrated to image mode, Easy Tier automatic data placement mode is no longer active on that volume.

The Easy Tier function does support evaluation mode for image mode volumes.

Easy Tier reports

The management GUI supports monitoring Easy Tier data movement in graphical reports. The data in these reports helps you understand how Easy Tier manages data between the different tiers of storage, how tiers within pools are used, and the workloads among the different tiers. Charts for data movement, tier composition, and workload skew comparison can be downloaded as comma-separated value (CSV) files.

To view Easy Tier data and reports in the management GUI, select one of the following paths:

- From the management GUI, select **Monitoring** > **Easy Tier Reports**
- From the management GUI, select **Pools** > **View Easy Tier Reports**

Data movement statistics

The data movement chart displays the migration actions that are triggered by Easy Tier.

The data movement chart displays the data for the previous 24-hour period, in one-hour increments. You can change the span of time display, and the increments that you would prefer, such as hourly.

To view Easy Tier data and reports in the management GUI, select one of the following paths:

- From the management GUI, select **Monitoring** > **Easy Tier Reports**
- From the management GUI, select **Pools** > **View Easy Tier Reports**

Access the data movement statistics chart by selecting the **Data Movement** tab.

The first time that you open the Easy Tier reports page, the Data Movement chart is automatically displayed. After you modify this chart, or view a different chart, your last action is retained until a new set of data is acquired through the Easy Tier Statistics tool. (The tool processes the Easy Tier data and generates the HTML reports and CSV files).

The following list describes the type of data that is displayed in the chart:

Extent capacity
The Y-axis displays the number of moved capacity extents. Promoted (rebalanced) extents display above the line, and demoted (swapped) extents are displayed below the line.

Selected viewport time
Displays the time and date that is rendered in the viewport selection.

Viewport X-axis time selection
Displays the time selection in hours. Values that can be chosen are in hours, and the options are 24, 12, 6, 4, 3, 2, or 1 hour.

Restriction: Displaying large amounts of data, such as a 24-hour period in 5-minute increments, is limited to the display area. There is not enough display space to see all 288 data points in the viewport. Use the horizontal scroll bar to view all the data. You can also use the viewpoint shifter to move to the next set of data that is not being displayed in the current viewpoint.

X-axis increment selection
Displays the time that is selected in increments. You can modify the increments to choose the time that is displayed in the viewport. Values are 1 hour, 30 minutes, 15 minutes, 10 minutes, or 5 minutes, depending on the increments available for that section.

Data movement key
Displays all the data movement actions that take place. The types of actions are:

Promote
When active data is moved from a lower tier of storage to a higher tier to improve the overall system performance.

Warm Promote
When an MDisk becomes overloaded, active data is moved from a lower tier to a higher tier to reduce the workload on the MDisk.

Balancing Move
When data is moved within the same tier from an array with a higher workload, to one with a lower workload to balance the workload within the tier.

Balancing Swap
When data is moved within the same tier from an array with higher workload, to one with a lower workload to balance the workload within the tier.

Promote Swap
When active data is moved from a lower tier of storage to a higher tier, other less-active data must first be moved to make space.

Warm Demote
When an MDisk becomes overloaded, active data is moved from a higher tier to a lower tier to reduce the workload on the MDisk.

Cold Demote
When inactive, or less-active data is moved from a higher tier of storage to a lower tier to free up space on the higher tier.

Movement details
Displays a textual and visual description of the different data movement actions that are listed in the key.

The detailed format is described in the following table:

Table 20. Data move statistic descriptions

Data	Data format	Definition	Comments
index	Integer	Row number	
dss	String	Cluster ID	IBM.2145-169933
volume_id	Integer	Volume ID (VDisk ID)	Example: 0
copy_id	Integer	cdisk ID	Example: 0
extent	Integer	The ID of the extent in this volume	

Table 20. Data move statistic descriptions (continued)

Data	Data format	Definition	Comments
time	ISO 8601 format	The time stamp when the movement happened	2016-10-26T21:37:05+05:00
type	Integer	Type of movement operation	Possible Value: 1~10
typedescription	String	Description of the type of data movement	Possible values: • PROMOTE(type=1) /* Move an extent to a more capable tier */ • SWAP(type=2) Move an extent to a less capable tier which frees space on a more capable tier. Also move an extent to the more capable tier. • DEMOTE(type=3) Move an extent to a less capable tier. • MOVE(type=4) Move an extent to another MDisk in the same tier • RB_SWAP(type=5) Swap extents between MDisks in the same tier • WARM_DEMOTE (type=6) Move an extent to a more capable MDisk in another tier. • WARM_PROMOTE(type=10)
source	Integer	The MDisk that extents are moved from	MDisk number
target	Integer	The MDisk that extents are moved to	MDisk number

Tier composition statistics

The Tier Composition chart displays the distributed workload between the top tier, middle tier, and bottom tier for the selected pool.

To view Easy Tier data and reports in the management GUI, select one of the following paths:

- From the management GUI, select **Monitoring** > **Easy Tier Reports**
- From the management GUI, select **Pools** > **View Easy Tier Reports**

Select the **Tier Composition** tab to display an interactive chart of the tiers that are within the selected pool.

The following list describes the type of data that is displayed in the chart:

Tier types
Displays the type of tiers that are in each category.

MDisk Mapping
Displays a table for the selected pool that describes which MDisks are mapped to which tier.

Extent Capacity
X-axis of the chart

Key
Displays the breakdown of the potential type of workload that might be in each tier category: The following values are available:

Active
Data with more than 0.1 IOPS / Extent access density for small IOPS

Inactive
Data with zero IOPS / Extent access density (no recent activity)

Active large
Data that is not classified

Low activity
Data with more than zero IOPS / Extent access density

Composition details
Displays more detailed information about the workload that is listed in the chart key.

The detailed format is described in the following table:

Table 21. Workload categorization descriptions

Data	Data format	Definition	Comments
index	Integer	Row number	
dss	String	Cluster ID	IBM.2145-169933
datetime	ISO 8601 format	The time stamp when the data is collected	2011-11-14T13:24:49+08:00
extpool	Integer	Pool ID	
tier	Integer	The tier of the MDisk	Possible tiers are in the range 0 - 2.
extentcount	Integer	Number of extents	
smallwps	Float	Small write IOPS targeted against the data	
smallrps	Float	Small read IOPS targeted against the data	
smallmbs	Float	Small I/O bandwidth targeted against the data	MBps
largewps	Float	Large write IOPS targeted against the data	
largerps	Float	Large read IOPS targeted against the data	

Table 21. Workload categorization descriptions (continued)			
Data	**Data format**	**Definition**	**Comments**
largembs	Float	Large I/O bandwidth targeted against the data	MBps
smallms	Float	Average response time on small I/O against the data	ms
largems	Float	Average response time on large I/O against the data	ms

Workload skew statistics

The Workload Skew Comparison chart displays the percentage of I/O workload compared to the total capacity.

To view Easy Tier data and reports in the management GUI, select one of the following paths:

- From the management GUI, select **Monitoring** > **Easy Tier Reports**
- From the management GUI, select **Pools** > **View Easy Tier Reports**

Access the Workload Skew Comparison chart by selecting the **Workload Skew Comparison** tab.

The following list describes the type of data that is displayed in the chart:

Tier Composition
Displays the usage of the different tiers in comparison to the I/O workload. The tiers are broken down to: Top Tier, Middle Tier, and Bottom Tier.

Key
Displays the sum of small I/O and the sum of MBps.

Skew Chart Details
Displays more details about the key and the make up of each tier.

Exporting Easy Tier report data by using the management GUI

You can export the Data Movement, Tier Composition, and Workload Skew Comparison data to your local system as a comma-separated value (CSV) file for viewing in other applications such as a spreadsheet.

To export the Easy Tier data by using the management GUI, complete these steps:

1. Access the Easy Tier statistics and charts by using one of two pathways:
   - From the management GUI, select **Monitoring** > **Easy Tier Reports**
   - From the management GUI, select **Pools** > **View Easy Tier Reports**

   Note: The first time that you open the page, the Data Movement chart is displayed. After you modify this chart, or view a different chart, your last action is retained until a new set of data is acquired through the STAT tool. (The STAT tool processes the Easy Tier data and generates CSV files).
2. Select the pool that contains the data that you want to export.
3. Click **Export** on the top of the GUI pane.
4. In the **Export Easy Tier Reports for Pool** page, select the type of data that you want to export for that pool. All data from the last 24 hours is exported. You can choose from the Data Movement, Tier Composition, and Workload Skew Comparison charts. You can select one, two or all three types of data.
5. Specify a destination that you want to save the file, and click **Save**.

Note: A compressed file with the default name of easytierstats_*date_timestamp*.zip is downloaded. You can modify the file name to suit your needs.

6. Locate the saved file, and extract the data to view in a spreadsheet or text editor. The default file names are data movement.csv, skewcurve.csv, and workload.csv.

Extracting and viewing performance data with the IBM Storage Tier Advisor Tool

You can use the IBM Storage Tier Advisor Tool, which is also referred to as advisor tool, to view performance data that is collected by IBM Easy Tier over a 24-hour operational cycle. The advisor tool creates Hypertext Markup Language (HTML) files that you use to view the data when you point your browser to the file.

About this task

To obtain the Storage Tier Advisor Tool, search for "Storage Tier Advisor" at the www.ibm.com/support website and download the software from IBM Fix Central.

To extract the summary performance data, follow these steps by using the command-line interface (CLI).

Procedure

1. Enter the following command to find the current dpa_heat.*node_name.date.time*.data files in the clustered system.

```
lsdumps -prefix /dumps/easytier
```

2. If necessary, enter the following command to copy the summary performance data-files to the current configuration node.

```
cpdumps -prefix /dumps/easytier/dpa_heat.node_name.date.time.data node_id | node_name
```

3. Use Secure Copy (SCP) or SSH File Transfer Protocol (SFTP) to copy the summary performance data in a binary format from the configuration node to a local directory.
4. From a Microsoft Windows command prompt, use the advisor tool to analyze the binary files. The results are stored in HTML files that are put into a local directory.
5. Point your browser to the HTML files in your local directory.

Results

The advisor tool displays the following statistical reports.

System summary report

- The number of storage pools that are monitored
- The number of volumes that are monitored
- The total capacity that is monitored
- The estimated total hot data capacity
- List of storage pools that are monitored, including pool ID, capacity, tier configuration, tier status, and data management status.

System recommendation report

- A sorted list of suggested flash MDisks to add to the storage pools (sorted by estimated performance improvement)
- A sorted list of suggested Enterprise MDisks to add to the storage pools (sorted by estimated IOPS improvement)
- A sorted list of suggested Nearline MDisks to add to the storage pools (sorted by estimated Bandwidth improvement)

- For each suggested MDisk: The estimated improvement

Storage pool recommendation report

- For each storage pool: list MDisk's current IOPS utilization and projected IOPS utilization
- For each storage pool: show the workload distribution on each tier
- For each storage pool: a list of suggested flash MDisks to add to the storage pools
- For each storage pool: a list of suggested Enterprise MDisks to add to the storage pools
- For each storage pool: a list of suggested Nearline MDisks to add to the storage pools
- For each suggested MDisk: The estimated improvement

Volume heat distribution report

- The distribution of hot data and cold data for each volume copy
- The configured capacity of a volume copy, along with the volume ID, copy ID
- The percentage of each volume copy's IO workload within the storage pool
- The capacity of each volume copy already on each tier

What to do next

You can view this information to analyze workload statistics and evaluate which logical volumes might be candidates for Easy Tier management. If you did not enable the Easy Tier function, the usage statistics that are gathered by the monitoring process can help you determine whether to use Easy Tier to enable potential performance improvements in your storage environment.

Easy Tier settings for storage pools and volumes

This table describes the storage pool Easy Tier settings.

Table 22. Easy Tier settings for storage pools and volumes

Storage pool Easy Tier setting	Number of tiers in the storage pool	Volume copy Easy Tier setting	Volume copy Easy Tier status
Off	One	Off	inactive (see note “1” on page 74)
Off	One	On	inactive (see note “1” on page 74)
Off	Two	Off	inactive (see note “1” on page 74)
Off	Two	On	inactive (see note “1” on page 74)
Measure	One	Off	measured (see note “2” on page 74)
Measure	One	On	measured (see note “2” on page 74)
Measure	Two	Off	measured (see note “2” on page 74)
Measure	Two	On	measured (see note “2” on page 74)
Auto	One	Off	measured (see note “2” on page 74)

Table 22. Easy Tier settings for storage pools and volumes (continued)

Storage pool Easy Tier setting	Number of tiers in the storage pool	Volume copy Easy Tier setting	Volume copy Easy Tier status
Auto	One	On	balanced (see note “3” on page 74)
Auto	Two	Off	active (see note “4” on page 74)
Auto	Two	On	measured (see note “2” on page 74)
On	One	Off	measured (see note “2” on page 74)
On	One	On	balanced (see note “3” on page 74)
On	Two	Off	measured (see note “2” on page 74)
On	Two	On	active (see note “4” on page 74)

Note:

1. When the volume copy status is `inactive`, no Easy Tier functions are enabled for that volume copy.
2. When the volume copy status is `measured`, the Easy Tier function collects usage statistics for the volume but automatic data placement is not active.
3. When the volume copy status is `balanced`, the Easy Tier function enables performance-based pool balancing for that volume copy.
4. When the volume copy status is `active`, the Easy Tier function operates in automatic data placement mode for that volume.

If the volume copy is in image or sequential mode or is being migrated, the volume copy Easy Tier status is `measured` instead of `active`.

The default Easy Tier setting for a storage pool is `auto`, and the default Easy Tier setting for a volume copy is `on`. These settings mean that Easy Tier functions except pool performance balancing are disabled for storage pools with a single tier, and that automatic data placement mode is enabled for all striped volume copies in a storage pool with two or more tiers.

Chapter 4. Volumes

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

Note: Single volume type systems are designed to support single copy volumes, with mirrored and stretched configurations available in multi-volume environments.

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.

If you delete a volume that has a snapshot assigned, the state of the volume changes to **Deleting** and the name of the volume is prefixed with *del_* and suffixed with a timestamp. The actual deletion of the volume occurs only after all snapshots associated to the volume are either deleted or expired.

Note: You cannot perform any action on a volume that is in the **Deleting** state.

Volumes can be assigned to an ownership group. An ownership group defines a subset of users and objects within 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.

© Copyright IBM Corp. 2025

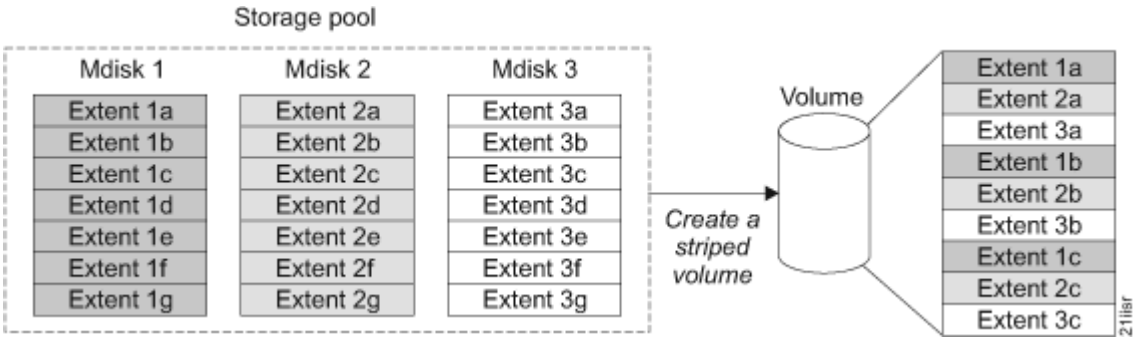


Figure 8. Storage pools and volumes

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.

Image mode MDisks are members of a storage pool, but they do not contribute to free extents. Image mode volumes are not affected by the state of the storage pool because the storage pool controls image mode volumes through the association of the volume to an MDisk. Therefore, if an MDisk that is associated with an image mode volume is online and the storage pool of which they are members goes offline, the image mode volume remains online. Conversely, the state of a storage pool is not affected by the state of the image mode volumes in the storage pool.

Image mode volumes can be used to migrate existing data from an external storage system. However, they cannot be used to export the replicated volumes to an external storage system. To configure replication on image mode volumes, add volumes to a volume group and then assign a replication policy. Image mode volumes cannot be added to volume groups in a storage partition. A volume group that contains volume with a cache mode other than **`readwrite`**, the replication direction cannot be changed.

To maintain replication on an external storage system until a new recovery copy is established, perform the following:

- Create an image mode volume with the cache mode set to **none**.
- Once the new recovery copy is established, change the cache mode to **`readwrite`**.

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. 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 23. 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. 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. **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: • **expandvdisksize** • **migratevdisk** • **rmvdiskcopy** • **rmvolumecopy** • **shrinkvdisksize** • **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 24. 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 that is 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 the 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.

Adding a copy to a basic volume

The **addvdiskcopy** command adds a copy to an existing basic volume, which changes a nonmirrored volume into a mirrored volume. In the management GUI, select **Volumes** and right-click the basic volume and select **Add Volume Copies**.

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.

You can enable the storage pool capacity allocated to the volume to automatically expand as the volume is written to. The provisioned capacity does not change. When disabled, the available buffer capacity decreases as the used capacity increases. The volume goes offline if the buffer capacity is fully used. You must use the CLI or REST API to expand the buffer capacity. To modify the auto-expand setting for a thin-provisioned volume, go to the **Volumes** view, select the volume and modify the volume properties.

Converting thin-provisioned volumes

You can convert thin-provisioned volumes into standard-provisioned volumes.

You can non-disruptively convert a thin-provisioned volume into a standard-provisioned volume by using the following volume mirroring procedure:

1. Select a single-copy, thin-provisioned volume.
2. Add a standard-provisioned copy to the volume. You can choose to have the system automatically remove the thin-provisioned copy after the standard-provisioned volume is synchronized or you can manually remove it.

3. Wait while the volume mirroring feature synchronizes.
4. Remove the thin-provisioned copy from the standard-provisioned volume.

Converting standard-provisioned volumes

You can convert standard-provisioned volumes to thin-provisioned volumes.

You can non-disruptively convert a standard-provisioned volume into a thin-provisioned volume by following this procedure:

1. Select a single copy, standard-provisioned volume.
2. Add a thin-provisioned copy to the volume.

   Use a small real capacity and the autoexpand feature. You can also choose to have the system automatically remove the standard-provisioned copy after the thin-provisioned volume is synchronized or you can manually remove it.
3. Wait while the volume mirroring feature synchronizes the copies.
4. Remove the standard-provisioned copy from the thin-provisioned volume.

Any grains of the standard-provisioned volume that contain all zeros do not cause any real capacity to be allocated on the thin-provisioned copy. Before you create the mirrored copy, you can fill the available capacity on the volume with a file that contains all zeros.

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**.

Note: Compressed volumes are not supported on systems that do not support data reduction pools.

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.

Deduplication is a type of data reduction that eliminates duplicate copies of data. On storage systems that support data reduction pool, deduplicated volumes must be created in data reduction pools for added capacity savings. 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**

To change the volume protection setting and disable volume protection, complete the following steps:

1. To change the system-wide value and disable volume protection on the system, see **chsystem** command.
2. To change the pool-wide value and disable volume protection on the pools, see **chmdiskgrp** command.

Expand or shrink volumes

You can use the management GUI or the command-line interface (CLI) to expand or shrink the capacity of a volume.

Expand volumes

The system supports expanding the size of a volume concurrent with I/O operations, when supported by the host operating system.

To expand a volume by using the CLI, see the **chvolume** command.

Note: If the volume being expanded is configured for high availability, and HA is currently *suspended*, then some host operating systems might fail to detect the resized device. In this case, remove the host paths to the offline copy of the volume, or wait for HA to be re-established before continuing.

Shrink volumes

Volumes can be reduced in capacity, if it is necessary. However, if the volume contains data, do not shrink the size of the disk. The system disables shrinking a volume if the selected volume is performing quick initialization. After the quick initialization completes, you can shrink the volume.



Attention:

1. It is difficult to anticipate how an operating system or file system uses the capacity in a volume. When you shrink a volume, capacity is removed from the end of the disk, whether or not that capacity is in use. Even if a volume has available capacity, do not assume that only unused capacity is removed when you shrink a volume.
2. If the volume contains data that is being used, *do not attempt under any circumstances to shrink a volume without first backing up your data*.
3. For performance reasons, some operating systems or file systems use the outer edge of the disk.

You can use the **shrinkvdisksize** command or the management GUI to shrink the usable capacity that is provisioned to the particular volume by the specified amount. You can also shrink the provisioned capacity of a thin-provisioned volume without altering the usable capacity that is assigned to the volume.

You cannot shrink the following:

- Volumes in a data reduction pool.
- A target volume in FlashCopy mappings. However, you can shrink a source volume, if the starting capacity of the FlashCopy mapping is not greater than the proposed capacity for the source volume. The management GUI filters target volumes that can be shrunk and disables the shrink action for those volumes that do not meet the criteria. If you are using the command-line interface, use the **lsfcmap**

command to display the starting capacity (*start_capacity* parameter) of the FlashCopy mapping. This parameter indicates the capacity that was used the last time the FlashCopy operation was completed.

- The capacity of replicated volumes.

To shrink a volume by using the CLI, see the **shrinkvdisksize** command.

Dependent volumes

You can use the command-line interface (CLI) to list the volumes that depend on the status of a node.

Before you begin

If a node goes offline or is removed from a system, all volumes that depend on the node go offline. Before you take a node offline or remove a node from a system, run the **lsdependentvdisks** command to identify any node-dependent volumes.

About this task

By default, the **lsdependentvdisks** command also checks all available quorum disks. If the quorum disks are accessible only through the specified node, the command returns an error.

Various scenarios can produce node-dependent volumes. The following examples are common scenarios in which the **lsnodedependentvdisks** command returns node-dependent volumes:

1. The node contains flash drives with the only synchronized copy of a mirrored volume.
2. The node is the only node that can access an MDisk on the SAN fabric.
3. The other node in the I/O group is offline (all volumes in the I/O group are returned).
4. Pinned data in the cache is stopping the partner node from joining the I/O group.

To resolve (1), allow volume mirror synchronizations between Flash drive MDisks to complete. To resolve (2-4), bring any offline MDisks online and repair any degraded paths.

Note: The command lists the node-dependent volumes at the time the command is run; subsequent changes to a system require running the command again.

Procedure

1. Issue the **lsdependentvdisks** CLI command.

   The following example shows the CLI format for listing the volumes that depend on node01:

   ```
   lsdependentvdisks -enclosure -delim : 0:1
   ```

   The following example shows the output that is displayed:

   ```
   vdisk_id:vdisk_name
   4:vdisk4
   5:vdisk5
   ```

2. If the command **lsdependentvdisks** returns an error, you must move your quorum disks to MDisks that are accessible through all nodes.

   Rerun the command until no errors are returned.

3. Reissue the **lsdependentvdisks** command.

   When the command returns no volumes, the system is free from any node-dependent volumes.

   The following example shows the command syntax for listing the volumes that depend on node01:

   ```
   lsdependentvdisks -delim : -node node01 :
   ```

If there are no node-dependent volumes in the system, the following example shows the command output:

```
vdisk_id       vdisk_name
```

Managing throttles for volumes

You can create, modify, and remove throttles for volumes by using the management GUI or the command-line interface.

Throttling is a mechanism to control the amount of resources that are used when the system is processing I/Os on a specific volume or volume group. If a throttle is defined, the system either processes the I/O, or delays the processing of the I/O to free resources for more critical I/O.

Managing throttles with the command-line interface

To create a new throttle or edit an existing throttle for a volume with the command-line interface, enter the following command:

```
mkthrottle -type vdisk -bandwidth bandwidth_limit_in_mb -iops iops_limit -vdisk
vdisk_id_or_vdisk_name
```

where *bandwidth_limit_in_mb* is the maximum amount of bandwidth that the system can process for the specified volume before the system delays processing of I/O for that volume. The *iops_limit* is the maximum I/O operations per second that the system can process before the system delays the processing for this volume. The *vdisk_id_or_vdisk_name* is the ID or name of the volume that the throttle applies to.

To change an existing throttle for the volume, enter the following command:

```
chthrottle -type vdisk -bandwidth bandwidth_limit_in_mb -iops iops_limit throttle_name_or_id
```

where *bandwidth_limit_in_mb* is the maximum amount of bandwidth that the system can process for the specified host before the system delays processing of I/O for that volume. The *iops_limit* is the maximum I/O operations per second that the system can process before the system delays the processing for this volume. The *throttle_id_or_name* is the ID or name of the throttle that is being changed.

To remove a throttle for volume with the command-line interface, enter the following command:

```
rmthrottle throttle_name_or_id
```

where *throttle_name_or_id* is the ID or name of the throttle that is being removed.

Volume groups

A volume group allows for many related volumes to be grouped so that they can be managed together. A volume group also can ensure consistency across all volumes in the group when used with snapshot or replication functions.

Volume groups can be used with the following functions:

Safeguarded Copy

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 are backed up collectively with the Safeguarded Copy function. Safeguarded snapshots can be configured through an internal scheduler that is defined in a snapshot policy or can be configured with an external snapshot scheduling application such as IBM Copy Services Manager.

Snapshot

Snapshots are point-in-time copies of a volume group that cannot be directly accessed by hosts. To access the snapshot contents, you can create a clone or thin clone. You can schedule snapshots using a snapshot policy or an external snapshot scheduling application. For more details, see .

High availability and disaster recovery

Policy-based high availability provides an active-active high availability solution that uses two independent systems in different locations. The storage will automatically remain accessible to hosts if there is an event that makes one of the systems unavailable. When using policy-based high availability, everything in a storage partition is highly available; all volume groups in a partition are replicated for high availability.

Policy-based replication for disaster recovery provides synchronous or asynchronous replication to another system for disaster recovery. Using this type of replication, a replication policy is assigned to individual volume groups to manage groups that are replicated.

You can move volumes from one volume group to another. If a volume group is within a storage partition, the new volume group must be in the same storage partition.

Transparent cloud tiering

Transparent cloud tiering enables volume data to be copied and transferred to cloud storage.

The system supports creating connections to cloud service providers to store copies of volume data in private or public cloud storage. With transparent cloud tiering, administrators can move older data to cloud storage to free up capacity on the system. Point-in-time snapshots of data can be created on the system and then copied and stored on the cloud storage. An external cloud service provider manages the cloud storage, which reduces storage costs for the system. Before data can be copied to cloud storage, a connection to the cloud service provider must be created from the system.

A cloud account is an object on the system that represents a connection to a cloud service provider by using a particular set of credentials. These credentials differ depending on the type of cloud service provider that is being specified. Most cloud service providers require the host name of the cloud service provider and an associated password, and some cloud service providers also require certificates to authenticate users of the cloud storage. Public clouds use certificates that are signed by well-known certificate authorities. Private cloud service providers can use either a self-signed certificate or a certificate that is signed by a trusted certificate authority. These credentials are defined on the cloud service provider and passed to the system through the administrators of the cloud service provider.

Once the system is authenticated, it can then access cloud storage to either copy data to the cloud storage or restore data that is copied to cloud storage back to the system. The system supports one cloud account to a single cloud service provider. Migration between providers is not supported.

Each cloud service provider divides cloud storage into segments for each client that uses the cloud storage. These objects store only data specific to that client. The names of the objects begin with a prefix that you can specify when you create the account for the system. A prefix defines system-specific content that the object stores and supports multiple independent systems to store data to a single cloud account. Each cloud service provider uses different terminology for these storage objects.

The following requirements apply to transparent cloud tiering:

1. When a cloud account is created, it must continue to use the same encryption type, throughout the life of the data in that cloud account. Even if the cloud account object is removed and remade on the system, the encryption type for that cloud account cannot be changed while back up data for that system exists in the cloud provider.
2. Performing rekey operations on a system with an encryption enabled cloud account, perform the commit operation immediately after the prepare operation. Remember to retain the previous system master key (on USB or in Key server) as this key can still be needed to retrieve your cloud backup data when performing a T4 recovery or an import.
3. Avoid the use of the **`Restore_uid`** option when backup is imported to a new system.

Supported cloud service providers

The system supports connections to various cloud service providers. Some cloud service providers require connections over external networks, and others can be created on a private network.

Each cloud service provider requires different configuration options. The system supports the following cloud service providers:

OpenStack Swift
OpenStack Swift is a standard cloud computing architecture from which administrators can manage storage and networking resources in a single private cloud environment. Standard application programmable interfaces (APIs) can be used to build customizable solutions for a private cloud solution.

Microsoft Azure
Microsoft Azure is an object storage solution that allows large volumes of unstructured data to be stored in the Azure public cloud. It provides an API and web-based interface to help you manage your cloud solution. Microsoft Azure is suitable for archiving, disaster recovery, and backup data storage.

Amazon S3
Amazon Simple Storage Service (Amazon S3) provides programmers and storage administrators with flexible and secure public cloud storage. Amazon S3 is also based on object storage standards and provides a web-based interface to manage, back up, and restore data over the web.

The following Amazon S3 (AWS) regions are supported by this code-level:

- us-east-2
- us-east-1
- us-west-1
- us-west-2
- af-south-1
- ap-east-1
- ap-south-2
- ap-southeast-3
- ap-southeast-4
- ap-south-1
- ap-northeast-3
- ap-northeast-2
- ap-southeast-1
- ap-southeast-2
- ap-northeast-1
- ca-central-1
- eu-central-1
- eu-west-1
- eu-west-2
- eu-south-1
- eu-west-3
- eu-north-1
- eu-south-2
- eu-central-2
- sa-east-1
- me-south-1

- me-central-1
- il-central-1

Security considerations for cloud accounts

Whenever the system accesses outside networks, the potential for unintentional or intentional exposure of sensitive data is a risk. When you are connecting the system to a cloud service provider over a public network, you can use encryption to protect data that is transferred to the cloud service provider.

The first level of encryption-based security provides secure communications between the system and the cloud service provider. The standard protocol, Transport Layer Security (TLS), protects these connections by encrypting data that is transferred between the system and the cloud service provider. Secure communications are mandatory for these connections and requires that public certificates are exchanged between the cloud service provider and the system. To configure certificates for secure communications in the management GUI, go to **Settings** > **Security** > **System Certificates**. You can also use the **chsystemcert** command to create system certificates. With secure communications, data is encrypted while it is transferred to the cloud, but might be stored on the cloud decrypted. Each cloud service provider has its own security measures to protect data once it is located in cloud storage; however, breaches can still occur and data can be compromised. Clients that use cloud service providers can add extra encryption methods to protect their data after it is stored on the cloud.

Since the system supports encryption of at-rest data, you can optionally configure encryption key management to further protect data that is stored on the cloud storage. If key management is configured on the system, data is encrypted before it leaves the system and is stored on the cloud. The system supports key management through either a USB flash drive or an encryption key server. When encryption is configured, a master encryption key is created and is stored separately on either a USB flash drive or key server. When you create snapshots of data to send to the configured cloud service provider, each volume and each cloud account have separate encryption keys. The encryption key that is used by the cloud account protects encryption keys for the volumes. The master encryption key protects the encryption key that is used by the cloud account. Because the master encryption key is physically present on the USB flash drive or key server, you must ensure that security measures are implemented to protect the master encryption key from theft or loss. When the data is transmitted between the system and the cloud service provider, the data is also encrypted by certificates that are configured for secure communications. The master encryption key also protects the data in transit and the data remains encrypted while it is stored on the cloud storage. Data also remains encrypted with the encryption master key when it transferred back to the system from the cloud during restore operations. Finally, data can be decrypted when it arrives at the system or it can be stored on an encrypted volume on the system.

When a connection to a cloud service provider is configured, you must decide whether to encrypt data at rest in the cloud for this account. After you decide, the encryption setting for the account cannot be changed without restoring all data from the cloud, reconfiguring the account, and re-creating cloud snapshots for the data.

Encryption protects the data, but also requires careful configuration and management of keys. Depending on how key management is configured on the system, access to the keys is required when data is copied to or restored from the cloud service provider.

Configuring transparent cloud tiering

With transparent cloud tiering, you can configure a cloud account on the system to be able to create and restore cloud snapshots of system volumes.

Creating a DNS server

Before you create a connection to a cloud service provider or connect to cloud storage, ensure that you specify at least one DNS server to manage hostnames. You can have up to two DNS servers that are configured on the system. To configure DNS for the system, enter a valid IP address and name for each server. Both IPv4 and IPv6 address formats are supported.

To create a DNS server for host name resolution, use the **Settings** > **Network** > **DNS** panel in the management GUI.

To create a DNS server using the command-line interface, see **mkdnsserver** command.

Enabling a cloud connection

To configure transparent cloud tiering on the system, you must enable a cloud connection to a supported cloud service provider.

A cloud account is an object on the system that represents a connection to a cloud service provider by using a particular set of credentials. These credentials differ depending on the type of cloud service provider that is being specified. Most cloud service providers require the host name of the cloud service provider and an associated password, and some cloud service providers also require certificates to authenticate users of the cloud storage. Public clouds use certificates that are signed by well-known certificate authorities. Private cloud service providers can use either a self-signed certificate or a certificate that is signed by a trusted certificate authority. These credentials are defined on the cloud service provider and passed to the system through the administrators of the cloud service provider.

Enabling a cloud connection to OpenStack Swift

To enable a cloud connection between the system and OpenStack Swift, use the **Settings** > **System** > **Transparent Cloud Tiering** panel in the management GUI. Expand the **OpenStack account** and enter the following information:

Tenant
Enter the name of the OpenStack tenant that the system uses to access the cloud storage.

User name
Enter the user name that is associated with the account. The user name identifies clients that can access content that is stored by the cloud service provider.

Password
Enter the password that is associated with the user name for the cloud account.

Endpoint
Enter the URL that the system uses to access the cloud service provider.

Container prefix
Enter a short string that is added to the beginning of the container name to identify which containers are associated with the system. Prefixes can also be used to identify multiple systems that use the same cloud account to store snapshots of volume data.

SSL Certificate
Select the SSL certificate that is used to authenticate the system to the cloud service provider.

Keystone
Keystone authentication system is required for all OpenStack Swift cloud accounts. The management GUI automatically detects whether Keystone is enabled. If it is not enabled, a warning message displays. Click **Enable** to configure Keystone authentication for the account.

Encryption
If encryption is configured on the system, this value is preselected. If encryption is not enabled on the system, this option is unavailable.

Bandwidth
Determine how bandwidth is used for downloads to and uploads from the cloud service provider. Specify **No Limit** to allow the connection to use unlimited bandwidth, regardless of other operations that occur on the system. If you do not want to give unlimited bandwidth to these operations, you can specify the maximum amount of bandwidth that the connection can use for either downloads or uploads.

To enable a cloud connection between the system and OpenStack Swift interfaces using the command-line interface, see **mkcloudaccountswift** command.

Enabling a cloud connection to Amazon S3

To enable a cloud connection between the system and Amazon S3 , use the **Settings** > **System** > **Transparent Cloud Tiering** panel in the management GUI. Expand **Amazon S3 account** and enter the following information:

Access key ID
The public part of the access key credential for the Amazon user that the system uses to access the cloud storage.

Secret access key

Enter the private part of the Amazon credential that the system uses to authenticate the system to cloud storage.

Region
Select the Amazon region that is used to both access and store data on cloud storage.

Bucket prefix
Enter the prefix for the Amazon S3 bucket names that the system will use or create. Prefixes can also be used to identify multiple systems that use the same cloud account to store snapshots of volume data. If you want to import or export snapshots between systems, use the same bucket prefix as the other system.

Encryption
If encryption is configured on the system, this value is preselected. If encryption is not enabled on the system, this option is unavailable.

Bandwidth
Determine how bandwidth is used for downloads to and uploads from the cloud service provider. Specify **No Limit** to allow the connection to use unlimited bandwidth, regardless of other operations that occur on the system. If you do not want to give unlimited bandwidth to these operations, you can specify the maximum amount of bandwidth that the connection can use for either downloads or uploads.

To enable a cloud connection between the system and Amazon S3 using the command-line interface, see **`mkcloudaccountawss3`** command.

Enabling a cloud connection to Microsoft Azure

To enable a cloud connection between the system and Microsoft Azure, use the **Settings** > **System** > **Transparent Cloud Tiering** panel in the management GUI. Expand **Cloud Provider Settings** page for Microsoft Azure, enter the following information:

Storage Account Name
Enter the name of the Microsoft Azure storage account that the system uses to authenticate to cloud storage.

Storage Account Key
Enter the account access key that the system uses to access the cloud storage.

Container Prefix
Enter the container prefix to identify the objects that the system creates and uses on the cloud storage.

Encryption
If encryption is configured on the system, this value is preselected. If encryption is not enabled on the system, this option is disabled by default.

Bandwidth
Determine how much bandwidth is used by download or upload operations for the cloud service provider. Specify **No Limit** to allow the connection to use unlimited bandwidth, regardless of other operations that occur on the system. If you do not want to give unlimited bandwidth to these operations, you can specify the maximum amount of bandwidth that the connection can use for either downloads or uploads. The maximum bandwidth limit for downloads and uploads is 10240 Mbps.

To enable a cloud connection between the system and Microsoft Azure using the command-line interface, see **mkcloudaccountazure** command.

Managing cloud accounts

A cloud account is an object on the system that represents a connection to a cloud service provider by using a particular set of credentials. These credentials differ depending on the type of cloud service provider that is being specified. Most cloud service providers require the host name of the cloud service provider and an associated password, and some cloud service providers also require certificates to authenticate users of the cloud storage. Public clouds use certificates that are signed by well-known certificate authorities. Private cloud service providers can use either a self-signed certificate or a certificate that is signed by a trusted certificate authority. These credentials are defined on the cloud service provider and passed to the system through the administrators of the cloud service provider. You can manage cloud accounts with the management GUI or the command-line interface.

Verify cloud account

To verify the cloud account state is online, use the **Settings** > **System** > **Transparent Cloud Tiering** panel in the management GUI. If the cloud account is not online, the connection to the cloud service provider is not available. Use the **Monitoring** > **Events** panel to investigate the cause of the problem. You might need to access the cloud service provider to determine the connection problem.

Changing settings on cloud account

Certain settings are controlled by the cloud service provider and cannot be changed. However, you can update credential information for a cloud account if necessary.

When settings are changed on the cloud account, any snapshot versions of volumes that are stored on the cloud becomes unavailable. Before you change any settings that are related to the cloud account, ensure that you restore data from the cloud storage or delete the cloud snapshots from cloud storage. After the changes are completed to the cloud account, new cloud snapshots can be created of volume data. Some situations might require that you update certain settings on the cloud account. For example, if a cloud service provider updates credential information, then changes to the cloud account information on the system are required. For some settings, the system restricts any updates.

OpenStack Swift account

For a cloud account with OpenStack Swift, the following settings cannot be changed:

Tenant
A tenant is a container that is used to store objects that are created for a particular account. Tenant names are defined by the cloud service provider and represent the underlying storage that is used by the account. Tenant names must remain consistent or data that is stored on the cloud cannot be accessed.

Container prefix
Container prefixes are added to the beginning of the tenant name to identify which tenants are associated with the system. Prefixes can also be used to identify multiple systems that use the same cloud account to store snapshots of volume data. Like with tenant names, these prefixes must remain consistent to access data on the cloud.

Encryption
After encryption is set for a cloud account, it cannot be changed.

Use the **Settings** > **System** > **Transparent Cloud Tiering** panel in management GUI to change settings for an OpenStack Swift account. Expand **Account information** to view the settings that are related to the account. The following settings can be updated:

User name
Enter the new user name that is associated with the account. The user name is used to identify clients that can access content that is stored by the cloud service provider.

Password
Enter the new password that gives access to the cloud storage for the specified user name. To update the password, click **Edit**.

Endpoint
Enter the new URL that the system uses to access cloud storage on the cloud service provider.

SSL Certificate
Select the new SSL certificate that is used to authenticate the system to the cloud service provider.

Keystone
Keystone authentication system is required for all OpenStack Swift cloud accounts. The management GUI automatically detects whether Keystone is enabled. If it is not enabled, a warning message displays. Click **Enable** to configure Keystone authentication for the account.

To change settings using the command-line interface for the OpenStack Swift storage account, see the **chcloudaccountswift** command.

Amazon S3 account

For a cloud account with Amazon S3, the following settings cannot be changed:

Region
Enter the user name that is associated with the account. The user name is used to identify clients that can access content that is stored by the cloud service provider.

Bucket Prefix
Enter short string that is added to the name of the bucket to identify which storage objects are associated with the system. Prefixes can also be used to identify multiple systems that use the same cloud account to store snapshots of volume data. A bucket is a container that Amazon S3 creates to hold objects that are associated with a particular client.

Encryption
If encryption is configured on the system, this value is preselected. If encryption is not enabled on the system, this option is unavailable.

Use the **Settings** > **System** > **Transparent Cloud Tiering** panel in management GUI to change settings for an Amazon S3 cloud account. Expand **Account information** to view the settings that are related to the account. The following settings can be updated:

Access key ID
Enter the new public part of the credential that the system uses to authenticate to cloud services. The access key ID is essentially a user name.

Secret access key
Enter the new private part of the credential that the system uses to authenticate the system to cloud services. To update the secret access key, click **Edit**.

To change settings using the command-line interface for the Amazon S3 storage account, see the **chcloudaccountawss3** command.

Microsoft Azure account

For a cloud account with Microsoft Azure, the following settings cannot be changed:

Container Prefix
Prefixes are added to the beginning of the container names. Container prefixes help to identify the objects that the system creates and uses on the cloud storage.

Encryption
If encryption is configured on the system, this value is preselected. If encryption is not enabled on the system, this option is not available.

Use the **Settings** > **System** > **Transparent Cloud Tiering** panel in management GUI to change settings for the Microsoft Azure storage account. Expand **Account information** to view the settings that are related to the account.

For Storage account key, you can add an access key or update the existing access key that the system uses to authenticate the cloud account.

To change settings using the command-line interface for the Microsoft Azure storage account, see the **chcloudaccountazure** command.

Testing a cloud account

The command-line interface supports diagnostics and status reports for network connectivity, authentication, and the ability to use cloud storage. You can use this information for troubleshooting and monitoring the status of your cloud account.

To test a cloud account, enter the command **testcloudaccount cloud_account_id_or_name**. For more information, see **testcloudaccount**.

Removing a cloud account

If a cloud account is no longer needed, it can be removed from the system. You can use the management GUI or the command-line interface to delete a cloud account.

If cloud snapshots are not currently in the cloud storage that is associated with the account, the cloud account settings and any containers and metadata are deleted from the cloud service provider. If cloud snapshots do exist on the cloud, only the cloud account information is deleted, but the cloud snapshots are retained on the cloud service provider. If at some point the data on the cloud is needed, you can re-create a new cloud account that uses the same settings and you can access the data that is stored on the cloud service provider. For the remove operation to succeed, no active cloud snapshots or restore operations can be in progress.

Use the **Settings** > **System** > **Transparent Cloud Tiering** panel in management GUI to remove the cloud account. Expand **Account information** to view the settings that are related to the account.

To remove a cloud account using the command-line interface, see the **rmcloudaccount** command.

Cloud volumes and snapshots

A cloud volume is any volume that is enabled for transparent cloud tiering. After transparent cloud tiering is enabled on a volume, point-in-time copies, or snapshots, can be created and copied to cloud storage that is provided by a cloud service provider. These snapshots can be restored to the system for disaster recovery purposes. Before you create cloud volumes, a valid connection to a supported cloud service provider must be configured.

With transparent cloud tiering, the system supports connections to cloud service providers and the creation of cloud snapshots of any volume or volume group on the system. Cloud snapshots are point-in-time copies of volumes that are created and transferred to cloud storage that is managed by a cloud service provider. A cloud account defines the connection between the system and a supported cloud service provider and must be configured before data can be transferred to or restored from the cloud storage. After a cloud account is configured with the cloud service provider, you determine which volumes you want to create cloud snapshots of and enable transparent cloud tiering on those volumes.

In the management GUI, volumes that do not have restrictions are enabled for cloud snapshots when a snapshot is created for that volume. If you are using the command-line interface, you must enable snapshots on the volume before you can create a snapshot of the volume.

A volume cannot be enabled for cloud snapshots if the volume is used in any of the following capacities:

- The volume is a VMware vSphere Virtual Volumes.
- A volume has a copy in a different storage pool.
- A volume is being migrated between different storage pools.
- A volume cannot be enabled for cloud snapshots if the cloud storage is set to import mode.
- A volume cannot be enabled for cloud snapshots if the maximum number of cloud volumes already exists. The maximum number of cloud volumes on the system is 1024.

Note: If the system exceeds this limit, you can disable cloud snapshots on an existing cloud volume and delete its associated snapshots from the cloud storage to accommodate snapshots on new cloud volumes.

- The volume is associated with user-owned legacy FlashCopy mappings.

After the volume is enabled for the function, cloud snapshots of those volumes can be created and stored on cloud storage. Multiple snapshots can be created at different times and each version is timestamped and stored in the cloud storage. Each snapshot version on the cloud storage represents the data at a specific point in time, so different versions can be restored if necessary. The system supports both snapshots of local system data and also data on different systems. Some of these cloud snapshots can have corresponding cloud volumes on the local system, but other snapshots might not. You can restore data from these different snapshots of the cloud volume if necessary.

The system supports two types of cloud snapshots:

Full snapshots
When the cloud snapshot is created for the first time for a volume, all data is copied to the cloud storage. Any subsequent cloud snapshots of the volume include only the changed data from the initial snapshot. Each snapshot version of the volume is timestamped and stored on the cloud storage. Therefore, full plus incremental snapshots ensure that all data is copied to the cloud storage and is available if a restore operation is necessary. If the volume contains a large amount of data, full snapshots can be time-consuming.

Incremental snapshots
The version contains changed data from the last time the cloud snapshot was created for a selected volume. Incremental-only snapshots are quicker to complete than a full snapshot.

Individual volumes can be added to a single volume group for creating snapshots for all the volumes in that group. If volumes have similar content or are used by a specific host or application, they can be added to volume group to simplify snapshot operations. Separate snapshots are started for each volume in the group and the group must be created before the snapshots are created. Volume groups can be configured by using the **mkvolumegroup** command. The management GUI lists all cloud volumes and specifies the name of the volume group in which the volumes belong.

If you want to limit access of users to specific volume groups that are used in creating cloud snapshots, you can assign these volume groups to an ownership group. An ownership group defines a subset of users and objects within the system.Volumes that are within the volume group must be within a child pool.

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. A user cannot change the ownership group of the resource, but can change the ownership group of the parent object. The following rules apply to volume groups that are defined in ownership groups:

- Volumes can be added to a volume group if both the volume and the volume group are within the same ownership group or if both are not in an ownership group. When migrating existing resources to ownership groups, the volume group and its volumes can belong to different ownership groups temporarily until the migration completes.
- The ownership of a volume group does not affect the ownership of the volumes it contains. If a volume group and its volumes are owned by different ownership groups, then the owner of the child pool that contains the volumes can change the volume directly. For example, the owner of the child pool can change the name of a volume within it. The owner of the volume group can change the volume group itself and indirectly change the volume, such as deleting a volume from the volume group. Neither the ownership group of the child pools nor the owner of the volume group can directly manipulate the resources that are not defined in their ownership group.

Individual cloud volumes and volume groups have different statuses to indicate whether the volume is available for snapshot operations. These statuses can be used to monitor in-progress snapshots and error states for the snapshot. Cloud volumes have the following snapshot statuses:

Off
Transparent cloud tiering is disabled. Snapshots are not enabled for the volume or volume group.

Ready
Transparent cloud tiering is enabled. A new snapshot can be started for the volume or volume group.

Copying
Transparent cloud tiering is enabled. A snapshot operation is in progress and data is being transferred to the cloud.

Copying error
Transparent cloud tiering is enabled. An existing snapshot operation is in-progress and cannot be completed because of an error. Causes for copying errors include losing the network connection to the cloud service provider or exceeding the amount of capacity on the cloud storage.

Not ready
Transparent cloud tiering is enabled. A new snapshot operation cannot be started because a restore is in progress for the volume.

In addition to creating snapshots of volume data to transfer to cloud storage, Transparent cloud tiering also supports restoring volume data from cloud storage back to the system. Like with snapshots, versions of each copy operation are timestamped and stored on the cloud, which allows for specific point-in-time recovery operations to the original system or to another system. Older versions can be deleted and removed from the system and the cloud storage simultaneously to manage space. Snapshots can be restored only for individual volumes. You can restore multiple volumes simultaneously, but each restore operation needs to be started separately for each volume.

Certain recovery scenarios can require the restoration of an older version of the cloud snapshot. However, all newer versions on the cloud storage are deleted, including any in-progress snapshot operations. The management GUI verifies the version of cloud snapshot and issues a warning if the selected snapshot is not the most recent version. To proceed with the restore operation, you must confirm the deletion of these other versions.

You restore volume data with any snapshot version that was copied to the cloud, including restoring from a snapshot version that is still in progress. The snapshot version and volume that the snapshot version is being restored to must be the same size. A version cannot be deleted while it is being used for a restore. You can restore a snapshot version with following methods:

Restore to the production volume
If the snapshot version is restored to the production volume, which is the original volume from which the snapshots were created. The snapshot version replaces the current data that exists on production volume with the data that is stored on the cloud storage. The production volume goes offline during the restore operations. Data is not fully restored to the production volume until the changes are committed.

Restore to a new volume
When the snapshot version is restored to a new volume, you can use the restored data independently of the original volume from which the snapshot was created. If the new volume exists on the system, then the restore operation uses the unique identifier (UID) of the original volume. If the new volume does not exist on the system, you need to choose whether to use the UID from the original volume or create a new UID. If you plan on using the new volume on the same system, use the UID that is associated with the snapshot version that is being restored. You want a unique UID if you are restoring a version of a volume that exists on another system to the current system.

Individual cloud volumes have different statuses to indicate whether the volume is available for restore operations. These statuses can be used to monitor in-progress restore operations. Cloud volumes have the following restore statuses:

None
No snapshot version for this volume available for restore.

Available
Indicates that the selected volume has available snapshot versions on cloud storage that can be restored.

Restoring
Indicates that the restore is in progress for this volume.

Restoring error
Indicates that the restore operation is in progress, but an error occurred. Causes for restoring errors include losing the network connection to the cloud service provider or exceeding the amount of capacity on the pools in which the volume belongs.

Managing cloud volumes

Cloud volumes are any volume on the system that is enabled for cloud snapshots. You can use the management GUI or the command-line interface (CLI) to manage the cloud volumes and cloud snapshots.

Enable cloud snapshots on volumes

In the management GUI, when you create a snapshot, cloud snapshots are automatically enabled for the volume when the volume is not restricted from cloud snapshots. If you are using the command-line interface, you must enable cloud snapshots for the volume separately.

Cloud snapshot restrictions
A volume cannot be enabled for cloud snapshots if the volume is used in any of the following capacities:

- The volume is a VMware vSphere Virtual Volume.
- A volume has a copy in a different storage pool.
- A volume is being migrated between different storage pools.
- A volume cannot be enabled for cloud snapshots if the cloud storage is set to import mode.
- A volume cannot be enabled for cloud snapshots if the maximum number of cloud volumes already exists.

Note: If the system exceeds this limit, you can disable cloud snapshots on an existing cloud volume and delete its associated snapshots from the cloud storage to accommodate snapshots on new cloud volumes.

- Volume is associated with user-owned legacy FlashCopy® mappings.

Note: Cloud snapshots are supported with the system's new FlashCopy management model based on snapshot function. For more information about snapshot function, see Snapshots.

To enable snapshots on a volume using the command-line interface, see **chvdisk** command.

Disable cloud snapshots for a volume

If you no longer need to create snapshots for a cloud volume, you can disable the function for the volume with the management GUI or the command-line interface.

When cloud snapshots are disabled on a volume, the volume is disconnected from the cloud account. Existing snapshot versions of the volume that are located in cloud storage are not deleted. A full snapshot of the volume is required if a cloud snapshot is needed later. You cannot disable cloud snapshots on a volume if a current snapshot operation is in progress. You must either wait for the snapshot operation to complete or cancel the snapshot. The cloud snapshots can be used only with volumes that have a capacity greater than zero bytes.

To disable cloud snapshots for a volume, use the **Volumes** > **Cloud Volumes** panel in the management GUI. Right -click the volume and select **Disable cloud snapshots**.

To disable snapshots on a volume using the command-line interface, see **chvdisk** command.

Create volume groups

For cloud volumes, you can create a group of volumes of which you can create snapshots and transfer to cloud storage. Volumes groups can be created to easily manage cloud snapshots for volumes. It is convenient to provide a single action to take a snapshot of a preset group of volumes.

Volume groups provide consistent snapshot operations for cloud volumes. You can view cloud volumes and display whether the volumes are part of a group in the management GUI, but you must configure the volume group in the command-line interface.

To create volume group, see the **mkvolumegroup** command, and to create or add cloud volumes to the volume group, see the **chvdisk** command.

Create cloud snapshots

After you create a cloud account, you can specify an individual volume or groups of volumes to be copied and transferred to cloud storage by creating cloud snapshots.

Cloud snapshots are point-in-time copies of volumes that are created and transferred to cloud storage that is managed by a cloud service provider. You can select individual or multiple volumes to create cloud snapshots of the volume data. You can use the command-line interface to configure volume groups to create snapshots for multiple volumes for consistent management.

When a snapshot is created for multiple volumes or volumes in a group, a separate snapshot operation is started for each individual volume.

In the management GUI, volumes that do not have restrictions are enabled for cloud snapshots when a snapshot is created for that volume. If you are using the command-line interface, you must enable snapshots on the volume before you can create a snapshot of the volume.

To create a cloud volume, use the **Volumes** > **Cloud Volumes** panel in the management GUI and click **Add Volumes**. This enables cloud snapshots on the added volumes before you can create cloud snapshot of the volumes. User can choose the following options:

- **Let the system decide the type of snapshot**
- **Complete a full snapshot**

To enable snapshots on a volume using the command-line interface, see the **chvdisk** command and to create cloud snapshots on a volume, see the **backupvolume** command. If you want to complete a full snapshot where all data on the cloud volume is copied to the cloud storage use **backupvolume -full volumen_id_or_name**.

Restore cloud snapshots

With transparent cloud tiering enabled, snapshot versions of volumes are stored on the cloud storage that is provided by a supported cloud service provider. You can restore data from the cloud back to the system from any snapshot version to recover from data loss on the system.

A volume cannot be used for restoring data from the cloud if the volume is a VMware vSphere Virtual Volume.

To restore a snapshot generation from cloud storage to the system, use the **Volumes** > **Cloud Volumes** panel in the management GUI. Right-click the volume to display all snapshot versions that are available to restore and select **Restore** and snapshot version that you want to restore to the system.

Select the target volume from the following options:

Restore a production volume group with the contents of one of the volume groups snapshot
The contents of the production system is overwritten with a previously saved version of the data.

Restore to the production volume
The snapshot version is restored to the production volume, which is the original volume from which the snapshots were created. After the restore operation completes, the snapshot version completely replaces the current data that exists on production volume. During the restore operation, the production volume goes offline until it completes. If you do not want to have the production volume offline for the restore, you can restore a cloud snapshot to a new volume. The production volume remains online and host operations are not disrupted.

Restore to a new volume
When the snapshot version is restored to a new volume, you can use the restored data independent of the original volume from which the snapshot was created. If the new volume exists on the system, then the restore operation uses the unique identifier (UID) of the new volume. If the new volume does not exist on the system, you need to choose whether to use the UID from the original volume or create a new UID. If you plan on using the new volume on the same system, use the UID that is associated with the snapshot version that is being restored.

You can use the command-line interface to either restore a snapshot directly to the production volume or to a new volume. To determine the available snapshot versions, see the **lsvolumebackupgeneration** command.

Restore to a production volume
To restore the most recent or old version of the cloud snapshot to the production volume use **restorevolume**. To restore an older version of the cloud snapshot, you must specify the **-deletelatergenerations** parameter to confirm the deletion of all later versions of the snapshot.

Restore to a new volume
You can also restore a snapshot version for a volume that might not exist on the local system. In this case, the unique identifier of the volume snapshot data is used. You can also restore the entire snapshot by replacing the volume UID with the UID of the volume that was originally used for the snapshot. Ensure that the UID is not currently used on any other volume in the system. To restore a snapshot version to the selected volume use the **restorevolume** command. You must specify **-restoreuid** to retain the UID for that volume, otherwise it restores data on new volume without replacing volume UID of original volume.

Delete a version of a cloud snapshot

A cloud volume can have multiple versions of snapshots on the cloud storage. If capacity on the cloud becomes limited, you can delete versions of the volume that are no longer necessary.

Before you delete a version of a snapshot, ensure that no snapshot or restore operations are in progress for the selected volume. In addition, ensure that the connection to the cloud service provider is online, so the deletion can be validated and completed on the cloud storage. Wait until these operations complete before you delete any snapshots that are associated with the volume. Deleting a backup generation (as opposed to canceling an in-progress backup) requires the connection to the cloud account to be available when the task is submitted in order to validate user input before you start an asynchronous delete operation.

Before you delete a version of a cloud snapshot, ensure that the connection to the cloud service provider is online. To delete a version of a cloud snapshot from the cloud storage, use the **Volumes** > **Cloud Volumes** panel in the management GUI. Right-click the cloud volumes and select **Manage Cloud Snapshots** and select the version of the snapshot.

To delete a single or all snapshot version for a volume using the command-line interface, see the **rmvolumebackupgeneration** command.

Import cloud snapshots

You can restore a cloud snapshot for a volume that was backed up from another system to the same container by importing a snapshot version of that volume to the local system.

When you import a cloud snapshot of a volume that is not on the current system, you can choose to use the unique identifier of the snapshot to replace the existing UID on the volume. In this case, the UID of the snapshot replaces the existing UID for the volume that is selected to receive the restored data. Ensure that the new UID is not currently used on any other volume in the system. The management GUI does not support UID replacement when a snapshot is imported from another system.

Note: The cloud account on the target system must use the same login credentials and the same container prefix as the cloud account that was used to back up the volume initially.

To import a cloud snapshot from another system,

1. Determine the import system ID for the system that the cloud snapshot is being imported from by running CLI **lscloudaccountimportcandidate** on the target system.
2. To change the mode of the cloud account, you use different commands based on the type of cloud account that is configured, use CLI **chcloudaccountswift**, **chcloudaccountawss3**, **chcloudaccountazure**. You must specify -mode import and import system ID.
3. Use the **restorevolume** command to restore the cloud snapshot for the volume that was backed up.
4. To confirm that the cloud account is in import mode, use CLI **lscloudaccount**. The target system can now access the data that is in the cloud snapshot of the source system.
5. To return the cloud account to normal mode, you use different commands based on the type of cloud account. You must specify -mode normal.

Transparent cloud tiering restrictions

Transparent cloud tiering on the system is defined by configuration limitations and rules.

The system restrictions specific to transparent cloud tiering for this release follow.

- Cloud volumes cannot be expanded or shrunk.
- A volume cannot be configured for a cloud snapshot if any of the following conditions are valid:
  - The volume is part of a volume group that is replicated, or configured for high availability.
  - The volume is a VMware vSphere Virtual Volumes volume, including FlashCopy owned volumes that are used internally for Virtual Volumes restoration functions.
  - The volume is a mirrored volume with copies in different storage pools.
  - The volume is being migrated between storage pools.
  - The volume is associated with any user-owned legacy FlashCopy mappings.

    Note: Cloud snapshots are supported with the system's new FlashCopy management model based on snapshot function. For more information about snapshot function, see related information section.
- A volume cannot be enabled for cloud snapshots if the cloud storage is set to import mode.
- A volume cannot be enabled for cloud snapshots if the maximum number of cloud volumes exists. If the system exceeds this limit, you can disable cloud snapshots on an existing cloud volume and delete its associated snapshots from the cloud storage to accommodate snapshots on new cloud volumes.
- A volume cannot be used for a restore operation if the following conditions are valid:
  - A Virtual Volume, including FlashCopy volumes that are used internally for Virtual Volumes restoration functions.
- A volume that is configured for backup or is being used for restoration cannot be moved between I/O groups.
- Only one operation (cloud snapshot, restore, or snapshot deletion) is allowed at a time on a cloud volume.
- Cloud volume traffic is allowed only through management interfaces (1 Gbps or 10 Gbps).

Snapshots

A volume group snapshot is a consistent point-in-time copy of a volume group.

A volume group snapshot can be used to restore the parent volume group or a subset of the parent volumes of the snapshot, or to create clones of the snapshot. For details, see “Restore or refresh from a snapshot” on page 103.

A volume group snapshot can be safeguarded, which adds protection to the snapshot. For details, see Safeguarded snapshots.

Snapshots policies manage creation and retention of snapshots of volumes in a volume group at specific intervals. For details, see “Snapshot policies” on page 186.

Highly available snapshots can be enabled when using high availability. For details, see “Highly available snapshots” on page 197.

Managing snapshots

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

You can add or delete a snapshot by using the **Volumes** > **Volume groups** panel in the management GUI.

Rename snapshots
You can change the name of an existing non-HA snapshot by using the **Rename** action. HA snapshots cannot be renamed.

Snapshot properties

Properties include basic properties of a snapshot and specific requirements that must be met.

The basic properties of a snapshot are:

- Snapshot states:
  - **ready:** A snapshot is created but not triggered.
  - **active:** Maintaining the point-in-time image.
  - **deleting:** Process of deleting.
  - **dependent_deleting:** The snapshot is instructed to be deleted but is being used to populate a clone, thin-clone, or is being used in a restore operation. Once the snapshot is no longer needed for these tasks, it automatically deletes.
- Snapshots are immutable, meaning that their content is protected from manipulation by the host.
- Snapshots consume space in the parent volume’s storage pool. By default, they are in the same pool as the parent volume, but they can be added to child pool of the parent volume’s pool. Or, if the parent volume is in a child pool itself, then they can be added to a child pool that has the same parent pool as the parent volumes pool.
- Expiration time value is assigned at the time a snapshot is created based on the source volume being part of a volume group with a snapshot policy.

  Note: On detection of an anomaly that leads to a Threat detection event, by default the latest volume group snapshot expiration time is extended to 7 days if it would otherwise expire beforehand. Snapshots created after the anomaly link to the anomaly fix procedure.

To add snapshots to a volume group or set of volumes, the following requirements must be met:

- Snapshots cannot be added to the volumes that are target volumes of legacy FlashCopy® mappings.
- 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 permitted except in the migration use-case.
- Volumes of zero size cannot have snapshots.

Cloning snapshots

A new volume group can be created and prepopulated with the contents of a snapshot. The new volumes and volume group inherit properties from the parent volumes and volume group of the snapshot.

The new volume groups can be of type thin-clone or clone. A thin-clone volume maintains space efficiency whereas a clone volume becomes an independent volume.

Administrators can use the management GUI or the command-line interfaces to create a clone or thin-clone volume group. In addition, the command-line interface supports the ability to create single volume clones or thin-clone outside of a volume group.

Thin-clone

A thin-clone is a host-accessible copy of a snapshot. By default, it is created in the same storage pool and I/O group as the source volume. A thin-clone can be used like any other volume or volume group with a few exceptions:

- A thin-clone is always dependent on the snapshot from which it was created.
- The capacity of a thin-clone can be expanded or shrunk. However, it cannot be shrunk to a size that is smaller than the snapshot from which it was created.
- If a snapshot is a local snapshot, a new thin clone volume group must be created outside the partition of the source volume group if the partition has a high availability replication policy assigned.
- If a snapshot is a highly available snapshot, a new highly available thin clone volume can be created inside a volume group within the partition of the source volume group. For details, see "Highly available snapshots" on page 197.

Clone

A clone is a host-accessible copy of a snapshot. Once all the data from the source snapshot is copied to the clone volume, it becomes independent of the snapshot from which it was created.

During the copy process, the volume displays as type 'clone', which changes once the copy is complete and the volume becomes a standard volume. By default, it is created in the same storage pool and I/O group as the source volume. Alternative pools or I/O groups can be specified.

A clone volume can be used like any other volume during the cloning process except it cannot be shrunk to a size smaller than the snapshot from which it was created.

Convert thin-clone to clone

A user can convert a thin-clone volume to a clone volume. When doing this any uncopied data is copied from the snapshot volume to the volume being converted and once that copy has completed the converted volume becomes a stand-alone volume. The system supports converting entire thin-clone volume groups and individual thin-clone volumes.

Restore or refresh from a snapshot

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

Restoring from a snapshot

This type of restore operation overwrites the contents of the production system with a previously saved version of the data.

You can restore either an entire volume group or subset of volumes within that group, which includes restoring clones and thin-clones. Additionally, snapshots of thin-clones can be used to restore the source volume group or volumes from which the thin-clone was created.

You can restore a volume group from a snapshots by using the **Volumes** > **Volume groups** > **Snapshots** option on the management GUI.

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

- Snapshots can be used to restore to the parent volume group or volumes, or to restore to the source of the parent volume group or volumes if the parent volume group is a thin-clone. The restore to source only works if the parent is a thin-clone.
- The volumes in a restore operation must all be present in the volume group associated with the snapshot. To restore a whole group then all volumes must be in the volume group. To restore a subset of

volumes that subset of volumes must be in the volume group, and it does not matter if there are other volumes present or missing in the volume group.

- A restore operation is permitted if the volumes to be restored have been deleted, but only using the command line interface (CLI).
- For restoring to the source of thin-clones, the thin-clone must be the same at time of restore as it was when the snapshot was taken. Also, ensure that none of the thin-clone volumes are deleted. A thin-clone, which is created from a subset of snapshot volumes, can be used to restore that subset of volumes to the source of the thin-clone.
- The volumes being restored must be the same virtual capacity as when the snapshot was added.
- You cannot restore volumes if they are transparent cloud tiering (TCT) enabled.

Refresh from a snapshot

Any existing data on the set of volumes is 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.

Safeguarded snapshots

Safeguarded snapshots are snapshots with an expiration time and restrictions regarding their premature deletion, that is the deletion of a snapshot before its expiration time. Only a Security administrator can delete a safeguarded snapshot prematurely.

Like standard snapshots, safeguarded snapshots can be created with either an internal scheduler or an external scheduler.

To safeguard snapshots using the internal scheduler, you can assign a snapshot policy and specify that the snapshot schedule should be safeguarded. This schedule then adds safeguarded snapshots to the volume group.

Note: You cannot configure safeguarded snapshots with the external scheduler if the volume group has a safeguarded backup policy assigned to it.

Safeguarded snapshot user roles

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.

Chapter 5. Hosts

A host system is a server that is connected to the storage system through supported connection protocols.

The storage system supports host connections through Fibre Channel, Serial Attached SCSI (SAS), NVM Express (NVMe) over Fibre Channel (FC-NVMe), NVM Express (NVMe) over RDMA, NVM Express (NVMe) over TCP, or an IP network. See the System Overview page on the relevant Hardware Guide to understand which compatible adapters can be used for each host attach protocol.

A host object is a logical object that represents a collection of host ports. Depending on host protocol, the host ports will be one of the following formats:

- Worldwide port names (WWPNs) for Fibre Channel and SAS hosts
- NVMe qualified names (NQNs) for FC-NVMe, NVM Express (NVMe) over RDMA and NVM Express (NVMe) over TCP hosts
- iSCSI qualified names (IQNs) or extended unique identifiers (EUIs) for iSCSI hosts

A typical configuration has one host object for each server that is attached to the storage system. If a cluster of hosts accesses the same storage, you can add the hosts into a host cluster object to make configuration simpler. For details, see .

To create and edit host objects in the management GUI, select **Hosts** > **Hosts**.

The system can detect WWPNs that have been logged into the system but are not yet configured in host objects, such that these can be allocated to host objects on creation (or added to existing host objects). These are visible in the management GUI in **Hosts** > **Hosts** > **Unconfigured hosts**. Some HBA device drivers do not leave ports logged in if no disks are detected on the network.

If you are configuring a host, such as hosts that are running the VMware operating system, define each host in the system as a separate host object. Do not combine multiple worldwide port names (WWPNs) from multiple hosts as one host object.

To create host objects by using the CLI, use the **mkhost** command. To add ports to an existing host object, use the **addhostport** command. To delete an existing host object, use the **rmhost** command.

The Fabric Device Management Interface (FDMI) enables any storage endpoint to register itself to the Fibre Channel (FC) fabric and query the HBA and port details of the entire fabric. For details, see .

You can manually enter WWPNs, NQNs and iSCSI names when creating hosts or adding ports to host objects.

Hosts can be assigned to an ownership group. Users in an ownership group are restricted to viewing and managing objects within their ownership group. For details, see .

Mapping volumes to a host

After host objects and volumes have been created, you must map the volumes to the host object (or hostcluster object) to provision storage to the host server.

Host mapping is similar in concept to logical unit number (LUN) mapping or masking. LUN mapping is the process of controlling which hosts have access to specific logical units (LUs) within the disk controllers. LUN mapping is typically done at the storage system level. Host mapping is done at the software level.

The act of mapping a volume to a host makes the volume accessible to the WWPNs or iSCSI names, such as iSCSI qualified names (IQNs) or extended-unique identifiers (EUIs) that are configured in the host object.

Note: If a volume is mapped to more than one host then those hosts must all use the same storage protocol (SCSI or NVMe). It is not supported to map a volume to the same server using different transport

© Copyright IBM Corp. 2025

protocols (e.g. Fibre Channel or Ethernet), however a volume can be mapped to different servers using different transport protocols.

Volumes and host mappings

Mappings can either be shared or private. Shared mappings are volume mappings that are shared among all the hosts that are in a host cluster. When a host cluster is created, any common volume mappings become shared among all the hosts within the host cluster. If a mapping is not common, it remains a private mapping for that host only. Private mappings are mappings that are associated with an individual host. To create, edit, or remove mappings, select **Volumes** > **Host mappings**. The Mappings view displays all private mappings by default. From the View menu, select either **Shared Mappings** to view shared mappings for hosts or **All Host Mappings** to display all configured private and shared mappings. Without a multipathing device driver, most operating systems present each path to a volume as a separate storage device. The multipathing software manages the many paths that are available to the volume and presents a single storage device to the operating system. If there are multiple paths, the system requires that the multipathing software run on the host or host cluster.

When you map a volume to a host, you can optionally specify a SCSI ID for the volume. This ID controls the sequence in which the volumes are presented to the host. Check the host software requirements for SCSI IDs because some require a contiguous set. For example, if you present three volumes to the host, and those volumes have SCSI IDs of 0, 1, and 3, the volume that has an ID of 3 might not be found because no disk is mapped with an ID of 2. The system automatically assigns the lowest available SCSI ID if none is specified.

Figure 9 on page 108 and Figure 10 on page 109 show two volumes, and the mappings that exist between the host objects and these volumes.

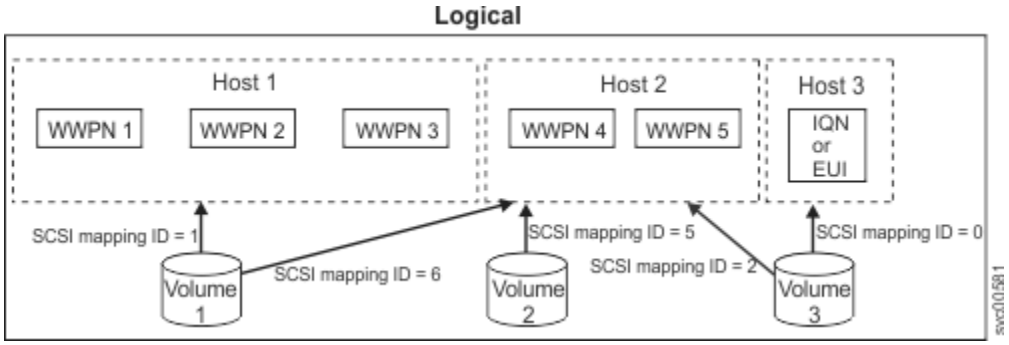


Figure 9. Hosts, WWPNs, IQNs or EUIs, and volumes

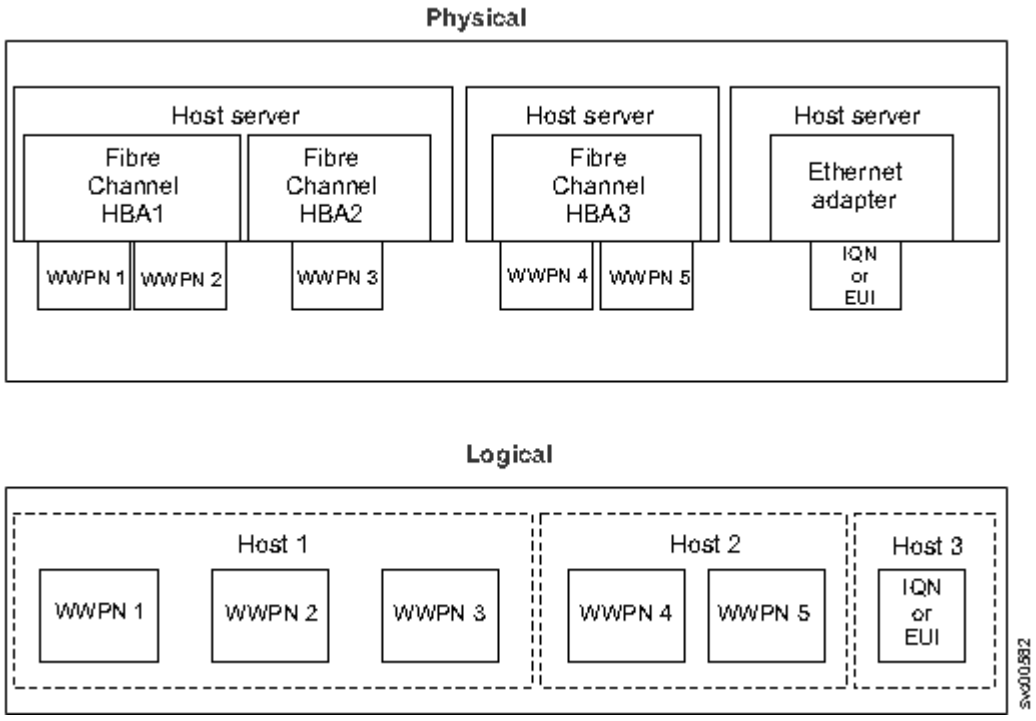


Figure 10. Hosts, WWPNs, IQNs or EUIs, volumes, and SCSI mappings

LUN masking is usually implemented in the device driver software on each host. The host has visibility of more LUNs than it is intended to use, and device driver software masks the LUNs that are not to be used by this host. After the masking is complete, only some disks are visible to the operating system. The system can support this type of configuration by mapping all volumes to every host object and by using operating system-specific LUN masking technology. However, the default, and recommended, system behavior is to map only those volumes that the host is required to access.

Host mappings inherit the ownership group of the host or host cluster and volume in the mapping. For details, see "Ownership groups" on page 211.

Determining the volume name from the device identifier on the host

You can use the command-line interface (CLI) to determine the volume name from the device identifier on the host.

About this task

Each volume that is exported by the system is assigned a unique device identifier. The device identifier uniquely identifies the volume and can be used to determine which volume corresponds to the volume that the host detects.

Complete the following steps to determine the volume name from the device identifier:

Procedure

1. Find the device identifier.
   For example, if you are using the native multipathing plug-in (NMP), the disk identifier is referred to as the virtual path (vpath) number. You can issue the following NMP command to find the vpath serial number:

   ```
   datapath query device
   ```

   For other multipathing drivers, refer to the documentation that is provided with your multipathing driver to determine the device identifier.

2. Find the host object that is defined to the system and corresponds with the host that you are working with.
   a) Find the worldwide port numbers (WWPNs) by looking at the device definitions that are stored by your operating system.
   For example, on AIX® the WWPNs are in the ODM and if you use Windows you must go into the HBA BIOS.
   b) Verify which host object is defined to the system for which these ports belong. The ports are stored as part of the detailed view so that you must list each host by issuing the following CLI command:

      ```
      lshost id | name
      ```

      Where *id* | *name* is the name or ID of the host.
   c) Check for matching WWPNs.
3. Enter the following command to list the host mappings:

   ```
   lshostvdiskmap hostname
   ```

   Where *hostname* is the name of the host.
4. Find the volume UID that matches the device identifier and record the volume name or ID.

Volume reservations

The SCSI **Reserve** command and the SCSI **Persistent Reserve** command are specified by the SCSI standards. Servers can use these commands to prevent ports in other servers from accessing particular LUNs.

Reserving LUNs prevents accidental data corruption that is caused when a server overwrites data on another server. The **Reserve** and **Persistent Reserve** commands are often used by clustered-system software to control access to system volumes.

If a server is not shut down or removed from the server system in a controlled way, the standard and persistent reservations of the server are maintained. Reservations prevent other servers from accessing data that is no longer in use by the server that holds the reservation. In this situation, you might want to release the reservation to allow a new server to access the volume.

When possible, have the server that holds the reservation explicitly release the reservation to ensure that the server cache is flushed and that the server software is aware that access to the volume was lost. In circumstances where this action is not possible, you can use operating-system-specific tools to remove reservations. Consult the operating system documentation for details.

When you use the **rmvdiskhostmap** CLI command or the management GUI to remove host mappings, system nodes can remove the standard reservations and persistent reservations that the host has on the volume.

Host portsets

A host portset defines the set of node ports that a host can use to communicate with the system. Portsets apply to both Ethernet and Fibre Channel (FC) environments. Use portsets to control host access, support isolation between workloads, and provide predictable connectivity.

A host object can be assigned to only one portset. A portset can include multiple hosts. Hosts in a cluster can use different portsets. If you do not assign a portset, the host uses the default portset. NVMe-FC hosts cannot use the default portset. They must use a user-defined Fibre Channel portset.

For more information and requirements about portsets, see .

Fibre Channel host portsets

Fibre Channel portsets define the FC target ports that a host can access. You can use these portsets with automated switch zoning tools to create consistent zoning rules.

For more information about Fibre Channel portset, see “Configuring Fibre Channel portsets” on page 335.

Ethernet host portsets

Ethernet portset requirements

- A host portset can include up to four IP addresses per node.
- A host portset must include at least one IP address before you add a host.
- Host attach IP addresses within a single portset can be on different VLANs or subnets
- Each physical port supports one unique routable IP address with a gateway. The same routable IP can be referenced in multiple portsets.
- For multipathing support, configure IP addresses symmetrically by assigning IPs from the same port number on both nodes. Example: node1 port3 and node2 port3.
- Symmetric IP configuration across both nodes ensures proper failover, failback, and load balancing for host connectivity.

For more information about Ethernet portsets in general, see “Configuring Ethernet portsets” on page 334.

Host connectivity

The system reports various information to help the user understand host connectivity and troubleshoot any issues.

For Fibre Channel, FC-NVMe, NVMe over RDMA, and NVMe over TCP-attached hosts, the system reports the node login count, which is the number of nodes that can detect each WWPN or NQN.

To view connectivity for Fiber Channel-attached hosts in the management GUI, select **Settings** > **Network** > **Fibre Channel Connectivity** or use the **`lsfabric`** command. To view connectivity issues for FC-NVMe, NVMe over RDMA, and NVMe over TCP-attached hosts in the management GUI, select **Settings** > **Network** > **NVMe Connectivity** or use the **`lsnvmefabric`** command.

For Ethernet-attached hosts, the number of logged-in nodes refers to sessions that are created between hosts and nodes. The number of logged-in nodes can be greater than the current number of nodes on the system.

- You can display the login information for each host in the management GUI to determine which ports for the logged-in nodes are connected to the Ethernet-attached hosts. To display this information, select **Hosts** > **Hosts** and right-click the host and select IP Login Information.
- To view login information in the command-line interface, use the **`lshostiplogin`** to list login session details for a specified host that is logged in to a configured Ethernet port on the system.
- You can also set a status policy and site status to determine how host statuses are displayed. For example, if you select Complete, the system displays an Online status only when all host ports are logged in to all nodes and hosts can access all mapped volumes. If a single host or node loses connectivity, the status for the host displays as Degraded. If you select Redundant for the Status Policy, an Online status is displayed when a sufficient number host ports to node logins exist to retain access to all its mapped volumes, even when a host port or single node is removed or lost.
- If hosts are part of a multiple site configuration, each host needs to be assigned to a specific site and you can specify how host statuses are displayed for these hosts. In the management GUI, you can specify All for Site Status to indicate that host connectivity to the nodes on all sites determine the status of the host. If you select Local for Site Status, the host status relates only to connectivity to nodes on the local site. Host status for other nodes mapped to different sites is not included.

Host clusters

A *host cluster* is a group of logical host objects that can be managed together. For example, you can create a volume mapping that is shared by every host in the host cluster.

The system uses internal protocols to manage access to the volumes and ensure consistency of the data. Host objects that represent hosts can be grouped in a host cluster and share access to volumes. A host can only be in one host cluster at a time. New volumes can also be mapped to a host cluster, which simultaneously maps that volume to all hosts that are defined in the host cluster.

Each host cluster is identified by a unique name and ID, the names of the individual host objects within the cluster, and the status of the cluster. A host cluster can contain up to 128 hosts. However, a host can be a member of only one host cluster.

A host cluster can have one of the following statuses:

Online
All hosts in the host cluster are online.

Host degraded
All hosts in the host cluster are either online or degraded.

Host cluster degraded
At least one host is offline and at least one host is either online or degraded.

Offline
All hosts in the host cluster are offline (or the host cluster does not contain any hosts).

Host clusters can be assigned to an ownership group. An ownership group defines a subset of users and objects within 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. Host clusters can be owned if they are assigned an ownership group explicitly or by inheritance from the user creating them. New or existing hosts that are defined in a host cluster inherit the ownership group that is assigned to the host cluster.

Volume mapping

By default, hosts within a host cluster inherit all shared volume mappings from that host cluster, as if those volumes were mapped to each host in the host cluster individually. Hosts in a host cluster can also have their own private volume mappings that are not shared with other hosts in the host cluster.

With *shared mapping*, volumes are mapped on a host cluster basis. The volumes are shared by all of the hosts in the host cluster, if there are no Small Computer System Interface (SCSI) LUN conflicts among the hosts. Volumes that contain data that is needed by other hosts are examples of a shared mapping.

If a SCSI LUN conflict occurs, a shared mapping is not created. SCSI LUN conflicts are caused when multiple volumes are mapped to the same SCSI LUN ID. SCSI LUN conflicts can also be caused when a single volume is mapped to multiple SCSI LUN IDs.

With *private mapping*, individual volumes are directly mapped to individual hosts. These volumes are not shared with any other hosts in the host cluster. A host can maintain the private mapping of some volumes and share other volumes with hosts in the host cluster.

For example, Figure 11 on page 113 shows two hosts that are members of a host cluster. The volume that contains data is a shared mapping with Host A and Host B. However, Host A and Host B also have a private mapping to their respective boot volumes (LUN A and LUN B).

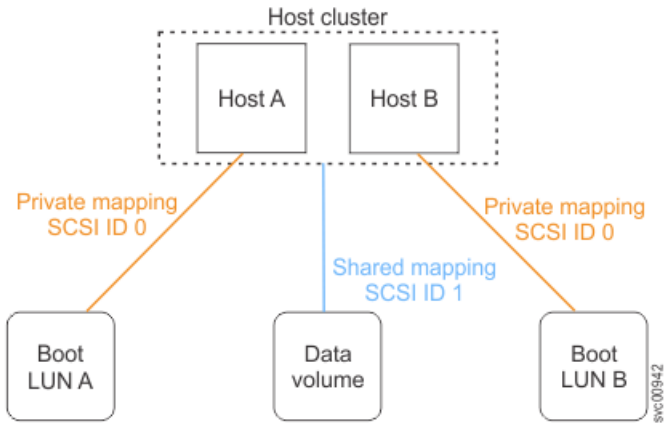


Figure 11. Host cluster that contains private and shared volumes

Managing host clusters

A *host cluster* is a group of logical host objects that can be managed together. For example, you can create a volume mapping that is shared by every host in the host cluster.

The system uses internal protocols to manage access to the volumes and ensure consistency of the data. Host objects that represent hosts can be grouped in a host cluster and share access to volumes. A host can only be in one host cluster at a time. New volumes can also be mapped to a host cluster, which simultaneously maps that volume to all hosts that are defined in the host cluster.

When you create a host cluster, you can specify the hosts to include in the cluster. You can also create an empty host cluster and add hosts later. If you select hosts to add to the host cluster, all of the common private volume mappings for the hosts become shared mappings for the host cluster. However, you can also exclude specific volumes from being shared. The private volume mappings are not shared with any other hosts in the host cluster. A host can be a member of only one host cluster at a time.

With *shared mapping*, volumes are mapped on a host cluster basis. The volumes are shared by all of the hosts in the host cluster, if there are no Small Computer System Interface (SCSI) LUN conflicts among the hosts. Volumes that contain data that is needed by other hosts are examples of a shared mapping.

When adding hosts to a host cluster, if you selected hosts with volumes that have SCSI ID conflicts, the system does not add these mappings to the host cluster. A SCSI LUN ID conflict occurs when multiple hosts are mapped to the same volume but with different SCSI IDs. In this case, a shared mapping is not created because the system does not allow a volume to be mapped more than once to the same host. The **Summary** page lists all volumes that contain conflicts and the system retains these mappings as private mappings to the original hosts.

If throttles are configured on any of the selected hosts, the throttles must be removed to be included in the host cluster. Throttles can be applied only to the host cluster and not individual hosts within the cluster. If you choose not to remove the throttles from the host, the host is excluded from the host cluster.

When a host is removed from a host cluster, you can convert the shared mappings of the host cluster to private mappings or remove all shared mappings from the host. If shared mappings are converted to private mappings, volume access is not removed since the volume remains mapped to the host and its SCSI ID remains unchanged. If you choose to remove the shared mappings when you remove a host, only the private mappings for that host are retained when it is removed from the host cluster. When you remove the shared mapping, hosts cannot access those volumes that were mapped as shared mappings to the host cluster.

You can use several tools to manage host clusters, including the management GUI, and the command-line interface (CLI). To work with host clusters in the management GUI, select **Hosts** > **Host Clusters**.

Managing WWPN changes of replaced host HBAs

You can use the command-line interface (CLI) to record a change to a defined host object.

Before you begin

Sometimes it is necessary to replace the host-bus adapter (HBA) that connects the host to the SAN. You must inform the system of the new worldwide port names (WWPNs) that the replacement HBA contains.

Ensure that your switch is zoned correctly.

Procedure

To inform the system of a change to a defined host object, complete the following steps.

1. Enter the following CLI command to list the candidate HBA ports.

   ```
   lsfcportcandidate
   ```

   or

   ```
   lssasportcandidate
   ```

   You see a list of the HBA ports that are available for addition to host objects. One or more of these HBA ports corresponds with one or more WWPNs that belong to the new HBA port.
2. Locate the host object that corresponds with the host in which you replaced the HBA. The following CLI command lists all the defined host objects:

   ```
   lshost
   ```

3. Enter the following CLI command to list the WWPNs that are currently assigned to the host object.

   ```
   lshost hostobjectname
   ```

   Where *hostobjectname* is the name of the host object.
4. Enter the following CLI command to add the new ports to the existing host object.

   ```
   addhostport -fcwwpn one or more existing port names
    separated by : hostobjectname/ID
   ```

   Where *one or more existing port names separated by :* is the WWPNs that are currently assigned to the host object and *hostobjectname/ID* is the name or ID of the host object.
5. Enter the following CLI command to remove the old ports from the host object.

   ```
   rmhostport -fcwwpn one or more existing port names
    separated by : hostobjectname/ID
   ```

   Where one or more existing WWPNs separated by a colon (`:`) are the WWPNs that are currently assigned to the host object and *hostobjectname/ID* is the name or ID of the host object.

   Note: If the following conditions are met when volume protection is enabled for the system, the deletion of the specified host port fails.

   - It is the last active port on the host.
   - It is mapped to any volume that received I/O within the specified volume protection interval.

   If volume protection is enabled, and the host port being deleted is the last port for a host, which is mapped to any volume that received I/O within the defined volume protection time period, then the command fails. If multiple hosts are mapped to the same active volume, the system deletes the port if the host is offline.

Results

Any mappings that exist between the host object and the volumes are automatically applied to the new WWPNs. Therefore, the host sees the volumes as the same SCSI LUNs as before.

What to do next

See the *IBM Multipath Subsystem Device Driver User's Guide* or the documentation that is provided with your multipathing driver for additional information about dynamic reconfiguration.

Managing throttles for hosts

You can create, modify, and remove throttles for hosts by using the management GUI or the command-line interface.

Throttles are a mechanism to control the amount of resources that are used when the system is processing I/Os on a specific host or host cluster. If a throttle is defined, the system either processes the I/O, or delays the processing of the I/O to free resources for more critical I/O.

Using the management GUI

To create a new throttle for a specific host, complete these steps:

1. In the management GUI, select **Hosts** > **Hosts**.
2. Right-click a host or multiple hosts and select **Edit Throttle...**.
3. On the **Edit Throttle** page, enter values for the following throttles:

   Bandwidth Limit
   Enter the maximum amount of bandwidth the host can process before the system delays processing for this host. Click **Create**.

   IOPS Limit
   Enter the maximum I/O operations per second the host can process before the system delays processing for this host. Click **Create**.

To change an existing throttle, complete these steps:

1. In the management GUI, select **Hosts** > **Hosts**.
2. Right-click a host cluster and select **Edit Throttle...**.
3. On the **Edit Throttle** page, enter values for the following throttles:

   Bandwidth Limit
   Enter the maximum amount of bandwidth the host can process before the system delays processing for this host. Click **Save**.

   IOPS Limit
   Enter the maximum I/O operations per second the host can process before the system delays processing for this host. Click **Save**.

To remove throttles for a specific host, complete these steps:

1. In the management GUI, select **Hosts** > **Hosts**.
2. Right-click a host or multiple hosts and select **Edit Throttle...**.
3. On the **Edit Throttle** page, click **Remove** by the throttle that you want to delete.

Using the command-line interface

To create a new throttle for bandwidth and IOPS for a host with the command-line interface, enter the following command:

```
mkthrottle -type host -bandwidth bandwidth_limit_in_mb -iops iops_limit -host
host_id_or_host_name
```

where *bandwidth_limit_in_mb* is the maximum amount of bandwidth that the system can process for the specified host before the system delays processing of I/O for that host. The *iops_limit* is the maximum I/O operations per second that the system can process before the system delays processing for this host. The *host_id_or_host_name* is the ID or name of the host that the throttle applies to.

To change an existing throttle for the host, enter the following command:

```
chthrottle -type host -bandwidth bandwidth_limit_in_mb -iops iops_limit throttle_name_or_id
```

where *bandwidth_limit_in_mb* is the maximum amount of bandwidth that the system can process for the specified host before the system delays processing of I/O for that host. The *iops_limit* is the maximum I/O operations per second that the system can process before the system delays processing for this host. The *throttle_id_or_name* is the ID or name of the throttle that is being changed.

To remove a throttle for a host with the command-line interface, enter the following command:

```
rmthrottle throttle_name_or_id
```

where *throttle_name_or_id* is the ID or name of the throttle that is being removed.

Chapter 6. Host attachment

This topic describes configuration activities to enable hosts to access volumes and use host-based solutions to manage the storage system.

The IBM Storage FlashSystem and IBM SAN Volume Controller product family supports the following protocols for host attachment:

- Fibre Channel
- NVMe over Fibre Channel (FC-NVMe)
- NVMe over RDMA
- NVMe over TCP
- Serial-attached SCSI (SAS)
- Ethernet iSCSI
- Small Computer System Interface-Fibre Channel (SCSI-FCP)

Refer to the **System overview** page of the product hardware documentation to determine which protocols are supported on each hardware configuration.

For more information about host compatibility, see IBM System Storage Interoperation Center (SSIC).

Hosts that use the Fibre Channel connections are attached to the system either directly or through a switched Fibre Channel fabric. Each port on a node is identified by a worldwide port name (WWPN).

Host connections through NVMe over RDMA use specialized network switches and require fewer resources than FC-NVMe. RDMA allows higher throughput and better performance with low latency. In addition, RDMA requires less expertise at the storage networking level than the Fibre Channel implementation, potentially reducing overall costs.

The advantage of using NVMe over TCP data transfer is that unlike RDMA data transfer, NVMe over TCP uses the existing Ethernet adapters on the host and network infrastructure.

Refer to the following pages for information about host configuration.

Using VMware

There are a number of integrations available between IBM Storage Virtualize and VMware that whilst not essential, will improve the overall functionality and provide a better user experience.

IBM Storage Virtualize plugin for vSphere

The IBM Storage Virtualize plugin for vSphere is based on VMware's Remote plugin architecture, and facilitates provisioning and management of IBM Storage from within the vCenter UI. The plugin and associated workflows are integrated into the vCenter user interface and allows a vSphere administrator to remotely perform administrative or maintenance actions across multiple storage systems. For more information on vSphere plugin, see https://www.ibm.com/docs/en/svpfv.

IBM Storage Virtualize Family - Storage Replication Adapter (SRA)

The IBM® Storage Virtualize Family Storage Replication Adapter (SRA) is a software add-on that integrates with the VMware Site Recovery Manager (SRM) solution to run Disaster Recovery failover scenarios with supported storage systems. For more information on SRA, see https://www.ibm.com/docs/en/spectrumvirtual-sra.

© Copyright IBM Corp. 2025

Using OpenStack

The IBM Storage Virtualize family volume management driver (formerly known as the Storwize® Cinder driver) offers various block storage services.

For further details about OpenStack, see the following links:

- IBM Storage Virtualize family volume driver documentation
- Cinder driver support matrix
- Cinder release notes

The following OpenStack versions are supported with this release:

- Wallaby
- Xena
- Yoga
- Zed
- Antelope
- Bobcat
- Caracal
- Dalmatian
- Epoxy
- Flamingo

Using Kubernetes and OpenShift

Red Hat® OpenShift® Data Foundation (ODF) uses the ODF FlashSystem Operator and ODF FlashSystem Driver to integrate with your IBM Storage FlashSystem. When using the OpenShift Data Foundation with your storage system ensure that all configuration needs are met by following the Red Hat ODF documentation set.

For more information on supported Red Hat OpenShift Data Foundation (ODF) version support ODF FlashSystem Operator and ODF FlashSystem Driver and FlashSystem storage systems, see ODF FlashSystem driver support summary.

- For more information about Red Hat OpenShift Data Foundation, a software-defined storage for containers, see the OpenShift Data Foundation page.

The ODF FlashSystem Operator and ODF FlashSystem Driver enable IBM Storage Virtualize family products to be used with Red Hat OpenShift Data Foundation.

For more information and documentation, see the following:

- User information and release notes documentation for the Red Hat OpenShift Data Foundation: OpenShift Data Foundation documentation.
- General information about the OpenShift Data Foundation, a software-defined storage for containers: OpenShift Data Foundation

Copy offload by using ODX on Microsoft Windows

Offloaded data transfer (ODX) is a feature in Microsoft Windows Server. ODX function frees up hosts and speeds the copy process by offloading the entire orchestration to a storage array.

For virtual machine (VM) operations and large data transfers, Microsoft ODX offloads the heavy lifting of data movement to the host storage array instead of doing normal read and write operations. Examples include copy operations such as:

- VM creation

- VM migration
- VM cloning
- Microsoft Windows Virtual Hard Disk (VHD) creation
- VHD conversion
- VM backup and recovery
- File copy

ODX is most relevant in a Hyper-V environment for VM heavy lifting operations. Standard buffered copy is run by reading the data from the storage controller into the host, buffering it, and then writing it to another volume. The ODX function frees up hosts and speeds the copy process by offloading the entire orchestration to a storage array. This offloading is done by using tokenization for read operations and write operations, and it avoids buffering, which can ultimately cut down on processor cycles.

Planning for offloaded data transfer (ODX)

For a copy workload in a Microsoft Windows environment, you can use Microsoft offloaded data transfer (ODX) to offload the copy workload to the storage controller to speed job completion. When you decide on an ODX implementation, review the planning considerations to help you determine whether ODX is appropriate for your environment.

It is important if you have a Windows environment to know whether existing Windows workload consists of copy workload, and if it does, what is its extent. By default, Windows offloads any copy greater than 256 K to the controller if the volumes are published as ODX capable. Two important points exist.

1. With ODX enabled, Windows offloads any copy greater than 256 K to the system controllers. The offload frees up the hosts CPU and network. Depending upon the type of copy, the system can even experience faster copy in certain situations. However, for multiple Windows hosts in the environment, all Windows hosts start pumping copy offload jobs to the storage controllers. If you use ODX, plan for this impact.
2. Workloads other than copy might experience higher latencies if the copy workload is high. Since the offloaded jobs tend to “Copy offload by using ODX on Microsoft Windows” on page 118 faster, they require more controller resources in a specified time. The result might be higher latencies with non-copy workloads.

If you have a Hyper-V environment, the benefit is immense through copy offload ODX technology. With ODX enabled, all Windows jobs like file copy, VM migration, cloning, and VHD creation would be offloaded to the storage controller. It is advisable to check the impact of copy offload on other types of workloads in such an environment.

If you decide to implement ODX, it is advisable that you review existing workloads and make sure that capacity planning is completed to meet the demand. If you are not interested in the benefits of copy offload, you can disable system-wide ODX by using the CLI. Not using ODX prevents existing workloads from experiencing extra latency due to copy offload.

SCSI Fibre Channel host attachments

Small Computer System Interface (SCSI) Fibre Channel hosts with different operating systems can be attached to the system.

The system supports configurations that use N_port ID virtualization (NPIV) in the host bus adapter when they are connected to a SAN switch. However, these configurations are not supported when they are directly attached to the Fibre Channel ports on the system.

IBM Power Systems AIX

You can attach IBM Power Systems AIX hosts via Fibre Channel.

For more information about host compatibility, see IBM System Storage Interoperation Center (SSIC).

Refer to AIX Documentation for further configuration information.

Configuring the AIX operating system

You must configure the AIX operating system to use IBM Power Systems hosts with the system.

Configuring for fast fail and dynamic tracking

You must enable the fast fail and dynamic tracking attributes for hosts systems that run an AIX 5.2 or later operating system.

Before configuring your host system to use these attributes, ensure that the host is running the AIX operating system version 5.2 or later.

To configure your host system to use the fast fail and dynamic tracking attributes follow these steps.

Note: These steps need to be taken only on systems running AIX 6.1 or 5.2. The commands require the fscsi*X* device to be in the defined state.

1. Issue the `chdev -l fscsi0 -a fc_err_recov=fast_fail` command to set the Fibre Channel SCSI I/O Controller Protocol Device event error recovery policy to fast_fail for each Fibre Channel adapter. This example command is for adapter fscsi0.

   Note: By default, dynamic tracking is enabled on all systems running AIX 7.1.
2. Issue the `chdev -l fscsi0 -a dyntrk=yes` command to enable dynamic tracking for each Fibre Channel device. This example command is for adapter fscsi0.

Multipath support for IBM Power Systems hosts

You must install multipathing software on all IBM Power Systems AIX hosts that are attached to the system.

On IBM Power Systems hosts, the AIX Path Control Module (AIXPCM) provides multipathing support.

For more information about the multipath drivers, see www.ibm.com/support.

Configuring system devices with multiple paths per LUN

- The system supports multiple LUNs on an IBM Power Systems AIX host.
- The system supports multiple path configurations for a LUN. You can have multiple hdisks (logical hard disks) available on the host for each physical LUN.
- To configure multiple paths for all LUNs, run the **cfgmgr** command to scan for new devices.
  - If you are using the AIX Path Control Module (AIXPCM) and want to check whether the multipath device is available, run the **lsdev** command and search for the hdisk# line with the `MPIO IBM 2###### FC Disk` label against it.

For more information about which multipath driver to use, see http://www.ibm.com/support.

SAN boot support for IBM Power Systems hosts

You can reload the operating system from a boot disk on the SAN over Fibre Channel. You can initiate this process from the system volume if your IBM Power Systems hosts use AIX operating system version 5.3 or later.

You must use the AIX Path Control Module (AIXPCM) to boot from the SAN. Create an installation volume and map it to the host. Proceed with the installation per the AIX installation instructions and, when you are prompted, select the previously defined volume as the target installation disk.

Dynamically increasing volume size

If your IBM Power Systems AIX hosts use AIX 5.2 or a later AIX operating system version, the system supports the ability to dynamically increase volume size.

The **chvg** command options make it possible to increase the size of a physical volume that the Logical Volume Manager (LVM) uses, without interruptions to the use or availability of the system. See the AIX publication *System Management Guide: Operating System and Devices* for more information.

Virtual input/output for IBM Power Systems hosts

The system provides both single and dual Virtual input/output (VIO) server configurations on IBM Power Systems hosts that support VIO.

You can present the system volumes to the VIO server host bus adapters (HBAs) by using the same method as a standard AIX installation. For single VIO server configurations, volumes can be split up into logical volumes by the VIO server and mapped to the VIO clients. For dual VIO server configurations, volumes cannot be split into logical volumes, and must instead be mapped intact through both servers to the VIO clients.

Known AIX issues and limitations

The restriction of Direct Memory Access (DMA) resources is one of the known issues and limitations with the system and an IBM Power Systems AIX host.

On a heavily loaded system, the following symptoms can indicate that the host is low on direct memory access (DMA) resources:

- You might see errors that indicate that the host bus adapter (HBA) was unable to activate an I/O request on the first attempt.
- You might see lower-than-expected performance with no errors being logged.

To reduce the incidence of these messages, you can increase the resources by modifying the maximum transfer size attribute for the adapter as follows:

1. Type the `lsattr -El` *HBA* `-a max_xfer_size` command to view the current setting, where *HBA* is the name of the adapter that is logging the error. For this example, the HBA is fcs0.
2. Type the `chdev -l fcs0 -P -a max_xfer_size=0x1000000` command to increase the size of the setting.

   Note: To view the range of allowable values for the attribute, type `lsattr -Rl fcs0 -a max_xfer_size`.
3. Restart the host to put these changes into effect.

IBM Power Systems with Virtual I/O Server

Successful Fibre Channel attachments for IBM Power Systems with Virtual I/O Server require knowledge of related configuration details and known issues and limitations with the server.

Configuring the IBM i operating system

You must configure the Virtual I/O Server and the IBM i client operating system before you can use IBM i as a host.

Before you configure the host operating systems, the following tasks must be completed:

- Your system must be installed.
- You must install the appropriate host bus adapter and driver on your host.

After the prerequisite tasks are complete, use the following general steps to configure your host operating system:

1. Zone the Virtual I/O Server to the system on the Fibre Channel network.
2. Install the AIX Path Control Module (AIXPCM) on the Virtual I/O Server to enable the management of multiple paths to the system volumes.
3. Create the host on the system by using the worldwide port names (WWPNs). Map the volumes to the host as required.

   Note: The IBM i virtual I/O client supports boot from a SAN with a volume for the load source of 20 GB or larger.

4. Run the **cfgdev** command on the Virtual I/O Server to configure the new disks. Use the **lspath** command to ensure that disks are listed for all paths to the system.
5. Map the system LUNs (hdisks on the Virtual I/O Server) uniquely to the IBM i virtual SCSI client adapters by using the **mkvdev** command.
6. Use the Hardware Management Console (HMC) to verify that the I/O tagging for the IBM i virtual I/O client partition is correctly set for the load source and the alternate restart device.
7. Install IBM i version 6 release 1 or later.

IBM z Systems that run the Linux operating system

You can attach IBM z Systems® that run the Linux operating system via Fibre Channel.

For more information about host compatibility, see IBM System Storage Interoperation Center (SSIC).

Refer to IBM z Systems Linux documentation and the following pages for further configuration information.

Configuring the Linux operating system for IBM z Systems® hosts

You must configure the Linux operating system before you can use z Systems hosts with the system.

Multipath support for IBM z Systems hosts

You must install multipathing software on all z Systems hosts that are attached to the system.

On z Systems hosts that are running the Linux operating system, the following software provides multipathing support:

Linux Kernel 2.4
The logical volume manager.

Linux Kernel 2.6

- The multipath-tools package on SUSE Linux Enterprise Server.
- The device-mapper-multipath package on Red Hat Enterprise Linux.

For more information about using the logical volume manager or the multipathing tools package with z Systems hosts that run Linux kernel 2.4 or Linux kernel 2.6, see the documentation at the following website:

www.ibm.com/developerworks/linux/linux390/index.html.

Installing and configuring the host bus adapter on an IBM Z® host system

The host bus adapters (HBAs) for an IBM Z® host system must be ordered as features. These features are either factory-installed when you order a new system or installed into an existing system by a service representative.

To configure the HBA:

- Ensure that the appropriate features (FICON®, FICON Express, FICON Express2, or FICON Express4) are installed on your IBM Z® host system.
- Configure the HBA to run in FCP mode.

For more information about FCP connectivity, see www.ibm.com/support.

SAN boot support on IBM z Systems hosts

On a SCSI LUN, you can complete an initial program load (IPL) process.

IPL processes on SCSI LUNs can fail because there is no path failover capability during the boot process. See the latest *How to use FC-attached SCSI devices with Linux on z Systems* publication for more information about using IPL processes with z Systems hosts.

Defining the number of disks on IBM z Systems hosts

When you define the number of disks on z Systems hosts that are running the Linux operating system, you are provisioning capacity for configured disks.

On the Linux operating system, a device node is used for each path to the disk and a second device node is used for the path grouped device. The maximum number of devices depends on your Linux configuration.

System configuration for IBM z Systems hosts

Each attached system LUN has a set of device files in the `/dev` Linux directory.

These device files are intended to be used only by the multipath tools. Allowing them to be used directly by your system or applications can produce data inconsistencies. Instead, configure your system and applications to use the device nodes that are created by the multipath tools.

Known issues and limitations for IBM z Systems hosts

There are some restrictions for z Systems hosts that are running the Linux operating system.

The following website provides currently known restrictions for the latest Linux on z Systems releases:

www.ibm.com/developerworks/linux/linux390/index.html.

IPL restrictions

Device Mapper Multipath (DM-Multipath) is not available on either the root or boot devices.

For information about DM-Multipath configuration and administration, see Device Mapping Multipath. System re-IPL (`shutdown -r`) is supported on z/VM® guests only. It is not supported in LPAR mode.

Changing the timeout value for the FC transport class

You can change the timeout value for the Fibre Channel (FC) transport class when the I/O operation stalls.

The FC transport class keeps all outstanding I/O for a certain time until it determines that the remote port is not responding. This time period is controlled by the timeout value. The default time is 60 seconds. This situation causes a temporary I/O stall in the case of FC link incidents or when paths are varied offline.

Change the timeout value.

- Set the timeout value for each remote port by using the `sysfs` file system.

```
/sys/class/fc_remote_ports/rport-0:0-0/dev_loss_tmo
```

- Preset the timeout value for all remote ports when you load the `scsi_transport_fc` module by using the `dev_loss_tmo` parameter.

SCSI device state changes from running to offline

In some instances, the Small Computer System Interface (SCSI) mid layer changes the SCSI device state from **running** to **offline**.

I/O that is later sent to that SCSI device is rejected by the mid layer and never reaches the SCSI lower-level device drivers for processing. To solve this problem, you must increase the SCSI timeout value to at least 60 seconds as shown in the following example:

```
SUBSYSTEM=="scsi", ACTION=="add", ATTR{vendor}=="IBM",
ATTR{model}==", ATTR{timeout}="60"
```

IBM z Systems running the IBM z/VSE operating system

You can attach IBM z Systems servers (hosts) that are running the IBM z/VSE® operating system. The z/VSE operating system can operate in a logical partition (LPAR) and in an IBM z/VM guest environment.

For more information about host compatibility, see IBM System Storage Interoperation Center (SSIC).

The following IBM website provides general z/VSE information about supported z/VSE releases, and supported IBM z Systems servers: www-03.ibm.com/systems/z/os/zvse/about/status.html.

You need a FICON Express (FICON Express, FICON Express2 or higher) adapter (also named host bus adapter (HBA)) on your z Systems server.

You must configure the FICON® Express adapter or HBA in FCP mode. FCP mode means that the HBA is configured with channel path ID (CHPID) type FCP in the input/output control data set (IOCDS) of your z Systems server.

Note: With N_Port ID Virtualization (NPIV), multiple virtual FCP adapters can be defined. Each virtual FCP adapter can have its own unique WWPN. You can use the WWPN of a virtual FCP adapter to map volumes to the host. This action can be used to define multiple hosts that share the same physical FCP adapter. The NPIV support was introduced with System z9®.

Configure your z/VSE system to use FCP-attached SCSI disks. After you configure a path to the SCSI disk, the SCSI disks are recognized by z/VSE user programs and system programs as fixed block architecture (FBA) disks only. You can use the SCSI disks with existing FBA interfaces. The z/VSE operating system does not provide SCSI command support for applications. See the IBM z/VSE Administration publication for additional information.

SAN boot support for IBM z Systems servers running the z/VSE operating system

If you have installed the z/VSE operating system on SCSI disks, you can complete the initial program load (IPL) process from the SCSI disks.

See the *IBM z/VSE Administration* publication for information regarding what specifications are required to initiate the IPL process in an LPAR environment and in a z/VM guest environment.

Hosts that run the Linux operating system

You can attach hosts running the Linux operating system via Fibre Channel.

For more information about host compatibility, see IBM System Storage Interoperation Center (SSIC).

Refer to Linux Documentation for further configuration information.

Requirements for attaching systems to hosts running the Linux

Ensure that your system meets the requirements for attaching to a host that is running the Linux operating system.

The following list provides the requirements for attaching the system to the host that is running the Linux operating system:

- Check the LUN limitations for your host system.
- Ensure that you have the documentation for your host and access to the hardware installation information for the correct model of your system. All system publications are available from the following website:www.ibm.com/support.
- Ensure that you install the correct operating systems and are running a supported kernel of Linux.
- When you attach the system to a BladeCenter blade server, see the BladeCenter documentation for SAN configuration details.

You must ensure the following configuration settings to restore the path if they are lost during an upgrade or other hardware maintenance that results in a node being offline:

- SCSI INQUIRY TIMEOUT

- SCSI COMMAND TIMEOUT
- Multipath settings in multipath.conf

Note: If you configure your system to use NPIV, there are fewer chances that the paths get lost.

Table 25. Parameter and recommended settings for linux hosts

Parameter	Recommended Settings
lpfc_max_luns	2048

- Use the following command to check the current value of maximum number of luns supported.

```
cat /sys/class/scsi_host/hostXX/lpfc_max_luns
```

  Note: The default value of maximum number of luns that are supported is 255.
- To configure the **lpfc_max_luns** parameter, you can change the `modprobe` configuration `/etc/modprobe.d/lpfc.conf` by using the following command:

```
options lpfc lpfc_max_luns=2048
```

  Then, recreate the `initramfs` and restart the host to reflect the configuration changes.

To set SCSI INQUIRY TIMEOUT, use these steps:

- Add **scsi_mod.inq_timeout=70** to the kernel boot command-line through *GRUB* configuration. By adding the **scsi_mod.inq_timeout=70** parameter, the change in the parameter is persistent from a server restart. Linux hosts can also regain system node paths when lost. This change can be done by completing the following steps.

  For SLES12 and SLES15 servers, follow these steps:

  1. To make the permanent change, edit `/etc/default/grub` and add to the `GRUB_CMDLINE_LINUX_DEFAULT` line:

```
scsi_mod.inq_timeout=70
```

  2. Run the following command to rewrite the boot record:

```
grub2-mkconfig -o /boot/grub2/grub.cfg
```

- For RHEL6, RHEL7, and RHEL8 servers, follow these steps:

  1. To make the change permanent, edit `/etc/sysconfig/grub` and add to the GRUB_CMDLINE_LINUX line:

```
scsi_mod.inq_timeout=70
```

  2. Run the following command to rewrite the boot record:

```
grub2-mkconfig -o /etc/grub2.cfg
```

  For RHEL9 and RHEL10 servers, follow these steps:

  1. To ensure the SCSI inquiry timeout survives subsequent reboots, append it to the kernel command line configuration by using GRUB.
  2. Add following parameter to all configured kernels:

```
bash
grubby --update-kernel=ALL --args="scsi_mod.inq_timeout=70"

Note: You can specify the specific kernel to update instead of ALL kernels
Verify the argument has been properly appended to the default boot configuration:

bash
grubby --info=DEFAULT
```

```
Look for scsi_mod.inq_timeout=70 within the args= line of the output.
```

The previous steps will be effective after rebooting. In RHEL6, RHEL7, RHEL8, RHEL9, RHEL10, SLES12, and SLES15, you can choose to change the **inq_timeout** parameter temporarily without rebooting. This method will not keep the parameter value persistent if the system ever restarts in the future and if you do not edit the GRUB configuration by following these steps.

Use the following command to change the **inq_timeout** parameter temporarily without rebooting:

```
echo 70 > /sys/module/scsi_mod/parameters/inq_timeout
```

In RHEL6, RHEL7, RHEL8, RHEL9, RHEL10, SLES12, and SLES15, enter the following command to view that the change was made:

```
systool -m scsi_mod -A inq_timeout
```

The output of the command shows that the value is changed to 70:

```
Module = "scsi_mod"
```

```
inq_timeout         = "70"
```

Note: It is best to perform both tasks in case the server is rebooted in the future.

To set the SCSI COMMAND TIMEOUT

- Set the *udev* rules for SCSI command timeout to 120s. This is the recommended setting for all versions of Linux.

Udev rules file creation

1. To increase the SCSI command timeout for the system, create the following udev rule:

```
# Set SCSI command timeout to 120s (default == 30 or 60) for IBM 2145 devices
SUBSYSTEM=="block", ACTION=="add", ENV{ID_VENDOR}=="IBM",ENV{ID_MODEL}=="2145", RUN+="/bin/sh -c 'echo 120 >/sys/block/%k/device/timeout'"
```

2. After you set up your volumes, confirm that they are set for 120 seconds. Locate the block device paths by running `multipath -ll | grep sd` from the command prompt. Then, run `cat /sys/block/sdX/device/timeout` (where X is each 2145 block device path).
3. To reload the udev rules without rebooting (or dynamically), you can run the following commands:

```
udevadm control -R
/sbin/udevadm trigger --type=devices --action=add
```

Settings for Linux hosts

To ensure path recovery in failover scenarios, certain Device Mapper Multipath (DMMP) settings and udev rules for the attachment of Linux hosts to the system are recommended. These settings are valid for IBM System x, all Intel or AMD-based servers, and Power® platforms.

You must restart your host after you complete the following two steps:

- Editing the multipath settings in `/etc/multipath.conf`
- Editing the udev rules for SCSI command timeout

For each Linux distribution and releases within a distribution, refer to the default settings under `[/usr/share/doc/device-mapper-multipath.*]` for Red Hat and `[/usr/share/doc/packages/multipath-tools]` for Novell SuSE. Ensure that the entries added to `multipath.conf` match the format and syntax for the required Linux distribution. Use the `multipath.conf` only from your related distribution and release. Do not copy the `multipath.conf` file from one distribution or release to another.

For some operating system levels, the `polling_interval` needs to be located under defaults instead of under **device settings**. If `polling_interval` is present in the device section, comment out `polling_interval` by using a # key.

For example:

```
Under Device Section
#         polling_interval 5,

Under Defaults Section
defaults {
        user_friendly_names yes
        polling_interval 5
}
```

Multipath settings for specific Linux distributions and releases

Edit `/etc/multipath.conf` with the following parameters and confirm the changes by entering:

```
multipathd -k
multipathd> show config
```

Note: Oracle Linux versions are the same as Red Hat Linux versions.

Red Hat Linux version 7.x, 8.x, and 9.x

```
     vendor "IBM"
     product "2145"
     path_grouping_policy "group_by_prio"
     path_selector "service-time 0" # Used by Red Hat 7.x
     prio "alua"
     path_checker "tur"
     failback "immediate"
     no_path_retry 5
     rr_weight uniform
     rr_min_io_rq "1"
     dev_loss_tmo 120
```

SUSE Linux Versions 12 and 15

```
     vendor "IBM"
     product "2145"
     path_grouping_policy "group_by_prio"
     path_selector "service-time 0"
     prio "alua"
     path_checker "tur"
     failback "immediate"
     retain_attached_hw_handler "yes"
     no_path_retry 5 # or no_path_retry "fail"
     fast_io_fail_tmo 5
     rr_min_io 1000
     rr_min_io_rq 1
     rr_weight "uniform"
```

Ubuntu

```
     vendor "IBM"
     product "2145"
     path_grouping_policy "group_by_prio"
     path_selector "service-time 0"
     prio "alua"
     path_checker "tur"
     failback "immediate"
     no_path_retry 5 # or no_path_retry "fail"
     retain_attached_hw_handler "yes"
     fast_io_fail_tmo 5
     rr_min_io 1000
     rr_min_io_rq 1
     rr_weight "uniform"
```

DM-MPIO for dev_loss_tmo

After a problem is detected on an FC port and it set to infinity, the SCSI layer can wait until 2147483647 seconds (68 years) before removing it from the system. The default value is determined by the OS.

All Linux hosts should have a **dev_loss_tmo** setting, but the value in seconds is how long to wait for the device/paths to be pruned. The suggested duration is around 120-150 seconds, but extended duration is also supported.

Care needs to be taken if it is too low since if paths are pruned, then they also need to be rediscovered and if too low, that may require manual rescan later. If inquiry timeout is right, the host should be able to re-add the paths when the SVC nodes are restored.

If the inquiry is too short such as 20 seconds, then the inquiry may timeout before the paths are ready.

Multipathing driver

If you lose paths and are not automatically restored, you can manually get them back with the following process.

- If you are using Linux *dm-multipath* as the `multipathing` software on the host, and have NPIV disabled on the SVC, it may be necessary to rescan ports after each node restores its paths.
- If you have NPIV disabled on the SVC, additional configuration is required. During the upgrade, an SVC with NPIV disabled shuts down the ports for an extended period of time, which may cause Linux to remove the ports.
- This setting may be applied to a running system, however it must also be applied to the GRUB configuration to ensure that it persists over restarts.
- It is possible that, even with this setting, an SVC upgrade may keep the ports down longer than the timeout setting allows. It may be necessary to rescan the ports once the SVC node restores operation.
- Use the **multipath** command to check path status once the SVC node completes the upgrade before you update the next node:

```
# multipath -ll /dev/mapper/mpatha
(360050768028211d8b000000000000061) dm-11 IBM,2145
size=10G features='1 queue_if_no_path' hwhandler='1 alua' wp=rw
|-+- policy='service-time 0' prio=50 status=enabled
| |- 4:0:1:0  sder    129:48  failed faulty running
| `- 5:0:1:0  sdfb    129:208 failed faulty running
`-+- policy='service-time 0' prio=10 status=active
  |- 4:0:0:0  sdem    128:224 active ready running
  `- 5:0:0:0  sdew    129:128 active ready running#
```

Choose a device name, instead of mpatha in the example, to match a device being utilized by the SVC that is performing the upgrade. The failed and faulty status, indicates that paths are still down for Linux. You can use **multipath -ll** command to list all devices name, and then scan the output for failed paths, confirming that failed paths are on devices that are associated with the SVC upgrade.

To rescan all SCSI targets and the ports on Red Hat Enterprise Linux or SUSE Linux Enterprise Server, use the **rescan-scsi-bus.sh** command, which is part of the `sg3_utils` package.

If your Linux distribution does not include the **rescan-scsi-bus.sh** command, use the **SCSI-rescan** command to rescans all the SCSI targets.

Hosts that run the Microsoft Windows Server operating system

You can attach hosts running the Microsoft Windows Server operating system via Fibre Channel.

For more information about host compatibility, see IBM System Storage Interoperation Center (SSIC).

Use the Microsoft Device Specific Module (MSDSM) multipathing driver for the Windows Server operating system. MSDSM is included in the Windows Server operating system and Hyper-V.

Refer to Windows documentation and the following pages for further configuration information.

Changing the disk timeout on Microsoft Windows Server

To change the disk timeout value on Windows Server operating systems, follow steps in the Windows Server host and in the registry browsing tool.

On your Windows Server hosts, change the **disk I/O timeout** value to `60` in the Windows registry, as follows:

1. In Windows, click **Start** > **Run**.
2. In the dialog text box, type `regedit` and click **Enter**.
3. In the registry browsing tool, locate the `HKEY_LOCAL_MACHINE\System\CurrentControlSet\Services\Disk\TimeOutValue` key.
4. Confirm that the value for the key is `60` (decimal value). If necessary, change the value to `60`.

Configuring the Brocade HBA for hosts that run a Windows Server operating system

After you install the Brocade host bus adapter (HBA) and the driver on a Microsoft Windows Server operating system host, you must configure the HBA.

For more information about supported Brocade models, see the following website:

www.ibm.com/support

To configure the HBA BIOS, either use the Brocade Host Connectivity Manager (HCM) software or reboot into the Brocade configuration tool. Configure the following settings:

- Host Adapter BIOS: Disabled (unless the machine is configured for SAN Boot)
- Queue depth: 4

You can disable the BIOS by using the command-line configuration tool:

- `bcu bios –disable 1/0`
- `bcu bios –disable 1/0`

You can set the queue depth by using the command-line configuration tool:

- `bcu fcpim –qdepth 1/04`
- `bcu fcpim –qdepth 1/14`

You can query the settings by using the command-line configuration tool:

- `bcu port –query 1/0`
- `bcu port –query 1/1`

Configuring the QLogic HBA for hosts that run the Windows Server operating system

After you install the QLogic HBA and the device driver on hosts that run a Microsoft Windows Server operating system, you must configure the HBA.

About this task

For more information about supported QLogic models, see the following IBM website:

www.ibm.com/support

To configure the HBA BIOS, either use the QLogic HBA manager software or reboot into the Fast!UTIL tool. Configure the following settings:

- Host Adapter BIOS: Disabled (unless the machine is configured for SAN Boot)
- Adapter Hard Loop ID: Disabled

- Connection Options: 1 - point to point only
- LUNs Per Target: 0
- Port Down Retry Count: 15

Set the execution throttle to a suitable queue depth for your environment, for example, a value of 100. For more information about the execution throttle and queue depth, see QLogic documentation.

If you are using Native Multipathing Plug-In (NMP), set Enable Target Reset to No. See Table 26 on page 130 to include the required parameters for the registry key.

Table 26. Registry key parameters for QLogic models

Key	Required parameters
HKEY_LOCAL_MACHINE > SYSTEM > CurrentControlSet > Services > ql2*xxx* > Parameters > Device > DriverParameters	**`Buschange=0;FixupInquiry=1`** **Note:** If you are using QLogic driver version 9.1.2.11 or later, **Buschange** cannot be set to zero. See your device driver documentation for details.

Configuring the Emulex HBA for hosts that run a Windows Server operating system

After you install the Emulex host bus adapter (HBA) and the driver on hosts that run a Microsoft Windows Server operating system, you must configure the HBA.

For the Emulex HBA StorPort driver, accept the default settings and set topology to 1 (1=F_Port Fabric). For the Emulex HBA FC Port driver, use the default settings and change the parameters that are given in Table 27 on page 130.

Note: The parameters that are shown in parentheses correspond to the parameters in HBAnywhere.

Table 27. Configuration file parameters for the Emulex HBA

Parameters	Recommended Settings
Query name server for all N-ports (BrokenRSCN)	Enabled
LUN mapping (MapLuns)	Enabled (1)
Automatic LUN mapping (MapLuns)	Enabled (1)
Allow multiple paths to SCSI targets (MultipleSCSIClaims).	Enabled
Scan in device ID order (ScanDeviceIDOrder)	Disabled
Translate queue full to busy (TranslateQueueFull).	Enabled
Retry timer (RetryTimer)	2000 milliseconds
Maximum number of LUNs (MaximumLun)	Equal to or greater than the number of the system LUNs that are available to the HBA.

Configuring Microsoft Device Specific Module (MSDSM)

Use the Microsoft Device Specific Module (MSDSM) multipathing driver for the Windows Server operating system. MSDSM is included in the Windows Server operating system and Hyper-V.

The Microsoft Multipath I/O (MPIO) feature must also be enabled when using MSDSM .

Complete the following steps to install MSDSM on the Windows Server operating system.

1. To use MSDSM, install the Microsoft Multipath I/O (MPIO) feature by using the **Add Roles and Features** wizard in the Windows Server Manager.

2. Open the **MPIO Properties** window from the MPIO application in Windows Administrative Tools.
3. Select **IBM 2145** on the **Discover Multi-Paths** tab in the **MPIO Properties** window and click **Add**.
4. Click **OK** to restart the host when prompted. Wait for the host to restart.
5. Open Windows PowerShell and enter the **`mpclaim -s -d`** command.

   Ensure that MSDSM is installed correctly and can claim the volumes correctly. Use the default MSDSM settings.

Migrating existing SAN boot images

If you have a host that runs the Microsoft Windows Server operating system and existing SAN boot images that are controlled by storage controllers, you can migrate these images to image-mode volumes that are controlled by the system.

About this task

To migrate your existing SAN boot images, complete the following steps:

Procedure

1. Shut down the host.
2. Complete the following configuration changes on the storage controller:
   a) Remove all the image-to-host mappings from the storage controller.
   b) Map the existing SAN boot image and any other disks to the system.
3. Zone one port of each host bus adapter (HBA) to one of the system ports that is associated with the I/O group for the target image-mode volume.
4. Complete the following configuration changes on the system:
   a) Create an image-mode volume for the managed disk (MDisk) that contains the SAN boot image.

      Use the MDisk unique identifier to specify the correct MDisk.
   b) Create a host object and assign it to the HBA port that you zoned to the system port in step “3” on page 131.
   c) Map the image mode volume to the host.

      For example, you might map the boot disk to the host with SCSI LUN ID 0.
   d) Map the swap disk to the host, if required.

      For example, you might map the swap disk to the host with SCSI LUN ID 1.
5. Change the boot address of the host by completing the following steps:
   a) Restart the host and open the BIOS utility of the host during the booting process.
   b) Set the BIOS settings on the host to find the boot image at the worldwide port name (WWPN) of the node that is zoned to the HBA port.
6. Boot the host in single-path mode.
7. Uninstall any multipathing driver that is not supported for system hosts that run the applicable Windows Server operating system.
8. Install a supported multipathing driver.
9. Restart the host in single-path mode to ensure that the supported multipath driver was properly installed.
10. Zone each HBA port to one port on each system node.
11. Add HBA ports to the host object that you created in step “4.b” on page 131.
12. Configure the HBA settings on the host by using the following steps:
    a) Restart the host and open the host's BIOS utility during the booting process.
    b) Ensure that all HBA ports are boot-enabled and can see both nodes in the I/O group that contains the SAN boot image. Configure the HBA ports for redundant paths.

c) Exit the BIOS utility and finish booting the host.

13. Map any additional volume to the host as required.

Hosts that run the VMware ESXi operating system

Fibre Channel attachments between the system and various guest host operating systems that run on the ESXi operating system have specific requirements and other guidelines that need to be followed.

Note: The system supports VMware vSphere Virtual Volumes that allow VMware vCenter to manage system objects like volumes and pools. Virtual Volumes provide storage that is used by specific virtual servers, and also correspond to single volumes on the system. The system administrator can assign ownership to VMware administrators to create the pools and volumes that are used exclusively for Virtual Volumes. For more information, see the information about configuring Virtual Volumes.

Configuring the ESXi operating system

You must configure the VMware ESXi operating system before you can connect hosts that run VMware ESXi with the system.

Before you begin

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

- Your system must be installed.
- You must install the appropriate host bus adapters and ensure that the latest drivers and firmware are applied.

Procedure

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

1. Define the host object by using the worldwide port name identifiers (WWPNs) for each physical ESXi host in the VMware vSphere cluster.
2. If more than one physical ESXi host is present in the vSphere cluster, create a host cluster and add each ESXi host object as a member.

   When you map volumes to ESXi hosts, ensure that you create shared host-cluster mappings for all shared storage to ensure consistent SCSI IDs across the ESXi hosts for the same volume. If the ESXi host is SAN booted from a volume on the system, use private mappings for these volumes.
3. Review the multipath settings by running the following CLI command:

   ```
   # esxcli storage nmp device list -d naa.<UUID of volume>
   ```

   The output is similar to the following example:

   ```
   naa.600507640084011070000000000008c
      Device Display Name: IBM Fibre Channel Disk (naa.600507640084011070000000000008c)
      Storage Array Type: VMW_SATP_ALUA
      Storage Array Type Device Config: {implicit_support=on; explicit_support=off;
   explicit_allow=on; alua_followover=on; action_OnRetryErrors=on; {TPG_id=16,TPG_state=AO}
   {TPG_id=17,TPG_state=ANO}}
      Path Selection Policy: VMW_PSP_RR
      Path Selection Policy Device Config: {policy=rr,iops=1,bytes=10485760,useANO=0;
   lastPathIndex=1: NumIOsPending=0,numBytesPending=0}
      Path Selection Policy Device Custom Config:
      Working Paths: vmhba4:C0:T0:L250, vmhba3:C0:T0:L250
   ```

   Verify the following settings:

   a. The Storage Array type plug-in for IBM volumes is set to VMW_SATP_ALUA.

   b. The path selection policy is set to RoundRobin.

c. The Round Robin IOPS is set to 1 to evenly distribute I/Os across as many ports on the system as possible. For information about how to make this change, see the VMware information at this website: Adjusting Round Robin IOPS limit from default 1000 to 1 (2069356) .

4. Enable ATS hardware accelerated locking and set to ATS-only public.To verify the current settings, use the following CLI command:

```
# vmkfstools -Ph -v1 /vmfs/volumes/<datastore_name> | grep Mode
Mode: public ATS-only
```

If the datastore Mode is not reporting public ATS-only, consider either of the following two options:

a. Creating a LUN and creating a VMFS5 or VMFS6 datastore and use Storage vMotion to migrate virtual machines to the new datastore.

b. Enforcing ATS-only on all affected datastores. This process requires stopping all I/O to the datastore first. For more information, see Change VMFS Locking to ATS-Only. This CLI command provides the UUID of the device:

```
# vmkfstools -Ph -v1 /vmfs/volumes/<datastore_name>  | grep -A 1 Partitions
Partitions spanned (on "lvm"):        naa.60050764008401107000000000000008b:1
```

Using the UUID of the device, the following CLI commands can be used to change the Lock Mode to ATS-only:

```
# vmkfstools --configATSOnly 1 /vmfs/devices/disks/naa.60050764008401107000000000000008b:1
```

5. Set the XCopy transfer size to 4096.

```
# Get-VMHost  | Get-AdvancedSetting -Name DataMover.MaxHWTransferSize | select Entity,
name, value


Entity                         Name                         Value
------                         ----                         -----
vmlab11c2.ssd.hursley.ibm.com DataMover.MaxHWTransferSize  4096
```

6. Use VAAI, which is a set of SCSI commands that are used by VMware ESXi hosts to offload operations to the storage array. Enable **HardwareAcceleratedMove**, **HardwareAcceleratedLocking**, and **HardwareAcceleratedInit**.

```
# Get-VMHost  | Get-AdvancedSetting -name *HardwareAccelerated* | select Name, value
Name                              Value
----                              -----
DataMover.HardwareAcceleratedMove     1
VMFS3.HardwareAcceleratedLocking      1
DataMover.HardwareAcceleratedInit     1
```

7. Wherever possible, set the guest virtual machines to have SCSI timeouts set to at least 60 seconds although 120 seconds is recommended. Refer to the host vendor documentation for instructions on which settings to change.

8. The number of paths per volume is recommended not to exceed 8.

However, during operations such as a nondisruptive volume move, if the Fibre Channel target port mode is transitional for the I/O group, up to 16 paths per volume is supported. See the VMware documentation for the maximum number of devices and paths for the version of VMware ESXi that you are using.

9. If you use Datastore Clusters where vSphere Storage vMotion is used to migrate virtual machines based on I/O latency, it is important to understand the system requirements.

a. Ensure that each LUN has independent and isolated performance characteristics:

i) Where possible, present each volume from a different pool, especially when you use arrays or managed disks that are created from traditional spinning (non-SSD or flash) media.

ii) Review the preferred node on each LUN in the Datastore Cluster to ensure an even distribution between the nodes within the I/O group.

b. Given VMs can migrate between datastores within the Datastore Cluster. Be aware of dynamic I/O workloads and how Easy Tier will respond. Do one of the following tasks:

i) Disable automatic storage migrations when you use Easy Tier, or

ii) Use a single tier of storage between all LUNs that are used in a Datastore Cluster.

10. If you have dual-site topologies, consider the following requirements:

- If the system topology is set to `stretched`, consider using Host and Virtual Machine DRS Groups to restrict a given VM to move only to other hosts within the same site.
- VMware ESXi host objects should have a site set to avoid unnecessary intersite traffic.

VMware ATS heartbeating

VMware ESXi uses the SCSI Compare and Write command (VMware refers to this command as Atomic Test and Set - ATS) to "heartbeat" periodically to datastores.

Note: The use of ATS heartbeating is not supported on ESXi hosts that run version 5.5 update 2 or later or ESXi version 6.0 before update 3.

Due to the low timeout value for heartbeat I/O that uses ATS, host disconnects can occur if delays of 8 seconds or longer are experienced in completing individual heartbeat I/Os on backend storage systems or the SAN infrastructure. Under certain circumstances, such as if the ATS heartbeats are ended due to delays of 8 seconds or more, the ESXi host might retry the ATS heartbeat command and it might fail with a miscompare or lost access to datastore message that is logged in the vmkernel.log on the host. In many cases, ATS heartbeat miscomparisons during verification and lost access to datastore messages are results of ATS heartbeats that work as expected when I/O delays are present.

ATS heartbeating increases the load on the system and can lead to access issues on busy systems, particularly during maintenance procedures. To reduce this load, ATS heartbeats can be disabled.

For VMFS5 and VMFS6 datastores, run these commands:

- To disable ATS heartbeats, run the following CLI command:

```
# esxcli system settings advanced set -i 0 -o /VMFS3/UseATSForHBOnVMFS5
```

- To enable ATS heartbeats, run the following CLI command:

```
# esxcli system settings advanced set -i 1 -o /VMFS3/UseATSForHBOnVMFS5
```

For VMFS3 datastores, run the following command to disable ATS heartbeats:

```
# esxcli system settings advanced set -i 0 -o /VMFS3/useATSForHBOnVMFS3
```

SAN boot support for hosts running ESXi operating systems

The system can be used as a boot device for the ESXi operating system.

For SAN boot support for the ESXi operating system, you must meet the following requirement:

- The ESXi operating system must be on a SAN disk.

See the software restrictions page on the following IBM support website for any other restrictions for SAN boot support: www.ibm.com/support.

Hosts that run the Microsoft Hyper-V operating system

You can attach hosts running the Microsoft Hyper-V operating system via Fibre Channel. This allows you to run multiple guest operating systems on the server.

For more information about host compatibility, see IBM System Storage Interoperation Center (SSIC).

Refer to the Windows Documentation and the following pages for further configuration information.

Configuring the QLogic HBA for a host that runs the Microsoft Hyper-V operating system

After you install the QLogic host bus adapter (HBA) and the device driver on hosts that are running the Microsoft Hyper-V operating system, you must configure the HBA.

About this task

To configure the HBA BIOS, either use the QLogic HBA manager software or reboot into the Fast!UTIL tool. Configure the following settings:

- **Host Adapter BIOS:** Disabled (unless the machine is configured for SAN Boot)
- **Adapter Hard Loop ID:** Disabled
- **Connection Options:** 1 - point to point only
- **LUNs Per Target:** 0
- **Port Down Retry Count:** 15

Set the execution throttle to a suitable queue depth for your environment, for example, a value of 100. For more information about the execution throttle and queue depth, see QLogic documentation.

If you are using Native Multipathing Plug-In (NMP), set **Enable Target Reset** to No. See Table 28 on page 135 to include the required parameters for the registry key.

Table 28. Registry key parameters for QLogic models

Key	Required parameters
HKEY_LOCAL_MACHINE > SYSTEM > CurrentControlSet > Services > ql2*xxx* > Parameters > Device > DriverParameters	**Buschange=0;FixupInquiry=1** **Note:** If you are using QLogic driver version 9.1.2.11 or higher, **Buschange** cannot be set to zero. See your device driver documentation for details.

Configuring the Emulex HBA for hosts that run the Microsoft Hyper-V operating system

After you install the Emulex host bus adapter (HBA) and the driver on hosts that are running the Microsoft Hyper-V operating system, you must configure the host bus adapter.

Before you begin

For the Emulex HBA StorPort driver, accept the default settings and set topology to 1 (1=F_Port Fabric).

Using Microsoft ODX on Microsoft Hyper-V hosts

Ensure that you are familiar with the Microsoft Hyper-V host attachment requirements to verify and troubleshoot your configuration.

ODX is disabled by default so that you can enable the MPIO feature. After enabling the MPIO feature, you can enable the ODX by entering the **chsystem -odx on** CLI command.

Although Microsoft ODX is supported, the volumes that are mapped to the host are not *ODX capable* in Windows until they are unmapped and then mapped again.

To unmap and map the volumes, follow these steps:

1. Bring down the applications by using the volumes so that the applications stop using the volumes.
2. Disconnect or logout the host from the storage.

3. Rediscover the volumes from the Server Manager (the volumes go away).
4. Reconnect or log in the host to the storage.
5. Rediscover the volumes from the Server Manager (the disks reappear).
6. Start the applications.

Note: If these LUNs are formatted with an allocation unit size that is less than 4 KB, ODX might not work on these LUNs. A subpage-aligned I/O (less than 4 KB) is not allowed.

You can add new LUNs to a Windows host for consumption of an ODX-aware application. The system array works best with ODX when LUNs that are exported to the Windows host are formatted with 32 KB NTFS cluster size. Therefore, ensure that the allocation unit size for the New Technology File System (NTFS) is greater than or equal to 32 KB. The default on the Windows Server 2012 release 2 is 4 KB, which might not result in the performance that the system can deliver.

Citrix XenServer hosts

You can attach Citrix XenServer hosts via Fibre Channel. .

Refer to Citrix documentation at http://www.citrix.com and the following page for further configuration information.

Note: From the **Citrix** home page, follow this path: **Products & Solutions** > **XenServer/Essentials** > **Essentials for XenServer** > **Specifications**.

Configuring the Citrix XenServer host

You must configure the Citrix XenServer host before you can use the host with the system.

The following Citrix XenServer website provides the most current information about configuring the Citrix XenServer:

www.citrix.com

Linux operating system

Linux virtual machines (VMs) are always supported in a paravirtualized mode.

Windows operating system

Windows VMs are always supported as hardware-based virtual machines (HVMs) with paravirtualized drivers for storage and network. The drivers are installed from the Citrix XenServer installation CDs.

If you plan to use Windows VMs running under Citrix XenServer, ensure that the server processors have hardware virtualization support, such as Intel Virtualization Technology (VT)-based processors or AMD-V-based processors.

Multipath support for Citrix XenServer hosts

The multipath driver is installed as part of the Citrix XenServer software.

The following website provides the most current information about the IBM configuration file (`multipath.conf`) for the system:

```
www.ibm.com/support
```

To find the appropriate configuration file, complete the following steps:

1. Search for the following term: `subsystem device driver for Linux`
2. Scroll down the list until you find the device-mapper multipath configuration file for the system.
3. Select and download the file for Red Hat Enterprise Linux 5.

Clustering support on Citrix XenServer hosts

Clustering for guest operating systems either within or across Citrix XenServer hosts, or clustering of Citrix XenServer hosts is not supported.

Zoning and portsets

Ensure that you are familiar with these zoning details and host zones. More details are included in the SAN configuration and zoning rules summary.

IBM FlashSystem auto zoning feature provides a complete solution for Fibre Channel zone management that integrates seamlessly with the Brocade Target Driven Zoning (TDZ) technology. This feature is designed to address a wide range of challenges typically faced by SAN administrators, including:

- Creating, modifying, and deleting zones within the Fibre Channel fabric
- Identifying optimal zone definitions to achieve optimal multi-path configurations for host servers
- Managing zoning in accordance with SAN best practices
- Backing up and restoring zone configurations to ensure reliability and recoverability
- Automatically adapting zone definitions by tracking changes in host or storage connectivity

  With auto zoning feature, storage target ports manage zoning operations on the FC fabric. Auto zoning creates peer zones, which provide an optimized zoning model:

  Peer zone contents

  - One principal device is typically a storage array target port.
  - Multiple non-principal devices are typically host initiator ports.

  Peer zone communication rules

  - The principal device can communicate with all nonprincipal devices, and non-principal devices can communicate with the principal device.
  - Nonprincipal devices cannot communicate with each other.

  Peer zones replace traditional single-initiator zoning and maintain consistent and predictable zoning patterns.

Supported Switch Platforms

The auto zoning feature is supported by Brocade FC switches that meet the following conditions:

- Fabric OS (FOS) version 9.1.1 or later
- The TDZ feature must be enabled on the switch ports where the storage target ports are connected. The feature can be enabled through the switch CLI. Refer to the switch documentation for detailed steps to enable this capability.
- The TDZ feature does not need to be enabled on host ports.

The TDZ feature must be explicitly enabled by the fabric administrator for each relevant port. The auto zoning feature initiates automatic zoning actions only after the TDZ is enabled.

Fabric Connectivity Requirements

For auto zoning to operate correctly and predictably, the fabric connectivity of the storage target ports and connected host ports must meet specific requirements. These considerations can ensure consistent zone formation, proper path distribution, and stable host-storage communication across fabrics.

- Auto zoning is supported only for user-defined FC portsets.
  - All FC I/O port IDs included in an auto zoning enabled portset must demonstrate symmetric fabric connectivity. The same FC I/O port ID from every node in the cluster must be connected to the same fabric. For example, FC IO Port ID 1 from all nodes must connect to Fabric A.
  - Ideally, a portset should contain an equal number of target FC I/O ports from each redundant fabric.

- Host initiator ports must be connected to the same fabric as the target ports defined within the portset that you want to associate with host object.

Auto zoning creates peer zones between host initiator ports and the target ports that share fabric.

Auto zoning Workflow

To start using auto zoning when adding new hosts in an existing configuration:

1. Ensure that fabric requirements are met as described above.
2. Create an auto zoning enabled Fibre Channel portset.
3. Create the host by using the FDMI hostname or by specifying a list of host port WWPNs.
4. Associate host with the auto zoning enabled portset.

Zoning policies

Auto zoning supports configurable zoning policies that determine how initiator WWPNs are matched with target port WWPNs when forming peer zones.

Zoning policies are applied at the portset level, and different portsets can use different policies based on deployment requirements.

Available zoning policies are:

one_to_one connectivity

This policy enforces a single path between every initiator port and one target port on each canister within the same fabric.

- Each initiator WWPN is zoned with only one FC IO Port ID from the portset. Each FC IO Port ID represents a single target port from each canister.
- Zoning is performed between the initiator port and the portset target port, provided both ports are within the same fabric.
- Initiator ports are evenly distributed across all target ports from the portset to achieve better workload distribution.
- After zoning, that initiator WWPN is not zoned with any additional target ports from the portset, even if multiple suitable target ports exist on that fabric.

This policy is optimal for environments requiring strict path control or host configurations where minimal connectivity is preferred.

one_to_one_all_fabrics

This policy also creates one-to-one zoning, but it is designed for host environments where initiator WWPNs can not be visible on all fabrics during initial discovery.

- Each initiator WWPN is zoned with one target port WWPN from each fabric but will have a single host path from fabric where Host port is connected .
- This policy is useful where initiator WWPNs do not appear upfront in the fabric
  - SAN-boot scenarios
  - NPIV (N_Port ID Virtualization) WWPN configurations

This policy ensures reliable boot-time zoning even when initiator WWPN visibility is limited or dynamic.

Zone Naming Convention

Zones are created that uses the following format, `IBM_SV_<Target Port WWPN>_00`, for example, IBM_SV_50050768101655c4_00.

This predictable naming scheme allows easy identification of IBM Storage managed zones in the switch fabric.

Zone Name Prefix Configuration

Storage administrator has an option to specify a zone name prefix for all automatically created zones. This prefix can be configured by using the chsystem command.

- The zone name prefix must be configured before the first auto zoning enabled portset is created on the system.
- Once the prefix is set on a system that contains any auto-zone-enabled portset, the prefix cannot be modified.
- If no prefix is provided, the default naming scheme is used, for example, `IBM_SV_<target-WWPN>_00`.

This mechanism ensures naming consistency across all auto-generated zones in the fabric.

If the user configures the prefix FC_Infinity, it is automatically prepended to the zone name. For example, if **`Prefix`** is FC_Infinity and **`Target Port WWPN`** is 50050768101655c4, then the resulting auto-generated zone name becomes: FC_Infinity_IBM_SV_50050768101655c4_00.

This example demonstrates how user-defined prefixes are applied to the default zone naming format.

For information on auto zoning, see .

Paths to hosts

The total supported number of paths for a host to a volume must not exceed 4 per node canister, for a total of 8 per system. If high availability is configured between systems, then the limit remains 4 per node canister for a total of 16 paths across both systems. Configurations in which this number is exceeded are not supported.

Portset avoids wrong zoning configuration and limits the number of paths though which hosts can access external storage devices. Zoning host ports to too many ports of external system can create redundant paths. Also, unnecessary host login to too many Fibre Channel ports of external system consumes resources and leads to non-uniform distribution of host login on ports. Portsets limit the number of paths that the host can access to external storage system and distributes the host login evenly across multiple ports. For more information on Fibre Channel portsets, refer to the “Portsets” on page 333.

Fibre Channel portset is mapped with Fibre Channel port and host object. The host port WWPNs must be zoned only to external system Fibre Channel ports that are associated with the portset. If host ports are zoned to any other Fibre Channel port, the system notifies non-porset member login. Such logins can be removed by switching zones. An administrator can identify such logins by using the **`lsfabric`** command. The **`mkhost`** command shows host login counts on ports and enable the administrator to identify less loaded portsets. The command also shows the WWPNs that an administrator can use with host ports for zoning.

Adding WWPN to portsets provides greater flexibility and can be used for specific functional use cases. For example, SCSI host WWPN on a Fibre Channel port can be added to Portset1, whereas NVMeFC host WWPN on that port can be added to different portset.

To find the worldwide port names (WWPNs) that are required to set up Fibre Channel zoning with hosts, use the **`lstargetportfc`** command. This command also displays the current failover status of host I/O ports.

Portset avoids wrong zoning configuration and limits the number of paths though which hosts can access external storage devices. Zoning host ports to too many ports of external system can create redundant paths. Also, unnecessary host login to too many Fibre Channel ports of external system consumes resources and leads to non-uniform distribution of host login on ports. Portsets limit the number of paths that the host can access to external storage system and distributes the host login evenly across multiple ports. For more information on Fibre Channel portsets, refer to the **`Fibre channel portset`**.

Host zones

The configuration rules for host zones are different depending upon the number of hosts that access the system. For configurations of fewer than 64 hosts per system, the system supports a simple set of zoning rules that enable a small set of host zones to be created for different environments. For configurations of more than 64 hosts per system, the system supports a more restrictive set of host zoning rules.

Zoning that contains host HBAs must ensure host HBAs in dissimilar hosts or dissimilar HBAs are in separate zones. Dissimilar hosts means that the hosts are running different operating systems or are different hardware products; thus different levels of the same operating system are regarded as similar.

You can map each volume that you create on the system to multiple HBA Fibre Channel ports in a specific host. There can also be multiple paths across the SAN. For these reasons, each host must run multipathing software, such as the Microsoft Device Specific Module (MSDSM). The multipathing software manages the many paths that are available to the volume and presents a single storage device to the operating system. The system supports various multipathing software packages. The specific multipathing software that the system supports depends on the host operating system that you use the software with.

The number of paths through the network from the system nodes to a host must not exceed eight. Configurations that exceed this number are unsupported.

- Each system node has multiple ports and each I/O group has two system nodes. Therefore, without any zoning, the number of paths to a volume is the number of system ports in the I/O group times the number of host ports.
- This rule exists to limit the number of paths that must be resolved by the multipathing device driver.

If you want to restrict the number of paths to a host, zone the switches so that each HBA port is zoned with one system port for each node in the clustered system. If a host has multiple HBA ports, zone each port to a different set of system ports to maximize performance and redundancy.

To obtain the best overall performance of the system and to prevent overloading, the workload to each port must be equal. This can typically involve zoning approximately the same number of host Fibre Channel ports to each Fibre Channel port.

Systems with fewer than 64 hosts

For systems with fewer than 64 hosts that are attached, zones that contain host HBAs must contain no more than 40 initiators, including the ports that act as initiators. A configuration that exceeds 40 initiators is not supported. A valid zone can be 32 host ports plus 8 ports. When it is possible, place each HBA port in a host that connects to a node into a separate zone. Include exactly one port from each node in the I/O groups that are associated with this host. This type of host zoning is not mandatory, but is preferred for smaller configurations.

Note: If the switch vendor recommends fewer ports per zone for a particular SAN, the rules that are imposed by the vendor take precedence over system rules.

Systems with more than 64 hosts

Each HBA port must be in a separate zone and each zone must contain exactly one port from each node in each I/O group that the host accesses.

Fibre Channel port-name identification

The locations of port names are identified for Fibre Channel connections.

About Fibre Channel port-name identification

The format and content of the Fibre Channel (FC) port identifier are determined by the manufacturer of the link control facility for the applicable Fibre Channel port. The identifier is an eight-byte field, which the Fibre Channel protocols use to uniquely identify the Fibre Channel port.

The worldwide port name (WWPN) consists of 16 hexadecimal characters (0 - 9 and A - F). The system uses the WWPN to uniquely identify the Fibre Channel host interface adapter that is installed in your host.

The system automatically finds the WWPN for your host Fibre Channel host interface adapter when you attach your host to the system.

Note: Each Fibre Channel port has a unique WWPN. There are 1 or more ports per host interface adapter, and therefore 1 or more WWPNs per host interface adapter. Each host interface adapter port has a unique WWPN that must be added to the host object.

Locating the WWPN for an IBM System p, eServer, or an RS/6000 AIX host

You can locate the WWPN for an IBM System p, eServer, or an RS/6000 AIX host by completing the following steps:

1. Log in as root.
2. Type `lscfg -vl fcs`*x*, where *x* is the adapter number.

   The network address is the Fibre Channel adapter port WWPN value.

Locating the WWPN for a host that is running the Linux operating system

You can locate the WWPN for a host that is running the Linux operating system with a QLogic adapter by completing the following steps.

Procedure

1. Restart the server.
2. Press Alt+Q to get the **FAST!Util** menu.

   If you have more than one Fibre Channel host bus adapter (HBA) installed, all the Fibre Channel HBAs are displayed. Scroll down to the adapter you want. Press Enter.
3. From the **FAST!Util** menu, scroll down and select **Select Host Adapter**.
4. Scroll up and highlight **Configuration Settings**. Press Enter.
5. From the **Configuration Settings** menu, click **Host Adapter Settings**.
6. Write down the 16-digit alphanumeric string that is displayed.

Locating the WWPN for a host running the Microsoft Windows operating system

Determining the WWPN of a host that runs a Windows operating system depends on the type of HBA in your host server.

For QLogic, you can use the SANsurfer GUI/IBM FAStT Management Suite Java™ (MSJ) if you have it, or restart the host and enter Ctrl+Q to open the QLogic BIOS, where you can find the HBA WWPNs.

For Emulex hosts, use the elxcfg tool that is packaged with the firmware. This tool opens in the Windows operating system and does not require a restart.

Locating the WWPN for an Oracle SPARC host

You can locate the worldwide port name (WWPN) for an Oracle SPARC host by following the steps in this topic.

Procedure

1. After you install the adapter and you restart the host system, view the `/var/adm/messages` file.
2. Search for the line that contains the applicable phrase for your host bus adapter (HBA):
   a. For the JNI SBUS HBA, search for `fcaw`*x*`: Fibre Channel WWNN`, where *x* is the adapter number (0, 1, and so on). You can find the WWPN on the same line immediately after the worldwide node name (WWNN).
   b. For the JNI PCI HBA, search for `fca-pci`*x*`: Fibre Channel WWNN`, where *x* is the adapter number (0, 1, and so on). You can find the WWPN on the same line as the WWNN.

c. For the QLogic QLA2200F HBA, search for `qla2200-hba`*x*`-adapter-port-name` where *x* is the adapter number (0, 1, and so on).

Locating the WWPNs for a host that is running a VMware operating system

You can locate the worldwide port names (WWPNs) for a host that is running a VMware operating system.

Procedure

To locate the WWPNs for the host, complete the following steps:

1. Open the VMware Management Interface and click the **Options** tab.
2. Select **Storage Management**.
3. In the new window, click the **Adapter bindings** tab.

   The WWPN is listed at the end of each port heading line, which is the lines that start with `vmhba`. For example, in the following line, **`21:00:00:E0:8B:1A:E4:C6`** is the WWPN of the HBA port:

```
vmhba0: QLogic Corp QLA231x/2340 (rev 02) (21:00:00:E0:8B:1A:E4:C6)
```

Locating the WWPN for an SGI Origin host

You can locate the WWPN for an SGI Origin host that is running the IRIX operating system with a QLogic adapter by following the steps.

1. Restart the server.
2. Type the `scsiha -w [bus_number | device]` command.

   For example, type `scsiha -w 6 7 8 9`. shows an example of the command output.

```
# scsiha -w 6 7 8 9
6 Portname: 210000e08b05d207
7 Portname: 210000e08b04d539
8 Portname: 210000e08b050808
9 Portname: 210000e08b038fe6
#
```

Queue depth in Fibre Channel hosts

In Fibre Channel networks, the queue depth is the number of I/O operations (SCSI commands) that can be run in parallel on a device.

Usually a higher queue depth can lead to better performance. However, when too many concurrent I/Os are sent to a storage device, the device responds with an I/O failure message of queue full. This message is intended to cause the host to try the I/O again a short time later. However, not all operating systems correctly handle the queue-full failure, which can cause unnecessary I/O failures or delays to applications.

To avoid these delays or failures, configure your hosts so that they send only a finite number of I/Os to the storage system to avoid exhausting the resources of the storage system.

In configurations with multiple hosts, configure each host with a similar queue depth to maintain fairness between the hosts.

- For most hosts, set the HBA queue depth to 32.
- For hosts that are significantly busier or where not many hosts are configured on the storage system, set the HBA queue depth to 128.

Ensure that the queue depth total for all hosts that are connected to a single physical Fibre Channel port on the storage system does not exceed 2048.

If host I/O queue full failures regularly occur, consider reducing the queue depths or distributing the host zoning across more physical Fibre Channel ports on the storage system.

iSCSI Ethernet host attachment

iSCSI Ethernet hosts with different operating systems can be attached to the system.

If you are configuring a host object on the system that uses iSCSI connections, ensure that you complete the necessary host-system configurations and configure the system for iSCSI connections. For iSCSI attached hosts, you can also configure CHAP authentication settings, such as usernames and passwords.

For more information about host compatibility, see IBM System Storage Interoperation Center (SSIC).

iSCSI overview

iSCSI is an IP-based standard for transferring data that supports host access by carrying SCSI commands over IP networks. The iSCSI standard is defined by RFC 3720.

Connections from iSCSI-attached hosts to system nodes are supported. iSCSI connections from system nodes to IBM Storage FlashSystem, IBM Spectrum Accelerate, and Dell EqualLogic systems that are used as external storage systems are also supported. If the external storage is not in this list, then raise SCORE approval request.

Table 29 on page 143 shows that iSCSI and Fibre Channel terms have analogous components.

Table 29. Comparison of iSCSI and Fibre Channel components

iSCSI components	**Fibre Channel components**
iSCSI host bus adapter	Fibre Channel host bus adapter
Network interface controller (NIC) and iSCSI software initiator	Fibre Channel host bus adapter
IP switch	Fibre Channel switch
IP router	–
iSCSI name, such as IQN (iSCSI qualified name) or EUI (extended-unique identifier)	WWNN (worldwide node name)

iSCSI initiators and targets

In an iSCSI configuration, the iSCSI host or server sends requests to a node. The host contains one or more initiators that attach to an IP network to initiate requests to and receive responses from an iSCSI target. Each initiator and target are given a unique iSCSI name such as an iSCSI qualified name (IQN) or an extended-unique identifier (EUI). An IQN is a 223-byte ASCII name. An EUI is a 64-bit identifier. An iSCSI name represents a worldwide unique naming scheme. This scheme identifies each initiator or target in the same way that worldwide node names (WWNNs) are used to identify devices in a Fibre Channel fabric.

To transport SCSI commands over the IP network, an iSCSI driver must be installed on the iSCSI host. The driver is used to send iSCSI commands and responses through a network interface controller (NIC) or an iSCSI HBA in the host.

iSCSI host connection options

Figure 12 on page 144 shows an iSCSI host that connects to the system over an Ethernet network.

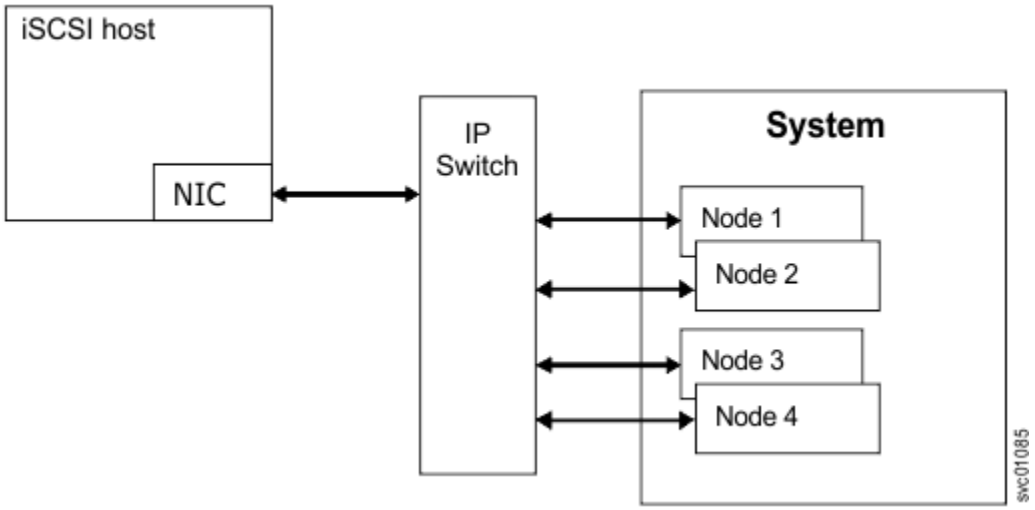


Figure 12. Transmitting SCSI over TCP/IP

iSCSI multisession support

Multisession support, or host-level multipath support, provides multiple paths between iSCSI initiators and targets, which is useful for high availability and load balancing.

Some hosts require that an iSCSI initiator login with a specific IP address of the target versus logging in with a target name. Logging in with a target name in those hosts causes the iSCSI initiator to log in to all of the IP addresses in the target and causes session reinstantiation that results in session loss of the previous login.

The system offers single subnet multi-session configurations and multiple and dual subnet multisession configurations that remove the restriction against logging in to a target name. iSCSI initiators use Internet storage name service (iSNS) servers to log in, which enables the discovery of the iSCSI targets and their IP addresses.

- An iSCSI session is a TCP relationship between an iSCSI initiator node port and an iSCSI target node port. When established, iSCSI control, data, and status messages are communicated over the session.
- Each session is identified with multiple iSCSI initiator names and the target name plus the iSCSI initiator session ID (ISID) plus the target portal group tag (target-side ID).
- The system multi-session approach is to have relationships between multiple initiator instances (separate ISIDs) to a single target portal group in an iSCSI target. This approach is similar to having multiple SCSI initiator ports from the same iSCSI initiator device have relationships with a single SCSI target port.

iSCSI storage support

The system supports iSCSI connections to IBM Storage FlashSystem and Dell EqualLogic systems that are used as external storage systems. For specific configuration information, see “iSCSI configuration details for host connections” on page 146.

Discovery and path configuration for iSCSI controllers is done as follows:

- Specify only a target IP and issue a discovery request to the iSCSI target controller
- The target controller returns a list of all configured iSCSI qualified names (IQNs) on that controller
- Establish connections to the discovered target IQNs

Example iSCSI configurations

Configure your system in a network that uses Internet Small Computer System Interface (iSCSI).

This figure shows a two-node clustered system that is connected to a single subnet. Each node has two Ethernet ports, each of which is used for iSCSI data transfers. One node in the system also acts as the system configuration node. In this example, port 1 on the configuration node provides the system management IP interface.

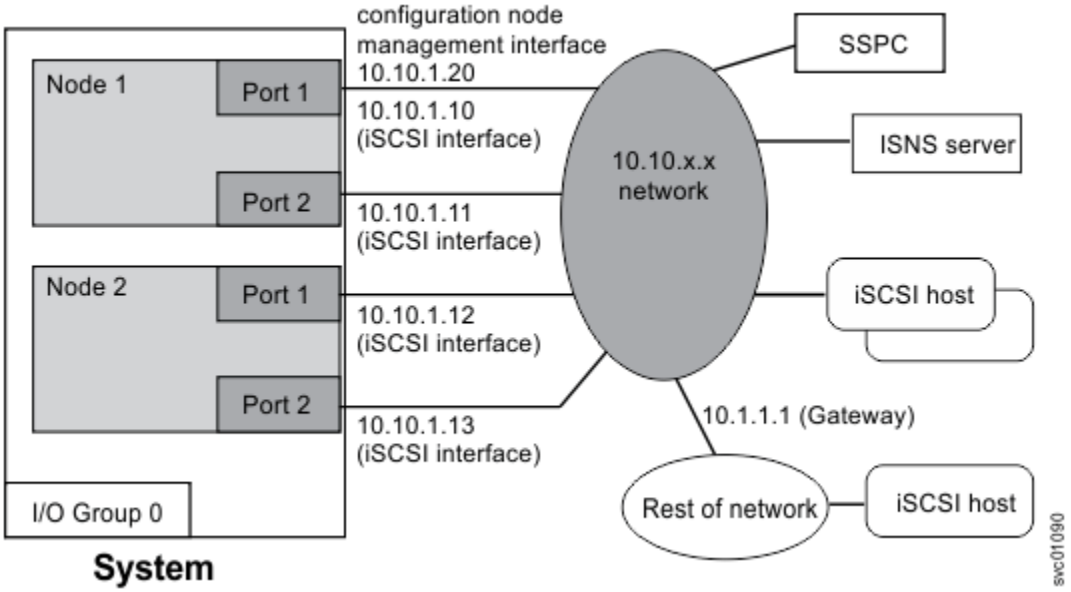


Figure 13. iSCSI configuration with a single subnet

This figure shows a two-node system that is connected to multiple subnets. Each node has two Ethernet ports (port 1 and port 2) that are connected to different IP subnets. One node in the system also acts as the configuration node. This configuration provides alternate IP interfaces on different subnets for the system management interface.

Note: With a multiple subnet configuration, if one subnet fails and the iSNS server is part of that subnet, the nodes and hosts are not able to access the iSNS server. The I/O is not affected and the other subnet is still active.

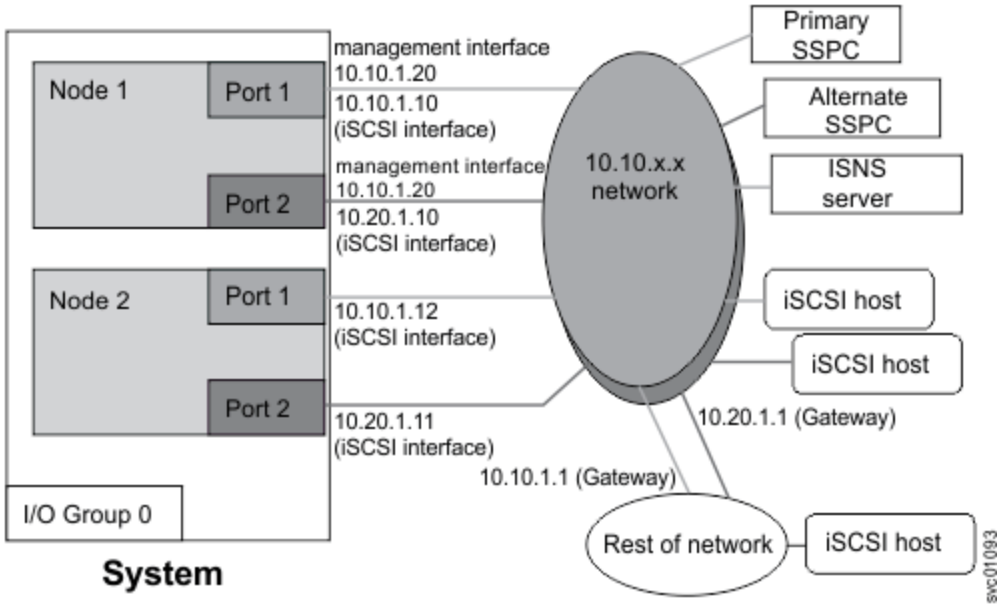


Figure 14. iSCSI configuration that uses multiple subnets and provides alternative configuration interfaces

This figure shows the same two-node system as Figure 14 on page 145. However, it illustrates the benefits of using multipathing and redundant host interface adapters in the host systems.

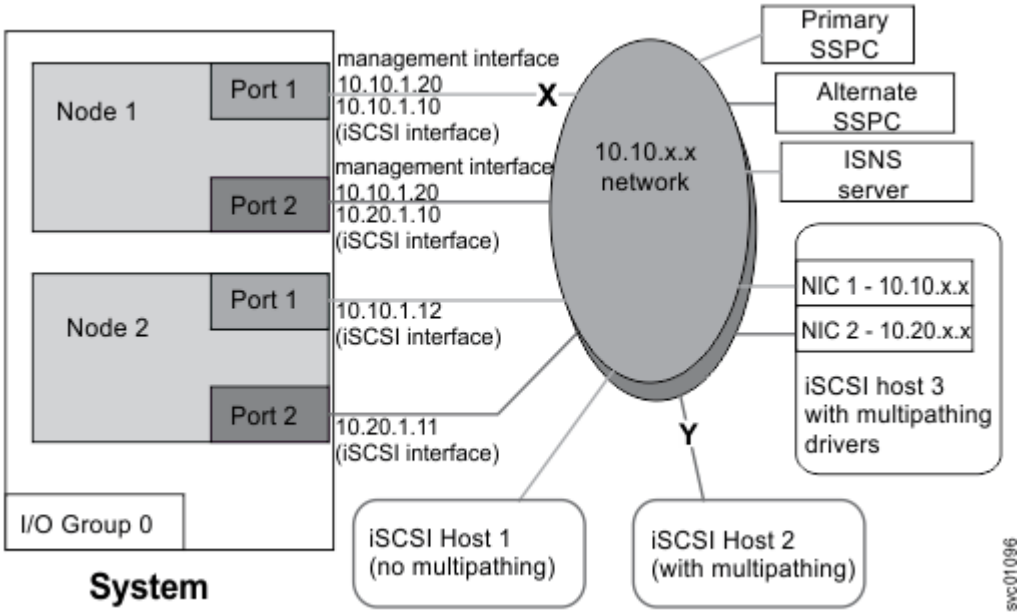


Figure 15. iSCSI configuration with host multipathing

In this example, host 1 does not use multipathing. A volume in the system's I/O group appears as four separate devices in host 1. The host selects one device to perform I/O to the volume, which corresponds to a particular IP address at a system node port, 10.10.1.10. If a connection between the host and this system port is broken (the link at X is broken), an I/O error is recorded on host 1 for that volume if I/O is in progress. No system state changes or IP failover occurs.

Host 2 uses multipathing. A volume in the system I/O group appears as a single device to the applications on host 2, even though the multipathing driver can detect four separate devices for each volume. The multipathing driver selects one or more of these devices during I/O. If the connection between the host and one system node port is lost, the multipathing driver can select an alternative path to the system I/O group. The I/O between the host and system continue without error. However, host 2 has only one host interface adapter. Therefore, host 2 reports I/O errors (such as the link at Y is broken) if the connection between that host interface adapter and the network is lost.

Host 3 uses multipathing and redundant host interface adapters. If a host interface adapter fails, the multipathing driver can still find paths from the host to a volume in the system I/O group, and the application I/O can continue without error. Because the host interface adapters are connected to different IP networks, the overall configuration can tolerate a single network failure without I/O errors occurring on host 3.

Multipathing drivers are not required to do system maintenance when system nodes are removed or replaced in an I/O group. However, multipathing host drivers are required for load balancing and for surviving host interface adapter, link, or network failures.

iSCSI configuration details for host connections

The following guidelines describe the iSCSI configuration requirements for host connections.

Attach the system to iSCSI hosts by using Ethernet ports on systems.

Note: The system supports SAN devices that bridge iSCSI connections into a *Fibre Channel* network.

iSCSI connections route from hosts to systems over the LAN. Follow these configuration rules:

- Hardware iSCSI initiators are not supported on IBM Storage FlashSystem.

- Validate the supported software iSCSI initiator by using the Server Support Interoperability Center (SSIC). If the software iSCSI initiator is not listed in the SSIC, submit a SCORE request to enable support.
- Ensure that SLP is not supported for discovery.
- Header and data digest support is provided only if the initiator is configured to negotiate.
- Only one connection per session is supported.
- Only `ErrorRecoveryLevel 0` (session restart) is supported.

iSCSI hosts connect to the system by using the node-port IP addresses. If a node fails, the iSCSI name, iSCSI alias, and the node-port IP addresses of the failed node transfer to the partner node. During the takeover, the iSCSI initiator logs out of the failed node and establishes a new session with the partner node by using the failed node's IP address. After the failed node recovers, the node-port IP addresses and the iSCSI name and alias return to the original node.

Note: 100 Gbps adapters support iSCSI. However, the performance is limited 25 Gbps per port.

IP requirements for iSCSI

Note: For general IP requirements, see "iSCSI Ethernet host attachment" on page 143.

IP addresses are used to discover storage volumes and to use I/O to access volumes. Node iSCSI IP addresses are also used to access a remote Internet Storage Name Service (iSNS) server, if configured.

- Each node Ethernet port can be configured on the same subnet with the same gateway. Or, you can have each Ethernet port on separate subnets and use different gateways. Network separation between node ports can result in better performance. Example: Ensure that the node Ethernet ports X and Y are on different VLAN or subnet for traffic separation.
- If you configure a system to use node Ethernet ports 1 and 2 for iSCSI I/O, ensure that the overall configuration meets the system IP requirements.
- To ensure iSCSI IP failover operations, connect nodes in the same I/O group to the same physical network segments on the same node ports. However, you can configure node Ethernet ports in different I/O groups to use different subnets and different gateways.
- Each iSCSI target can support both IPv4 and IPv6 concurrently.
- Ensure that all IP addresses, such as service and configuration, that are associated with a clustered-system Ethernet port are on the same subnet. However, IP addresses associated with a node Ethernet port that is used for iSCSI traffic can be configured to belong to different subnets for traffic isolation.
- Each connection between a host initiator and a system Ethernet port can use IPv4 or IPv6 (where supported). Concurrent use of both IPv4 and IPv6 for the same initiator to the system Ethernet port is not supported.

Port IP addresses enable iSCSI-attached hosts to process I/O. Provide a portset when you assign an IP address to a port. The portset ensures that a host does not discover more than four ports per node. For more information about portsets, see "Portsets" on page 333.

- Use the **mkip** command-line interface (CLI) command to configure a new port IP address to a specified node Ethernet port with an IPv4 or IPv6 address. To view a list of ports, use the **lsportethernet** command.

  The optional **-vlan** parameter sets the virtual local area network (VLAN) ID for an IP address that is configured for iSCSI host attachment. For more information about the **-vlan** parameter, see the information about configuring VLAN for iSCSI using the CLI.
- To remove an IP address from a node Ethernet port, run either of the following CLI commands:
  - To delete an IP configuration for the specified Ethernet port, run the **rmip ip_address_id** command.
  - To list the valid values for the IP address ID, run the **lsip** command.

Configuring or modifying an iSCSI alias by using the CLI

You can use the CLI to create or change the iSCSI alias for the selected node. An iSCSI alias is a user-assigned name that identifies a node to the iSCSI-attached host. To configure or modify an iSCSI alias, refer to the following steps.

You can use the **chnode** CLI to specify that:

1. The name or iSCSI alias to be configured is the name or alias of the node in the I/O group, and
2. The name or iSCSI alias to be configured is the name or alias of the partner node in the I/O group.

   If the system does not have a partner node, the values that you set are applied to the partner node when it is added to the system. If this parameter is used when a partner node already exists, the name or alias of that node changes.

After you create iSCSI aliases, you can configure the address for the system iSNS server.

Configuring the iSNS server address by using the CLI

If you are using iSCSI attached hosts with a system, you can use the CLI to configure the address of the iSNS server. Host systems use the iSNS server for target management and discovery. The system supports either IPv4 or IPv6 that is used by the iSNS server at a time. Example - If you configure an IPv6 iSNS IP address when you already configured an IPv4 iSNS IP address, the new IPv6 IP address becomes the iSNS IP address. You cannot use the old IP address for iSNS function.

- Use the **chsystem** CLI to specify an IPv4 address or an IPv6 address for the iSCSI storage name service (iSNS).

Host mapping

The following information specifies iSCSI host mapping. For general host mapping information, see “Ethernet connectivity” on page 299.

- iSCSI uses either iSCSI qualified names (IQN) (223 bytes) or extended unique identifier (EUI) (64-bit) names.
- Hosts use iSCSI names to connect to the node. Each node has an IQN in which the system name and node name are used as part of that IQN.
- iSCSI hosts that have a multipath driver installed can be moved in a nondisruptive manner.
- Multiple sessions from iSCSI hosts, with a maximum of four sessions from one iSCSI host to each system iSCSI target (FlashSystem node).
- A host can discover and log in only to IP addresses that are mapped to it by using portsets.
- The system rejects any login attempt that is not associated with a configured host object. A host object with a valid IQN must exist for the system to accept logins from a host.
- A host can log in to "n" SVC nodes only if its portset contains the port IP addresses from all "n" nodes.
- If a host that is associated with portset x initiates discovery through portset y, the discovery succeeds, but returns only the IP addresses that are mapped through portset x.
- Modification of IP addresses that are mapped to a portset causes login disruptions from hosts that are mapped to the portset.
- Host objects can be associated to different host attach portsets by using **mkhost** and **chhost** commands. IP addresses can be configured on different portsets by using **mkip** command. A host object can login through IP addresses belonging to the same portset. For more details, see the limits and considerations for the portset.
- A volume can be mapped to either a Fibre Channel host or an iSCSI host. You cannot map a volume to a host by using both iSCSI and Fibre Channel is not supported.
- The system does not support I/O from both, Fibre Channel and iSCSI, initiators in the same hosts to the same volumes.
- AIX® does not support iSCSI with IPv6.

- The system supports various host machines. Refer to the documentation for specific instructions and requirements for a particular host. For a list of supported hosts, see http://www.ibm.com/support.

Host operating systems, such as Linux, which use multiple Ethernet interfaces can experience Address Resolution Protocol (ARP) Flux. ARP Flux occurs when a host responds to ARP requests for interfaces on the same or different subnet from any interface on that same or different subnet. Although this behavior is typically harmless, in specific cases, the ARP Flux generates unexpected behavior of applications due to an incorrect mapping between IPv4 addresses and MAC addresses.

- To avoid ARP Flux on Linux, use the following setting on the host:

```
echo 1 > /proc/sys/net/ipv4/conf/all/arp_filter
```

- To make this behavior persistent, add a `net.ipv4.conf.all.arp_filter=1` entry in `/etc/sysctl.conf`.

  If you use other operating systems, such as VMware ESX, check your settings to avoid ARP Flux on those hosts as well.
- Enter the following command to disable iWARP Port Mapper (iwpmd) service on RHEL 7.5 and higher host machine:

```
$ systemctl stop iwpmd
```

  Note: With the IBM Storage Virtualize software, the iWARP Port Mapper (iwpmd) service causes host system hang or restart.

For more information about host mapping, see “Mapping volumes to a host” on page 107.

iSCSI authentication

iSCSI supports two types of authentication through the Challenge Handshake Authentication Protocol (CHAP):

- Unidirectional (one-way) authentication: The iSCSI target authenticates the iSCSI initiators.
- Bidirectional (two-way) authentication: The iSCSI target and the iSCSI initiators authenticate each other.
- Unidirectional (one-way) CHAP authentication, where you can specify an authentication username with a maximum of 255 characters, or you can use the initiator's IQN as the username.
- Bidirectional (two-way) authentication with initiators that use the *storage_username* and *storage_secret* for username and password fields.
- CHAP algorithms MD5, SHA1, SHA2, and SHA3-256 for unidirectional (one-way) and bidirectional (two-way) CHAP authentication.
- Only SHA3-256 and SHA256 hashing algorithms are supported when iSCSI host authentication mode is enabled.



Attention: The iSCSI initiator supports separate passwords for discovery and I/O sessions. The target (IBM Storage FlashSystem) requires the same password for both. You can configure no CHAP, unidirectional CHAP, or bidirectional CHAP. For bidirectional CHAP, use credentials that are different from the unidirectional CHAP credentials.

Configuring system iSCSI authentication by using the CLI

You can use the CLI to configure the system to authenticate with iSCSI-attached hosts by using the Challenge-Handshake Authentication Protocol (CHAP). After the CHAP is set for the system, all attached hosts must be configured to authenticate. When you are troubleshooting a problem, you can delay your configuration of the CHAP authentication until after you configure the first one or two hosts and test their connectivity.

To configure authentication between the system and iSCSI-attached hosts, refer to the following steps:

1. Use the **chhost** CLI to configure unidirectional (one-way) CHAP authentication for an iSCSI host. If you do not specify the username, the IQN of the initiator is taken as the username for unidirectional (one-way) CHAP authentication.
2. To set the bidirectional (two-way) CHAP authentication, set the unidirectional (one-way) CHAP first and then run the **chhost** CLI. The specified CHAP secret cannot begin or end with a space.

   Note: You can also set *storage_username* when you configure the storage secret or by default *clustername* is taken as *storage_username* by using the **chhost** CLI.
3. Use the **-nohostsecret** parameter of the **chhost** CLI to clear CHAP secrets for host authentication that were previously set.
4. Run the **lshost** command to check the CHAP configuration status.

After you configure the CHAP secret for the system, ensure that the system CHAP secret is added to each iSCSI-attached host. On all iSCSI-attached hosts, specify a CHAP secret that the hosts use to authenticate to the system.

iSCSI protocol supported session parameters

The following Ethernet session parameters are supported:

```
initial_r2t = 1
immediate_data = 0
max_connections = 1
Max_recv_segment_data_length = 32k
max_xmit_data_length = 32k
max_burst_length = 32k
first_burst_length = 32k
default_wait_time = 2
default_retain_time = 20
max_outstanding_r2t = 1
data_pdu_inorder = 1
data_sequence_inorder = 1
error_recovery_level = 0
header_digest = CRC32C,None
data_digest = CRC32C,None
ofmarker = 0
ifmarker = 0
ofmarkint = 2048
ifmarkint = 2048
```

Note: When the host object on the system side is created or re-created for an iSCSI host, log out of any existing sessions from that host and then, log in again.

iSCSI host object status

After you configure the iSCSI host object, use the **lshost** CLI to display the status of the configured host. An iSCSI host can have one of the following statuses:

Online
The host is logged in to both FlashSystem nodes and has a vdisk mapped, which indicates full connectivity and vdisk access.

Degraded
The host is logged in to one or both FlashSystem nodes without a vdisk mapping, which limits access to vdisks and results in degraded functionality.

Offline
The host has no active sessions with either node, which results in no connectivity. The vdisk mapping does not apply in this state.

Setting up a host server as an iSCSI initiator

Setting up a host server as an iSCSI initiator depends on the host type and operating system.

To configure a host, select a software-based iSCSI initiator, such as the Microsoft Windows iSCSI software initiator.

Note: Hardware iSCSI initiator is not supported. Submit a SCORE request to enable support.

You can use either the management GUI or the command-line interface (CLI) to set up a host server. To configure the host server as an iSCSI software-based initiator, complete the following steps.

Using the management GUI

1. To set up your system for iSCSI, complete the following steps:
   a. In the management GUI, select **Settings** > **Network** > **Ethernet Ports**.
   b. Expand the I/O group, right-click the node's corresponding port that you want to change and select **Manage IP Addresses Settings**.
   c. On the **Manage IP Addresses Settings** page, select **Add IP Address**. Specify the details for the IP address that you want to add to the selected port. Select the IP version (IPv4 or IPv6), enter the subnet mask or prefix, and optionally provide the gateway and VLAN. Select the portset to associate with this IP address. You can choose the default portset or create a new one from the **Portset** tab, which is used for the host attach operation. Then click **Save**.
   d. Verify that you configured the node and the system Ethernet ports correctly. The **Ethernet Ports** page displays all changed nodes as configured and **Host Attach** is set to **Yes**.
   e. To create hosts, select **Hosts** > **Hosts** and select **Add Host**. Select **iSCSI Hosts** and enter an iSCSI initiator name in the **iSCSI port** field. Check the host server to determine the iSCSI initiator name. Enter more details about the host and click **Add Host**.
   f. To create volumes, select **Volumes** > **Volumes** and select **Create Volumes**. Select the type of volume that you want to create and select the pool, size, and other details for the volume. Click **Create and Map to Host**.
   g. After the volume is created, click **Continue** to map the host to the volume. Select the host that you create and click **Map Volumes**.
2. To set up your host server, complete the following steps:
   a. Ensure that you configured your IP interfaces on the server.
   b. Install the software for the iSCSI software-based initiator on the server.
   c. On the host server, run iSCSI configuration commands, which enable the iSCSI initiator to log in to the system and discover volumes. The host then creates host devices for volumes.
3. After the host devices are created, you can use them with your host applications.

Using the CLI

1. To set up the system for iSCSI, complete the following steps:
   a. Select a set of IPv4 or IPv6 addresses for node Ethernet ports, which are in the I/O groups that use iSCSI volumes.
   b. Verify that the Host flag is enabled for the node Ethernet ports by using the **lsportethernet** command. If the flag is not enabled, enable it by using the **chportethernet** command.
   c. Configure the node Ethernet ports on each node in the system by using the **mkip** command. When configuring IP addresses, you can select the default portset or create a new portset by using the **mkportset** command. The portset is required for the host attach operation.
   d. Review the output of **lsip** and **lsportethernet** commands to verify that the node and the system Ethernet ports correctly configured.
   e. Use the **mkvdisk** command to create volumes on the system.
   f. Use the **mkhost** command to create a host object on the system that describes the iSCSI server initiator. Then, map volumes to the host object.
   g. Use the **mkvdiskhostmap** command to map the volume to the host object in the system.
2. To set up a host server, complete the following steps:

a. Ensure that you configured your IP interfaces on the server.

b. Install software for the iSCSI software-based initiator on the server.

c. On the host server, run the configuration methods for iSCSI so that the host server iSCSI initiator logs in to the system and discovers the volumes. The host then creates host devices for volumes.

3. After the host devices are created, you can use them with your host applications.

Linux iSCSI host attachment

You can use iSCSI Ethernet to attach a Linux host to the system.

Before you begin host attachment procedures:

- Read the information about setting up the host server.
- Ensure that you are familiar with the command-line interface (CLI) commands. Information about the CLI is available in this product information.

Installing the Linux software iSCSI initiator

To install the Linux software iSCSI initiator for iSCSI host attachment, set the iSCSI qualified name, target discovery, and authentication. Also, enable multipathing for the Linux hosts.

Before you begin

The system supports the use of the iSCSI-initiator software that is provided by your Linux distribution vendor. Ensure that you have the latest release of the initiator software that is installed on your host before you continue with the instructions.

The initiator software on RHEL systems is packaged as *iscsi-initiator-utils*. The initiator software on SLES systems is packaged as *open-iscsi*.

Procedure

1. Configure the iSCSI-initiator daemon by editing the `/etc/iscsi` file.

```
[root@R64SRV6 ~]# cd /etc/iscsi/

[root@R64SRV6 iscsi]# ls

initiatorname.iscsi  iscsid.conf

[root@R64SRV6 iscsi]#
```

2. Manage the software by using the **chkconfig** command and the **service** command. The service on RHEL systems is named *iscsid*. The service on SLES systems is named *open-iscsi*. Consider the examples for RHEL and SLES systems.

   For RHEL systems:

```
[root@R64SRV6 iscsi]# chkconfig --list | grep "iscsid"
iscsid          0:off   1:off   2:off   3:off   4:off   5:off   6:off

[root@R64SRV6 iscsi]# chkconfig iscsid on

[root@R64SRV6 iscsi]# chkconfig --list | grep "iscsid"
iscsid          0:off   1:off   2:on    3:on    4:on    5:on    6:off
[root@R64SRV6 iscsi]#
```

   For SLES systems:

   Yas1 and Yast2 provide an appropriate method to properly configure iSCSI services and initiator. From a command line for iSCSI services in SLES:

```
#/etc/init.d/open-iscsi start

#chkconfig open-iscsi on
```

3. Use the default settings for the iSCSI initiator. These settings are described in `/etc/iscsi/iscsid.conf`.

Setting the iSCSI qualified name for the Linux host

To correctly specify the iSCSI qualified name (IQN) for the Linux host, the initiator name that you choose must match the iSCSI name that is given to the host object for this system.

The iSCSI initiator IQN that you configure on the host side should never be the same as the target IQN on the system node.

1. Edit the `/etc/iscsi/initiatorname.iscsi` file to specify the IQN name. The following example shows a sample name:

```
InitiatorName= iqn.1986-03.com.ibm:iscsihost.0
```

2. After you change the initiator name, you must start the iSCSI-initiator service again.

Configuring more than one Linux host Ethernet port for iSCSI

If you want to configure more than one Linux host Ethernet port for iSCSI, methods are available to improve routing and maintain optimal redundancy in cases of dynamic discovery.

About this task

You can use one of the following two methods:

- Configure different subnet IP addresses for the host ports and the system ports. Each target IP of a system node is accessible only from one host IP of the corresponding subnet. As a result, pairs of host and system ports are in separate subnets.
- In the `/etc/iscsi/ifaces` subdirectory, create a separate *iface* configuration file for each network object that you want to bind a session to. Each iface config file must have a unique name that is less than or equal to 64 characters in length.

Procedure

The following *iface* example shows a simple case with two network devices (eth0 and eth1):

1. Edit the `/etc/iscsi/ifaces/eth0` file to specify the iface.net_ifacename:

```
iface.net_ifacename = eth0
```

2. Similarly, edit the `/etc/iscsi/ifaces/eth1` file:

```
iface.net_ifacename = eth1
```

iSCSI target discovery for Linux hosts

Two possible methods can be used to discover iSCSI target portals. The following list includes the methods:

- Send Targets
- Internet storage name service (iSNS)

After discovery, you use the **`iscsiadm`** command to log in to or log out of a specific portal on a discovered target.

Discovering iSCSI targets by using send targets

You can discover iSCSI target portals by using the send targets method. The **`iscsiadm`** command returns the IQN and target portals that are associated with the system node.

Run the discovery command:

```
iscsiadm --mode discovery --type sendtargets --portal x.x.x.x
```

where *x.x.x.x* is the IP address of a node Ethernet port on the system.

Tip: You can substitute `-m` for `--mode`, `-t` for `--type`, and `-p` for `--portal` in the discovery command.

Command output similar to the following example is displayed:

```
x.x.x.x:3260, 1 iqn.1986-03.com.ibm:2145.china5.hlnc111874
```

Discovering iSCSI targets on your Linux host by using iSNS

You can discover iSCSI target portals by using the internet storage name service (iSNS) method. Before you use the iSNS discovery, you must do the following tasks:

- Configure the system to use iSCSI.
- Configure the system to use your iSNS server.
- Verify that the system nodes registered their target portals with the iSNS server. If the node iSCSI addresses are visible to the iSNS server, this registration occurs automatically.

To discover iSCSI targets by using iSNS, complete the following steps:

1. Modify the iSNS server settings in the `/etc/iscsi/iscsid.conf` file by entering the IP address and the port number that is used by your iSNS server:

   ```
   isns.address = <iSNS server IP address>
   isns.port = <iSNS server port>
   ```

2. Make the iSNS configuration changes effective by restarting the iSCSI initiator service.
3. Run the iSNS discovery command to generate a list of all iSCSI target portals that are registered with the iSNS server:

   ```
   iscsiadm --mode discovery --type isns
   ```

Logging in to discovered targets

You can connect to a specific portal on a discovered target by using the **`iscsiadm`** command. The system supports only one iSCSI session between an initiator and a target. Ensure that you do not attempt to connect to the same target node more than once.

Log in the iSCSI initiator to your selected system targets by using the following appropriate command examples:

For SLES:

```
iscsiadm --mode node --target target_IQN --portal x.x.x.x --login
```

For example: `iscsiadm --mode node --target iqn.1986-03.com.ibm:2145.cluster10g1.node1 --portal 192.168.64.191 --login`.

For RHEL:

```
iscsiadm --mode node --targetname target_IQN --portal x.x.x.x --login
```

For example: `iscsiadm --mode node --targetname iqn.1986-03.com.ibm:2145.cluster10g1.node1 --portal 192.168.64.191 --login`.

When the login command completes, the host can access the volumes that are mapped to the host from the target after the SCSI devices are rescanned.

Logging out from targets

You can log out of a specific system target or you can log out of all established sessions.

Note: After you log out of a system target, the iSCSI volumes that it serves are no longer available to the host.

1. To log out of a specific system target, enter the following command:

```
iscsiadm --mode node --target <IQN> --portal x.x.x.x --logout
```

2. To log out of all your established sessions, enter the following command:

```
iscsiadm --mode node --logoutall=all
```

iSCSI timeout configuration settings for the system

You can set the timeout value for disks. The disk timeout is set in the `/sys/block/<device>/device/timeout` file. Depending on operating system and OS version, increase the default value of 30 or 60 seconds to 120 in most situations.

For example, if the device is `/dev/sdc`, you would use:

```
[root@R64SRV6 ~]#echo -n 60 > /sys/block/sdc/device/timeout
```

To increase the SCSI command timeout for the system, create the following *udev* rule so that it is not necessary to set the timeout manually for each mapped disk:

```
udev rules file
cat /etc/udev/rules.d/99-ibm-2145.rules

# Set SCSI command timeout to 120s (default == 30 or 60) for IBM 2145 devices
SUBSYSTEM=="block", ACTION=="add", ENV{ID_VENDOR}=="IBM",
ENV{ID_MODEL}=="2145", RUN+="/bin/sh -c 'echo 120 > /sys/block/%k/device/timeout'"
```

Reconfirm the settings after your restart the system. After you set up your volumes, use `multipath -ll | grep sd` to locate the block device paths. Also, check with `cat /sys/block/sd<x>/device/timeout` for 120 seconds, where <x> is a 2145 path. For example:

```
edam:~ # multipath -ll | grep sd | tail -1
  `- 10:0:0:4 sdl 8:176 active ready running
edam:~ # cat /sys/class/block/sdl/device/timeout
120
edam:~ #
```

Setting up authentication for Linux hosts

You can set up unidirectional (one-way) CHAP authentication for Linux hosts. After you configure unidirectional (one-way) authentication that is working for your host, you can optionally set up bidirectional (two-way) authentication.

Before you begin

The system supports two Challenge Handshake Authentication Protocol (CHAP) methods:

- Unidirectional (one-way) CHAP authentication (the system authenticates the host iSCSI initiator).
- Bidirectional (two-way) CHAP authentication (both the system and the initiator authenticate each other).

Note: CHAP secrets that you select for unidirectional (one-way) authentication and bidirectional (two-way) authentication must be different.

Procedure

To set up authentication for a Linux host, follow these steps:

1. Open `/etc/iscsi/iscsid.conf` or `/etc/iscsid.conf` by using an appropriate editor.
2. Go to the CHAP settings paragraph.

   The following example shows the output of CHAP settings for a Linux host:

```
#*************
#CHAP Settings
```

```
#**************

#To enable CHAP authentication set node.session.auth.authmethod
#to CHAP. The default is None.
#node.session.auth.authmethod = CHAP

#To set a CHAP username and password for initiator
#authentication by the target(s), uncomment the following lines:
#node.session.auth.username = username
#node.session.auth.password = password
node.session.auth.username = rhel_username
node.session.auth.password = xxxxxxxxxxxxx
#To set a CHAP username and password for target(s)
#authentication by the initiator, uncomment the following lines:
#node.session.auth.username_in = username_in
#node.session.auth.password_in = password_in
node.session.auth.username_in = clustername
node.session.auth.password_in = yyyyyyyyyyyy
#node.session.auth.chap_algs = SHA3-256,SHA256,SHA1,MD5
#To enable CHAP authentication for a discovery session to the target
#set discovery.sendtargets.auth.authmethod to CHAP. The default is None.
#discovery.sendtargets.auth.authmethod = CHAP
#To set a discovery session CHAP username and password for the initiator
#authentication by the target(s), uncomment the following lines:
#discovery.sendtargets.auth.username = username
#discovery.sendtargets.auth.password = password

#To set a discovery session CHAP username and password for target(s)
#authentication by the initiator, uncomment the following lines:
#discovery.sendtargets.auth.username_in = username_in
#discovery.sendtargets.auth.password_in = password_in
```

3. Set up authentication.
   - Set up unidirectional (one-way) authentication:
     a. Set a CHAP username and password to your initiator name.
        i) `node.session.auth.authmethod = CHAP`
        ii) `node.session.auth.username = <initiator's user name>`
        iii) `node.session.auth.password = <CHAP secret for host>`
        iv) (Optionally) `node.session.auth.chap_algs = SHA3-256,SHA256,SHA1,MD5`

        Note: When iSCSI host authentication mode is enabled, only SHA3-256 and SHA256 hashing algorithms are accepted.

     b. Set a discovery session CHAP username and password to your initiator name.
        i) `discovery.sendtargets.auth.authmethod = CHAP`
        ii) `discovery.sendtargets.auth.username = <initiator's user name>`
        iii) `discovery.sendtargets.auth.password = <CHAP secret for host>`
     c. Save these settings. Log out of any current sessions and rediscover the system iSCSI target for the CHAP secret to be effective.

     Note: In the previous example, *xxxxxxxxxxxx* is the CHAP secret for the host, and the *rhel_username* is the IQN name of the initiator. This username must be the same value that you set with the **chhost** command (*hostusername* field) for this host.
   - Set up bidirectional (two-way) authentication.

     Note: It is not mandatory to set up two-way authentication. Before you configure for bidirectional (two-way) authentication, ensure that unidirectional (one-way) authentication is configured and is working for your host.

     a. Edit the `password_in` to CHAP secret that you set up with the **chhost** command on the system. Set the CHAP username as *storage_username* which can be found by **lshost** command.
        i) Set a CHAP username and password for the target or targets.
           – `node.session.auth.username_in = <storage_username>`

- node.session.auth.password_in = *<CHAP secret for clustered system>*

ii) Set a discovery session CHAP username and password for the target or targets.

- discovery.sendtargets.auth.username_in = *<storage_username>*
- discovery.sendtargets.auth.password_in = *<CHAP secret for clustered system>*

b. Save these settings. Log out of any current sessions and rediscover the system iSCSI target for the CHAP secret to be effective.

Note: Before upgrading to 8.5.3.0, if the customer has configured two-way chap authentication, they must first switch to unidirectional (one-way) chap and then back to bidirectional (two-way) chap once the upgrade is complete. This is required because `clustername` as a username is not supported for two-way chap secret in earlier releases.

Enabling multipathing for Linux hosts

You can enable the system to support multipathing for Linux hosts.

Before you begin

The system supports the Linux device-mapper multipathing driver for the iSCSI host attachment. After the iSCSI devices are discovered, you can enable multipathing.

Procedure

To enable multipathing for a Linux host, complete the following steps:

1. If the daemon is already running, stop the multipathing daemon with the `service multipathd stop` command.
2. Add the entry for the system by editing the multipath configuration file, `/etc/multipath.conf`. In the DMMP configuration file, `/etc/multipath.conf`, ensure that the default settings include the values that are described in Settings for Linux hosts.

   Note: Prevent blacklisting of the IBM 2145 devices by either disabling blacklisting of devices entirely or by adding an appropriate entry in the `blacklist_exceptions` section.
3. Save the file.
4. Restart the multipathing daemon with the `service multipathd start` command.

   If you are using SUSE Linux Enterprise Server 10 Service Pack 2 or later, replace the line `prio "/sbin/mpath_prio/alua /dev/%n"` with `prio "alua"` to disable the `Using deprecated prio` message.

   Note: Linux usually enables multipathing automatically. However, in some cases such as the RH 6.3 SAN boot installation, the multipathing is not configured properly. If this situation occurs, you must manually edit the `/etc/multipath.conf` file. For information about DM-multipath, see the Linux *Storage Administration Guide* or *DM-Multipathing Installation and Configuration Guide* for your kernel from the Linux provider. For more information about Linux system configuration, see the following website:

   https://www.ibm.com/docs/en/linux-on-systems

Configuring multisession on the Linux operating system

Hosts with Linux operating systems support multisession connections between the host and target.

About this task

Discover the target, and then complete the following steps to configure a session between the host and the target:

Procedure

1. To create a session between the host and the target, run one of the following commands:
   - `iscsiadm –m node -l`

     This command establishes a session on all of the portals reported by the iSCSI target at the time of discovery.
   - `iscsiadm –m node –T` *`iSCSI target name`* `-l`

     This command establishes a single session between the host and the target at the specified target portal.
2. To remove a session between the host and the target, run one of the following commands:
   - `iscsiadm –m session –u`

     This command removes sessions on all of the portals reported by the iSCSI target at the time of discovery.
   - `iscsiadm –m node –T` *`iSCSI target name`* `-u`

     This command removes a single session between the host and the target at the specified target portal.

Microsoft Windows iSCSI host attachment

You can use iSCSI Ethernet to attach a Microsoft Windows host to the system.

Before you begin host attachment procedures:

- Read the information about setting up the host server.
- Ensure that you are familiar with the command-line interface commands.

Setting the iSCSI qualified name for the Microsoft Windows host

The iSCSI qualified name (IQN) can be viewed or changed on the Windows host.

To view or change the iSCSI IQN, click the **Configuration** tab on the Windows Server operating system. The initiator name that you choose must match the iSCSI name that is given to the host object for the system.

iSCSI target discovery for Windows hosts

The supported methods for discovering iSCSI target portals are:

- Send Targets
- Internet storage name service (iSNS)

Discovering iSCSI targets using Send Targets

You can discover iSCSI target portals by using the Send Targets method. The Send Targets method is configured from the **Discovery** tab of the iSCSI Initiator control panel.

This form of discovery requires that you to type the address of one or more of your system node Ethernet ports. For each node that you want to complete discovery on, follow these steps:

1. Click the **Add** button (click **Discover Portal** on Windows Server 2008) that is associated with the list of target portals.
2. Type the IP address or DNS name of the node Ethernet port on which you want to complete discovery. Leave the IP port name at the default value of 3260. If CHAP authentication is configured for this host on the system, see “Authentication for Windows hosts” on page 160.
3. Click **OK**. The system node port address is added to the list of target discovery portals.

Discovering iSCSI targets on your Windows host by using iSNS

You can discover iSCSI target portals by using the internet storage name service (iSNS) method. Before you use the iSNS discovery, complete the following tasks:

- Configure the clustered system to use your iSNS server.
- Verify that the system nodes registered their target portals with it.

The iSNS method is configured from the **Discovery** tab of the iSCSI Initiator control panel.

1. Click the **Add** button that is associated with the list of iSNS servers.
2. Type the IP address or DNS host name of your iSNS server.

Connecting to discovered targets

After you discover the iSCSI targets, you connect to a specific portal on a target.

The system supports only one iSCSI session between an initiator and a target. Ensure that you do not attempt to connect to the same target node more than once. Open the **Targets** tab of the **iSCSI Initiator** control panel. The list of discovered targets shows an entry for each discovered node.

To connect to a specific portal on a discovered target, complete the following steps:

1. Select the node that you want to connect to from the list and click **Log on...** (click **Connect** on Windows Server 2008). You can select whether to restore the connection automatically at boot or to enable multipathing for the iSCSI target from the **Connect to Target** window.
2. To connect to a specific portal on the target, such as to perform load balancing between the node Ethernet ports or to configure CHAP authentication, click **Advanced...** to access these settings.
3. Click **OK** to change the status for the selected targets from **Inactive** to **Connected**.
4. Select the node to disconnect from the list and click **Log Off** to stop the I/Os to the iSCSI disks. Use this option to log off from the iSCSI target.

Viewing and managing the discovered disks

You can use the Windows host to view the discovered disks and bind the devices to the **Microsoft iSCSI Initiator Service**. The system volumes that are mapped to the iSCSI host are now visible to the Windows disk management services. The system volumes can be initialized, formatted, and mounted.

Ensure that the iSCSI devices are available before dependent applications or services are started. You can bind these devices to the **Microsoft iSCSI Initiator Service** from the **Bound Volumes/Devices** tab on the **iSCSI Initiator** control panel.

If the system is rebuilt or restored from backups, you must log off and log back in to any previously connected iSCSI targets that are supplied from the system.

From the Windows host, you can view information to help you differentiate each of the discovered disks. To do so, complete the following steps.

1. From the Windows **Command Prompt**, type `diskpart`.
2. Next, type `list disk`.
3. Next, select the disk for which you want to view the details. Type `select disk` *x* where *x* is the number of the disk that you want to view.
4. Type `detail disk` to see the type of disk and other information. The following screen shows a sample output.

```
DISKPART> list disk
  Disk ###  Status      Size     Free     Dyn  Gpt
  --------  ----------  -------  -------  ---  ---
  Disk 0    Online       149 GB    78 GB   *
  Disk 1    Online       149 GB    78 GB   *
  Disk 2    Online       565 MB   565 MB
  Disk 3    Online       337 MB   337 MB

DISKPART> select disk 2

Disk 2 is now the selected disk.

DISKPART> detail disk

IBM      2145             SCSI Disk Device
Disk ID: 00000000
Type   : iSCSI
Bus    : 0
Target : 2
LUN ID : 0

There are no volumes.

DISKPART>
```

Changes to system registry to optimize the initiator

You can change the system registry to optimize the iSCSI operation.

For reliable iSCSI operation during system maintenance, you can update these values in the system registry. See Table 30 on page 160. In addition, you can also update the system registry to change the delayed acknowledgment settings to address slow performance with iSCSI operations between systems. .

Table 30. System registry

Registry key	Type	Value (decimal)
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Control\Class\{4D36E97B-E325-11CE-BFC1-08002BE10318}*bus id*\Parameters\LinkDownTime	DWORD	120
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Control\Class\{4D36E97B-E325-11CE-BFC1-08002BE10318}*bus id*\Parameters\MaxRequestHoldTime	DWORD	120
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Services\Disk\TimeOutValue	DWORD	60
HKEY_LOCAL_MACHINE\SYSTEM\CurrentControlSet\Control\Class\{4D36E97B-E325-11CE-BFC1-08002BE10318}*bus id*\Parameters\ MaxPendingRequests	DWORD	2048

Note: You must restart the computer for these changes to take effect.

Authentication for Windows hosts

The system supports both unidirectional (one-way) and bidirectional (two-way) Challenge Handshake Authentication Protocol (CHAP) methods for Microsoft Windows host attachments.

Discovery sessions or normal sessions use unidirectional (one-way) CHAP authentication where the system authenticates the host iSCSI initiator. In a bidirectional (two-way) CHAP authentication, the system and the initiator authenticate each other.

Setting up authentication for discovery sessions for Windows hosts

You can set up authentication for discovery sessions for Microsoft Windows hosts at the same time that you are connecting to the system volumes. The CHAP authentication information is on the same **Advanced Settings** panel.

To set up unidirectional (one-way) authentication for a Windows host, complete the following steps:

1. From the Control Panel, select the **iSCSI Initiator** option.
2. From the **iSCSI Initiator Properties** panel, click the **Discovery** tab.
3. After you click the **Discovery** tab, click **Add** under the **Target Portals** section to see the **Add Target Portal** dialog box.

   Note: For Windows Server 2012, select **Discover Portal** and then enter the IP address under **IP address or DNS name**.
4. Click **Advanced...** to see the **Advanced Settings** panel.
5. Select **CHAP logon information**.

   Note: For Windows Server 2012, select **Enable CHAP logon**.
6. Type in a value for the **User name**. The user name must be the same value that you set with the **chhost** CLI command (*hostusername* field) for this host.
7. Type in a value for the **Target secret**. The target secret must be a value of 12 characters and is the same value that you set with the **chhost** command on the system for this host. Click **OK**.

Tasks in the management GUI:

1. Select **Hosts** > **Hosts**.
2. Right-click the host to configure CHAP authentication and select **Properties**.
3. On the **Properties** page, ensure that **Show Details** is selected and click **Edit**.
4. In the **iSCSI CHAP secret** field, enter the CHAP secret that you use for the host.
5. In the **iSCSI user name** field, enter the host user name that you use for the host. Click **Save**.

Setting up authentication for normal sessions for Windows hosts

You can set up authentication for normal sessions for Microsoft Windows hosts at the same time that you are connecting to a target or volume. The CHAP authentication information is on the same **Advanced Settings** panel.

To set up unidirectional (one-way) authentication for a Windows host, complete the following steps:

1. From the Control Panel, select the **iSCSI Initiator** option.
2. From the **iSCSI Initiator Properties** panel, click the **Discovery** tab.
3. After you click the **Targets** tab, click **Log On...** under the **Targets** section to see the **Log On to Target** panel.

   Note: For Windows Server 2012, highlight your target, select **Connect**, and then select **Advanced** from the pop-up screen.
4. Click **Advanced...** to see the **Advanced Settings** panel.
5. Select **CHAP logon information**.

   Note: For Windows Server 2012, select **Enable CHAP logon**.
6. Type in a value for the **User name**. The user name must be the same value that you set with the **chhost** CLI command (*hostusername* field) for this host.
7. Type in a value for the **Target secret**. The target secret must be a value of 12 characters and is the same value that you set with the **chhost** command on the system for this host. Click **OK**.

For tasks in the management GUI, refer to the Setting up authentication for discovery sessions for Windows hosts section.

Setting up bidirectional (two-way) authentication for Windows hosts

You can set up bidirectional (two-way) authentication for Microsoft Windows hosts.

Verify that you have a working unidirectional (one-way) authentication and that you configured a system iSCSI Challenge Handshake Authentication Protocol (CHAP) secret on the system.

To set up bidirectional (two-way) authentication for a Windows host, complete the following steps:

1. From the Control Panel, select the **iSCSI Initiator** option.
2. From the **iSCSI Initiator Properties** panel, click the **General** tab.

   Note: For Windows Server 2012, select **Configuration**.
3. Click **Secret** to see an **iSCSI Initiator** panel from which you can type your CHAP secret. Enter the system CHAP secret.

   Note: For Windows Server 2012, click **CHAP...** and then enter the system CHAP secret.

   Click **OK**.

   Notes:

   a. This setting applies to both the discovery session and normal session.

   b. The CHAP secrets for the system and the host initiator cannot be the same.
4. To set up the bidirectional (two-way) authentication, repeat the previous steps, but in this instance, select **Perform mutual authentication** from the **Advanced Settings** panel.
5. Select **iSCSI Initiator Properties** > **iSCSI Initiator Mutual CHAP secret**. The initiator username value must be *storage_username* and chap secret must match with what is configured with **chhost**. The *storage_username* can be found by **lshost** command.

Tasks in the management GUI:

1. Select **Settings** > **Network** > **iSCSI**.
2. On the **iSCSI Configuration** page, select **Modify CHAP Configuration**.
3. On the **Modify CHAP Configuration** page, enter a *storage_secret* and select **Use for iSCSI-attached hosts**. Click **Modify**.

   Note: Before upgrading to 8.5.3.0, if the customer has configured two way chap authentication, they must first switch to one way chap and then back to two way chap once the upgrade is complete. This is required because `clustername` as a user name is not supported for two way chap secret in the earlier releases.

Configuring multisession on the Microsoft Windows operating system

The Microsoft Windows operating system supports multisession connections.

About this task

Discover the iSCSI targets, and then complete the following steps to configure a multisession connection.

Note: The following Microsoft Multipath I/O (MPIO) timer key values are recommended when you use multisession on Microsoft Windows and the Microsoft iSCSI Multipath DSM:

- PathVerifyEnabled:

  This flag enables path verification by MPIO on all paths every *N* seconds (where *N* depends on the value set in the PathVerificationPeriod panel). The type is boolean and must be filled out with either 0 (disable) or 1 (enable). The value is set to Disabled by default. You must change the value to Enabled.
- PathVerificationPeriod:

  This setting indicates the period, in seconds, within which MPIO is expected to complete path verification. This field is only honored if PathVerifyEnabled is True. This default value is 30 seconds. The maximum value that is allowed is MAXULONG. You must set this value to 60 seconds.

- PDORemovePeriod:

  This setting controls the amount of time, in seconds, that the multipath pseudo-LUN will remain in system memory, even after losing all paths to the device. The default is 20 seconds. The maximum value that is allowed is MAXULONG. You must set this value to 60 seconds.

  Detailed information about the procedure to set these values is available by searching for "Microsoft iSCSI Software Initiator Version 2.08" from the following website:

  www.microsoft.com/downloads

To create a multisession connection, complete the following steps.

Procedure

1. Select the **Targets** tab from **iSCSI Initiator Properties**, and click **Log On**.
2. From **Log On to Target**, select **Enable multi-path**, and then select one of the following two options:
   - Click **OK**.

     Repeat this procedure to create multiple sessions. Each time that you repeat this procedure, a new session is established between the initiator and the target.
   - Click **Advanced**, then make the following selections:
     - At **Local Adapter**, click **Microsoft iSCSI Initiator**.
     - At **Source IP**, click the host IP from which your target portal can be pinged.
     -
     - At **Target Portal**, click **Default**.

     Click **OK**.

What to do next

To delete a multisession session, select the **Targets** tab from **iSCSI Initiator Properties**, and click **Details**. Place a check mark in the box of each identifier connection to be terminated, and then click **Log off** and **OK**.

Note: Multiple sessions are not supported for Microsoft Clustering with iSCSI.

IBM AIX iSCSI host attachment

You can use iSCSI Ethernet to attach an IBM AIX host to the system.

Before you begin host attachment procedures:

- Read the information about setting up the host server.
- Ensure that you are familiar with the command-line interface commands.

Configuring the AIX iSCSI software initiator

Ensure that you follow all the steps to set up the host server and configure the AIX iSCSI software initiator.

Install the AIX software initiator on your AIX host if it is not already installed.

1. To verify that the file set `devices.iscsi_sw.rte` is installed, use the **`lslpp -l`** command.
2. If necessary, install the interim fix for the system Object Data Manager (ODM) stanzas.

The software initiator is configured by using the System Management Interface Tool (SMIT). To configure the software initiator, follow these steps:

1. Select **Devices**.
2. Select **iSCSI** and select **iSCSI Protocol Device**.

3. Select **Change / Show Characteristics of an iSCSI Protocol Device**. Select the device that is associated with the software iSCSI initiator, such as `iscsi0`.
4. Verify that the **Initiator Name** value is correct. The initiator name value is used by the iSCSI target during login. The name must match the iSCSI name that was set for the system host object that is associated with this host.

   Note: A default initiator name is assigned when the software is installed. You can change the initiator name to match the local network naming conventions.
5. The **Maximum Targets Allowed** field corresponds to the maximum number of iSCSI targets that can be configured. If you reduce this number, you also reduce the amount of network memory that is preallocated for the iSCSI protocol driver during configuration.

Adding the iSCSI targets for AIX hosts

Add the iSCSI targets for AIX hosts so that they can be discovered through a properly configured network interface.

To add the iSCSI targets during the device configuration, follow these steps:

Note:

The system does not support Ethernet port aggregation. Therefore, host-based multipath driver support for system LUNs is necessary to provide active-active controller use. Operating systems such as Linux, Windows, and VMware can configure iSCSI with multiple paths to system LUNs by using iSCSI interconnects. However, AIX MPIO does not support multiple paths to system LUNs by using iSCSI interconnects.

This statement does not mean that LUN redundancy across controller failure is not supported over iSCSI for AIX hosts. If a controller fails, the surviving controller takes over all the IP addresses and the iSCSI target that is configured on the failing controller. The AIX host then reconnects with the failed over iSCSI targets on the surviving controller and continues with its IO operations. Therefore, controller failure does not disrupt operations when AIX hosts communicate to the system through iSCSI.

1. Edit the `/etc/iscsi/targets` file to add the iSCSI target portal details for all system nodes that the host must log on to.

   Each uncommented line in the file represents an iSCSI target and is in the following format.

   ```
   192.168.1.7     3260      iqn.1986-03.com.ibm:2145.sahyadri.node1
   ```

   Note: As part of the "no iSCSI MPIO" restriction, it is not valid to have the same target IQN multiple times in the targets file, even with different IP addresses. If the same target exists multiple times, comment out the duplicate targets in order for the process to run.
2. After you edit the `/etc/iscsi/targets` file, type the following command:`cfgmgr -v -l iscsi0`

   This command causes the driver to attempt to log on to each of the targets that are listed in the `/etc/iscsi/targets` file. It also defines a new hard disk (hdisk) for each LUN on the targets that are found.
3. Review the discovered hard disks by entering the following command:

   ```
   lsdev -c disk
   ```

   The system volumes appear in the output in the following format:

   ```
   hdisk2   Available               IBM 2145 iSCSI Disk Drive
   ```

   Note: If the appropriate disks are not discovered, review the configuration of the initiator, the target, and any iSCSI gateways to ensure that they are correct. Run the **`cfgmgr`** command again.

Configuring iSCSI host disk timeouts

To configure the iSCSI host disk timeouts for the system, use the **chdev** CLI command to set the value and then verify your changes.

Use the value of 60 seconds for the rw_timeout attribute for the system iSCSI host disks. To set the value, complete the following steps:

1. Set the value with the **chdev** command as shown in the following example for each iSCSI device:

```
chdev -l hdisk10 -a rw_timeout=60
```

2. Verify the changes by using the **lsattr** command:

```
lsattr -E -l hdisk10 -a rw_timeout
```

Setting up authentication for AIX hosts

AIX hosts can be correctly set up for authentication on the system by following certain guidelines and tasks.

Although the system supports both one-way authentication and two-way authentication for iSCSI, the AIX software initiator currently supports only one-way authentication. The system target authenticates the initiator.

CHAP settings are defined in the `/etc/iscsi/targets` file on the host. The AIX initiator or host bus adapter (HBA) always uses its iSCSI qualified name (IQN) as the CHAP user name.

To set up authentication on an AIX host, complete the following steps:

1. Open the `/etc/iscsi/targets` file with any editor.
2. For each line that contains a target definition, append the CHAP secret of the initiator in quotation marks:

```
192.168.1.7      3260     iqn.1986-03.com.ibm:2145.sahyadri.node1 "secret"
```

   The CHAP secret value that you set here must match the value that was configured on the system for the host object that is associated with this host. Because the system authenticates on a per-initiator basis, the CHAP secret is the same for all the targets on a particular clustered system.

   Note: When setting up authentication for two way chap username and password, the `node.session.auth.username_in` value must be `clustername`. The `clustername` can be found by **lssystem** command.

   An example of CHAP settings for an AIX host for the `/etc/iscsi/targets` file is shown below.

```
#ChapSecret             = %x22*( any character ) %x22
#                       ;   "                     "
#                       ; ChapSecret is a string enclosed in double quotes. The
#                       ; quotes are required, but are not part of the secret.
#
#EXAMPLE 1: iSCSI Target without CHAP(MD5) authentication
#      Assume the target is at address 192.168.3.2,
#      the valid port is 5003
#      the name of the target is iqn.com.ibm-4125-23WTT26
#The target line would look like:
#192.168.3.2 5003 iqn.com.ibm-4125-23WWT26
#
#EXAMPLE 2: iSCSI Target with CHAP(MD5) authentication
#      Assume the target is at address 10.2.1.105,
#      the valid port is 3260
#      the name of the target is iqn.com.ibm-K167-42.fc1a
#      the CHAP secret is "This is my password."
#The target line would look like:
#10.2.1.105 3260 iqn.com.ibm-K167-42.fc1a "This is my password."
#
#EXAMPLE 3: iSCSI Target with CHAP(MD5) authentication and line continuation
#      Assume the target is at address 10.2.1.106,
#      the valid port is 3260
#      the name of the target is iqn.com.ibm:00.fcd0ab21.shark128
```

```
#      the CHAP secret is "123ismysecretpassword.fc1b"
#The target line would look like:
#10.2.1.105 3260 iqn.2003-01.com.ibm:00.fcd0ab21.shark128


192.168.1.41 3260 iqn.1986-03.com.ibm:2145.pahar.dvt110702
192.168.2.43 3260 iqn.1986-03.com.ibm:2145.moscow.dvt110706 "svcchapsecret"
```

The two targets in the previous example are members of different clustered systems. One target is configured to authenticate the initiator, and the other target is not configured to authenticate the initiator.

Note: Before upgrading to 8.5.3.0, if the customer has configured two-way chap authentication, they must first switch to one-way chap and then back to two-way chap once the upgrade is complete. This is required because `clustername` as a user name is not supported for two-way chap secret in earlier releases .

Updating ODM stanzas for system iSCSI devices

Updating the Object Data Manager (ODM) stanzas for the system iSCSI devices requires that you follow certain guidelines.

An interim fix is available to update the AIX ODM stanzas to recognize iSCSI system volumes. Ensure that you have installed the official PTFs. The following website provides information about the available PTFs: www.ibm.com/support

When you install the interim fix package, the AIX host can recognize iSCSI system volumes. The system volumes are shown in the below example with the following attributes for `lsdev -C -1 hdisk1`.

```
lsdev -C -l hdisk1

hdisk1     Available              IBM 2145 iSCSI Disk Drive


# lsattr -E -l hdisk1
clr_q          no                            Device CLEARS its Queue on error True
host_addr      9.71.43.106                   Hostname or IP Address           False
location                                     Location Label                   True
lun_id         0x0                           Logical Unit Number ID           False
max_transfer   0x40000                       Maximum TRANSFER Size            True
port_num       0xcbc                         PORT Number                      False
pvid           none                          Physical volume identifier       False
q_err          yes                           Use QERR bit                     True
q_type         simple                        Queuing TYPE                     True
queue_depth    8                             Queue DEPTH                      True
reassign_to    120                           REASSIGN time out value          True
reserve_policy no_reserve                    Reserve Policy                   True
rw_timeout     60                            READ/WRITE time out value        True
start_timeout  60                            START unit time out value        True
target_name    iqn.1986-03.com.ibm:2145.china6.hlcn111890   Target NAME      False
unique_id      3521360050768019102968800000000000000204214503IBMiscsi Unique device
identifier      False

lscfg -v -l hdisk1


  hdisk1             IBM 2145 iSCSI Disk Drive

        Manufacturer.................IBM
        Machine Type and Model......2145
        ROS Level and ID............30303030
        Serial Number...............
        Device Specific.(Z0)........0000043268101002
        Device Specific.(Z1)........
        Device Specific.(Z2)........
        Device Specific.(Z3)........
```

VMware iSCSI host attachment

iSCSI Ethernet hosts for the VMware operating system can be attached to FlashSystem.

Before you begin host attachment procedures:

- Read the information about setting up the host server.
- Ensure that you are familiar with the command-line interface (CLI) commands.

Installing the VMware initiator software

The following example shows the VMware ESXi 8.0 host configuration. Ensure that the supported release of the initiator software is installed on your host before you proceed.

Complete the following steps to install the VMware initiator software:

1. Enter the IP address of the ESX server, user ID, and password to log in to the system.

   Note: Use the default initiator SCSI settings.

   The following functions are not supported:

   - Guest OS support
   - SAN boot support

   For more details on support metrics, see IBM System Storage Interoperation Center (SSIC).
2. Prepare an iSCSI network: Create a dedicated VLAN or subnet and ensure end-to-end MTU consistency if you are using jumbo frames. For details, see Best practices for running VMware vSphere on iSCSI network.
3. Configure a vSwitch for iSCSI: Create a dedicated standard vSwitch with one physical NIC per iSCSI path. For details, see Create a vSphere standard Switch.
4. Create iSCSI VMkernel adapters: Create a VMkernel per iSCSI path and enable iSCSI traffic. For details, see Create a VMkernel adapter.
5. Enable VMware iSCSI software adapter: Add and enable the software iSCSI adapter on the ESXi host. For details, see Add or remove the software iSCSI adapter.
6. Bind VMkernel adapters to the iSCSI adapter: Bind all iSCSI VMkernels for multipathing. For details, see Configure port binding for iSCSI or iSER on ESXi.
7. Disable iSCSI delayed ACK: Disable DelayedAck on the iSCSI adapter to avoid storage latency issues. Host restart is required after you change the DelayedAck settings on the ESXi host. For details, see ESX/ESXi hosts might experience read or write performance issues with certain iSCSI storage arrays.
8. Add IBM Storage FlashSystem iSCSI targets: Add FlashSystem target IPs by using the **Dynamic Discovery**. For details, see Configure dynamic or static discovery for iSCSI and iSER on ESXi host.

   Note: If you remove the target IP on the **Dynamic Discovery** tab, then all target portals are removed.
9. Set up CHAP for iSCSI Storage Adapters (optional): You can set up unidirectional (one-way) and bidirectional (two-way) CHAP authentication for VMware hosts.

   To set up unidirectional (one-way) and bidirectional (two-way) CHAP authentication on a VMware host, complete the following steps:

   a. From the **vSphere Client** page, select **Inventory** tab.
   b. Select **Storage Adapters** > **iSCSI storage adapter**, and then right-click and select **Properties**.
   c. Select **CHAP authentication**.
   d. Select the preferred settings and enter the password:
      i) From the **CHAP target authenticates host** menu, select **Use CHAP**.
      ii) Select **Use initiator name** to use IQN of the initiator as the username for CHAP authentication or enter the username that you set with the **`chhost`** CLI command **`hostusername`** field on the system. Enter the CHAP secret.

   To set up the bidirectional (two-way) authentication, repeat the preceding procedure from this step and then complete the following steps:

   a. From the **Mutual CHAP** authentication menu, select **Use CHAP**.
   b. Select **Name** and enter the *storage_username* value that can be found by **`lshost`** command.

c. Enter **CHAP secret**. Ensure that the CHAP secret matches the value that configured with the **chhost** command.

Note: If you configure two-way CHAP authentication before you upgrade to software version 8.5.3.0, then switch to the one-way CHAP first and then back to two-way CHAP after the upgrade is complete. This switch is required because *clustername* as a username is not supported for the two-way CHAP secret in older releases.

10. Rescan storage adapters to discover iSCSI LUNs from FlashSystem.
11. Present storage to VMs: Create VMFS datastores or use RDM for direct LUN access. For details, see Best practices for running VMware vSphere on iSCSI.
12. Set a multipathing policy: Configure all iSCSI devices to use Round Robin for optimal load balancing. VMware supports different types of multiple paths.
    a. Round Robin: VMware accesses different paths available at fixed intervals.
    b. Fixed path: The I/O initiates access from the path that is initially selected, and switches to a different path if problems are encountered.
    c. Most recently used path: VMware uses the most recently used path, and switches to a different path if problems are encountered.

Viewing available paths

To view available paths in VMware, complete the following steps:

1. In the **vSphere Client**, go to the **ESXi host** > **Configure**.
2. Under **Storage** > **Storage Devices**, select the storage device whose paths you want to view.
3. On the **Properties** tab and review the module that owns the device, such as NMP or HPP. Under **Multipathing Policies**, you can also view the **Path Selection Policy** and the **Storage Array Type Policy** assigned to the device, if applicable.
4. On the **Paths** tab. Review all paths available for the storage device and the status of each path. The following path status information appears:

Table 31. Path status information

Status	Description
Active (I/O)	Indicates a working path or multiple paths that currently transfer the data.
Standby	Indicates paths that are not active. If the active path fails, standby paths become operational and start I/O transfer.
Disabled	Indicates paths that are deactivated by the administrator.
Dead	Indicates paths that are no longer available for processing I/O. A physical medium failure or array misconfiguration can cause this status.

Note: If you are using the **Fixed** path policy, you can see the preferred path, which is marked with an asterisk (*) in the **Preferred** column.

For more details, see Viewing and managing storage paths on ESXi hosts.

Queue depth in iSCSI hosts

In iSCSI networks, the queue depth is the number of I/O operations that can be run in parallel on a device.

Configure your servers to limit the queue depth on all of the paths to the volumes in configurations that contain many servers or volumes.

Where a number of servers in the configuration are idle or do not initiate the calculated quantity of I/O operations, queue depth might not need to be limited.

Queue depth limits in iSCSI networks

The queue depth is the number of I/O operations that can be run in parallel on a device. If you are designing a configuration for an iSCSI network, you must estimate the queue depth for each node to avoid application failures.

If a node reaches the maximum number of queued commands, it returns error codes to the host such as `Resource unavailable`. Many operating systems cannot recover if the situation persists for more than 15 seconds. This inability to recover can result in one or more servers that present errors to applications and application failures on the servers.

The formula for queue depth calculation considers the following factors:

- The maximum number of queued commands is calculated on a per-node basis. An I/O group contains two nodes. The system must continue to function when one of the nodes in an I/O group is not available.
- If a volume is mapped to be seen by more than one server, each of the servers can send commands to it.
- If a device driver times out of a command, it immediately reissues the command. The commands are stored in the command queue.

After you calculate the queue depth limit, you must apply it. Each operating system has a particular method of limiting the queue depth on a per-volume basis; see your host operating system documentation to change queue depth.

Homogeneous queue depth calculation in iSCSI hosts

Set the homogeneous queue depth to a calculated value in iSCSI hosts.

Set the queue depth for each volume on the servers by using the calculation:

```
q = (n/(v * c)) * 0.75
```

where

- n = the allowed command pool per I/O group; this value is fixed at 4000
- v = number of volumes per server for the particular I/O group under consideration
- c = number of servers in an I/O group

Therefore, `q = (4000/(v * c)) * 0.75`.

Example

Consider a system with 20 host servers `(c = 20)`, each with 20 volumes mapped `(v = 20)`. For this configuration, the value of the queue depth is `q = (4000/(20 * 20)) * 0.75 = 7.5` (or the lower boundary limit of 7).

Use the minimum value of 2 if the value of derived queue depth exceeds the maximum allowed value on a particular host operating system.

Some common host server configurations are summarized in Table 32 on page 170, considering 20 volumes mapped to each host server.

Table 32. Common queue depth configurations			
Number of iSCSI hosts per I/O group	**<20**	**20-40**	**64**
Recommended queue depth	Seven (or the maximum number that is allowed by the host OS, whichever is smaller)	Three to four (or the max number that is allowed by the host OS, whichever is smaller)	Two

Nonhomogeneous queue depth calculation in iSCSI hosts

After you understand certain criteria for nonhomogeneous queues, you can calculate their queue depth in iSCSI hosts.

Nonhomogeneous queues meet one of the following criteria:

- One or more servers must be allocated more resources so that they can queue more commands.
- Volumes must not be distributed evenly among the I/O groups in the clustered system.

For each volume, consider each server to which that volume has a mapping in an I/O group. This mapping results in a set of server/volume pairs. If the sum of the server and volume queue depth for all of the pairs is less than 4000, the server does not experience problems due to a full queue.

Multisession

Multisessions for iSCSI hosts can be configured by using different operating systems. Several characteristics must be considered when you configure an iSCSI multisession.

The following points apply to multisession configuration:

- Check the multisession support for respective host operating system.
- The use of the supported iSCSI multipathing is recommended when running the I/O with multisession.
- If you use CHAP authentication under the Microsoft Windows operating system, use the same secret across each session that is initiated from that host to a particular target.
- If the Microsoft Windows clustered system is rebuilt or restored from a backup, hosts must log off and then log back on to any previously connected iSCSI targets that are supplied from the system.
- If the host application uses SCSI 3 Persistent Reservation, you must run only a single iSCSI session per node.
- Running with multiple iSCSI sessions is not supported for applications that use Microsoft Windows Failover clustering over an iSCSI interface. You must run only a single session per node in these situations.

For more information on session limit maximums for an iSCSI target, an I/O group, and a system, see Configuration limits.

Single subnet configurations

Some multisession configurations can support a single subnet.

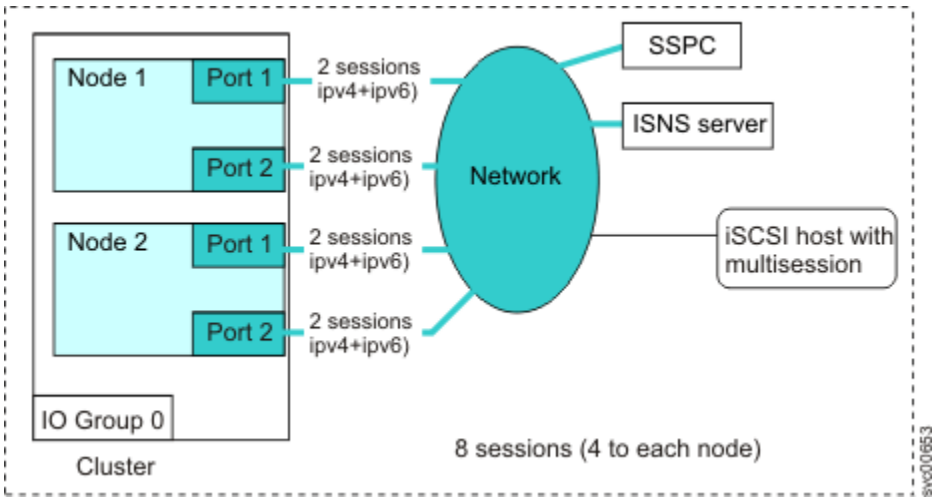


Figure 16. Four two-session iSCSI multisession configurations for a single subnet

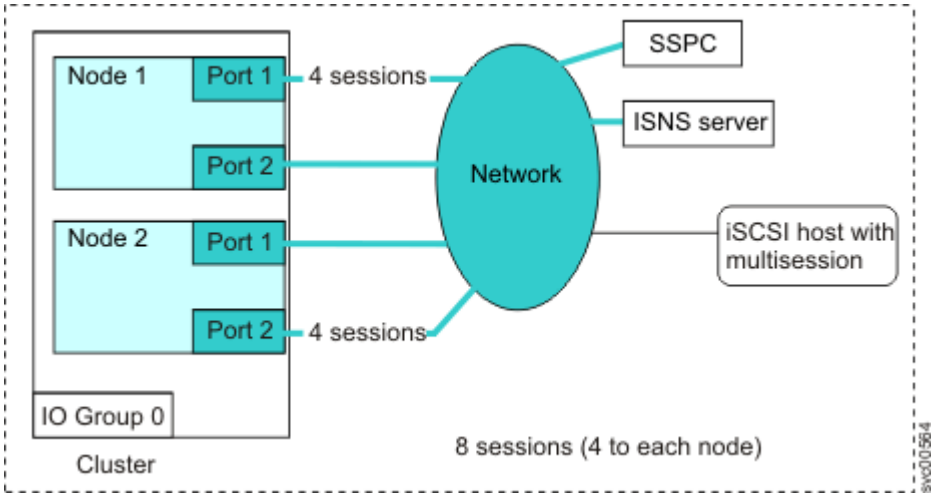


Figure 17. Two four-session iSCSI multisession configurations for a single subnet

Multiple and dual subnet configurations

Some multisession configurations support multiple and dual subnets.

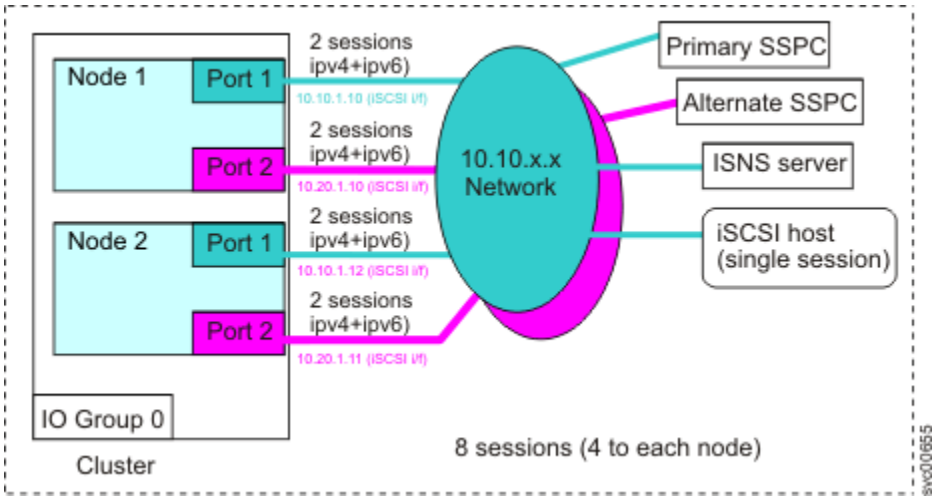


Figure 18. Four two-session iSCSI multisession configurations for multiple subnets and dual configuration

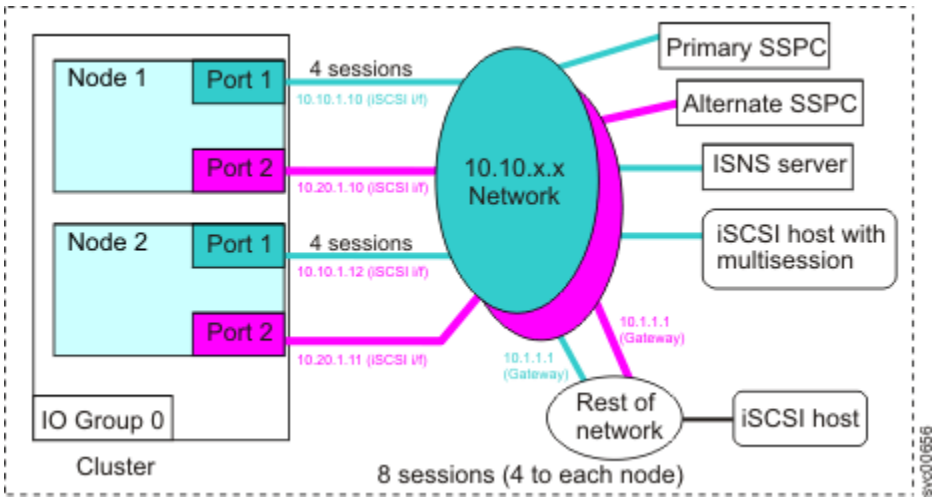


Figure 19. Two four-session iSCSI multisession configurations for a multiple subnets and dual configuration

NVMe over Fibre Channel host attachments

A host system can be attached to the storage controller using the NVM Express (NVMe) over Fibre Channel (FC) protocol.

FC-NVMe configuration guidelines

Consider NVMe over Fibre Channel target configuration guidelines when you plan and configure the hosts.

Include the following details in your plan:

- A FC-NVMe host can connect to a maximum of four ports on a node. To allow failover, configure up to four additional ports on the partner node.

  Zone up to four ports from a single host to detect up to four ports on a node. To allow failover and avoid outages, zone the same or additional host ports to detect up to four ports on the second node in the I/O group. The recommended configuration is to include only those ports in a zone that are part of the respective portset.
- Ensure that the configuration does not exceed the number of sessions per FC port as well as the number of hosts per system. For more information, see Configuration limits.
- A new FC-NVMe host can be mapped to a user-defined Fibre Channel portset only.

NVMe hosts that run the Linux operating system

When you install and configure attachments between the system and a host that runs the Linux operating system, follow specific guidelines.

For other specific information about NVMe over Fibre Channel (FC-NVMe) such as interoperability requirements, see Configuration limits.

Attachment requirements for hosts that are running the Linux operating system

Ensure that your system meets the requirements for attaching to a host that is running the Linux operating system.

1. Follow the HBA vendor instructions to update to the right firmware and driver level.
2. Install the required NVM-Express user space tools, by using one of the following commands:
   - For SLES

     ```
     zypper install nvme-cli
     ```

- For RHEL

```
yum install nvme-cli
```

3. If you are working with an Emulex adapter, verify that Emulex auto-connect is installed. To verify, run the following command:

```
rpm -q nvmefc-connect
```

Note: For SLES 15, RHEL 8.3 and latest kernel levels, the command **nvmefc-connect** is no longer applicable.

Note: When using FC-NVMe on RHEL 9.0 , use nvme-cli version 1.16 or above.

Configuring the Linux operating system for FC-NVMe hosts

After you ensure that your system meets the requirements for attaching to a Linux host, configure the Linux operating system.

You must install the appropriate host bus adapters with the correct firmware and driver levels that support FC-NVMe.

1. Zone the host ports to the NVMe ports on the system.

   For more information about how to identify NVMe ports on the system, see the CLI host commands.
2. Find the host NVMe Qualified Name (NQN) address (under `/etc/nvme/hostnqn`).
3. On the system, create the NVMe host object by using the host NQN.

```
svctask mkhost -force -name fc-nvmehost -nqn
nqn.2014-08.org.nvmexpress:uuid:449f8291-9c1e-446c-95c1-0942f55fa208
-portset fcnvme -protocol fcnvme -type generic
```

4. Map relevant volumes to the NVMe host. The same volumes cannot be mapped concurrently to both NVMe and SCSI.
5. To discover and connect to NVMe targets, enter the following commands.
   a. The **NVMe Discover** command.

```
nvme discover --transport=fc --traddr=nn-0x$twwnn:pn-0x$twwpn
--host-traddr=nn-0x$wwnn:pn-0x$wwpn
```

   The return of this command is the **NVMe Discover** log page, which consists of the target's subsystem NQN.

```
Discovery Log Number of Records 1, Generation counter 0
=====Discovery Log Entry 0======
trtype:  fibre-channel
adrfam:  fibre-channel
subtype: nvme subsystem
treq:    not required
portid:  <>
trsvcid: none
subnqn:  nqn.1986-03.com.ibm:nvme:2145.<>.iogroup<>
traddr:  nn-$twwnn:pn-0x$twwpn
```

   b. The **NVMe Connect** command, which has the same syntax with the addition of the subnqn.

```
nvme connect --transport=fc --traddr=nn-0x$twwnn:pn-0x$twwpn
--host-traddr=nn-0x$wwnn:pn-0x$wwpn -n $subnqn
```

   c. **NVMe connect-all** includes the **Discover** and **Connect** commands within the same command. This command connects automatically to the subnqn that is presented in the discover command:

```
nvme connect-all --transport=fc --traddr=nn-0x$twwnn:pn-0x$twwpn
--host-traddr=nn-0x$wwnn:pn-0x$wwpn
```

d. Initiate, the **NVME disconnect** command, to disconnect any NVMe device:

```
nvme disconnect -d /dev/nvme-$x
```

e. Initiate the **NVME disconnect-all** command, to disconnect all NVMe devices:

```
nvme disconnect-all
```

Removing a LUN

To remove a LUN from the storage system, perform the following steps:

Attention: Before removing a LUN from the storage system, ensure that all processes and applications that access the LUN are stopped and nothing is accessing one or more storage systems. If any process is still accessing the removed LUN, an emergency shutdown may occur, resulting in loss of access.

1. Stop all files access to the LUN, which needs to be removed and unmount the files system, if applicable.

   If LVM was used on one or more storage systems from which the LUN is being removed, make sure that no Logical Volumes, Volume Groups or Physical Volumes are still present or in use of one or more relevant storage systems. To ensure the above function, run the following commands:

   - **man lvremove**
   - **man remove**
   - **man pvremove**

2. To verify that all files are closed, enter the **lsof/dev/nvme*n*** commands. The output must be empty.
3. Place the LUNs on the storage side. From the storage system, unmap the volume from the host that uses the GUI.

   a. From the left-side panel, open the **Hosts** tab.

   b. Right-click on the hostname and select **Modify Volume Mappings**.

   c. Select the volume to be unmapped and click **Remove Volume Mapping**.

Multipath configuration for FC-NVMe hosts

Follow FC-NVMe multipath configuration recommendations for successful attachment of Linux hosts to the system.

Turning Native NVMe Multipath on or off

Hosts can be configured to work with the traditional Device Mapper or with the Native NVMe Multipath as follows:

- On SLES15, Native NVMe Multipath is turned on by default. On SLES12SP4, Native NVMe Multipath is not turned on by default.

  To check whether Native NVMe Multipath is on, enter:

```
# systool -m nvme_core -A multipath
Module = "nvme_core"

    multipath           = "Y"
```

  If the multipath = "N", and you want to turn it on, enter:

```
echo "options nvme_core multipath=Y" > /etc/modprobe.d/50-nvme_core.conf
dracut -f
reboot
```

Note: On SLES12SP4, if the multipath is not turned on even after reboot, then follow these steps:

1. Edit the file `/etc/default/grub`.
2. Append **`nvme-core.multipath=on`** at the end of the GRUB_CMDLINE_LINUX_DEFAULT variable.
3. Save the file.
4. Run `grub2-mkconfig -o /boot/grub2/grub.cfg` to apply the new configuration.
5. Reboot.

If you choose to work with Device Mapper, enter:

```
echo "options nvme_core multipath=N" > /etc/modprobe.d/50-nvme_core.conf
dracut -f
reboot
```

- On Red Hat Enterprise Linux 8.0 or later, Native NVMe Multipath is not turned on by default.

  To check whether Native NVMe Multipath is enabled in the command-line, enter:

  ```
  # cat /sys/module/nvme_core/parameters/multipath

      Y
  ```

  If the output is "N", and you want to enable it, follow these steps:

  – Using command-line:

    1. Add the `nvme_core.multipath=Y` option on the command-line:

       ```
       # grubby --update-kernel=ALL --args="nvme_core.multipath=Y"
       ```

    2. On the 64-bit IBM Z architecture, update the boot menu:

       ```
       # zipl
       ```

    3. Reboot the system.

  – Using module configuration file:

    1. Create the `/etc/modprobe.d/nvme_core.conf` configuration file with the following content:

       ```
       options nvme_core multipath=Y
       ```

    2. Back up the `initramfs` file system:

       ```
       # cp /boot/initramfs-$(uname -r).img \
            /boot/initramfs-$(uname -r).bak.$(date +%m-%d-%H%M%S).img
       ```

    3. Rebuild the `initramfs` file system:

       ```
       # dracut --force --verbose
       ```

    4. Reboot the system.

  If you choose to work with Device Mapper, enter:

  1. Remove the `nvme_core.multipath=Y` option from the command-line:

     ```
     # grubby --update-kernel=ALL --remove-args="nvme_core.multipath=Y"
     ```

  2. On the 64-bit IBM Z architecture, update the boot menu:

     ```
     # zipl
     ```

  3. Remove the `options nvme_core multipath=Y` line from the `/etc/modprobe.d/nvme_core.conf` file, if it is present.
  4. Reboot the system.

For more details to enable native NVMe and Device Mapper on Red Hat Enterprise Linux 8.0 or later, see Enabling multipathing on NVMe devices.

Device Mapper or Native NVMe Multipath configuration

Two framework options are available to configure the multipath configuration file: Device Mapper or Native NVMe Multipath.

To configure Device Mapper, complete the following steps:

1. Edit the `/etc/multipath.conf` file to include the following code:

```
devices {
    device {
        vendor "NVME"
        product "IBM     2145"
        path_grouping_policy "multibus"
        path_selector "round-robin 0"
        prio "ANA"
        path_checker "none"
        failback "immediate"
        no_path_retry "queue"
        rr_weight uniform
        rr_min_io_rq "1"
        fast_io_fail_tmo 15
        dev_loss_tmo 600
    }
}
defaults {
    user_friendly_names yes
    path_grouping_policy    group_by_prio
}
```

2. Run the following commands to validate that the multipath daemon is running:

```
systemctl enable multipathd.service
systemctl start multipathd.service

# ps -ef | grep -v grep | grep multipath
root      1616     1  0 Nov21 ?        00:01:14 /sbin/multipathd -d -s
```

   After you enable the multipath service for SLES12SP4, rebuild `initrd` with multipath support:

```
dracut --force --add multipath
```

3. Run the following commands to apply configurations:

```
multipath -F
multipath
multipath -ll
```

If you use Native NVMe Multipath, add the following code to the multipath configuration file:

```
blacklist {
    device {
        vendor  "NVME"
        product "IBM\s+2145"
    }
}
```

Device Mapper performance

Testing shows that performance is superior when you work with multiqueue I/O scheduling as described in the following steps.

Note: Before you apply these steps, contact a SUSE representative to verify that this option is supported.

For performance improvements, apply multiqueue I/O scheduling with blk-mq (in SLES):



Warning: Do not change **`nvme-core.multipath`** value when Native NVMe is in use.

1. Edit the /etc/default/grub file and the following text:

```
GRUB_CMDLINE_LINUX_DEFAULT="BOOTPTimeout=20 BootpWait=20
biosdevname=0 powersaved=off resume=/dev/system/swap splash=silent
quiet showopts crashkernel=175M,high dm_mod.use_blk_mq=y
scsi_mod.use_blk_mq=1 transparent_hugepage=never"
```

2. Apply the new configuration:

```
swfc178:~ # grub2-mkconfig -o /boot/grub2/grub.cfg
Generating grub configuration file ...
Found theme: /boot/grub2/themes/SLE/theme.txt
Found linux image: /boot/vmlinuz-4.12.14-25.19-default
Found initrd image: /boot/initrd-4.12.14-25.19-default
Found linux image: /boot/vmlinuz-4.12.14-23-default
Found initrd image: /boot/initrd-4.12.14-23-default
done
```

3. Reboot.
4. Validate that the multiqueue feature is enabled in multipath:

```
mpatho (eui.880000000000000b0050760071c60044) dm-3 NVME,IBM     2145
size=1.5G features='3 queue_if_no_path queue_mode mq' hwhandler='0' wp=rw
```

VMware ESXi installation and configuration for NVMe over Fibre Channel hosts

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 SCSI 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 datastores and can read or write to them.

Configuring the ESXi operating system

You must configure the VMware ESXi 8.0 and up operating system before you can connect hosts that run VMware ESXi datastores with the system.

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

- Your system must be installed.
- You must install the appropriate host bus adapters and ensure that the latest drivers and firmware are applied.

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

1. Make sure you use Emulex Gen6 HBA minimal supported versions.
   - Driver version: 12.4.293.2
   - Firmware version: 12.4.243.17

   Enter the following command to verify:

   ```
   esxcli storage san fc list
   ```

2. Make sure that FC-NVMe is enabled by setting lpfc_enable_fc4_type to 3.

   To change the value, enter the following command:

   ```
   esxcli system module parameters set -m lpfc -p lpfc_enable_fc4_type=3
   ```

3. Make sure all port of the ESXi hosts and system ports are zoned correctly.

To see the WWPNs of FC-NVMe, enter the following command:

```
esxcli storage core adapter list
```

The NVMe WWPNs have the driver name `lpfc`.

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

   You can find the NQN by entering the following command:

   ```
   esxcli nvme info get
   ```

5. If more than one physical ESXi host is present in the vSphere cluster, create a host cluster and add each ESXi host object as a member.
6. After you create volumes on the system and map them to the host, do VMware storage rescan.

   If rescan does not work, try resetting the Fibre Channel ports. To reset the ports, enter the following command:

   ```
   esxcli storage san fc reset -A <HBA name>
   ```

7. Create VMware datastores based on the system mapped volumes.

NVMe hosts that run the Microsoft Windows operating system

Ensure that your system meets the requirements for attaching to a host that is running the Microsoft Windows operating system.

The following list provides the requirements for attaching the system to the host that is running the Windows operating system:

- Before you install the Emulex HBA Manager application, verify the correct level with Broadcom. At the time of publication, `OneInstall-Setup-12.6.240.40.exe`, which includes driver level 12.6.165.0, was used.
- Install the Emulex HBA Manager application from Broadcom. Under **Installation options**, select the **FC Driver** checkbox to include the driver in the installation. Also, select the correct **Emulex HBA Manager application** kit checkbox.

Configuring the Windows operating system for FC-NVMe hosts

After you ensure that your system meets the requirements for attaching to a Microsoft Windows host, configure the Windows operating system.

You must install the appropriate host bus adapters with the correct firmware and driver levels that support FC-NVMe.

1. Zone the host ports to the NVMe ports on the system.

   For more information about how to identify NVMe ports on the system, see the CLI host commands.

2. Find the Windows host NVMe Qualfied Names (NQNs).

   The NQNs of the Windows host displays as a string composed from a constant and the WWPN address of the host port. The constant is `nqn.2017-01.com.broadcom:ecd:nvmf:fc:` plus add the port WWPN.

3. On the system, create the Windows host by using the Windows ports NQN. The system supports up to 2 NQN addresses per host object, which are separated by a comma. For instructions about creating a host, refer to the host commands information.
4. On the system, map relevant volumes to the Windows NVMe host. The same volumes cannot be mapped concurrently to both NVMe and SCSI.
5. Open Emulex HBA Manager application.

   Select a port to configure and click **Driver Parameters**. Change the following defaults and repeat for each port:

a. For **EnableNVMe**, set to 1 to enable NVMe support.
b. For **NumNVMENode**, set to the maximum value of 512.
c. For **NumNVMENS**, set to the number of volumes that are mapped or to the maximum value of 255.

6. Apply the settings and restart the Windows host.
7. Either validate the correct firmware level or select **Download Firmware** to update.

Multipath configuration for FC-NVMe Windows hosts

Follow FC-NVMe multipath configuration recommendations for successful attachment of Microsoft Windows hosts to the system.

1. On **Windows server manager** > **Tools**, MPIO might need to be added. You can add MPIO either through **Manage** > **Add Roles/Features** or through the **powershell** command **Add-WindowsFeature -Name 'Multipath-IO'**.
2. In the **MPIO** window, select **Discover Multi-Paths**. Under **Others**, add `IBM NVMe 2145`. This process might take a while.
3. Restart the Windows host.
4. You can view the multipath device by entering the **powershell** command `mpclaim -s -d`.

IBM Power Systems AIX

You can attach an NVMe host to an IBM Power Systems AIX host by following specific requirements and instructions.

Attaching to IBM Power Systems AIX hosts

For interoperability information about attachments between the system and an IBM Power Systems AIX host, check the IBM support website: http://www.ibm.com/support

Note: The IBM Power Systems information applies to all AIX hosts that are listed on the interoperability support site.

Attachment requirements for IBM Power Systems hosts

Ensure that your system meets the requirements for attaching to a host that is running the AIX operating system.

Before you attach an IBM Power Systems host, ensure that you meet the following prerequisites:

- You installed the correct operating systems and version levels on your host, including any updates and APARs (Authorized Program Analysis Reports) for the operating system.
- You have the documentation for your host system and you have access to the Installing section in this product information.

  All system publications are available from the following website:

  IBM

  https://www.ibm.com/docs/en/aix

At the time of writing, the supported servers are POWER9 servers with minimum firmware version of 950.

AIX environments for IBM Power Systems hosts

Ensure that each IBM Power Systems host uses a supported operating system and level.

The following website provides current interoperability information about the supported operating system levels for IBM Power Systems hosts:

https://www.ibm.com/docs/en/aix

At the time of writing, the supported operating system is AIX 7.2 TL 5 from 2035_72V and above.

Drivers and firmware for IBM Power Systems hosts

Ensure that you use the correct host bus adapter device driver and firmware levels for yourIBM Power Systems AIX hosts.

For POWER9 hosts, the minimum firmware level is 950.

The following website provides current interoperability information about device driver and firmware levels:

https://www.ibm.com/docs/en/aix

Configuring the AIX operating system

You must configure the AIX operating system before you can use IBM Power Systems hosts with the system.

Before you configure the AIX host operating systems, the following tasks must be completed:

- The system must be installed.
- The appropriate host bus adapters must be installed.

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

1. Zone the host to the system nodes on the Fibre Channel (FC) SAN.
2. Verify that the FC ports appear in AIX. Each FC port appears through the **lsdev** command with a name starts with *fcs*.

   ```
   bash-4.4# lsdev | grep fcs
   fcs0       Available 00-00       PCIe3 2-Port 32Gb FC Adapter (df1000e314101506)
   fcs1       Available 00-01       PCIe3 2-Port 32Gb FC Adapter (df1000e314101506)
   ```

   Note: If *fcs* device is not found, you can run the **cfgmgr** command.
3. Check for FC-NVMe device availability.

   Note: AIX creates children devices of the *fcs* device for each protocol that AIX supports.

   a. The FC-NVMe protocol appears through the **lsdev** command with a name starts with *fcnvme*.

      ```
      bash-4.4# lsdev | grep fcnvme
      fcnvme0    Available 00-00-02    FC-NVMe Protocol Device
      fcnvme1    Available 00-01-02    FC-NVMe Protocol Device
      ```

   b. If the FC-NVMe device is not detected, you can enable it by using a **smitty** command for each port.

      ```
      smitty -> Devices -> FC Adapter -> FC-NVMe Device -> Enable a FC-NVMe Protocol Device
      ```
4. Find the host NVMe Qualified Name (NQN) from each FC-NVMe port.

   ```
   bash-4.4# lsattr -El fcnvme0
   attach     switch      How this adapter is connected  False
   autoconfig available   Configuration State            True
   host_nqn   nqn.2018-04.com.ibm:aix.fc.pn.0x100000109ba9d1c9.nn.0x0x200000109ba9d1c9   Host
   NQN (NVMe Qualified Name) True
   ```
5. Create FC-NVMe host type in target.

   Note: In AIX, each port has its own NQN. You can configure two NQNs to add a host to the target. If you have more ports, you can work on the host cluster, where each pair is defined as `host`.
6. Create a volume and map to the FC-NVMe host object.
7. Run the **cfgmgr** command without any parameters.
8. To check the FC-NVMe paths and disks, run the following command:

```
lsdev | grep -i nvme
```

The NVMe Discovery Controller and the NVMe Dynamic Controller appear for each path. The following example shows the four paths between the host and the target for the NVMe Discovery Controllers and the NVMe IO dynamic Controllers:

```
nvme0 Available 00-00-02 NVMe Discovery Controller
nvme1 Available 00-00-02 NVMe Discovery Controller
nvme2 Available 00-01-02 NVMe Discovery Controller
nvme3 Available 00-01-02 NVMe Discovery Controller
nvme4 Available 00-00-02 NVMe Dynamic Controller
nvme5 Available 00-00-02 NVMe Dynamic Controller
nvme6 Available 00-01-02 NVMe Dynamic Controller
nvme7 Available 00-01-02 NVMe Dynamic Controller
```

The following example shows the mapped volumes as hdisk:

```
hdisk1 Available 00-00-02 NVMe Flash Disk
hdisk2 Available 00-00-02 NVMe Flash Disk
hdisk3 Available 00-00-02 NVMe Flash Disk
hdisk4 Available 00-00-02 NVMe Flash Disk
hdisk5 Available 00-00-02 NVMe Flash Disk
hdisk6 Available 00-00-02 NVMe Flash Disk
hdisk7 Available 00-00-02 NVMe Flash Disk
hdisk8 Available 00-00-02 NVMe Flash Disk
hdisk9 Available 00-00-02 NVMe Flash Disk
hdisk10 Available 00-00-02 NVMe Flash Disk
```

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.

Note: NVMe over RDMA or NVMe over TCP is not supported for direct attach host connection.

Important: FlashSystem 9600, 7600, and 5600 support NVMe over TCP but do not support NVMe over RDMA for host connectivity.

For more information on NVMe support, see Configuration limits.

For more information about host compatibility, 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 compatible operating system and adapter cards. Refer to vendor documentation if updates are needed.

Note: For more information about host compatibility, 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 compatible operating system and adapter cards. Refer to vendor documentation if updates are needed.

For more information about host compatibility, see IBM System Storage Interoperation Center (SSIC).

Table 33. NVMe over RDMA minimum compatible 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 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.

Any single host NQN is only permitted to have a single connection to any target IP address. If a host is using multiple initiators, they should be configured on separate subnets with their intended target IP addresses.

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.

      Note: To optimize NVMe-TCP performance, disable the TCP delayed ack feature by using the following command:

      ```
      vsish -e set /net/tcpip/instances/defaultTcpipStack/sysctl/_net_inet_tcp_delayed_ack 0
      ```

      For more details, see VMware ESXi vendor guide.

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. Discover and connect to the storage controller.
   - Identify the discovered controller through command:

     ```
     esxcli nvme controller list
     ```

   - To discover and connect use commands:

     a. `esxcli nvme fabrics discover -i <ip_address_of_target_rdma/tcp> -a <vmhba> -p 4420`

     b. `esxcli nvme fabrics connect -i <ip_address_of_target_rdma/tcp> -a <vmhba> -p 4420 -s <subsystem-nqn> -c <controller-id> -t <keep-alive-timeout>`

     where <controller-id> is an optional parameter.

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.

Chapter 7. Policies

A policy defines a set of rules that can be used to simplify storage management. The system supports several different types of policies.

Provisioning policies

A provisioning policy is an entity that defines a set of rules for allocating capacity from a pool. Any volumes that are created in the pool adopt the capacity-saving methods that are defined in the provisioning policy.

An administrator or security administrator can use a provisioning policy to define a set of rules that are applied when volumes are created within a storage pool. When a provisioning policy is assigned to a storage pool, all the volumes that are created in the pool automatically get the capacity savings that are defined by the policy. Provisioning policies allow a simpler and consistent provisioning process that can be used with external automation software. The capacity savings for existing volumes do not change if the provisioning policy is removed from the storage pool or a different policy is applied. A provisioning policy defines only the capacity savings that are used when you create new volumes in the pool.

Provisioning policies support the following capacity-saving options:

None
Use the **none** option when you don't want volumes to use any additional capacity savings. With this option, volumes use only the built-in capacity saving functions of the storage pool. For example, pools containing certain drive types, such as FlashCore Module drives, are self-compressing.

Thin-provisioning
Use the thin-provisioning option to indicate that all volumes in a pool use thin-provisioning to save capacity. Thin-provisioning allocates capacity only when data is written to storage.

You can optionally define the following attributes:

- *Buffer size*: indicates the capacity in the storage pool that a thin-provisioned volume attempts to reserve as a buffer.
- *Warning threshold*: indicates the threshold at which a warning is shown for the used capacity of thin-provisioned volumes.

Compression
Use the compression option to indicate that all volumes in a pool use compression to save capacity. Compression reduces the capacity that data requires by removing unnecessary content from the data.

Deduplication
Use deduplication when you create a new provisioning policy to add deduplication to your capacity saving for a data reduction pool. Both compression and thin-provisioning capacity saving can be used with deduplication to increase capacity savings for data reduction pools. Standard pools do not support deduplication.

You can rename a provisioning policy after the policy is created. However, the capacity savings option of a provisioning policy cannot be changed. To change the capacity savings to be used for new volumes, create and assign a new provisioning policy to the storage pool, replacing any existing policy. The old policy can be deleted if it is no longer required. A provisioning policy can be associated with multiple pools. However, a pool can only have one policy.

The capacity savings for a volume can be changed after a volume is created. A provisioning policy for a storage pool is only applied when new volumes or volume copies are created.

After you add a provisioning policy to a pool, or assign a different provisioning policy, use the **addvolumecopy** command to convert existing volumes to use the capacity savings of the new provisioning policy. The **-autodelete** parameter can be used to automatically delete the original copy of the volume after the new copy is synchronized.

© Copyright IBM Corp. 2025

To configure provisioning policies in the GUI, select **Policies** > **Provisioning Policies**. Assign the provisioning policy to a pool using the panel **Pools** > **Pools**.

To configure provisioning policies on the CLI, use **mkprovisioningpolicy**. You can add a provisioning policy to a pool using **mkmdiskgrp** or **chmdiskgrp**.

Snapshot policies

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

A snapshot policy has the following properties:

Backup interval
The length of time interval between snapshot backups. The minimum value is 60 minutes.

Retention period
The length of time which snapshots are retained. The maximum length is 365 days.

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 can be assigned to one or more volume groups, whereas a volume group can only have one snapshot policy. When assigning a snapshot policy to a volume group, you can specify that the snapshots are Safeguarded.

The snapshot policy configures the internal scheduler, which creates and deletes snapshots. The system provides predefined polices that contain specific retention and frequency values for common use-cases. The system also allows creation of the user-defined policies.

The following predefined policies are supported and can be used:

- A policy for frequent snapshot creation and shortest retention. With this policy, snapshots are created every six hours and retained for one week. Use this policy for volume data that requires the highest recovery point objective (RPO). This policy is intended for data that is frequently updated and business-critical. Examples include customer account data, order data, or other proprietary information. For more information about RPO requirements, refer to your organization's business continuity plan.
- A policy for less frequent snapshot creation with medium retention. With this policy, snapshots are created weekly and retained for one month. Use this policy for application data that is updated frequently and requires a high RPO, but does not contain business-critical data.
- A policy for less frequent snapshot creation with long retention. With this policy, snapshots are created monthly and retained for one year. Use this policy for data that is not updated frequently but must be retained, such as historical customer data or employee records.
- A policy for cloud snapshots only. With this policy, cloud snapshots are created daily and retained for one month. Use this policy to store snapshots of volumes in cloud storage.
- A hybrid policy can be used to create schedules for local snapshots and cloud snapshots. With this policy, local snapshots are created every six hours and retained for one week. Cloud snapshots are created daily and retained for one month. Use this policy for volume data that requires the highest RPO and for storing snapshots of volumes in cloud storage.

Note: Snapshots can also be added and managed with external schedulers like IBM Copy Services Manager.

Managing snapshot policies

An administrator or security administrator can create, define, and assign a snapshot policy to a volume group.

Assigning snapshot policies

Policies can be created and assigned to volume groups from the **Policies** > **Snapshot policies** tab of the management GUI.

If the backup start time is in the past, it will be advanced according to the backup interval until a time in the future is reached. A monthly scheduled snapshot creates snapshot on the same day of each month for days 1-28 based on the start time. If the scheduled day is not available in that month, the last day of the month is considered.

Note: For example, if snapshot policy is specified for every month and the start date is 31 May, the next snapshot date would be 30 June, followed by 31 August. Daylight savings do not affect the backup and retention intervals that are specified by the schedule. For example, if snapshots are triggered daily and occur every day at 2:00 AM before daylight savings start, they are created at 1:00 AM when daylight savings pushes the time ahead by an hour. The backup interval remains constant at 24 hours in this example.

Note: If you change the system time or timezone (for example, when moving domain controllers or correcting time settings) after a schedule has been set up, the change will not affect these schedules. They are based on an internal tick-driven time delta, not the actual clock time delta.

The command-line interface provides the following command to assign snapshot policies:

- **mksnapshotpolicy** creates a snapshot policy.
- **mkvolumegroup** assigns a snapshot policy to a volume group at the time of creation.
- **chvolumegroup** changes or assigns the snapshot policies on an existing volume group.

If you assign a policy to a volume group that already has a policy associated with it, the newly assigned policy replaces the existing policy.

When a snapshot policy is added to a volume group, you can specify the schedule to create a Safeguarded snapshot. For example, `chvolumegroup -snapshotpolicy 0 -safeguarded MyGrp`

Deleting snapshot policies

Use the **rmsnapshotpolicy** command in the command-line interface to delete the user-defined snapshot policies.

The predefined policies cannot be deleted.

Security administration privileges are required to delete policies that are associated with volume groups.

Changing snapshot policies

Use the **chsnapshotpolicy** command in the command-line interface to change the name of the user-defined snapshot policy. The name can also be changed from the **Policies** > **Snapshot policies** tab in the management GUI.

Suspending snapshot policies

If the source volumes are compromised during a cyberattack or during a disaster recovery scenario, any snapshot that is taken during an attack can contain malicious data or software. The system supports suspending snapshot policies until the compromised data can be determined. This can be done at a volume group level by using the **Volumes** > **Volume groups** panel on the GUI, or the **chvolumegroupsnapshotpolicy** command.

Alternatively, all snapshot policies on the system can be suspended using the **chsystem** command. Snapshots can be used for cloning or restoring while any snapshot policies are suspended.

Replication policies

A replication policy defines how replication is configured between systems.

The topology of a replication policy specifies the type of replication to be used. The following topologies are available:

High availability

- This type of replication policy is used with storage partitions.
- The policy specifies the systems that are used for .

  The locations of the systems in the policy define the layout of the systems seen in the management GUI. The system in location 1 is shown on the left, and the system in location 2 is shown on the right. This layout remains fixed and is not affected by the direction of replication between systems, simplifying monitoring and management from either of the systems.
- This type of policy can be used for 2-site high availability and 3-site (high availability with disaster recovery) configurations.
- This type of replication policy will be configured with highly available snapshots enabled by default if both systems support this feature.
- You can start from the panel in the GUI to add or remove high availability for a partition.

Disaster recovery policies

policies can be categorized into two types: synchronous and asynchronous. Synchronous or asynchronous disaster recovery can be used with volume groups that are in a storage partition. The storage partitions are linked to define the systems that are used for disaster recovery replication. It can be used for 2-site disaster recovery or 3-site (high availability with asynchronous disaster recovery) configurations. You can use the **Volume groups** panel in a storage partition to add or remove replication for a volume group. You can also switch between synchronous and asynchronous policies.

Synchronous disaster recovery policies

Synchronous disaster recovery policy provides a high-performance solution for protecting data across sites. In synchronous replication, data is written to both the primary and secondary sites at the same time, ensuring consistency between locations. The storage partitions are linked to define the systems that are used for disaster recovery replication. Synchronous replication is not supported over TCP/IP connections.

Asynchronous disaster recovery policies

Asynchronous replication helps protect data across sites without requiring high-bandwidth connectivity. Unlike synchronous replication, data writes occur first at the primary site and then replicate to the secondary site with a slight delay. This design reduces network load while still meeting your **Recovery Point Objective (RPO)**.

2-site asynchronous disaster recovery

- This type of replication policy is used with volume groups that are not in a storage partition.
- This type of policy is suitable for environments that cannot use storage partitions for disaster recovery. For example, replication for hosts that are not currently supported in a partition. For more information, see .
- The policy specifies the systems and the I/O groups in those systems that are used for replication.
  - The locations of the systems in the policy define the layout of the systems seen in the management GUI. The system in location 1 is shown on the left, and the system in location 2 is shown on the right. This layout remains fixed and is not affected by the direction of replication between systems, simplifying monitoring and management from either of the systems.
- This type of policy can be used for 2-site replication only.
- You can use the **Policies** tab of the **Volumes** > **Volume groups** panel in the management GUI to add or remove this type of replication for a volume group.
- The volumes in the recovery volume group are offline and inaccessible if using a replication policy with a 2-site-async-dr topology.

Both types of replication policies for asynchronous replication define a maximum acceptable recovery point for the recovery copy. An alert is generated if the recovery point exceeds this value.

Replication policies cannot be changed after they are created. However, a new policy can be used to replace the existing one if changes are required.

Chapter 8. Replication

Policy-based replication can be used to implement high availability and disaster recovery solutions between systems in different locations.

The possible configurations include:

High availability
Active-active high availability between two independent systems.

Disaster recovery
Asynchronous or synchronous replication to another system for disaster recovery.

High availability and disaster recovery
Active-active high availability between two independent systems, with disaster recovery replication to a third system.

Note: On some platforms, software features, or combinations of software features are limited by available hardware. If applicable, the limitations are documented in a *Software feature limits* page in the relevant Hardware Guide.

High availability

Policy-based high availability provides an active-active high availability (HA) solution using two independent systems in different locations. The storage will automatically remain accessible to hosts if there is an event that makes one of the systems unavailable.

Storage partitions are used to implement policy-based high availability and disaster recovery solutions. Partitions contain volumes, volume groups, hosts, host clusters, and host-to-volume mappings. For more information, see Storage partitions.

When an HA replication policy is associated with a partition, all the objects contained in the partition are automatically configured across both systems that are defined in the replication policy. Hosts that are created in a partition will be able to discover paths to the mapped volumes through both systems. Hosts that are configured with a location prioritize the use of a co-located system for I/O.

Using policy-based high availability, everything in a storage partition is highly available. When required, both management and I/O failover between systems to maintain access.

Each partition that is associated with an HA replication policy has two properties:

- *Preferred management system*
- *Active management system*

The active management system is the system from which all configuration tasks must be performed. In the event of a failure that affects the current active management system, the active management system role will automatically failover to the other system in the partnership.

The preferred management system is the system that you would like to be the active management system under ideal conditions and is the system that is preferred to keep the partition accessible if the two HA systems are disconnected. The preferred management system can be changed from the current active management system. Different partitions can be configured with different preferred management systems.

If the active management system and the preferred management system are not the same system, the system will automatically failover the active management system role to the preferred management system when it is able to.

As part of the high availability solution, a partnership IP quorum application must be deployed on a server that is separate from the storage systems. The partnership IP quorum application is used to arbitrate between systems in a partnership, to determine which system will operate as the active management system in the event of a failure.

© Copyright IBM Corp. 2025

Related concepts

"IP quorum application" on page 504
The IP quorum application is a Java™ application that runs on a server that is separate from the storage system.

Related information

"Partnerships" on page 206
Partnerships are used to connect systems together to enable migration, data replication, and high-availability solutions.

"Policies" on page 185
A policy defines a set of rules that can be used to simplify storage management. The system supports several different types of policies.

"Storage partitions" on page 243
A storage partition groups the configuration for one or more applications that need to be managed together.

Planning high availability

When you plan for replication for high availability, check that your environment supports it and that you understand the limits of configuration.

The following configuration limits apply to policy-based high availability:

- A limit of one I/O group can be configured in the system.
- A maximum of one partnership can be used for high availability.
- A partnership cannot be used for replication for both high availability and disaster recovery.
- SCSI hosts connected through Fibre Channel or iSCSI can be configured in a storage partition that is configured for high availability.
- NVMe-TCP and NVMe-FC hosts can be configured in a storage partition that is configured for high availability.
- Host clusters can be configured in a storage partition that is configured for high availability if both systems are running IBM Storage Virtualize 9.1.0 or later.
- SCSI Persistent Reservations are supported for Fibre Channel hosts if both systems are running IBM Storage Virtualize 8.7.2 or later.
- SCSI Persistent Reservations are supported for iSCSI hosts if both systems are running IBM Storage Virtualize 8.7.3 or later.
- NVMe reservations are not supported for NVMe-FC and NVMe-TCP hosts.
- Volumes that are configured for high availability cannot be mapped to other IBM Storage Virtualize systems.

Policy-based high availability requires a partnership to be configured between the systems by using either Fibre Channel or Ethernet (short distances that use RDMA) connectivity.

The systems should be sufficiently far apart such that anticipated disruptions affect only one location. The systems must be 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 millisecond round-trip time (RTT) is supported for the inter-site link between systems.

IP connectivity is required between systems for management. Use of Network Address Translation (NAT) between systems is not supported. The following TCP ports are required to be accessible on the system management IP address:

- Port 7443 for REST API access between systems.
- Replication setup that uses the management GUI requires access to both systems from the host where the web browser is running. Ensure that any firewalls between the web browser and storage systems allow inbound traffic to port 7443 and 443 on the system management IP address.

In addition, both systems must have IP network connectivity to an external partnership IP quorum application to arbitrate if a site or system failure occurs. Multiple quorum applications can be deployed for redundancy.

Hosts must use an ALUA-compliant multipath policy. Policy-based high availability supports a reduced set of the operating systems supported by IBM Storage Virtualize systems. For more information, see Supported host operating systems for high availability.

Hosts can be configured with a non-uniform topology, where each host has connectivity to only one of the locations in the high availability configuration. To accurately report the high availability status of these hosts, specify a location for each host, and group those hosts in a host cluster with the uniform connectivity parameter set to *No*. Add the hosts and the host cluster to a storage partition with a replication policy that identifies the system at each location. Hosts should be defined in pairs, one for each location to provide redundant access by the host cluster in the case of an individual system outage.

The following restrictions apply to policy-based high availability:

- When a storage partition is configured for high availability, names of partitions, hosts, volume groups, and volumes in the partition can be changed only if all systems are running IBM Storage Virtualize 9.1.0 or later. You can rename an object on the active management system and the changed object name applies to all copies.
- Policy-based high availability cannot be used with volumes that are:
  - Image mode
  - Cache disabled
  - Configured to use Transparent Cloud Tiering (TCT)
- A volume in a storage partition with a replication policy assigned cannot be:
  - Shrunk. Volume expansion is supported if all systems are running IBM Storage Virtualize 9.1.0 or later.
  - Migrated to image mode, or have an image mode copy added.
- A volume can be moved between volume groups within a partition if all systems are running IBM Storage Virtualize 9.1.0 or later.

Implementing high availability

Policy-based high availability can be configured starting from the **Storage partitions** panel in the management GUI, navigating to the partition to be modified.

The GUI guides you through the required steps to:

- Configure a partnership between systems, or select an existing partnership.
- Link storage pools between systems.
- Configure a partnership IP quorum application.

Configuration and monitoring can also be performed by using the REST API, or the CLI.

Creating a new high availability configuration 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: “Partnerships” on page 336.
- If not already configured, follow the link to set up a linked storage pool: “Pool links” on page 207.
- If not already configured, follow the link to set up a partnership IP quorum application: “IP quorum application” on page 504.

Follow the steps to create a new policy-based high availability (HA) configuration by using command-line interface:

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 **mkpartition**.
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**.

   Note: To prioritize a colocated system for I/O, specify a location for the host.
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. You can monitor and change the active and preferred management systems using the **lspartition** and **chpartition** commands.

Adding high availability to an existing configuration by using the CLI

It is recommended that policy-based high availability (HA) is configured and managed by using the GUI. The GUI helps you create a new configuration or modify an existing configuration. You can also configure policy-based HA by using the REST API or the CLI.

Prerequisites to configure your systems to support the HA solution:

- If not already configured, follow “Partnerships” on page 336 to set up a partnership.
- If not already configured, follow “Pool links” on page 207 to set up a linked storage pool.
- If not already configured, follow “IP quorum application” on page 504 to set up an IP quorum application.

Follow these steps to add policy-based high availability to an existing configuration by using the CLI:

1. Volumes must be in a volume group to add high availability. If necessary, to create a new volume group, enter the following command:

```
mkvolumegroup
```

For more information, see **mkvoluemgroup command**.

2. If needed, use **lsvolumegroup** to see the ID of the volume group. For more information, see **lsvolumegroup** command.
3. To add the volumes that you want to be HA to the volume group, enter the following command:

```
chvdisk –volumegroup <volume group ID or name> <volume ID or name>
```

For more information, see **chvdisk command**.

4. To create a draft partition, enter the following command:

```
mkpartion –draft –name my_partition
```

For more information, see **mkpartition** command.

Note: lspartition shows **draft=yes** for a draft partition. A draft partition can be used to construct a storage partition, but has no effect until it is published.

5. To add a volume group to the draft partition, enter the following command:

```
chvolumegroup –draftpartition my_partition <volume group ID or name>
```

For more information, see **chvolumegroup command**. Adding the volume group to the draft partition automatically adds any dependent hosts and other volume groups to the draft partition. It does not automatically add volumes that are not in volume groups.

6. The draft partition must be published to make it into a storage partition that is eligible for use in a high availability configuration. The draft partition can be published only when all the required volumes are in volume groups. To publish a draft partition, enter the following command:

```
chpartition –publish my_partition
```

For more information, see **chpartition command**.

7. If required, you can create another draft partition and repeat the process. Only one draft partition is allowed for each system.
8. If required, partitions can be merged. To merge a second partition into the first partition, enter the following command:

```
mergepartition –targetpartition <name or ID of other storage partition> my_partition
```

For more information, see **mergepartition command**.

9. 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>
-location2system <name or ID of other system> -location1iogrp 0 -location2iogrp 0
```

If necessary, use **lspartnership** to get the name or ID of each system. For more information, see **mkreplicationpolicy command**.

10. To associate the published storage partition with the replication policy, enter the following command:

```
chpartition –replicationpolicy <HA policy ID or name> my_partition
```

Note: As long as the partition is associated with the replication policy, the volumes, volume groups, hosts and host-to-volume mappings that are contained in the partition are maintained across both systems.

11. Hosts can now discover paths to volumes at both locations. Use **lspartition** to monitor the health of the HA solution. Use **lspartnership** to monitor connectivity between the systems. Use **chpartition** to change the active and preferred management systems.

Note: To prioritize a colocated system for I/O, use **chhost** to specify a location for the host.

Host operating systems for high availability

Policy-based high availability has been verified for compatibility with a subset of the operating systems that are compatible with IBM Storage Virtualize.

Fibre Channel SCSI hosts

- IBM AIX (version 7.2 and later).
- IBM i 7.4 TR11 or later, IBM i 7.5 TR5 or later, IBM i 7.6 or later. Fabric-attached Fibre Channel only.
- Linux
  - Red Hat Enterprise Linux, including with KVM (all supported versions).
  - SUSE Linux Enterprise (all supported versions).
  - Linux on IBM Z (all supported versions and distributions).
- Microsoft Windows Server, including Hyper-V (all supported versions). For more information, see Storage tests on a failover cluster might not discover all shared LUNs.
- Oracle Solaris.
- Proxmox VE 9.x.
- HPE VM Essentials 8.1.1.
- VMware ESXi (all supported versions).

iSCSI hosts

- Linux
  - Red Hat Enterprise Linux, including with KVM (all supported versions).
  - SUSE Linux Enterprise (all supported versions).
  - Linux on IBM Z (all supported versions and distributions).
- Microsoft Windows Server, including Hyper-V (all supported versions). For more information, see Storage tests on a failover cluster might not discover all shared LUNs.
- Proxmox VE 9.x.
- VMware ESXi (all supported versions).

NVMe hosts

The VMware ESXi 8.0 operating system is supported only with NVMe/TCP or NVMe/FC. For information about NVMe/TCP support, see VMware ESXi installation and configuration for NVMe over RDMA or TCP hosts. For information about NVMe/FC support, see VMware ESXi installation and configuration for NVMe over Fibre Channel hosts.

Operating system specific information

Proxmox VE

Configure high availability with the following settings:

- Configure storage for volumes that are configured for high availability according to Proxmox VE recommendations. Proxmox VE recommends using plain LVM with dedicated volume groups for shared LUNs. For more information, see Proxmox VE Storage documentation.
- Configure a uniform storage configuration so that each host can connect to both storage systems.
- Remove the LVM system ID from all shared volume groups to enable cross-node activation during failover. Run **vgchange --systemid "" <vgname>**.
- Register storage pools with the `--shared 1` flag so that Proxmox automatically activates them on the surviving node.

Limitations:

- In a non-uniform configuration, Proxmox high availability does not automatically restart a virtual machine on another cluster node after a storage-induced failure.
- Proxmox high availability recovers only from quorum loss. It does not recover from storage loss.

Requirements and considerations

- Hosts that are configured for high availability must have the Generic host type. Change existing hosts that use deprecated host types to the Generic host type.
- Volumes configured for high availability cannot be mapped to other IBM Storage Virtualize systems.
- Compatibility with the listed operating systems requires the native multipath driver. Third-party multipath solutions are not qualified.
- For more information about host compatibility, see IBM System Storage Interoperation Center (SSIC).

Related concepts

"Hosts that run the Linux operating system" on page 124
You can attach hosts running the Linux operating system via Fibre Channel.

"IBM Power Systems AIX" on page 119
You can attach IBM Power Systems AIX hosts via Fibre Channel.

"IBM z Systems that run the Linux operating system" on page 122
You can attach IBM z Systems® that run the Linux operating system via Fibre Channel.

"Hosts that run the VMware ESXi operating system" on page 132
Fibre Channel attachments between the system and various guest host operating systems that run on the ESXi operating system have specific requirements and other guidelines that need to be followed.

"Hosts that run the Microsoft Hyper-V operating system" on page 134
You can attach hosts running the Microsoft Hyper-V operating system via Fibre Channel. This allows you to run multiple guest operating systems on the server.

Highly available snapshots

A highly available (HA) snapshot is a volume group snapshot that exists on both systems with the same data, and can be restored without suspending high availability (HA) for the partition.

A local snapshot is a volume group snapshot that exists on a single system. Restoring from a local snapshot results in suspension of high availability for the partition until a resynchronization of the restored data is completed.

HA snapshots can be used in any configuration that supports high availability, including as part of a 3-site configuration for high availability and disaster recovery.

If both systems support this feature, HA snapshots are enabled by default when configuring high availability on a partition by using the GUI, or when creating a replication policy for HA by using the CLI or REST API.

On the management GUI, the HA status is displayed on the **Snapshots** tab of the volume group. On the CLI or REST API, the **ha_status** field indicates whether a snapshot is HA or local.

You can use the management GUI to enable or disable HA snapshots from the replication tile on the partition. Enabling or disabling HA snapshots affects new snapshots only and any existing snapshots are unchanged. To enable or disable HA snapshots by using the CLI, create a new replication policy with the desired properties and then assign the new replication policy to the partition.

Using HA snapshots

With HA snapshots enabled:

- Snapshots can be taken on the active management system of an HA partition:

  - If HA is established, HA snapshots are created.
  - If HA is not established, local snapshots are created.
- The snapshot policy that is associated with the volume group on the active management system defines the scheduling of snapshots.
- Configuring a snapshot policy on both HA systems ensures that snapshots continue to be triggered if the active management system changes. Safeguarded and retention settings for snapshots are defined by the settings of the volume group on the active management system.

HA snapshots can be removed from the active management system of the HA partition and it results in removal of the snapshots on both systems.

An HA snapshot can be restored from the active management system of the HA partition, either the whole volume group or a single volume:

- If HA is established, the restore is applied on both systems without loss of HA.
- If HA is not established, the restore is applied on the active management system and gets synchronized with other system when HA is reestablished.

Clones or thin-clones can be created from an HA snapshot outside of the HA partition on either of the HA systems and are created only on the system where the command is run.

If a partition has HA snapshots, removing HA from that partition results in HA snapshots being converted to local snapshots.

HA thin clone volumes

An HA thin clone volume is a host-accessible copy of a specific volume from a HA snapshot. An HA thin clone volume exists on both systems with the same data.

HA thin clone volumes can be used in any configuration that supports high availability, including as part of a 3-site configuration for high availability and disaster recovery.

An HA thin clone volume can be created on the active management system of an HA partition:

- If the volume type is **thinclone**.
- If the source snapshot is an HA snapshot.
- If the thin clone volume is placed into a volume group within the same HA partition.

HA thin clone volumes can be refreshed. The refresh is applied on the active management system and gets synchronized with the non active management system when HA is reestablished.

HA thin clone volumes can be converted to clones. The convert to clone happens only on the system where the command was run.

Related information

“Snapshots” on page 101
A volume group snapshot is a consistent point-in-time copy of a volume group.

Disaster recovery

Policy-based replication can be used to implement replication between systems for disaster recovery.

In the event of a disaster that prevents access to the production storage, host access can be enabled to the recovery copy. When the original production system is accessible again, replication can be restarted in either direction to restore redundancy.

Replication for disaster recovery can be configured with or without the use of storage partitions. Using storage partitions allows high availability to be easily added later and will also allow the use of future FlashSystem grid functionality. Storage partitions can be used where the host environment meets the supported configurations. For more information, see Storage partitions.

Existing configurations using replication for disaster recovery without storage partitions can be migrated into a storage partition by using the CLI. This allows an existing 2-site disaster recovery (DR) configuration to be extended to a 3-site HA and DR configuration while continuing to support disaster recovery operations and without requiring resynchronization of the recovery copy.

Replication for disaster recovery offers the flexibility to use different replication policies for different volume groups to match business requirements for different recovery point objectives, even on the same partnership or within the same storage partition.

Planning replication for disaster recovery

Before implementing replication for disaster recovery, you must determine your recovery point objectives, application requirements, and consider the connectivity available between locations.

To implement replication for disaster recovery, consider the following:

- Determine the round-trip time (RTT) between the systems. As the RTT increases, the maximum throughput decreases.
- Determine the change rate of data for applications storing data on the system. Environments with high change rates and low bandwidth may be unable to achieve low RPOs.
- The use of replication protects data from a large-scale disaster on the site or storage systems. Consider using safeguarded snapshots in addition to replication to also provide protection against data corruption and cyber attacks.

Replication for disaster recovery requires a partnership to be configured between the systems. Partnerships for asynchronous replication can use any of the supported connectivity types. Partnerships can be configured with up to three other systems for disaster recovery.

IP connectivity is required between systems for management. Use of Network Address Translation (NAT) between systems is not supported. The following TCP ports are required to be accessible on the system management IP address:

- Ports 6443, 7443, and 443 for REST API and GUI access between systems on the default management IP addresses.
- Replication setup using the management GUI requires access to both systems from the host where the web browser is running. Ensure that any firewalls between the web browser and storage systems allow inbound traffic to port 443 on the system management IP address.
- Additionally, IP replication requires inbound and outbound access using port 3260 on the system management IP and port 3265 on the IP addresses for the data path connections.

Policy-based replication for disaster recovery can be used with all host operating systems that are compatible with IBM Storage Virtualize systems.

For more information about host compatibility, see IBM System Storage Interoperation Center (SSIC).

For the best performance, configure the host multipath drivers to use SCSI ALUA or NVMe ANA.

With policy-based replication, you can replicate thin-clones and FlashCopy target volumes to the recovery system.

- Snapshots cannot be replicated directly. Thin-clones can be refreshed from snapshots and the deltas are copied automatically to the recovery system.
- FlashCopy maps cannot be started if the FlashCopy target volume is replicated.

The following restrictions apply to policy-based replication for disaster recovery:

- The name of a volume group, and names of volumes in the volume group can be changed only if all systems are running IBM Storage Virtualize 9.1.0 or later. You can rename an object on the production system only and the changed object name applies to all copies.
- Policy-based replication cannot be used with volumes that are configured to use Transparent Cloud Tiering (TCT).

- Image mode volumes with cache disabled can be replicated to import data from external storage. Cache must be enabled after the replica is established. After cache is enabled in read/write mode, it cannot be disabled.
- A volume in a volume group with a replication policy assigned cannot be:
  - Shrunk. Volume expansion is supported if all systems are running IBM Storage Virtualize 9.1.0 or later.
  - Migrated to image mode, or have an image mode copy added.
  - Moved to a different I/O group.
- All volumes in a volume group must be in the same I/O group.
- A volume can be moved between volume groups if the current group and the new group are not both configured with replication for disaster recovery, and if all systems are running IBM Storage Virtualize 9.1.0 or later.

Implementing replication for disaster recovery

You can implement replication for disaster recovery with policy-based replication.

Policy-based replication uses volume groups and replication policies to deploy and manage replication for disaster recovery.

Configuring disaster recovery using storage partitions

From the **Storage partitions** panel, navigate to the partition to be configured.

The GUI guides you through the required steps to:

- Configure a partnership between systems, or select an existing partnership.
- Link storage pools between systems.
- Create and link a partition on the remote system for disaster recovery.

After configuring the partition, navigate to the **Volume group** view within the partition to create and assign replication policies on individual volume groups as required.

Configuring disaster recovery without using storage partitions

If not configured already, navigate to the **Settings** > **System** > **Partnerships** page to create a partnership between systems for replication.

From the **Volumes** > **Volume groups** panel, navigate to the volume group to be configured and select the **Policies** tab.

The GUI guides you through the required steps to:

- Select an existing partnership.
- Link storage pools between systems.
- Create a new replication policy or select an existing policy to assign to the volume group.

Managing replication for disaster recovery

You can manage replication for disaster recovery by using the management GUI, CLI, or REST API.

When a replication policy for disaster recovery is assigned to a volume group, the volume group operates in one of three modes:

Production
: The volumes in a production volume group are accessible for host I/O and configuration changes are allowed to the volume group. This system acts as the source for replication to the target volumes on the recovery system. The production volume group contains the most up-to-date copy of data for host access.

Recovery
The volumes in a recovery volume group are the target for replication. Read-only access from a host is permitted if the volume group has a valid recovery point and is replicated using a disaster recovery replication policy (used with volume groups that are in a storage partition), otherwise the volumes are inaccessible for host I/O; starting a recovery test or enabling independent access to a volume group will transition the volumes into an online and read-writeable state. While a volume is in a read-only state it is not guaranteed to be completely consistent at all times as the data will be changing as writes are replicated from the production copy, but it is sufficient for discovering the filesystem and mounting to a server. Configuration changes are not allowed on the recovery copy, but changes made to the production copy are copied to the recovery copy.

Independent
If access is enabled for the recovery copy of a volume group, such as during a disaster, each copy of the volume group is independently accessible for I/O and configuration changes are allowed on both. Replication is suspended in this state. While independent, any data or configuration changes that are made on either system are not copied to the other system. When replication is restarted, the system automatically updates the data and the configuration on the recovery volume group to match the production volume group.

Replication policies do not define a direction for replication. The initial direction of replication is determined when a replication policy is associated with a volume group or a volume group is created specifying a replication policy. The system where the replication policy is assigned starts as the production system.

Removing a replication policy from the production copy of a volume group causes the recovery volume group and its volumes to be deleted. Independent access must be enabled on the recovery copy first to keep the copy, and the replication policy can then be removed from the volume group without affecting the other copy.

The replication policy assigned to a volume group can be replaced with a different replication policy without requiring any resynchronization, for example, to use a policy with a different recovery point objective alert for the volume group. The new policy must specify the same systems in the same locations.

Managing and monitoring disaster recovery for a volume group

Using the management GUI
To configure disaster recovery, assign a replication policy to the volume group. Using the management GUI, navigate to the **Volume groups** page and select the desired volume group, then select the **Policies** tab. When a replication policy is assigned to the volume group, this page is used for managing and monitoring replication.

Using the command-line interface
To assign a replication policy to a volume group using the CLI, use the following command:

```
chvolumegroup -replicationpolicy <policy> <volumegroup>
```

To unassign a replication policy from a volume group, use the following command:

```
chvolumegroup -noreplicationpolicy <volumegroup>
```

To view a list of all the volume groups that are configured for replication, use the following command:

```
lsvolumegroupreplication
```

To view the detailed status of replication for an individual volume group, use the following command:

```
lsvolumegroupreplication <volumegroup>
```

Accessing the recovery copy of a volume group

Using the management GUI

1. Log in to the system with the recovery copy.
2. Navigate to the **Volume groups** page and select the desired volume group.
3. Select the **Policies** tab and review the recovery point for the copy.
4. If the recovery point is acceptable, click **Enable access** to enable independent access to the volume group.
5. If the production system is inaccessible, the production copy is made independent when connectivity is restored. The recovery copy will always become independent and available for use immediately.

Using the command-line interface

You can use the **chvolumegroupreplication** command to enable independent access when logged in to the system with the recovery copy:

```
chvolumegroupreplication -mode independent <volumegroup>
```

Restarting replication from the independent state

When the replication mode for a volume group is in the independent state, replication can be restarted in either direction.

The replication policy assigned to the volume group on the system you want to use as the production copy will be used to define the replication requirements, such as the RPO alert for the group. If the volume group was previously the recovery copy in a storage partition, it might not yet have a replication policy assigned. You must assign a replication policy to the volume group before you can restart replication from this system. If you use the management GUI to restart replication, the GUI guides you through the required steps to create a new replication policy or select an existing policy to assign to the volume group as part of the restart task.

Using the management GUI

1. Log in to the system with the volume group you want to use as the new production copy.
2. Navigate to the **Volume groups** page and select the desired volume group.
3. Select the **Policies** tab.
4. Click the **Restart replication** action to restart replication using the volume group on the local system as the production copy. Any changes made to the copy on the remote system will be overwritten when replication is restarted.

Using the command-line interface

You can use the **chvolumegroupreplication** command to restart replication when logged in to the system with the copy you want to use as the production copy:

```
chvolumegroupreplication -mode production <volumegroup>
```

Migrating replication for disaster recovery to use partitions

Existing configurations that use replication for disaster recovery without storage partitions can be migrated into a storage partition by using the CLI. This allows an existing 2-site DR configuration to be extended to a 3-site HA+DR configuration while continuing to support disaster recovery operations and without requiring resynchronization of the recovery copy.

This procedure requires all the volume groups that are to be migrated to the same partition to replicate in the same direction.

Using the command-line interface

1. On each system:
   a. Use **mkpartition -draft -name <partition_name>** to create a new draft partition.

      Note: If appropriate for your environment, consider naming the partitions differently to allow easier identification of which partition is typically used for disaster recovery during normal running.

   b. Use **chvolumegroup -draftpartition <partition> <volumegroup>** to move the volume groups to be migrated into the new draft partition.

      Note: Moving an object into a draft partition may automatically add associated objects to the partition too. Review the complete set of objects in the draft partition before you publish it.

   c. Publish the partition by using: **chpartition -publish <partition>**
2. On the system with the recovery volume groups:
   a. Identify the **uuid** value of the partition that will be used for recovery by using: **lspartition <recovery_partition>**
3. On the system with the production volume groups:
   a. Configure a link for disaster recovery between the partitions by using: **chpartition -makedrlink -remotedrlinkedpartitionuuid <recovery_partition_uuid> -remotesystem <recovery_system_id/name> <production_partition_id/name>**

      Note: The remote partition being linked to must contain all of the associated remote volume groups for all disaster recovery volume groups in the local partition.

   b. For each 2-site-async-dr policy that needs to be replaced, create and assign a new replication policy with topology **async-dr** in the partition:
      i) Use the **lsvolumegroupreplication** command to identify the current replication policy assigned to each volume group to be migrated.
      ii) Use the **lsreplicationpolicy** command to identify the RPO alert value for each policy.
      iii) Create the new replication policy by using: **mkreplicationpolicy -name <name> -topology async-dr -partition <production_partition_id/name> -rpoalert <rpoalert>**
      iv) Assign the new replication policy to each volume group by using: **chvolumegroup -replicationpolicy <policy> <volumegroup>**
      v) Optionally, remove any unused **2-site-async-dr** replication policies by using: **rmreplicationpolicy <policy>**

After the CLI procedure above has been completed, further management and monitoring can be performed using the GUI, the CLI, or the REST API.

High availability and disaster recovery (3-site replication)

Combining policy-based high availability and policy-based replication for disaster recovery enables a 3-site replication solution for the most mission-critical workloads. You can manage the high availability and disaster recovery replication from the **Storage partition** page on the management GUI.

Policy-based HA provides a zero recovery point objective (RPO) solution by using two independent storage systems with synchronous replication across metro-area distances and active-active access for hosts at both sites with seamless failover for zero recovery time objective (RTO). Asynchronous replication (RPO > 0) to a third system provides disaster recovery capability (RTO > 0) when a metro-area disaster.

Three-site replication for a storage partition can be configured in any of the following ways:

- Starting with an existing 2-site high availability configuration, then adding replication for disaster recovery at a third side.
- Starting with an existing 2-site configuration for disaster recovery using storage partitions, then adding high availability to either of the partitions.

- Configuring 3-site replication for high availability and disaster recovery in a single operation.

On systems participating in 3-site policy-based replication, highly available snapshots can be enabled. For details, see “Highly available snapshots” on page 197.

The planning, implementation, management, and monitoring information for high availability and disaster recovery, when used individually, also applies when used in a 3-site configuration.

- Planning high availability
- Planning replication for disaster recovery

Partnerships must be configured between all three systems. High availability requires either a Fibre Channel or an RDMA-based Ethernet partnership. The partnership for disaster recovery also supports the use of a long-distance TCP partnership. However, for all partnership types, the disaster recovery partnership must have enough bandwidth to handle the write rate on the production volume. The different partnerships between systems can use different types of connectivity.

The use of separate node ports is recommended to isolate the different types of traffic for:

- High availability partnerships
- Disaster recovery partnerships
- Hosts and external storage, which can also be separated further if required

Example error scenarios

This table lists examples of possible error scenarios and the expected behavior in a 3-site solution, starting from an initial healthy state.

Table 34. Example error scenarios

Scenario	High availability	Disaster recovery
Loss of connectivity between the non-preferred HA system and the DR system	No impact	No impact
Loss of connectivity between the preferred HA system and the DR system	The active management system will change to the non-preferred system to maintain replication to the DR system. The storage partition remains active-active with no interruption to host access. The preferred system will automatically become the active management system again when connectivity to the DR system is restored.	No impact
Loss of connectivity between the HA systems	If the connectivity issue lasts more than 3 seconds, HA will be suspended. The volumes within the storage partition will remain accessible through one of the two HA systems.	No impact. Replication will continue between the active management system and the DR system.

Table 34. Example error scenarios (continued)

Scenario	High availability	Disaster recovery
Both HA systems lose connectivity to the DR system	No impact	Replication is suspended. Replication will resume automatically when connectivity is restored.
One or both of the HA systems lose connectivity to all quorum applications	No impact. However, the loss of quorum will prevent the non-preferred system from taking over the partition if the HA systems lose connectivity, or if the preferred system becomes unavailable.	No impact
The preferred HA system loses connectivity to the non-preferred system and all quorum applications at the same time. The non-preferred system still has connectivity to the active quorum application.	The storage partition will continue to be accessible on the non-preferred system. HA will automatically re-establish when connectivity is restored.	Replication will continue between the active management system and the DR system.
The preferred HA system is unavailable (for example, power failure, hardware failure)	The storage partition will continue to be accessible on the non-preferred system. HA will automatically re-establish when the preferred system is available.	Replication will continue between the active management system and the DR system.
The non-preferred HA system is unavailable (for example, power failure, hardware failure)	The storage partition will continue to be accessible on the preferred system. HA will automatically re-establish when the non-preferred system is available.	Replication will continue between the active management system and the DR system.
The DR system is unavailable (for example, power failure, hardware failure)	No impact	Replication is suspended. Replication will resume automatically when the DR system is available.
Loss of management IP connectivity only between the HA systems	No impact on host I/O. If high availability is established for the partition, configuration changes that would result in the loss of high availability are blocked. If high availability is not established, configuration changes are permitted and HA will re-establish automatically when connectivity is restored.	No impact.

Table 34. Example error scenarios (continued)		
Scenario	**High availability**	**Disaster recovery**
Loss of management IP connectivity only between the active management system and the DR system	No impact.	Configuration changes made to the HA storage partition cannot be replicated to the DR system. Replication will be suspended if configuration changes are required on the DR system. If replication is suspended, it will resume automatically when connectivity is restored.
Access is enabled to the recovery copy of one or more volume groups in the storage partition	No impact.	Replication is suspended for the independent volume groups. Host I/O and configuration changes are allowed on the HA and DR partitions independently.
Replication is restarted to synchronize changes from the DR system to the HA systems	HA is suspended while replication is running from the DR system to the active management system. HA will re-establish automatically when there are no volume groups replicating from the DR system.	No impact.

Partnerships

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

You can create a partnership in the following ways:

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

In addition to the partnership for data transfer between systems, IP connectivity is also required between the management IP addresses of each system to access the REST API for management. Management traffic uses authentication certificates to ensure secure communications between systems.

When you use the management GUI to configure high availability or disaster recovery by using storage partitions, the GUI guides through the steps that are required to configure suitable partnerships between systems. It is not necessary to create partnerships between systems beforehand.

Using the management GUI

When you create a partnership between systems for high availability or disaster recovery by using GUI, certificates are automatically exchanged between systems to create a truststore on each system. No additional steps are required.

Using the command-line interface

When using the command-line interface, you must exchange certificates and create a truststore on each system as a separate activity after you create the partnership between systems.

1. On the first system, use the **`lssystemcertstore`** command to check with certificate to export.

This lists the enabled scope certificates. In case the (scope id = 3) internal_communication is present, then use following command to export the certificate:

```
chsystemcertstore -export -scope internal_communication
```

The certificate is exported to a file named `system_certificate_slot_3.pem`.

Otherwise, use the following command to export the certificate:

```
chsystemcert -export
```

or

```
chsystemcertstore -export -scope default
```

The certificate is exported to a file named */dumps/certificate.pem*.

2. Copy the certificate file from the system to your local workstation by using the Secure Copy (SCP) or SSH File Transfer Protocol (SFTP) client. Then, use the SCP or SFTP to upload the certificate file to the `/tmp/` directory on the second system.
3. On the second system, create the truststore by using the following command:

```
mktruststore -file /tmp/certificate.pem -restapi on -name <remote_cluster_name>
```

Naming the truststore is optional. But, in environments with multiple partnerships or multiple truststores, it helps to identify which truststores are used for which purpose.

4. Repeat the preceding steps to export the certificate from the second system and create a truststore on the first system.

You can use the **lssystemcert** command to view information about the current system certificate and the **lstruststore** command to list the truststores that exist on the system.

Pool links

High availability and disaster recovery replication require the storage pools that are used for creating replicated volumes to be linked between systems. When a volume is created on the active management system for HA, or on the production system for DR, pool links define the pool that is used on the remote system to create the replicated volume.

For SAN Volume Controller systems that use a stretched system topology with mirrored volumes, only the pools that are in site 1 need to be linked.

Use of the management GUI to configure pool links simplifies the process by allowing pools to be linked from a single system without needing to login to the other system. Configuration can also be performed by using the CLI or the REST API, connecting to each system with pools to be linked.

Additionally, if child pools currently exist on one 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 second system, displaying the properties like name, capacity, and provisioning policy from the first system.

When you configure pool links by using the GUI, you can optionally assign a provisioning policy to the pools that you are linking. Provisioning policies control how capacity is provisioned on new volumes that are created 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 desired. The system creates fully provisioned volumes if a provisioning policy is not configured. For more information, see Provisioning policy.

Using the management GUI

When adding high availability or disaster recovery, the management GUI guides you through the required steps to link storage pools between systems, and optionally assign provisioning policies on each linked pool.

To link extra pools later, or remove the links between pools, go to the **Pools** > **Pools** page in the management GUI and select the pool that you want to modify. Right-click on the pool and select the action-prompt to **Add Pool Link for Replication** or **Modify Pool Links for Replication** as required.

Using the command-line interface

Pools on partnered systems are linked if they have the same *replication_pool_link_uid* value. To identify pools that are linked, use the **`lsmdiskgrp`** command on each system.

To modify or remove the links between pools by using the command-line interface, use the **`chmdiskgrp`** command.

Portset links

When you create a host in a storage partition that is configured for high availability, the portset link defines the portset that is used to create the host on the remote system based on the local portset.

Only portsets used for Ethernet or Fibre Channel host attachment can be linked between systems.

Using the management GUI to configure portset links simplifies the process by allowing portsets to be linked from a single system without needing to log in to the other system. Portset links can also be configured by using the CLI or REST API, connecting to each system.

Using the management GUI

To configure high availability for a partition, the management GUI guides you through the steps to link portsets between systems.

To link more portsets later, or manage existing links, select **Settings** > **Networks** > **Ethernet Portsets** or (based on your network type) in the management GUI and select the portset that you want to add or modify.

Using the command-line interface

Portsets on partnered systems are linked if they have the same *replication_portset_link_uid* value. To identify portsets that are linked, use the lsportset command on each system.

To create, modify, or remove the links between portsets by using the command-line interface, use the chportset command.

Replication policies

A replication policy defines how replication is configured between systems.

The topology of a replication policy specifies the type of replication to be used. The following topologies are available:

High availability

- This type of replication policy is used with storage partitions.
- The policy specifies the systems that are used for .

  The locations of the systems in the policy define the layout of the systems seen in the management GUI. The system in location 1 is shown on the left, and the system in location 2 is shown on the right.

This layout remains fixed and is not affected by the direction of replication between systems, simplifying monitoring and management from either of the systems.

- This type of policy can be used for 2-site high availability and 3-site (high availability with disaster recovery) configurations.
- This type of replication policy will be configured with highly available snapshots enabled by default if both systems support this feature.
- You can start from the panel in the GUI to add or remove high availability for a partition.

Disaster recovery policies

policies can be categorized into two types: synchronous and asynchronous. Synchronous or asynchronous disaster recovery can be used with volume groups that are in a storage partition. The storage partitions are linked to define the systems that are used for disaster recovery replication. It can be used for 2-site disaster recovery or 3-site (high availability with asynchronous disaster recovery) configurations. You can use the **Volume groups** panel in a storage partition to add or remove replication for a volume group. You can also switch between synchronous and asynchronous policies.

Synchronous disaster recovery policies

Synchronous disaster recovery policy provides a high-performance solution for protecting data across sites. In synchronous replication, data is written to both the primary and secondary sites at the same time, ensuring consistency between locations. The storage partitions are linked to define the systems that are used for disaster recovery replication. Synchronous replication is not supported over TCP/IP connections.

Asynchronous disaster recovery policies

Asynchronous replication helps protect data across sites without requiring high-bandwidth connectivity. Unlike synchronous replication, data writes occur first at the primary site and then replicate to the secondary site with a slight delay. This design reduces network load while still meeting your **Recovery Point Objective (RPO)**.

2-site asynchronous disaster recovery

- This type of replication policy is used with volume groups that are not in a storage partition.
- This type of policy is suitable for environments that cannot use storage partitions for disaster recovery. For example, replication for hosts that are not currently supported in a partition. For more information, see .
- The policy specifies the systems and the I/O groups in those systems that are used for replication.
  - The locations of the systems in the policy define the layout of the systems seen in the management GUI. The system in location 1 is shown on the left, and the system in location 2 is shown on the right. This layout remains fixed and is not affected by the direction of replication between systems, simplifying monitoring and management from either of the systems.
- This type of policy can be used for 2-site replication only.
- You can use the **Policies** tab of the **Volumes** > **Volume groups** panel in the management GUI to add or remove this type of replication for a volume group.
- The volumes in the recovery volume group are offline and inaccessible if using a replication policy with a 2-site-async-dr topology.

Both types of replication policies for asynchronous replication define a maximum acceptable recovery point for the recovery copy. An alert is generated if the recovery point exceeds this value.

Replication policies cannot be changed after they are created. However, a new policy can be used to replace the existing one if changes are required.

Replication statistics

You can use the metrics that the system collects for policy-based replication to analyze performance issues. Statistics for policy-based replication help you troubleshoot and improve system performance.

Volume group statistics

Policy-based replication operates on volume groups rather than individual volumes. Therefore, most of the replication statistics are recorded against a volume group. The following statistics are logged for volume groups:

Table 35. Statistics collection for volume group

Statistic name	Description
rarp	Indicates the average recovery point in seconds for the volume group since the last time statistics were collected.
rnrw	Indicates the cumulative number of replication writes to the target sent by this node for this volume group.
rnrb	Indicates the cumulative number of blocks that this node has replicated for this volume group.
rwrp	Indicates the worst recovery point in seconds for the volume group since the last time statistics were collected.

Node statistics

For node-level values, policy-based replication adds to the existing node XML statistics file.

Table 36. Statistics collection for node statistics

Statistic name	Description
rjru	Indicates the percentage of journal resources that are used for replicating data, averaged across all threads. Value out of 10000 representing a percentage to two decimal places.
rjrw	Indicates the highest percentage of journal resources that are used across all threads. Value out of 10000 representing a percentage to two decimal places.

Volume statistics

Generally, all statistics regarding replication are reported on the volume groups but many existing tools focus on displaying statistics on the individual volumes.

Table 37. Statistics collection for volume statistics

Statistic name	Description
varp	Indicates the average recovery point in seconds for the volume group for this volume since the last time statistics were collected.
vwrp	Indicates the worst recovery point in seconds for the volume group for this volume since the last time statistics were collected.

Chapter 9. Access

The Access menu in the GUI maintains who can log in to the system, defines the access rights to the user, and tracks what was done by each privileged user to the system.

Ownership groups

An ownership group defines a subset of users and objects within the system.

Ownership groups restrict access for users in the ownership group to only those objects that are defined within that ownership group. An *owned object* can belong to one ownership group. Users in an ownership group are restricted to viewing and managing objects within their ownership group. These restricted users can access only this subset of objects. Unrestricted users, who are not in an ownership group, can continue to view or manage all the objects on the system based on their defined user role, including objects within ownership groups. The following objects can be assigned to ownership groups:

- Child pools
- Volumes
- Volume groups
- Hosts
- Host clusters
- Host mappings
- FlashCopy mappings
- FlashCopy consistency groups
- User groups
- Portsets

Two basic use cases can be applied to using ownership groups on the system. The first use case for using ownership groups is on a system where new objects are created within the ownership group. There can also be other existing objects on the system that are not in the ownership group. The second use case for using ownership groups is on a system where these supported objects are already configured and you want to migrate these objects to use ownership groups.

When a user group is assigned to an ownership group, the users in that user group retain their role, but are restricted to only those resources within the same ownership group. The role that is associated with a user group can define the operations on the system and the ownership group can further limit access to individual resources. For example, you can configure a user group with the Copy Operator role, which limits user access to FlashCopy operations. Access to individual resources, such as a specific FlashCopy consistency group, can be further restricted by assigning it to an ownership group.

Inheriting ownership

Generally two basic inheritance rules that apply to all objects that can be associated to an ownership group:

1. Owned objects inherit the ownership group of the owner who creates those objects.
2. Only users with Security Administrator role can create, manage, or assign ownership group to new objects or existing objects.

Depending on the type of resource, ownership groups can be defined explicitly or inherited from these explicitly defined objects.

The following graphic shows how different objects inherit ownership from ownership groups:

© Copyright IBM Corp. 2025

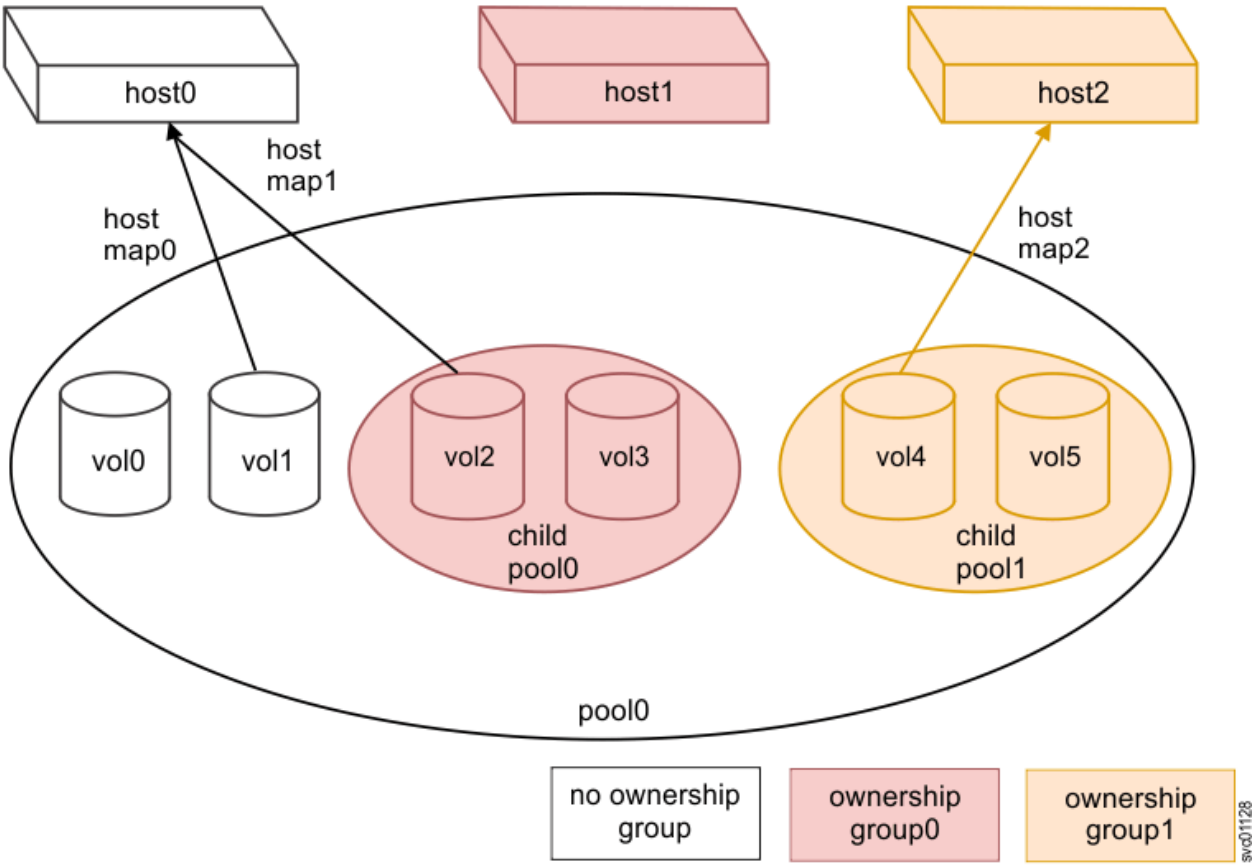


Figure 20. Ownership group inheritance

You can explicitly define ownership groups for the following objects:

Child pools
New or existing volumes that are defined in the child pool inherit the ownership group that is assigned for the child pool.

Host clusters
New or existing hosts that are defined in a host cluster inherit the ownership group that is assigned to the host cluster.

Hosts that are not part of a host cluster
If a host is not part of a host cluster, an ownership group can be associated with the host when the host object is created or changed.

Volume groups
Volume groups can be created to manage multiple volumes that are used with transparent cloud tiering support. The following rules apply to volume groups that are defined in ownership groups:

- Volumes can be added to a volume group if both the volume and the volume group are within the same ownership group or if both are not in an ownership group. When migrating existing resources to ownership groups, the volume group and its volumes can belong to different ownership groups temporarily until the migration completes.
- The ownership of a volume group does not affect the ownership of the volumes it contains. If a volume group and its volumes are owned by different ownership groups, then the owner of the child pool that contains the volumes can change the volume directly. For example, the owner of the child pool can change the name of a volume within it. The owner of the volume group can change the volume group itself and indirectly change the volume, such as deleting a volume from the volume group. Neither the ownership group of the child pools nor the owner of the volume group can directly manipulate the resources that are not defined in their ownership group.

Note: Volume groups can only be created by using the CLI.

FlashCopy consistency groups

FlashCopy consistency groups can be created to manage multiple FlashCopy mappings. The following rules apply to FlashCopy consistency groups that are defined in ownership groups:

- Users that are defined in ownership groups can add FlashCopy mappings to a consistency group if the volumes in the mapping and in the consistency group are within the same ownership group. For users that are not defined in an ownership group, no restriction for adding mappings to a FlashCopy consistency group. In this case, FlashCopy consistency groups are containers only and do not influence the ownership of the contents.
- When you are migrating existing resources to ownership groups, a FlashCopy consistency group and its resources can belong to different ownership groups temporarily until the migration completes.
- As with volume groups and volumes, the ownership of the consistency group has no impact in the ownership of the mappings it contains.

User groups

The following rules apply to user groups:

- Only users with Security Administrator role can create or manage ownership groups.
- Users with Security Administrator role cannot be assigned to an ownership group.

Portsets

The following rules apply to portsets:

- Users who are defined in an ownership group can view and create IP addresses or assign portset to host that are in their ownership group.
- Users who are defined in an ownership group can assign portset to host that are not assigned to an ownership group.
- When a user who are defined in an ownership group creates a portset, it is automatically assigned to the same ownership group as that restricted user.
- Users who are defined in an ownership group can view all IP addresses in portsets owned by them and portsets not assigned to an ownership group.
- Users who are defined in an ownership group cannot modify portsets associated with other restricted users.
- Users who are defined in an ownership group cannot view portsets that have been assigned to a different ownership group.
- Users who are defined in an ownership group cannot view IP addresses that are associated with a portset assigned to a different ownership group.
- Users who are defined in an ownership group cannot create or delete IP addresses from a portset that is part of a different ownership group.
- Users who are defined in an ownership group cannot own **storage** portsets.
- Users with the Security Administrator role can view, create, and manage all portsets on the system.
- Users with the Security Administrator role can view and create IP addresses to any portsets on the system.
- Users with the Security Administrator role can assign portsets that are not assigned to an ownership group to hosts in an ownership group.
- Users with the Security Administrator role can assign portset to hosts that are assigned the same ownership group.
- Users with the Security Administrator role cannot assign portset to hosts if the host and portsets are in different ownership groups.
- Users with the Security Administrator role cannot change the ownership of **storage** portsets.

Other objects inherit the ownership group from these explicitly assigned objects. Usually the ownership group is defined when an object is created. Ownership can also be changed when the explicitly defined

object from which the ownership is established is assigned to a different ownership group. The following objects inherit the ownership from the explicitly defined object:

Hosts that are a part of a host cluster

The following rules apply to hosts that are defined in ownership groups:

- The host inherits the ownership group of the host cluster that it belongs to.
- If a host is removed from a host cluster within an ownership group, the host inherits the ownership group of the host cluster it used to belong to.
- If a host is removed from a host cluster not within an ownership group, the host inherits no ownership group.
- Hosts can be added to a host cluster if the host and host cluster have the same ownership group.
- Changing the ownership group of a host cluster automatically changes the ownership group of all the hosts inside the host cluster.

Volumes

The following rules apply to volumes that are defined in ownership groups:

- 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.

Users

The following rules apply to users that are defined in ownership groups:

- A user inherits the ownership group of the user group it belongs to.
- Users that use LDAP for remote authentication can belong to multiple user groups, but belong to just one ownership group that is associated with one of the user groups.

Volume-to-host mappings

The following rule applies to volume-to-host mappings that are defined in ownership groups:

Volume-to-host mappings inherit the ownership group of the host or host cluster and volume in the mapping.

FlashCopy mappings

The following rules apply to FlashCopy mappings that are defined in ownership groups:

- FlashCopy mappings inherit the ownership group of both the volumes that are defined in the mapping.
- Like with FlashCopy consistency groups, it is possible for a consistency group and its mappings to belong to different ownership groups. However, the ownership of the consistency group has no impact in the ownership of the mappings that it contains.

Configuring ownership groups

You can configure ownership groups to manage access to resources on the system.

Before you create ownership groups and assign resources and users, it important to understand the following guidelines:

- Users within an ownership group can view or change resources within the ownership group in which they belong. Only users outside of an ownership group, usually users with Security Administrator roles, can view and manage all resources on the system, including users within an ownership group.
- Users within an ownership group cannot change any objects outside of their ownership group. This restriction includes global resources that are related to resources within the ownership group. Global resources cannot be assigned to ownership groups. For example, a user can change a volume in the ownership group, but not the MDisk or drive that provides the storage for that volume. The MDisk or drives are considered global resources and cannot be assigned to ownership groups.
- Users within an ownership group cannot view or change resources if those resources are assigned to another ownership group or are not assigned to any ownership group. However, users within ownership groups can view and display global resources. For example, if a volume is assigned to an ownership group, a user outside of that ownership group cannot display information for that volume. However, the same user can display information on drives on the system, since drives are a global resource that cannot be assigned to any ownership group.

Using the management GUI

These steps are for new systems where owned objects, such as child pools, hosts, and volumes are not configured yet.

To configure supported objects to use ownership groups on system, complete these steps:

1. In the management GUI, select **Access** > **Ownership Groups**.
2. On the **Ownership Groups** page, click **Create Ownership Group**. A new ownership group is created and displays on the **Ownership Groups** page.
3. After you have created an ownership group, create a user group to include users that are the owners of the ownership group. Select **Access** > **Users by Group** and click **Create User Group**.
4. On the **Create User Group** page, enter the following information:

   Name
   Enter the name of the user group.

   Role
   Select the role for all users within the user group. User groups that are assigned in ownership groups cannot use the Security Administrator role.

   Ownership Group
   Select the ownership group that you created earlier and assign it to this user group.
5. On the **Access** > **Users by Group** page, select existing users or create new users to assign to the user group. These users automatically inherit the ownership group that is assigned to the user group.
6. On the **Pools** > **MDisks by Pools** page, if a parent pool does not exist, click **Create Pool** and assign storage to that pool. Right-click the new parent pool and select **Create Child Pool**.
7. On the **Create Child Pool** page, enter the following information:

   Name
   Enter the name of the child pool.

   Capacity
   Select the amount of capacity that you are allocating from this child pool.

   Note: The **Capacity** field is not available when you are creating a data reduction child pool. A data reduction child pool has access to the complete free capacity of the parent pool without limit. A data reduction child pool cannot be created if the parent pool does not have free capacity.

   Ownership Group
   Select the ownership group that you created earlier and assign it to this child pool.
8. After child pools are configured and assigned to the ownership groups, users that you created in Step“5” on page 215 can create objects like hosts, volumes, and FlashCopy mappings.

Managing ownership groups

You can use the management GUI or command line interface to manage ownership groups.

Migrating to ownership groups

If you want to use ownership groups for objects that exist on your system, you need to reconfigure certain resources if you want to configure ownership groups.

An ownership group defines a subset of users and objects within the system. If you want to use ownership groups, you might need to reconfigure the following objects to implement ownership groups on the system:

- Child pools
- Volumes
- Volume groups
- Hosts
- Host clusters
- Host mappings
- FlashCopy mappings
- FlashCopy consistency groups
- User groups
- Portsets

If these objects are configured on your system and you want to use ownership groups to manage access, identify the objects that you want to assign to ownership groups and determine any dependencies these objects have to related objects outside of the ownership group. For example, if you have a volumes that are used in a FlashCopy mapping ensure that both volumes are assigned to the same ownership group, even if different child pools provide capacity. In this case, two separate child pools are configured with the same ownership group and their volumes inherit the ownership group from the pool. Objects like FlashCopy mappings and host mappings cannot be created with inconsistent ownership. Inconsistent ownership means that either an object is owned by different owners or dependent objects do not belong to the same ownership group. Some of these inconsistencies are can be temporary and resolve when all related objects are assigned to the correct ownership group. However some actions, such as adding a volume copy or migrating a volume can be created with inconsistent ownership intentionally. For example, when adding a volume copy or migrating a volume, you can have these objects be in different ownership groups. In the management GUI, this is done automatically when a volume is copied or migrated to a child pool in a different ownership group and warnings are generated. It is not recommended to leave objects, such as FlashCopy mappings or host mappings, in an inconsistent state. To ensure full information, only run system backup operations on the system when objects are in a consistent state.

Migration scenarios that require reconfiguration of these objects can be varied; however, two basic migration scenarios are available:

Migration Scenario 1: Child pools exist on the system

In this scenario, child pools exist on the system and a security administrator can create the ownership group and assign it to the existing child pool. All volumes that are associated with the child pool automatically inherit the ownership group of that child pool.

The following graphics show a configuration where child pools currently exist on the system and the process of assigning ownership groups to these child pools. The first graphic represents the current configuration before ownership groups are assigned to existing child pools. The second graphic shows the how some objects inherit the ownership group of the child pools and how others need to be added explicitly.

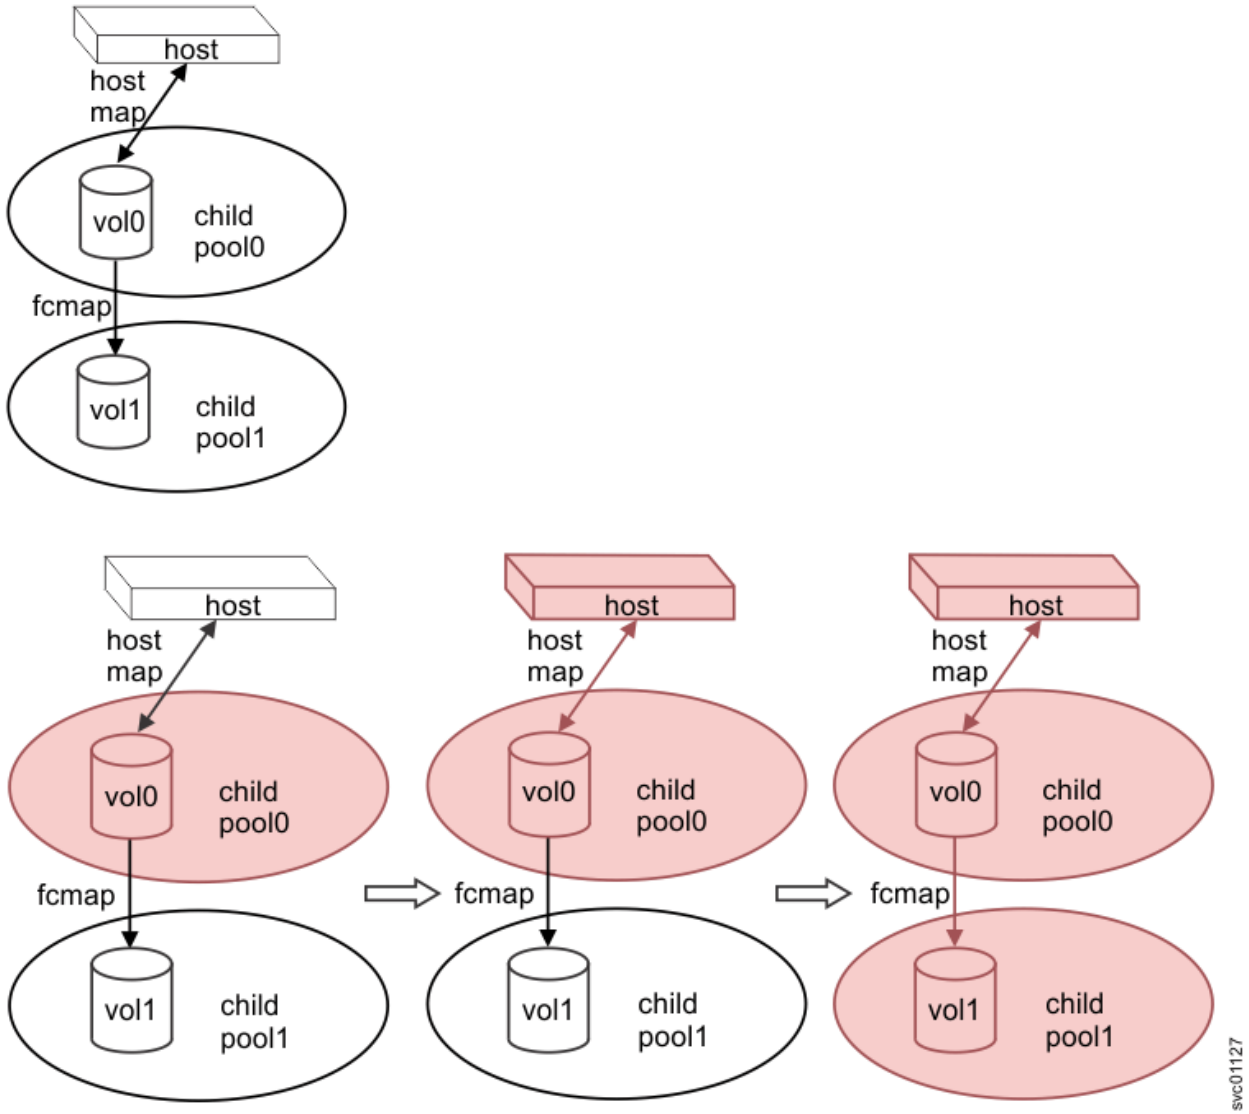


Figure 21. Example: Child pools are configured on the system

Using the management GUI

To use ownership group with existing child pools, complete the following steps:

1. In the management GUI, select **Access** > **Ownership Groups**.
2. On the **Ownership Group** page, select **Create Ownership Group**.
3. On the **Create Ownership Group** page, enter a name for the ownership group and select **Create**.
4. On the **Ownership Group** page, select **Assign Child Pools**.
5. On the **Assign Child Pools** page, select the child pool that you want to add to the ownership group and click **Assign**. In the example, *child pool0* is added to the ownership group. Because the volume inherits the ownership group of *child pool0* automatically.
6. If the system detects dependent resources for volumes that are within that child pool, the management GUI displays the **Additional Resources to Add** page where it lists hosts or host clusters that have mappings to volumes that inherit the ownership group. For example, in the preceding graphic, *child pool0* is added to the ownership group and *vol0* automatically inherits the ownership group of the child pool0. However, the host mapping between *vol1* and *host* does not inherit the ownership group from *vol0* because it also depends on the host. On the **Additional Resources to Add** page, select the host or host clusters to add to the ownership group. Click **Proceed**.

7. On the **Ownership Groups** page, verify that the volumes and hosts are assigned to the new ownership group.
8. In the example, *pool1* also needs to be assigned to the same ownership group as *pool0*. Repeat this process for all child pools that you want ownership groups to manage access.
9. After you have assigned ownership groups to all child pools, create a user group to include users that are the owners of the ownership group. Select **Access** > **Users** and click **Create User Group**.
10. On the **Create User Group** page, enter the following information:

    Name
    Enter the name of the user group.

    Role
    Select the role for all users within the user group. User groups that are assigned in ownership groups cannot use the Security Administrator role.

    Ownership Group
    Select the ownership group that you created earlier and assign it to this user group.
11. On the **Access** > **Users** page, select existing users or create new users to assign to the user group. These users automatically inherit the ownership group that is assigned to the user group. These users can now start creating objects like hosts and volumes within this ownership group, using capacity in the child pools that are assigned to the ownership group.

Using the command-line interface

Migration Scenario 2: Child pools do not exist on the system

In this scenario, child pools do not exist on the system, so they must be created from existing parent pools, and assigned to an ownership group. Volume mirroring can be used to create a copy of the volume in the new child pool. The volume copy inherits the ownership group that was assigned to the child pool.

The following graphics show a configuration where child pools do not exist on the system and the process of assigning ownership groups. The first graphic represents the current configuration before ownership groups are assigned.

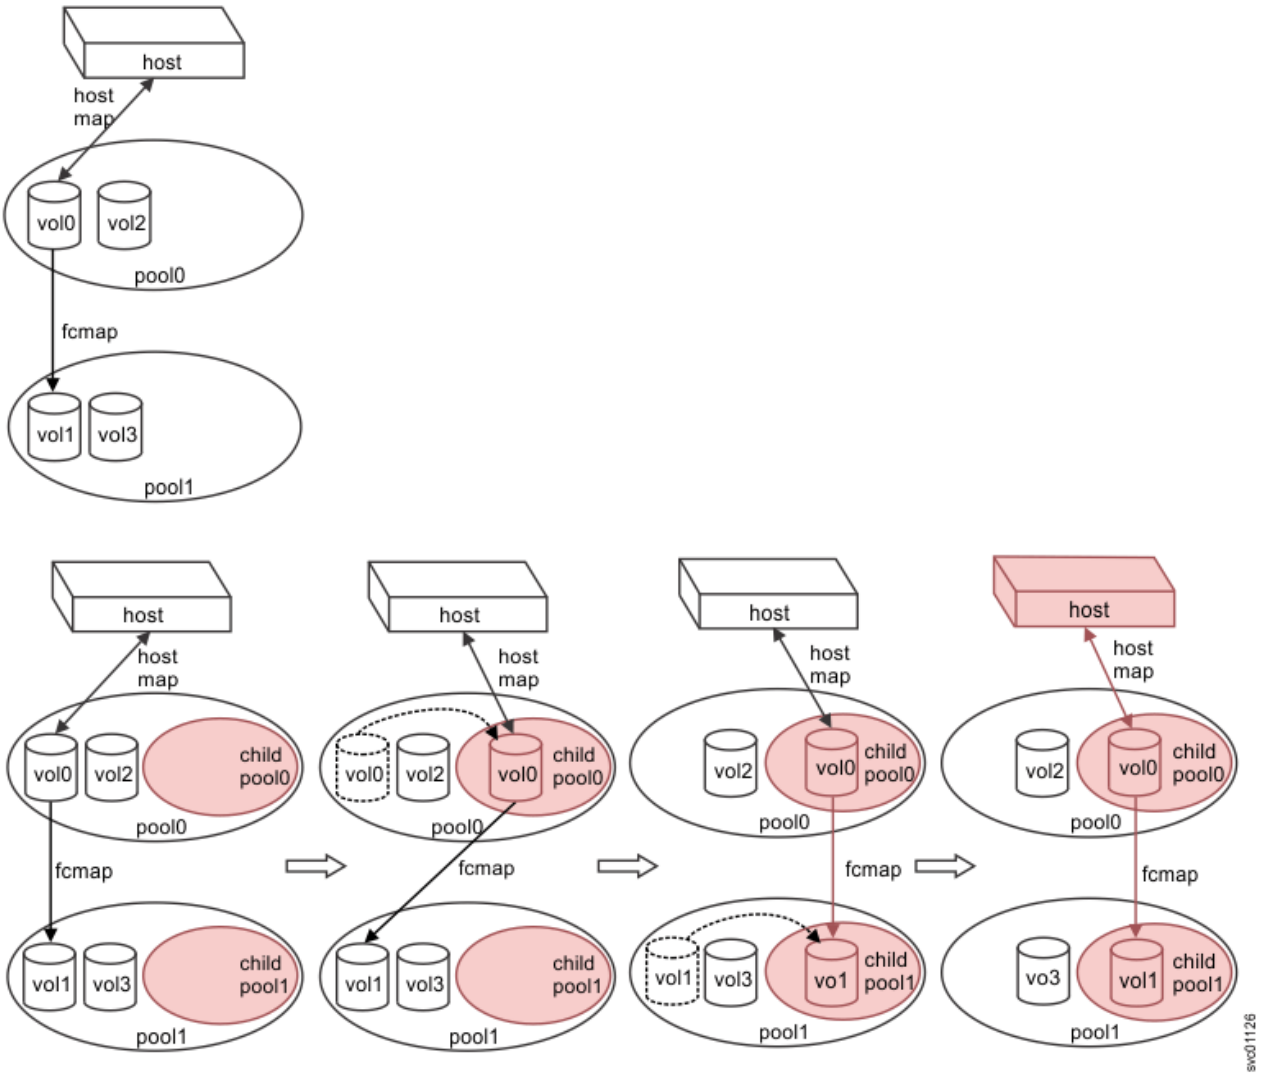


Figure 22. Example: Child pools are not configured on the system

In this example, the first graphic shows a system with two parent pools that contain two volumes each. Two volumes (*vol0* and *vol2*) use a parent pool (*pool0*) and two other volumes (*vol1* and *vol3*) use the other parent pool (*pool1*). The volume (*vol0*) is mapped to the host and is the source volume in a FlashCopy mapping with *vol0* in *pool1*. In this case, you want to limit access to *vol0* and *vol1* and its related resources.

The second graphic shows the required steps to use ownership groups in this environment. First, two child pools are created from the parent pools. After the child pools are created, a mirrored copy of each volume is created in the new child pools. After the volume copies are in the corresponding child pool, the child pools are assigned to an ownership group and the volumes with those child pools inherit that ownership group. The original copy of the pool can be deleted from the parent pool. More resources can be added to the ownership group, such as hosts and host clusters. The last graphic in the sequence shows the new configuration.

Using the management GUI

To migrate ownership groups on a system without child pools, complete the following tasks:

1. **Create child pools from existing parent pools:**
    a. In the management GUI, select **Pools** > **Pools**.

b. On the **Pools** page, right-click the pool that you want to create the child pool from and select **Create Child Pool**.

c. Repeat this step for all the pools that you want to use ownership groups. In the example, a separate child pool is created from both parent pools (*pool0* and *pool1*).

2. **Migrate volumes from parent pool to child pools:**

a. In the management GUI, select **Volumes** > **Volumes by Pools**.

b. Select the parent pool that you created the child pool from, and then right-click the volume and select **Add Volume Copy**.

c. On the **Add Volume Copy** page, select the child pool that you want to copy the volume to and click **Add**.

d. If necessary, delete the original copy from the parent pool.

e. Repeat steps "2.b" on page 220 to "2.c" on page 220 for each volume.

Note: This process uses volume mirroring to create a copy of an existing volume in the child pool. You can also use the **Migrate volume to another pool** action.

3. **Create ownership groups and assign child pools:**

a. In the management GUI, select **Access** > **Ownership Groups**.

b. On the **What is an Ownership Group?** page, enter a name of the ownership group and select **Create Ownership Group**.

c. On the **Create Ownership Group** page, expand **Child Pools** and select **Assign**.

d. On the **Assign Child Pools** page, select the child pools that you want to add to the ownership group and click **Assign**.

e. If the system detects dependent resources for volumes that are within that child pool, the management GUI displays the **Additional Resources to Add** page. The page lists hosts or host clusters that have mappings to volumes that inherit the ownership group. For example, in the preceding graphic, *child pool0* is added to the ownership group and *vol0* automatically inherits the ownership group of *child pool0*. However, the host mapping between *vol0* and *host* does not inherit the ownership group from *vol0* because it also depends on the host. On the **Additional Resources to Add**, select the host or host clusters to add to the ownership group. Click **Proceed**.

f. On the **Ownership Groups** page, verify that child pool is assigned to the new ownership group.

g. Repeat steps "3.c" on page 220 to "3.f" on page 220 for any other child pool.

h. After you assign ownership groups to all child pools, create a user group to include users that are the owners of the ownership group. Select **Access** > **Users** and click **Create User Group**

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

Name
Enter the name of the user group.

Role
Select the role for all users within the user group. User groups that are assigned in ownership groups cannot use the Security Administrator role.

Ownership Group
Select the ownership group that you created earlier and assign it to this user group.

j. On the **Access** > **Users** page, select existing users or create new users to assign to the user group. These users automatically inherit the ownership group that is assigned to the user group. These users can now start creating objects like hosts and volumes within this ownership group, using capacity in the child pools that are assigned to the ownership group.

Deleting an ownership group

You can use the management GUI or the command-line interface (CLI) to delete ownership groups.

An ownership group defines a subset of users and objects within the system.

Using the management GUI

1. Select **Access** > **Ownership Groups**. The **Ownership Group** page opens.
2. Click the menu button next to the ownership group that you want to delete, and then select **Remove Group**. The **Remove Ownership Group** page opens.
3. Select **Remove**.

Using the command-line interface

Use the **rmownershipgroup** command to delete an ownership group.

For example:

```
rmownershipgroup -keepobjects group_id_or_name
```

where *group_id_or_name* specifies ID or name of the ownership group to delete. The **keepobjects** parameter specifies to keep the objects that currently belong to the ownership group and to modify them to have no ownership group. You must specify the **keepobjects** parameter to delete an ownership group that still contains associated objects. You cannot delete an ownership group that still contains user groups, even if you specify the **keepobjects** parameter.

Users

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.

Security Administrators can create role-based user groups where any users that are added to the group adopt the role that is assigned to that group. Roles apply to both local and remote users on the system and are based on the user group to which the user belongs. A local user can belong only to a single group; therefore, the role of a local user is defined by the single group to which that user belongs. Roles are defined at the system level, which means that a user can be an administrator on one system, but a security administrator on another system. Remote users that use either LDAP or single sign-on require a user group on the system that matches a group defined on the remote authentication server. Security Administrators can also select to enable multifactor authentication per user group. With multifactor authentication, any users that are assigned to a user group must present a second factor to access the system. Multifactor authentication requires that an authentication server is configured to verify the identity of all users within the group. A user group is created automatically for each role. Create additional user groups only for remote authentication such as LDAP or to define access for tenant users.

You can assign the following roles to your user groups:

Monitor
: Users have access to all actions that are related to viewing objects and processes on the system. Monitor-role users cannot change the state of the system nor change the resources that the system manages. Monitor-role users can access all information-related GUI functions and commands, back up configuration data, and change their own passwords.

Service
: Users can set the time and date on the system, delete dump files, add and delete nodes, apply service, and shut down the system. Users can also complete the same tasks as users in the monitor role.

Administrator
: Users can manage all functions of the system except those functions that manage users, user groups, authentication, and encryption. Administrator-role users can run the system commands that the security-administrator-role users can run from the CLI, except for commands that deal with users, user groups, authentication, and encryption.

Security Administrator
Users can manage all functions of the system, including managing users, user groups, user authentication, and configuring encryption. Users with the Security Administrator role can run any system commands from the command-line interface (CLI). However, they cannot run the **satask** command from the CLI. Only the **superuser** ID can run **satask** command.

Restricted Security Administrator
Users with the security administrator role is changed to restricted security administrator when two person integrity (TPI) is enabled on the IBM Storage Virtualize. TPI requires two security administrators to work together to complete critical or risky tasks. For example, a restricted security administrator with an approved TPI request can remove Safeguarded snapshots.

User groups

Users with the Security Administrator role can organize users of the system by role through user groups.

The following user groups are configured by default:

SecurityAdmin
Users access all functions on the system, including managing users, user groups, and user authentication. Users can also configure encryption on the system.

Administrator
Users can complete most of the same tasks as users who are in the SecurityAdmin role. However, these users cannot access functions that deal with managing users, user groups, and authentication.

RestrictedAdmin
Users can complete the same tasks and run most of the same commands as administrator-role users. Support personnel can be assigned this role to help resolve errors and fix problems.

Service
Users can set the time and date on the system, delete dump files, add and delete nodes, apply service, and shut down the system. Users can also complete the same tasks as users in the monitor role.

Monitor
Users can view objects and the system configuration settings but cannot configure, modify, or manage the system or its resources.

Locking user accounts

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.

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. In addition, the following system-wide policies can be set:

- Automatic lock after (0-10) consecutive failed login attempts. A value of 0 means that the account does not lock because of failed login attempts.
- Set length of time auto-lock applies for (0-10080) minutes (7 days maximum). A value of 0 means that the account remains locked out indefinitely.
- Superuser account can be locked from general login.

User account can be manually locked using the **chuser** command, or automatically locked depending on the locking policy defined on the system. By default, user account locking is enabled for all users except the superuser. If a user fails to meet the password policy requirements, their account is automatically locked for 24 hours. For more information, see “Password policy” on page 487.

Locking the superuser account

For systems with a dedicated technician port, the system supports locking the superuser account for environments with strict security requirements.

The superuser account is the only account on the system that can complete these tasks:

- The superuser account is the only user that exists in service mode and can access the service assistant or run or view service-related command-line operations (or **satask sainfo** commands).
- The superuser account is the only user that can run the system recovery procedure and the backup and restore system configuration procedure.

Other Security Administrator accounts cannot complete these tasks if the superuser account becomes locked. If **Allow locking of the superuser account** is checked on the **Password Policy** page, these critical actions are disabled.

If you plan to allow the superuser account to be locked, ensure that you complete the following tasks on the system:

Enable Support Assistance
Support assistance enables support personnel to access the system either remotely from the support center, or locally from a console. If the superuser account is locked, a support person can log in to service the system when Support Assistance is enabled.

Ensure that password reset for the superuser account is enabled
You must ensure that the superuser password can be rest if you plan on allowing locks on the superuser account. This feature is enabled by default. If this feature is not enabled, unlocking the account by resetting the password on the service assistant GUI or service command-line interface is not possible. Use the command **svctask setpwdreset -enable** to enable this feature.

Ensure that an account with the role of SecurityAdmin exists
Ensure that a separate account with the user role of SecurityAdmin exists, either locally or on a remote LDAP server. This account can be used to unlock the superuser account.

The following table describes different use cases where the superuser account is locked and the possible recovery actions:

Locked superuser use case	Recovery action
Superuser account is locked after specified failed login attempts for a specified lockout time period.	Wait for the specified time for the lockout to expire and reattempt log in.
Superuser account is locked after the specified login attempts but an indefinite lockout was selected, or superuser account is manually disabled because it is not intended for use.	Complete one of these tasks to unlock the superuser account: • Manually unlock the superuser account that uses another account with the SecurityAdmin user role. • Create a remote support assistance request. • Use the technician port to access the service assistant GUI.
Superuser account is locked after two person integrity (TPI) is enabled.	Complete one of these tasks to unlock the superuser account: • Disable TPI. • Request remote support assistance. • Use the technician port to access the service assistant GUI.

Managing users and user groups

You can use the management GUI or the command-line interface to manage user and user groups on the system.

Changing user groups

You can use the command-line interface (CLI) to change user groups. User groups organize users of a system by role.

About this task

Roles apply to both local and remote users on the system and are based on the user group to which the user belongs. A local user can belong only to a single group; therefore, the role of a local user is defined by the single group that the user belongs to. Remote users can belong to one or more groups; therefore, the roles of remote users are assigned according to the groups that the remote user belongs to.

In addition to assigning roles to user groups, you can use user groups to control how user access different management interfaces and functions that are supported on the system. You can create a user group with specific roles and further control how users within that group interact with objects on the system by allowing or restricting the user group to functions or interfaces. With this capability, you can also control access from scripts or automated services, such as IBM Storage Insights that access the system automatically. The following interfaces and functions can be enabled or disabled with user groups:

Table 38. Supported access control for management interfaces and functions with user groups

Management interface or function	**Supported interface for configuration**
Management GUI access	Command-line interface
Command-line interface access	Command-line interface
REST API access	Command-line interface
Multifactor authentication	Management GUI and command-line interface
Ability to require a password and key for access	Command-line interface

Note: If any remote user groups on the system are configured with multifactor authentication, then it is not possible to use the BatchMode=yes setting when logging in using SSH for any remote users on the system.

You can use this feature to enable or disable access to interfaces from automated scripts or services, such as data collection for IBM Storage Insights or other automated processes. You can further secure the system by creating an ID for the automation service with only the role that they require to access the information necessary to the application or service. Generally, to configure a user group and access controls for an automation service, you need to complete the following steps:

1. Create a separate user group for automation IDs.
2. Create the automation users with appropriate role. For example, for data collection for IBM Storage Insights, the Monitor role allows the service to view objects and the system configuration settings. Other automated services or scripts can require different role-based access depending on the operations the service completes on the system
3. Disables interfaces for these automation IDs. Depending on the automated service, you can control which interfaces the automation service can use to complete tasks on the system. For example, with IBM Storage Insights data collection, you can limit these automation IDs to command-line access only, since other interfaces and functions are not necessary.
4. In cases where SSH keys are used to establish a connection between the system and the automated service, ensure that the key itself is secured and use these SSH keys with the automated service only.

To change a user group in the management GUI, select **Access** > **Users**. Select a user group and select **Properties** from the **Actions** menu.

To change a user group using the CLI, use the **chusergrp** CLI command.

Procedure

1. Use the **chusergrp** CLI command to change the attributes of an existing user group.
   For example, enter the following command:

   ```
   chusergrp -role role_name -remote yes | no group_id_or_name
   ```

   where *role_name* specifies the role that is associated with any users that belong to this group and *group_id_or_name* specifies the group to be changed. The **remote** parameter specifies whether the group is visible to the authentication server.
2. Issue the **lsusergrp** CLI command to display the user groups that were created on the system.
   For example, enter the following command:

   ```
   lsusergrp usergrp_id_or_name
   ```

   where *group_id_or_name* specifies the user group to view. If you do not specify a user group ID or name, all user groups on the system are displayed.

Example

If you plan to disable or enable interface access or functions, you can use the either the **mkusergrp** or **chusergrp** command. Both commands provide the following parameters to disable and enable access and functions:

-disablegui
Specifies if the user group has management GUI access disabled or not. Value can be yes or no.

-disablecli
Specifies if the user group has command-line interface access disabled or not. Value can be yes or no.

-disablerest
Specifies if the user group has REST-API access disabled or not. Value can be yes or no.

-passwordkeyrequired
Specifies if the user group requires a password and SSH key. Value can be yes or no.

-multifactor
Specifies if the user group has multifactor authentication enabled or not. Value can be yes or no.

The following examples show the command operations for some common uses cases for these parameters:

Limiting access to automated service
Create or change a user group with the appropriate role and restrictions:

```
mkusergrp -name automatedservice -role Monitor -passwordkeyrequired yes -disablecli no
-disablerest yes
Modifying the authentication setting for this user group will affect logins for all users
in the group.
Are you sure you want to continue? (y/yes to confirm)  yes
```

To change an existing user group, run the following command:

```
chusergrp -name automatedservice -role Monitor -disablegui yes -disablecli no -disablerest
yes
Modifying the authentication setting for this user group will affect logins for all users
in the group.
Are you sure you want to continue? (y/yes to confirm)  yes
```

In both of these examples, a user group was dedicated to an automated service, such as the data collector for IBM Storage Insights, that disabled access to the management GUI and REST APIs, but allows access through the command-line interface. Users in this group can only perform functions that are allowed by the Monitor role, such as viewing and collecting data that is related to operations on the system.

Enabling a password and SSH key for a user group

```
chusergrp -name serviceusergroup -role Service -passwordkeyrequired yes
Modifying the authentication setting for this user group will affect logins for all users
in the group.
Are you sure you want to continue? (y/yes to confirm)  yes
```

Enabling multifactor authentication for a user group

```
chusergrp -name mfagroup -role Administrator --multifactor yes
Modifying the authentication setting for this user group will affect logins for all users
in the group.
Are you sure you want to continue? (y/yes to confirm)  yes
```

Note:

- Before you can enable multifactor authentication on the user group, you must ensure all prerequisites are completed as well as configuring a multifactor authentication service on the system.
- Multifactor authentication is not supported on the Common Information Model (CIM). Therefore, user groups configured with multifactor authentication, or password and SSH key authentication are disabled for Common Information Model (CIM) interface.

Deleting user groups

You can use the command-line interface (CLI) to delete user groups. User groups organize users of a volume storage system by role.

About this task

Roles apply to both local and remote users on the system and are based on the user group to which the user belongs. A local user can belong only to a single group; therefore, the role of a local user is defined by the single group that the user belongs to. Remote users can belong to one or more groups; therefore, the roles of remote users are assigned according to the groups that the remote user belongs to.

To delete a user group in the management GUI, select **Access** > **Users**. Select a user group and select **Delete** from the **Actions** menu.

To delete a user group with the CLI, complete these steps:

Procedure

Issue the **rmusergrp** CLI command to delete a user group:
For example:

```
rmusergrp -force group_id_or_name
```

where *group_id_or_name* specifies the group to delete. The **force** parameter specifies to delete the group even if there are users in the user group.

Important: Using the **force** parameter might result in a loss of access. Use it only under the direction of the support center.

Changing users

You can use the command-line interface (CLI) or the management GUI to change users on the system.

Before you begin

System users must provide either a password, a Secure Shell (SSH) key, or both. Local users are authenticated through the authentication methods that are on the system.

You can create two categories of users that access the system. These user types are based on how they authenticate to the system:

- Some users must provide an SSH password (or if not possible an SSH key).
- If a 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), a valid password and SSH key can be used.
- Users must be in a user group that is defined on the system.

Remote users can also configure local credentials if they need to access the system when the remote service is down. Remote users have their groups that are defined by the remote authentication service.

To change a user in the management GUI, select **Access** > **Users**. Right-click the user and select **Modify** from the **Actions** menu. For local users, you can change the user name, password, and the SSH Public key file. For remote users, you can change the user name and the SSH Public key file. The password is maintained and authenticated through the remote authentication service.

About this task

To change a user in the CLI, follow these steps:

Procedure

1. Use the **chuser** CLI command to change the attributes of an existing user.
   For example, enter the following command:

   ```
   chuser -usergrp group_id_or_name user_id_or_name
   ```

   where the *group_id_or_name* specifies the new group for the user and *user_id_or_name* specifies the user to be changed.
2. Use the **chcurrentuser** CLI command to change the attributes of the current user.
   For example, enter the following command:

   ```
   chcurrentuser -nokey
   ```

   where the **nokey** parameter specifies that the SSH key of the user is to be deleted.
3. Use the **lscurrentuser** CLI command to display the name and role of the logged-in user.
   For example, enter the following command:

   ```
   lscurrentuser
   ```

   The name and the role of the user are displayed.

Deleting users

You can use the command-line interface (CLI) or the management GUI to delete users on the system.

About this task

To delete a user in the management GUI, select **Access** > **Users**. Select a user and select **Delete** from the **Actions** menu.

To delete a user from the CLI, complete these steps:

Procedure

1. Issue the **rmuser** CLI command to delete a user:
   For example:

   ```
   rmuser user_id_or_name
   ```

   where *user_id_or_name* specifies the user to be removed.
2. Issue the **lsuser** CLI command to display a list of users that were created on the system.

For example:

```
lsuser
```

If the user ID or name is not specified, all users on the system are displayed. Verify that the deleted user was removed from the system.

Locking and unlocking user account automatically

Accounts can be locked automatically by the system after a certain number unsuccessful login attempts. A user with Security Administrator access can unlock the account to provide access again.

The account can be locked for a specific amount of time, or the account can be locked indefinitely. A user with the Security Administrator access can unlock the account to provide access again

Enable automatic account locking

Enable automatic account locking by using management GUI or command-line interface.

Using the management GUI

To enable automatic account locking using the management GUI, complete these steps:

1. In the management GUI, select **Settings** > **Security** > **Password Policies**.
2. On **Password Policies** page, select **Password expiration and account lockout**.
3. Select the checkbox **After a number of failed login attempts, automatically lock accounts for an amount of time**.
4. Enter the number of consecutive failed login attempts before the account is locked (You can select 1-10 attempts).
5. Select either Define time period (minutes) or Indefinite lockout. If you choose to define the time period, enter the number of minutes that an account should remain locked for before automatically unlocking (You can select 1-10080 minutes). If you select indefinite lockout, accounts remain locked until a user with the Security Administrator access unlocks the account manually. See “Unlocking user accounts manually” on page 230.
6. Select **Allow locking of the superuser account** to enable automatic account locking for the superuser account, for more information see “Locking and unlocking superuser accounts” on page 230.

   Note: If **Allow locking of the superuser account** is not enabled, then the superuser account will not be automatically locked when the specified number of failed attempts is reached.
7. Click **Save**.

Using the command-line interface

To enable automatic account locking using the command-line interface, enter the following command:

1. Decide how many consecutive failed login attempts (1-10) should lock the account, and how long the account should remain locked for when it becomes locked (0-10080 minutes).

   Note: A time of 0 minutes means that the account will be locked indefinitely until manually unlocked by a user with the Security Administrator access.
2. Specify required values, enter the following command:

   ```
   svctask chsecurity -maxfailedlogins 5 -lockoutperiod 30
   ```
3. To enable the superuser locking, enter the following command:

   ```
   svctask chseurity -superuserlocking enable
   ```

Disable automatic account locking

Disable automatic account locking by using management GUI or command-line interface.

Using the management GUI

To disable automatic account locking using the management GUI, complete these steps:

1. In the management GUI, select **Settings** > **Security** > **Password Policies**.
2. On **Password Policies** page, select **Password expiration and account lockout**.
3. Deselect the checkbox **After a number of failed login attempts, automatically lock accounts for an amount of time**.
4. Choose whether to deselect the checkbox **Allow locking of the superuser account**. This setting should not be deselected if there is a requirement to manually lock (or disable) the superuser account.
5. Click **Save**.

Using the command-line interface

To disable automatic account locking using the command-line interface, enter the following command:

1. To disable automatic account locking for all users, enter the following command:

```
svctask chsecurity -maxfailedlogins 0
```

2. To disable the superuser locking, enter the following command:

```
svctask chseurity -superuserlocking disable
```

Locking and unlocking user accounts manually

A user with the Security Administrator role can lock or unlock user accounts. This allows you to prevent the log in of the specific user if there is suspicious activity that is associated with this account.

Note: You can manually lock the user account even if automatic locking is not enabled. You cannot manually lock your own user account.

Locking user accounts manually

You can use the management GUI or the command-line interface to lock user accounts manually.

Before you begin

You must be assigned with the Security Administrator role to lock user accounts manually.

Using the management GUI

To lock the user account, complete the following steps:

1. In the management GUI, click **Access** > **Users by Group**. The **Users by Group** window is displayed, which shows the **All Users** section.
2. From the **All users** section, right-click on the user and click **Lock** to lock the selected user account.

Using the command-line interface

To lock the user account manually, complete these following steps:

1. To view the ID and name list of the user accounts, enter the following command:

```
lsuser
```

2. To lock the user account, enter the following command:

```
svctask chuser -lock <name_or_id>
```

Where <name_or_id> specifies the name or ID of the user account. The user account remains locked until you manually unlock the user account.

Unlocking user accounts manually

You can use the management GUI or the command-line interface to unlock user accounts manually.

Before you begin

You must be assigned with the Security Administrator role to unlock user accounts manually.

Using the management GUI

To unlock the user account, complete the following steps:

1. In the management GUI, click **Access** > **Users by Group**. The **Users by Group** window is displayed, which shows the **All Users**section.
2. From the **All users** section, right-click on the locked user and click **Unlock** to unlock the selected user account.

Using the command-line interface

To unlock the user account manually, complete these following steps:

1. To view the ID and name list of the user accounts, run **lsuser** command.
2. To unlock the user account, enter the following command:

```
svctask chuser -unlock <name_or_id>
```

Where <name_or_id> specifies the name or ID of the locked user account.

Locking and unlocking superuser accounts

A 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 systems with a dedicated technician port, the system supports locking the superuser account for environments with strict security requirements.

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.

The superuser account is the only account on the system that can complete these tasks:

- The superuser account is the only user that exists in service mode. This user can access the service assistant to run service-related command-line operations, such as (**satask** and **sainfo** commands).
- The superuser account is the only user that can run the system recovery procedure and the backup and restore system configuration procedure.

Other Security Administrator accounts cannot complete these service tasks. If the superuser account is locked, then these service actions are unavailable until the superuser account is unlocked.

If you plan to allow the superuser account to be locked, consider completing the following tasks on the system to ensure that superuser access can be restored when needed:

Enable Support Assistance
Support assistance enables support personnel to access the system either remotely from the support center, or locally from a console. If the superuser account is locked, a support person can log in to service the system when Support Assistance is enabled.

Ensure that password reset for the superuser account is enabled
You must ensure that the superuser password can be reset if you plan on allowing locks on the superuser account. This feature is enabled by default. If this feature is not enabled, unlocking the account by resetting the password on the service assistant GUI or service command-line interface is not possible. Use the command **`svctask setpwdreset -enable`** to enable this feature.

Ensure that an account with the role of SecurityAdmin exists
Consider creating a separate account on the system with the user role of SecurityAdmin, either locally or on a remote LDAP server. This account can be used to unlock the superuser account.

Manually locking a superuser account

The system allows to lock or disable the superuser account by using both the management GUI and the command-line interface.

Note: Superuser locking must be enabled before the system-wide automatic locking settings are applied to the superuser account.

Using the management GUI

To lock or disable the superuser account using the management GUI, complete the following steps:

1. Log into the management GUI as a user with the Security Administrator role.
2. Select **Settings** > **Security** > **Password Policies**.
3. Select **Password expiration and account lockout** and then select **Allow locking of the superuser account**.
4. Click **Save** to enable superuser locking.
5. Select **Access** > **Users by Group**.
6. Right-click the superuser account under **All Users**.
7. Click **Lock** to manually lock the superuser account.

Using the command-line interface

To lock or disable the superuser account by using the command-line interface (CLI), complete the following steps:

1. Log in to the CLI as a user with the Security Administrator role.
2. To enable superuser locking, run the following command:

```
svctask chsecurity -superuserlocking enable
```

3. To manually lock the superuser account, run the following command:

```
svctask chuser -lock superuser
```

Unlocking a superuser account

The system allows to unlock or re-enable the superuser account by using both the management GUI and the command-line interface.

If the locked superuser account is required to complete a service or recovery action, the account must be re-enabled. Re-enabling the superuser account is possible under following circumstances:

1. If there is another user with the Security Administrator role, then that user can unlock the superuser account using the management GUI or the CLI.

2. If Remote Support Assistance is enabled, then remote support personnel can connect to the system to perform service actions.
3. If dedicated technician port is available, the system allows to unlock the superuser account using that port. This requires physical access to the system. Connect to the system using the superuser account. The service GUI provides an option to unlock the superuser account, or use the service CLI to run the *satask unlocksuperuser* command.

Using the management GUI

To unlock or re-enable the superuser account using the management GUI, complete the following steps:

1. Log in to the management GUI as a user with the Security Administrator role.
2. Select **Access** > **Users by Group**.
3. Right-click the superuser account under **All Users** and click **Unlock**.
4. Select **Settings** > **Security** > **Password Policies**.
5. Select **Password expiration and account lockout** and then deselect **Allow locking of the superuser account**.

   Note: Deselecting the Allow locking of the superuser account check box prevents both manual and automatic locking of the superuser account. Skip the above step to keep automatic locking of the superuser account enabled if automatic locking is also enabled system-wide.
6. Click **Save**.

Using the command-line interface

To unlock or re-enable the superuser account using the command-line interface (CLI), complete the following steps:

1. Log in to the CLI as a user with the Security Administrator role.
2. To unlock the superuser account, run the following command:

   ```
   svctask chuser -unlock superuser
   ```
3. To disable superuser locking, run the following command

   ```
   svctask chsecurity -superuserlocking disable
   ```

   Note: Disabling superuser locking prevents both manual and automatic locking of the superuser account. Skip the above step to keep automatic locking of the superuser account enabled if automatic locking is also enabled system-wide.

The following table describes different use cases where the superuser account is locked and the possible recovery actions:

Locked superuser use case	Recovery action
Superuser account is locked after specified failed login attempts for a specified lockout time period.	Wait for the specified time for the lockout to expire and reattempt log in.
Superuser account is locked after the specified login attempts but an indefinite lockout was selected, or superuser account is manually disabled because it is not intended for use.	Complete one of these tasks to unlock the superuser account: • Manually unlock the superuser account that uses another account with the SecurityAdmin user role. • Create a remote support assistance request. • Use the technician port to access the service assistant GUI.

Locked superuser use case	Recovery action
Superuser account is locked after two person integrity (TPI) is enabled.	Complete one of these tasks to unlock the superuser account: • Disable TPI. • Request remote support assistance. • Use the technician port to access the service assistant GUI.

Working with local and remote users

You can create either a local or a remote user to access a system.

Before you begin

You can create two categories of users that access the system. These types are based on how the users are authenticated to the system. Local users must provide the *username* and password, and if you require command line access without entering a password, a Secure Shell (SSH) key - or both. Local users are authenticated through the authentication methods that are located on the system.

If the local user needs access to management GUI, a password is needed for the user. Access to the command-line interface (CLI) is also possible with the same password or (alternatively) a valid SSH key can be used. An SSH password is required if a user is working with both interfaces. User groups define roles that authorize the users within that group to a specific set of operations on the system.

Local users must be part of a user group that is defined on the system.

A remote user is authenticated on a remote service that is provided by an LDAP server and does not need local authentication methods. The remote user does not need to be added to the list of users on the system. If the user requires command-line access without entering a password, the remote user can be added to the system and an SSH key can be specified for the user.

Remote users cannot access the system when the remote service is down. A local user account must be used until the remote service is restored. The user groups a remote user is a member of are defined by the remote authentication service. To define a remote user, create a user group on the local computer that is also defined on the remote authentication service.

You can connect to the system by using the same username with which you log in to the system.

About this task

Complete these steps to create either a local or remote user:

Procedure

1. Select **Access** > **Users** .
2. Select the appropriate user group.
3. Click **Create User**.
4. Enter the information on the new user and click **Create**.

Audit log

An audit log documents user actions that are submitted through the different management interfaces that run commands on the system, such as the management GUI.

You can use the audit log to monitor user activity on all supported interfaces.

The audit log entries provide the following information:

- The time and date when the action or command was submitted on the system.

Note: When logs are displayed in the command-line interface, the time stamps for the logs in CLI are the system time. However, when logs are displayed in the management GUI, the time stamps are converted to the local time where the Web browser is running.

- The name of the user who completed the action or command.
- The IP address of the system where the action or command was submitted.
- The name of source and target node on which the command was submitted.
- The parameters that were submitted with the command, except confidential information like clear-text passwords.
- The results of the command or action that completed successfully. Failed commands are not added to the audit log.
- The sequence number and the object identifier that is associated with the command or action.
- The origin of the command or action. Possible values are: GUI, CLI, or REST.
- For service-related actions that are completed by support personnel, the audit log displays the following information:

  Challenge
  Indicates the challenge with which the support user authenticated to the system. Support users connect to the system through a challenge-response authentication. If authentication succeeds, a session is opened and support users can conduct service-related actions on the system. This field identifies the support user at the support center. This field is blank for other users.

Login message

You can create or change a message that displays when users log on to the system.

About this task

When you log in to the system with the management GUI, command-line interface, or service assistant, the message displays before you log in. To create or change the login message, you can use the **chbanner** command or the management GUI. If you are using the command, you must create the message in a supported text editor and use the Secure Copy (SCP) or SSH File Transfer Protocol (SFTP) to copy the file to the configuration node on the system. In the management GUI, go to **Settings** > **GUI Preferences** > **Login** to enable and create a login message.

Procedure

To change the login message from a SAN administrator workstation, complete the following steps.

1. Use a suitable text editor to create a file that contains the text of the message.

   Note: The message cannot exceed 4 Kbytes.
2. Use a Secure Copy client to copy the file to the configuration node of the system to be configured. Specify the management IP address of the system to be configured.
3. Log on to the system to be configured.
4. In the command-line interface, type the following command to set the login message:

   ```
   chbanner -file filepath
   ```

   where *filepath* specifies the file that contains the text of the new message.

Example

```
echo "Access to this system is permitted only for authorized personnel." > loginmessage
 scp loginmessage admin@12.23.34.45:/tmp/loginmessage
   (enter password)
 ssh admin@12.23.34.45
   (enter password)
 chbanner -file /tmp/loginmessage
```

Chapter 10. FlashSystem grid

A FlashSystem grid can be used to scale-out a storage system by adding more independent systems to create a federated cluster.

A FlashSystem grid allows for greater scalability than a traditional scale-out system (which uses I/O groups) as it supports up to 32 systems per FlashSystem grid with linear increases in performance, capacity, and number of volumes, hosts, and snapshots. Systems in a FlashSystem grid can be updated independently and systems with different hardware models and generations can be included within a single FlashSystem grid. This capability allows for hardware replacement or tiering within the same FlashSystem grid. The maximum number of systems in a FlashSystem grid depends on the software version running on the coordinator system. If member systems are running an older software version, review the relevant limits page to understand how many systems will be visible from the member system. If the limits are different, some systems may display authentication errors. To resolve authentication errors, upgrade all systems to a code level that supports the configured number of systems and follow the fix procedure for the error.

You can manage and monitor all connected systems in a FlashSystem grid as a single system from a single interface. Storage partitions can be non-disruptively migrated between systems in a FlashSystem grid. For more information, see Migrating storage partitions between systems.

Systems in a FlashSystem grid have the following requirements:

- The system supports a maximum of a single I/O group and the compatibility mode for multiple I/O groups must be disabled.
- The system topology must be standard.
- The system layer must be replication.

All systems in a FlashSystem grid must use one of the following options for the system management IP addresses:

- IPv4
- IPv6
- IPv4 using fully qualified domain names (FQDN) with DNS server configured
- IPv6 using fully qualified domain names (FQDN) with DNS server configured
- Combination of IPv4 and FQDN with DNS server configured

If your storage system is running 8.7.2 or an older software version, and is part of a FlashSystem grid, you must update the existing FlashSystem grid configuration before upgrading the system to the current software version. For more information, see https://www.ibm.com/support/pages/node/7181817.

Systems in a FlashSystem grid exchange certificates to enable secure communication between members. System-signed certificates are the default.

A **coordinator system** is the system on which a user creates and manages the FlashSystem grid. A coordinator system also becomes the first member of the FlashSystem grid.

A **member system** can be managed by the FlashSystem grid to which it is added as a member.

The FlashSystem grid dashboard shows details about storage systems and storage partitions in a FlashSystem grid that include aggregated physical and logical capacity, health, and reachability.

The **Storage systems** page of the FlashSystem grid shows a list of storage systems that are part of the FlashSystem grid along with details.

The **Storage partitions** page of the FlashSystem grid shows a list of storage partitions in the FlashSystem grid along with information about partition location, replication topology, replication status, volume groups, migration status, and placement evaluation status.

© Copyright IBM Corp. 2025

The **AI suggestions** include a list of advisories generated by IBM Storage Insights for storage systems that are part of the FlashSystem grid. The AI-powered recommendations are generated by analyzing historical data across all storage systems within the FlashSystem grid.

Video

The following FlashSystem grid video demonstrates the use case for creating a new FlashSystem grid and adding a member system from the coordinator system by using the management GUI.

Related information

"Storage partitions" on page 243
A storage partition groups the configuration for one or more applications that need to be managed together.

Implementing a FlashSystem grid

You can create a FlashSystem grid and manage systems within a FlashSystem grid by using the management GUI or CLI.

Creating a FlashSystem grid

The system on which you create a FlashSystem grid becomes the coordinator system for that FlashSystem grid.

To create a new FlashSystem grid in the management GUI, go to **Settings** > **System** > **FlashSystem grid** and select **Create a new FlashSystem grid** > **Create**.

You can also create a FlashSystem grid by using the **managegrid**CLI.

Adding a system to a FlashSystem grid

You can add a system to a FlashSystem grid from the FlashSystem grid **Storage systems** page. On the FlashSystem grid dashboard, go to **FlashSystem grid** > **Storage Systems** > **Add a system**. Then, enter the IP address or the FQDN of the system that you want to add to the FlashSystem grid. The management GUI guides you to verify the certificate to allow the certificate exchange between the coordinator system and the system that is being added to the FlashSystem grid.

Alternatively, a system can join an existing FlashSystem grid from the management GUI.

To join an existing FlashSystem grid, go to **Settings** > **System** > **FlashSystem grid** > **Join an existing FlashSystem grid**. Then, enter the IP address or the FQDN of the coordinator system of the FlashSystem grid that you want to join. The management GUI guides you to verify the certificate to allow the certificate exchange between the coordinator system and the system that is being added to the FlashSystem grid.

You can also use the **managegrid** CLI to add a system to the FlashSystem grid. After you create a FlashSystem grid, you can add a system to the FlashSystem grid by performing the following steps in order:

1. Check the list of scope certificates by using **lssystemcertstore** CLI. In case the (scope_id = 3) internal_communication certificate is present then download the certificate by using the **chsystemcertstore** CLI. Otherwise, download certificate by using the **chsystemcertstore** CLI for scope default or use **chsystemcert** CLI. Copy the certificate of the FlashSystem grid coordinator-system to the new system, and add it to the truststore by using the **mktruststore** CLI.
2. Request to join the FlashSystem grid by using the **join** option. Provide the IP address or the FQDN of the coordinator system and the entry, which the coordinator system has in the truststore.
3. Copy the certificate from the new system to the FlashSystem grid coordinator-system, and add it to the truststore.

4. On the coordinator system, accept the request to join the FlashSystem grid by using the **accept** option. Provide the IP address or the FQDN of the new system and the entry, which the new system has in the truststore.

Viewing a FlashSystem grid

You can use the FlashSystem grid dashboard to view the total number of systems and partitions that are part of the FlashSystem grid. The FlashSystem grid dashboard also shows whether storage systems in the FlashSystem grid are Healthy or Unhealthy, and Reachable or Unreachable.

The FlashSystem grid dashboard offers several key functions:

- **Hardware health monitoring**: indicates the health status of the hardware and highlights any errors or alerts within the FlashSystem grid systems.
- **Storage-capacity usage overview**: provides a consolidated view of the storage capacity usage across the FlashSystem grid.
- **System reachability**: indicates whether the coordinator system in the FlashSystem grid can connect with each member system in the FlashSystem grid.
- **Migrations**: indicates the count and status of partition migrations only if any active partition migrations are happening in the FlashSystem grid.
- **Physical capacity**: indicates the aggregated physical capacity metrics across systems in a FlashSystem grid.
- **Logical capacity**: indicates the aggregated logical capacity across systems in a FlashSystem grid. **Logical capacity** is visible only if Flash Core Module (FCM) drives are used in all systems of the FlashSystem grid.
- **AI suggestions**: indicates a list of AI-powered suggestions from IBM Storage Insights for storage systems in a FlashSystem grid.

  The **View FlashSystem grid operations** option navigates you to the **Storage systems** page of the corresponding FlashSystem grid in IBM Storage Insights.

The FlashSystem grid dashboard also shows if any active migrations exist along with the migration status. For more information, see .

The following information about each system in the FlashSystem grid is available on the **Storage systems** page:

- Product name
- System code level
- Last contacted time
- Connectivity health: indicates whether connectivity errors or alerts exist on the member system.
- Hardware health: indicates whether hardware errors or alerts exist on the member system.
- Physical capacity

The **Storage systems** page of a FlashSystem grid also shows error data that indicates the type of error faced during visualization of other systems in the FlashSystem grid. The error data is indicated in the form of icons at applicable fields and it also suggests actions to resolve errors.

You can also use the **lsgrid**, **lsgridsystem**, **lsgridmembers**, and **lsgridpartition** CLIs to visualize the FlashSystem grid.

The **Storage partitions** page of a FlashSystem grid shows details about storage partitions in the FlashSystem grid.

Removing a system from a FlashSystem grid

You can remove a system from a FlashSystem grid from the FlashSystem grid **Storage systems** page. On the FlashSystem grid dashboard, go to **FlashSystem grid** > **Storage Systems**, click the overflow menu,

and select **Remove system**. If a partition migration is in progress to or from a system in a FlashSystem grid, that particular system cannot be removed from the FlashSystem grid until the migration is complete.

A member system can also be removed from a FlashSystem grid by using the **managegrid** CLI.

To remove a system from a FlashSystem grid by using the CLI, use the **remove** option on the coordinator system.

When you run the **leave** option from a coordinator system, the system leaves the FlashSystem grid and the FlashSystem grid gets deleted. When you run the **leave** option from a member system, it partially removes that particular member system from the FlashSystem grid. However, to completely remove the member system from the FlashSystem grid, use the **-remove** option for the member on the coordinator system.

The truststore entries that were created during adding a system to the FlashSystem grid are removed when a member system is removed from the FlashSystem grid.

Removing the last system from a FlashSystem grid results in deleting the FlashSystem grid. When a FlashSystem grid is deleted, the coordinator system cannot see the details, such as system health, capacity, and partitions of other systems that are removed from the FlashSystem grid. A member can remove only itself from the FlashSystem grid, whereas a coordinator can remove any system that is part of the FlashSystem grid.

If you run the **managegrid -leave** command from a member system, remove the member system from the **Storage systems** page of the FlashSystem grid to avoid an inconsistent FlashSystem grid experience.

If a FlashSystem grid member system is running 8.7.2 or an older software version, then use the **-remove** option on the coordinator system to remove the member system from the FlashSystem grid. Also, the member system that is removed from the FlashSystem grid needs to run the **-leave** option.

Set up a system through FlashSystem grid onboarding

This feature enables you to configure a new system from the FlashSystem grid. To begin, initialize the system from the , enter the **cluster management IP address** and optional DNS, save the settings, and **restart** the system.

After the system restarts, sign in to the **FlashSystem grid coordinator** and select **Add a system** to add a new unconfigured system.

The FlashSystem grid detects that the system is unconfigured and prompts you to reset the password. This step ensures that the system is not added with default credentials.

Next, the unconfigured system imports the configuration from the FlashSystem grid coordinator. You can review the imported settings and **deselect** configuration groups before applying them. The system applies the selected settings, records progress for audit and resume operations, and then prompts you for **system-specific information** (such as, Call Home, Storage Insights, or encryption settings) to complete the setup.

Note: Do not perform any initial configuration if the setup is being completed through the FlashSystem grid.

Configuring a new system setup by using the management GUI

To add and configure a new system in the FlashSystem grid, complete the following steps:

1. **Add the new system from the FlashSystem grid user interface**
   From the coordinator system:
   - Specify the target system by using an **IP address or a fully qualified domain name (FQDN)**.
   - Reset the password when prompted to remove default credentials.
2. **Get configuration from the FlashSystem grid**
   Import the configuration from the FlashSystem grid coordinator.

3. **Review and apply the configuration on the new system**
   In the new system interface, review the prefilled settings. If needed, uncheck any **Imported settings** that are not required before applying the configuration.

   Note:

   - **Required settings** are mandatory and cannot be deselected.
   - If the imported **Required settings** do not include Call Home, Storage Insights (SI) configuration, the setup wizard displays Call Home/Storage Insights in **System-specific settings** so you can provide the missing information.

   Required settings
   Configure required system parameters needed to complete the system setup. You can validate the imported system settings such as DNS details, date and time, system location, user contact information, Call Home, and IBM Storage Insights configuration.

   Notifications
   Notification settings help the new system communicate events information, such as SNMP server configuration and syslog server details.

   Password policies
   Password policies define rules such as password complexity, reuse limits, and expiration behavior to enforce strong security and consistent password management across the system.

   Other security settings
   Additional controls ensure consistent security level across the grid. These controls may include inactivity logout, SSH rules, security protocol level, and volume protection settings.
4. **Complete system-specific settings**
   Complete any system-specific settings, such as encryption. If a setting fails, the interface displays an error message next to the setting so you can manually correct it.

Changing coordinator of a FlashSystem grid

FlashSystem grid enables you to change the existing coordinator system and make another member system from the grid as the new coordinator.

Changing FlashSystem grid coordinator

To change the coordinator of a FlashSystem grid by using the management GUI, go to **FlashSystem grid** > **Storage systems** page. From the overflow menu of the member system that you want to set as the new coordinator, select **Assign as coordinator**.

Note: Changing a member system to coordinator in a FlashSystem grid by using the management GUI is possible only from the current coordinator system of the grid.

You can also use the **`managegrid`** command to change the coordinator of a FlashSystem grid and set another member system from the grid as coordinator. To change the coordinator of a FlashSystem grid, run the **`managegrid -makecoordinator`** command on the member system that you want to assign as new coordinator of the grid.

Note: If the current coordinator system in a FlashSystem grid is offline or unreachable, you can change the coordinator by using the **`managegrid`** command only.

Guidelines for FlashSystem grid coordinator change

Changing a FlashSystem grid coordinator when the existing coordinator system is online

- Ensure that all the member systems in the FlashSystem grid are online and reachable.
- Verify the code or firmware level of each member system in the FlashSystem grid to ensure that all member systems support the coordinator change operation.

- If any member systems in the FlashSystem grid do not support the coordinator change operation or are unreachable, remove such member systems from the grid.

Changing a FlashSystem grid coordinator when the existing coordinator system is offline

To change the coordinator system in a FlashSystem grid when the current coordinator is offline, ensure that:

- All online member systems in the grid support the coordinator change operation, and
- The remaining member systems in the grid are either in *invalid_grid_membership* state or unreachable.
- Select a member system with the latest FlashSystem grid information as the new coordinator. Use to the **lsgridsystem** command to check the values of the **grid_gen_num** parameter to verify the system with the latest grid information.
- If any online member systems in the FlashSystem grid do not support the coordinator change operation, then run the **managergrid -leave** command on such member systems.
- On the member system that is selected to be the new coordinator, run the **managegrid -makemecoordinator -removeinvalidmembers** command:
  - To remove other unreachable or invalid member systems from the FlashSystem grid, and
  - To assign the selected member system as new the new coordinator.

  This command also removes the old coordinator system that is unreachable.

To join a FlashSystem grid or to create a new FlashSystem grid from unreachable systems that were removed in the coordinator change operation, run the **managergrid -leave** command after the systems are back online.

Note: Do not initiate multiple coordinator change requests simultaneously as it can result in a forced new FlashSystem grid creation.

Multi-tenancy

Multi-tenancy in a FlashSystem grid enables service providers and teams with delegated administration to manage defined sets of resources independently. Multiple tenants, such as teams, customers, or logical groups, can share infrastructure with clear isolation, controlled access, and simple administration.

FlashSystem grid administrators define each tenant scope by grouping and assigning systems, storage partitions, and pools to that tenant. Tenant administrators then define and manage volume groups on their assigned pools and manage only the resources in their scope, without involvement at the grid level.

Multi-tenancy works with , , and other functions.

The system provides role-based views that separate FlashSystem grid administration from tenant-scoped operations. FlashSystem grid administrators see all systems and shared resources across the FlashSystem grid. Tenant administrators see only the systems and partitions that are assigned to their tenant.

FlashSystem grid administrator

Roles and responsibilities

FlashSystem grid administrators configure and maintain shared, cross-system resources.

Scope and visibility

- **Systems:** All storage systems in the FlashSystem grid, including product name, code level, hardware health, connectivity health, and pool availability.
- **Partitions:** All partitions associated with all tenants, including replication topology, high availability status, disaster recovery status, migration status, placement evaluation, and volume group counts.
- **User groups:** All tenant-scoped user groups defined across the FlashSystem grid.

Typical tasks

- Plan capacity and resource governance.
- Configure high availability and disaster recovery.
- Create and manage tenants and partitioning.
- Maintain system health and perform system-level maintenance.

Tenant administrator

Tenant administrators work within assigned systems and partitions.

Scope and visibility

- **Storage systems:** Only systems assigned to the tenant, with health indicators and pool counts.
- **Storage partitions:** Only partitions allocated to the tenant, with replication topology, HA/DR status, and volume group usage.

Typical tasks

- Create volume groups and volumes.
- Add hosts and mappings.
- Assign and manage protection policies within tenant scope.

Key concepts

Stable identity

Each tenant is represented on every system by an **Ownership Group** and identified consistently across the FlashSystem grid. This stable identifier supports cross-system aggregation and conflict detection.

Separation of duties

- **Grid administrator**: Owns partitions, HA pairings, DR topology, and cross-system governance.
- **Tenant administrator**: Owns volumes, hosts, mappings, and protection policy assignments within the tenant.

Pool model without quotas

The system creates **child pools** per tenant per system without fixed capacity reservations. **Partitions** in the FlashSystem grid determine the usable capacity that administrators control. This approach reduces capacity fragmentation and supports growth.

Before you begin (prerequisites)

Before configuring multitenancy on any system or adding any system into an existing tenant.

Ownership group capacity

Each system must have a free slot to create a tenant ownership group. For details, see . A system supports a defined maximum number of ownership groups.

Pool availability

Each system must have at least one parent pool. Tenant child pools are created from the parent pool. For details, see .

System reachability

Each system must be reachable and responsive to complete tenant operations (create ownership groups, child pools, and user groups). If a system becomes unreachable, the configuration state shows pending, and the UI displays an error. Restore connectivity and retry.

Configuring multi-tenancy in FlashSystem grid

Multi-tenancy in FlashSystem grid can be configured through the management GUI or the command-line interface.

In the management GUI, grid administrators create tenants by specifying a tenant name, selecting the systems to include, and creating a tenant-specific user group and user, ensuring the correct system and IP

address are chosen for the User location. Tenant details, assigned systems, and storage partitions can be viewed under **Tenants**, with clear, actionable error messages provided as needed.

Grid administrators oversee cross-system resources, while tenant administrators operate within isolated, tenant-specific Ownership Groups.

Configuring multi-tenancy in the management GUI

It is recommended to configure multi-tenancy in the FlashSystem grid by using the management GUI. As a FlashSystem grid administrator, you can create tenants, assign systems, create tenant user groups and users, and remove tenants when they are no longer required.

Creating a tenant

To create a tenant by using the management GUI, log in to the FlashSystem grid interface with FlashSystem grid administrator credentials. From the navigation panel, select **FlashSystem grid** > **Tenants**, and then select **Create tenant**.

Enter a name for the tenant, select the systems to assign, and create the tenant user group and user.

Select the **User location (system)** where the user resides to create the user credentials.

After reviewing the information, complete the workflow to create the tenant. The FlashSystem grid then creates the tenant, assigns the selected systems, and generates the required tenant-scoped user groups, users, and child pools. The tenant identity is consistent across systems through the Tenant UUID.

After tenant creation is complete, a tenant administrator can create storage partitions within the tenant's Ownership Group and configure replication policies for those partitions as required.

Viewing tenant details

To view a tenant, select **FlashSystem grid** > **Tenants** and select the **Tenant name** from the list.

The tenant details page shows the systems that are assigned to the tenant, the storage partitions that the FlashSystem grid created. If the configuration has an issue, the management GUI displays an actionable message to help the administrator identify and correct the problem.

The management GUI supports adding systems to a tenant and removing systems from a tenant. Systems can be added or removed by using the **Add system** or **Remove system** for the tenant. A system can be removed from a tenant only after all associated resources are removed. These resources include storage partitions, child pools, and user groups.

Removing a tenant

A tenant can be removed from the FlashSystem grid only after all associated systems are removed from the tenant. When the last associated system with a tenant is removed, the tenant is automatically deleted.

Viewing tenant information in the command-line interface

Note: You can manage tenant resources by using the command-line interface. However, this method requires multiple synchronization steps for ownership groups and UIDs across systems. To reduce manual effort and configuration complexity, use the FlashSystem grid management GUI.

After you create the tenant, you can view tenant information by using following commands.

-
-
-
-
-

Chapter 11. Storage partitions

A storage partition groups the configuration for one or more applications that need to be managed together.

A storage partition allows for simplified and scoped management of the resources within the storage system or a larger FlashSystem grid environment. A storage partition acts as a virtual storage system with a focus on the logical resources that are used by an application.

Within a storage partition, all volumes must be in volume groups, and hosts can be mapped to volumes in the same storage partition only. Host clusters can contain hosts only in the same storage partition. Volumes in a storage partition cannot be mapped to other IBM Storage Virtualize systems. These rules ensure that storage partitions can be nondisruptively migrated to another system in the FlashSystem grid, or be configured for high availability (HA) without conflicting with other resources that exist on either of the systems.

New hosts, host clusters, volumes, and volume groups can be created in a storage partition. However, existing storage resources within the system cannot be added directly to a partition. A draft partition can be used to construct a partition with existing objects that meets the requirements of a partition. Some of the requirements of partitions are not enforced while constructing a draft partition. This allows the addition of resources to the partition iteratively as the dependencies between application are identified and resolved. A draft partition can be published only when all the requirements of a partition are met. The GUI helps to add either volumes, or volume groups, or both to the draft partition and indicates which volumes, volume groups, and hosts are added to the partition automatically. A draft partition cannot be migrated or made highly available.

Management IP address for storage partitions

Each storage partition can be assigned a static management IP address, providing management access to the partition through CLI (SSH) or REST API. When a partition is migrated, the partition management IP address stays with the partition, ensuring continuous access for automation, monitoring, and external tools.

Use the partition panel in the management GUI to configure an IP address for a partition. When the partition management IP address is created, it is bound to the same Ethernet port as the system management IP. If multiple system management IPs are configured, it attaches to the port with the lowest number.

For more information, see Management access using partition IP address.

Partition-scoped and system-scoped resources

Storage partitions contain resources that are scoped to the partition, such as volumes, volume groups, hosts, and IP addresses. These resources move with the partition when it is migrated to another system. Some resources are scoped to the system and do not migrate with the partition. These resources include system users, user groups, and ownership groups. Partition users are user accounts that are scoped to a specific partition. Partition users move with the partition during migration and ensure that access to the partition is maintained on the target system.

Certificates for storage partitions

After you configure the management IP address, you can configure a partition certificate for a storage partition by using the CLI (SSH) or the REST API.

You can assign a certificate to each storage partition through the CLI (SSH) or the REST API. When a partition is migrated, the partition management certificate migrates with the partition. This behavior ensures continuous access for automation, monitoring, and external tools.

© Copyright IBM Corp. 2025

You can use the **Storage partitions** panel in the management GUI to configure a certificate for a storage partition.

Storage partitions and replication

Storage partitions can be configured to be highly available between two independent storage systems by associating a replication policy to the partition. Systems automatically ensure that the logical configuration, data, and management access are configured to be highly available by managing all provisioning activities and configuring the replication. You can assign a management IP address to a highly available partition, ensuring seamless failover for automation, monitoring and external tools. When a storage partition operates in High Availability, certificate management for that partition is performed only from the Active Management System (AMS) and is restricted from all non-AMS interfaces. For more information about HA, see "High availability" on page 191.

Two storage partitions in different systems can be linked together to allow replication for disaster recovery for one or more volume groups within the partition. Volume groups within the partition can be configured with different replication policies, allowing for different recovery point objectives, or selective replication of a subset of the volume groups.

The HA and disaster recovery features can be configured simultaneously allowing for disaster recovery replication to be configured on volume groups within a highly available storage partition. This configuration allows for HA between two production systems, with replication to a third system for disaster recovery. For more information, see "High availability and disaster recovery (3-site replication)" on page 203.

vCenter partitions

You can register multiple vCenters from VMware vSphere servers with IBM Storage FlashSystem or IBM Storage Virtualize, which helps VMware environments to achieve multitenancy. Each of these vCenters acts as a separate system and shares the resources on the storage system. Each vCenter is a self-contained unit (unaware of each other's resources) that belongs to an ownership group.

A vSphere deployment can use Virtual Machine File System (VMFS) datastores from the same storage system.

You can use a storage partition to manage the entire storage workload for a vSphere deployment. A common storage partition includes:

- All VMFS datastore volumes managed by a single vCenter server.
- The ESX server host-objects in that vSphere deployment.

All related storage resources corresponding to the vSphere deployment can be included in a single partition.

Deployment scenarios

VMFS-only deployment

If the deployment uses only VMFS datastores, then the following lightweight approach is used:

- *Legacy mode*: In this mode, storage partitions are not used for VMFS datastores and ESX hosts. However, features like policy-based high availability and partition migration are not available.
- *vCenter partition mode*: A partition is used for managing VMFS datastores and ESX server host-objects. The VMFS datastore does not require a child pool or an ownership group.

Managing storage partitions

You can combine storage partitions by using the management GUI.

You can also split a storage partition by creating a draft partition and specifying the original partition as the source. Resources from the original partition can be added to the draft partition while they remain

associated with the original partition. When you publish the draft partition, the resources move from the original partition to the new partition.

Chapter 12. Data migration

This topic describes methods to migrate data.

Data migration applications

Data migration has the following characteristics and applications:

- Storage partitions can be non-disruptively migrated between systems in a FlashSystem grid. For details, see “Migrating storage partitions between systems” on page 247.
- The system supports migrating existing volumes to a data reduction pool with volume mirroring. For details, see “Moving volumes to a data reduction pool” on page 251.
- For systems with encryption enabled, you can migrate existing volumes from non-encrypted pools to encrypted pools. For details, see “Migrating volumes to an encrypted pool” on page 253.
- Changes the MDisks to which the volume is allocated.
- Does not disrupt access to the volume data.
- Can be used when you are removing MDisks from a clustered system.
- Can move data to an MDisk that has appropriate performance for host requirements.
- Can move data from image mode volumes to MDisks when you are adding a system to an existing SAN infrastructure.
- Can migrate striped and sequential volumes to an image mode volume.

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

Migrating storage partitions between systems

IBM Storage Virtualize supports the non-disruptive migration of storage partitions between systems, which enables a partition to be migrated from one Storage Virtualize system to another without any application downtime.

Migrating storage partitions can be used for the following cases:

- Migrating to a new storage array before decommissioning the original system.
- Load-balancing by migrating storage partitions from overloaded systems to other systems.

Migration enables relocation of storage partitions from a source system to a target system location. As a consequence of which, the following flow of events occurs:

- All resources that are associated with the storage partition are created on the migration target system, and the copy process starts from the source system to the target system.
- Host I/O is served from the migration target storage system once the synchronization process is complete.
- At the end of the migration process, the storage partition and all of its resources are removed from the source system.

A storage partition that has replication configured for DR can be migrated.

If a storage partition contains volume groups with snapshot policies assigned, snapshot policies are created on the target system and assigned to volume groups as part of migration process.

Storage partition migration is supported on Fibre Channel, iSCSI, and NVMe-oF hosts.

Limits and restrictions

- Draft storage partitions and storage partitions that are configured for policy-based high availability are not supported for storage partition migration.

© Copyright IBM Corp. 2025

- Only one storage partition can be migrated at a time from a system. Any consecutive migrations that are attempted get queued, and are scheduled automatically as per the sequence of invocations when the earlier migrations complete.
- If a system has storage partitions that are configured for high availability, partition migration can only be used with the HA partner system.

Prerequisites

Before you can use non-disruptive storage partition migration, ensure that the following prerequisites are met:

1. Migrating a storage partition requires both Fibre Channel and IP connectivity between the source and target storage systems. The requirements are the same as that for configuring high availability (HA). Review the high availability requirements to ensure that the storage partition supports being migrated and the host operating systems support this feature.
2. Confirm that both systems are members of the same FlashSystem grid, or that neither system is a member of a FlashSystem grid, and that both systems meet the requirements for FlashSystem grid.
3. You must link storage pools between the systems. Establish a pool link between systems on the current source and target systems. To create linked child pools, use the **Pools** panel on the management GUI on the system that does not have the child pools created.
4. If a storage partition has replication that is configured for disaster recovery (DR), then the target system (used for the partition migration) and the DR-linked system must be:
   - In a partnership that is already configured, or
   - Configured for automatic selection of Fibre Chanel or short-distance Ethernet RDMA connectivity, or
   - Configured for automatic selection of IP connectivity, if the *replication* type portsets are configured on systems.

Procedure

Placement advisory

The FlashSystem grid GUI shows an AI-powered list of target storage systems for storage partition migration. This information is provided by IBM Storage Insights, which analyzes the historical data of all systems to find the best placement for partitions in the FlashSystem grid.

Ensure that the following prerequisites are met:

- All systems in the FlashSystem grid must be added to the same IBM Storage Insights tenant.
- The system must have an IBM Storage Insights Pro license.
- The IBM Storage Insights API key must be registered with the FlashSystem.

Using the GUI to migrate a storage partition

On the GUI, storage partition migration can be initiated from the **Storage partitions** page of the FlashSystem grid only. A storage partition can be migrated only to systems that are part of the same FlashSystem grid.

To migrate a storage partition, go to **FlashSystem grid** > **Storage partitions** and select **Migrate partition**. You can use the **Partition placement** option on the **Storage partitions** panel to begin AI-powered evaluation of systems.

If user intervention is needed to complete a partition migration, the **Migration status** link guides you to take the necessary action.

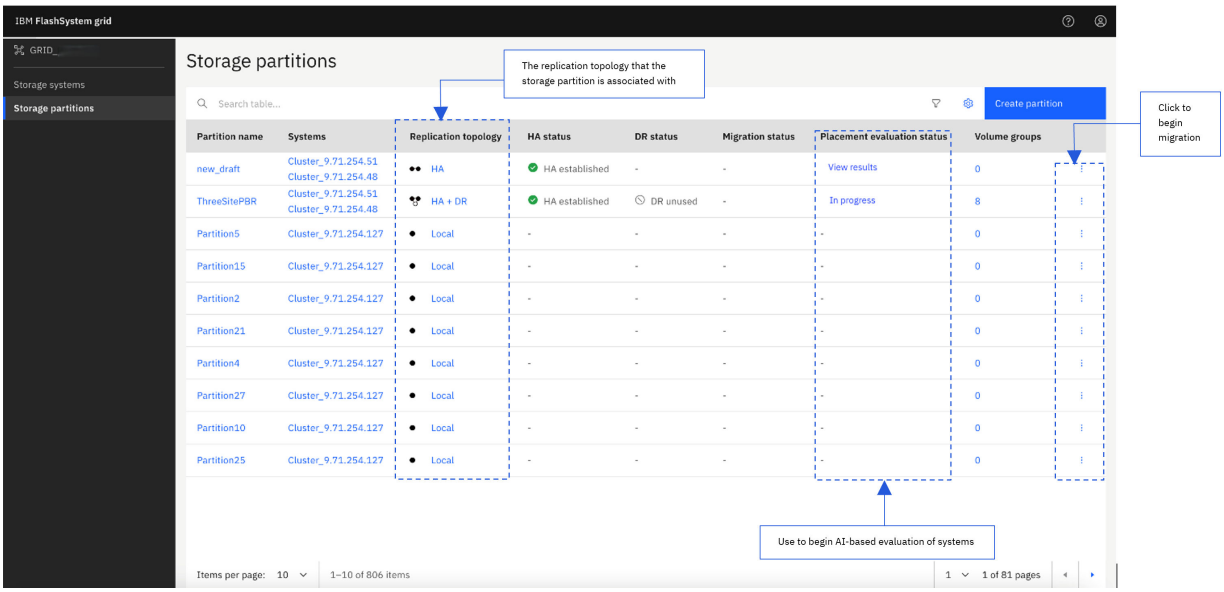


Figure 23. Storage partitions page of the FlashSystem grid.

Using the CLI to migrate a storage partition

1. Run the **chpartition** command with the `-location` or `-locationuuid` option to migrate the storage partition to a different system.
2. The following example initiates a migration of storage partition to the designated location system:

   ```
   chpartition -location <remote_system> mypartition1
   ```

   Note: If there are two or more systems with the same name, then specify remote system ID instead of remote system name while initiating the migration.
3. To check the migration status, run **lspartition**.

   Note: You can monitor the progress of the migration, including the amount of data remaining to be copied and the estimated time of completion, by using the **lsvolumegroupreplication** command.
4. Once the storage partition's data and configuration has been migrated to the target storage system, you must confirm that the affected hosts have discovered paths to the volumes in the target storage system.

   Note: The event for confirmation of path discoveries by the mapped hosts does not get raised if all the hosts that are mappedthat are to a partition are known to rescan the storage automatically after regular intervals, and the system detects the storage discoveries driven by the host. These host attributes can be managed by using the **autostoragediscovery** option in **mkhost** and **chhost** CLIs.
5. To complete the migration, remove the copy of the data and configuration of the storage partition from the source storage system. Before committing the migration, verify the host's multipath configurations, ensure path availability and redundancies, and assess IO performance at the target system.

   Note: The migration can be canceled and the partition rolled back to its original system before fixing the event if you observe any issues with the migrated storage partition.

   The rollback operation needs a confirmation at the source system for either continuing or canceling the migration operation. You can continue the migration if you are able to resolve the issues at the target system. To continue the migration operation without continuing with rollback, run the following at source system:

   ```
   chpartition -migrationaction fixeventwithchecks <partition_id/name>
   ```

   To continue the rollback and cancel the partition migration, run the following at source system:

```
chpartition -override -location <source_system> <partition_id>
```

6. If the storage partition has replication configured for DR, you might need to temporarily pause or resume DR replication.

   Note: When you migrate a storage partition configured for DR, if the partition contains recovery volume groups, DR replication must be paused temporarily by setting all recovery volume groups in the partition as independent. Before committing or aborting a migration, you can restart replication for DR on the volume groups that were made independent as part of the migration process. You must restart replication from the system that you want to use as the production copy for each volume group. To restart replication in the same direction as before the migration was started, with the volume groups on the local system configured as recovery copies, you must log into the remote system (in the DR-linked partition) and restart replication from that system.
7. If the storage partition on the source system contains volume groups with snapshot policies assigned, go to the target storage system and create snapshot policies and assign them to volume groups in the partition on the target system.
8. An informational event at the target storage system marks the completion of the storage partition migration.

   Note:

   You can monitor the migration by using `migration_status` field that is shown by **`lspartition`** command indicating that there is no migration activity active or queued for that storage partition.

   If you specify a different target location in an ongoing storage partition migration, the ongoing migration is aborted and the partition migration to the new target location is queued. You can use the **`-override`** option in the **`chpartition`** command to abort an ongoing partition migration and specify a new target location.

   You can abort an active or queued migration operation by specifying the source system location as the specified location. This needs an `-override` option since the abort operations remove the data that is replicated until that moment.

Completing the migration

If any user intervention is needed to complete a partition migration, click the status links in the **Migration status** column, which guides you to take the necessary action. You can do the following tasks:

Stop disaster recovery (DR)
When you migrate a storage partition configured for DR, if the partition contains recovery volume groups, DR replication must be paused temporarily by setting all recovery volume groups in the partition as independent.

Automated host rescan
Hosts mapped to a partition are enabled to rescan the storage automatically at regular intervals and the system detects the discovery status and the I/O activity. Host rescan of the storage systems recognizes and updates all the paths.

To initiate a rescan of hosts on the storage systems, go to **Storage partition** > **Host**. Then return to the migration wizard and click **Mark as complete**.

Restart disaster recovery (DR)
Before committing or aborting a migration, you can restart replication for DR on the volume groups that were made independent as part of the migration process. You must restart replication from the system that you want to use as the production copy for each volume group. To restart replication in the same direction as before the migration was started, with the volume groups on the local system that is configured as recovery copies, you must log in to the remote system (in the DR-linked partition) and restart replication from that system.

Roll back a migration
The migrated storage partition that serves I/Os from the target (after host health checks at the source) can be rolled back to its source location. You can roll back a migration from the target system.

Abort a migration
To discontinue the migration process, you can abort the migration from the source system.

Commit a migration
After migration is almost complete, the configuration and data is copied to the target system. You can commit the migration to finalize the changes. Committing the migration on the target system completes the migration process.

Rolling back a migration

Rollback is needed if the migrated storage partition does not perform as expected on the target system after migration. The rollback operation switches the I/O operations to the storage partition at the source location.

The migrated storage partition serving I/Os from the target location (after you check and fix the host health at the source) can be rolled back to its source location.

Moving volumes to a data reduction pool

Data reduction pools reduce the amount of data that is stored on internal drives and external storage systems through data reduction techniques including compression, deduplication, and thin provisioning. A volume in a data reduction pool can be configured to use different capacity savings methods simultaneously. Data reduction pools reclaim previously used capacity that is no longer needed by host systems. Support for the host SCSI unmap command is enabled by default.

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

Note: The Data Reduction Estimation Tool also provides some analysis of potential compression savings for volumes. However, use the management GUI or the command-line interface to run the integrated Comprestimator Utility to gather data for potential compression savings for volumes in data reduction pools.

Support for the host SCSI unmap command is enabled by default.

Using the management GUI

To migrate the volume to a data reduction pool, complete these steps:

1. Create a data reduction pool by completing these steps:
   a. In the management GUI, select **Pools** > **Pools**.
   b. On the **Pools** page, click **Create**.
   c. On the **Create Pool** page, enter a name of the pool and select **Data Reduction**.
   d. Click **Create**.
2. Add storage to the data reduction pool by completing these steps:
   a. In the management GUI, select **Pools** > **Pools**.
   b. Right-click the data reduction pool that you created and select **Add Storage**.
   c. Select from the available storage and reserve capacity to the pool. Click **Assign**.

      Note: You cannot select from available storage when you are creating a data reduction child pool. A data reduction child pool has access to the complete available capacity of the parent pool without limit. A data reduction child pool cannot be created if the parent pool does not have available capacity.
3. Create a copy of the volume in the standard pool by completing these steps:
   a. In the management GUI, select **Pools** > **Volumes by Pools**.

b. Select the existing pool that contains the volumes you want to migrate to the data reduction pool. Right-click the volume and select **Add Volume Copy**.

c. On the **Add Volume Copy** page, select the data reduction pool to create a copy of the volume. Copy 1 is the original volume in the original pool, while copy 2 is created in the data reduction pool. Ensure that either thin-provisioning or compressed is selected for the capacity savings. For the option that you select, you can also select to use deduplication for the volume that you create. For example, you can create a thin-provisioned volume that also uses deduplication to remove duplicate data.

d. Click **Add**.

e. Ensure that the copies are synchronized before you proceed to the next steps. On the **Volumes** page, ensure that **Yes** is displayed for copy 2 in the **Synchronized** column.

Note: The time synchronization takes to complete depends on the size of the volume and system performance. You can increase the synchronization rate by right-clicking the volume and selecting **Modify Mirror Sync Rate**.

4. To complete the migration, remove the original volume copy by completing these steps:

a. In the management GUI, select **Volumes** > **Volumes**.

b. Right-click the volume copy and select **Delete**.

c. Click **Yes** to confirm the deletion.

Migrating extents using the CLI

To improve performance, you can migrate extents using the command-line interface (CLI).

About this task

The system provides various data migration features. These features can be used to move the placement of data both *within* parent pools and *between* parent pools. These features can be used concurrently with I/O operations. You can use either of these methods to migrate data:

1. Migrating data (extents) from one MDisk to another (within the same parent pool). This method can be used to remove highly used MDisks.
2. Migrating volumes from one parent pool to another. This method can be used to remove highly used parent pools. For example, you can reduce the use of a pool of MDisks. Child pools that receive their capacity from parent pools, cannot have extents that are migrated to them.

Notes:

1. The source MDisk must not currently be the source MDisk for any other migrate extents operation.
2. The destination MDisk must not be the destination MDisk for any other migrate extents operation.

Migration commands fail if the target or source volume is offline, there is no quorum disk defined, or the defined quorum disks are unavailable. Correct the offline or quorum disk condition and reissue the command.

You can determine the use of particular MDisks by gathering input/output (I/O) statistics about nodes, MDisks, and volumes. After you collect this data, you can analyze it to determine which MDisks are used frequently. The procedure then takes you through querying and migrating extents to different locations in the same parent pool. This procedure can only be completed using the command-line interface.

If performance monitoring tools indicate that an MDisk in the pool is being overused, you can migrate data to other MDisks within the same parent pool.

Procedure

1. Determine the number of extents that are in use by each volume for the MDisk by issuing this CLI command:

```
lsmdiskextent mdiskname
```

This command returns the number of extents that each volume is using on the MDisk. Select some of these extents to migrate within the pool.

2. Determine the other MDisks that are in the same volume.
   a) To determine the parent pool that the MDisk belongs to, issue this CLI command:

   ```
   lsmdisk mdiskname | ID
   ```

   b) List the MDisks in the pool by issuing this CLI command:

   ```
   lsmdisk -filtervalue mdisk_grp_name=mdiskgrpname
   ```

3. Select one of these MDisks as the target MDisk for the extents. You can determine how many free extents exist on an MDisk by issuing this CLI command:

   ```
   lsfreeextents mdiskname
   ```

   You can issue the **lsmdiskextent** *newmdiskname* command for each of the target MDisks to ensure that you are not just moving the over-utilization to another MDisk. Check that the volume that owns the set of extents to be moved does not already own a large set of extents on the target MDisk.

4. For each set of extents, issue this CLI command to move them to another MDisk:

   ```
   migrateexts -source mdiskname | ID -exts num_extents
   -target newmdiskname | ID -threads 4 -vdisk vdiskid
   ```

   where *num_extents* is the number of extents on the *vdiskid*. The *newmdiskname* | *ID* value is the name or ID of the MDisk to migrate this set of extents to.

   Note: The number of threads indicates the priority of the migration processing, where **1** is the lowest priority and **4** is the highest priority.

5. Repeat the previous steps for each set of extents that you are moving.
6. You can check the progress of the migration by issuing this CLI command:

   ```
   lsmigrate
   ```

Migrating volumes to an encrypted pool

For systems with encryption enabled, you can migrate existing volumes from non-encrypted pools to encrypted pools. Both the management GUI and the command line interface can be used to migrate volumes to encrypted pools.

During system setup in the management GUI, you can activate and enable encryption licenses. Any pools that are created after encryption is enabled are assigned a key that can be used to encrypt and decrypt data. However, if encryption was configured after volumes were already assigned to non-encrypted pools, you can migrate those volumes to an encrypted pool by using child pools. When you create a child standard 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. You can use either the management GUI or the command line interface to migrate volumes to an encrypted pool.

When you create a child data reduction pool, the pool inherits the encryption status of the parent pool, and both pools use the same encryption key. You cannot use the **-encrypt** parameter when you create a child data reduction pool.

The system supports both internal storage and external storage. Internal storage consists of arrays that attached directly to the system through serial-attached SCSI connections. External storage consists of arrays that are attached to the system over a storage area network (SAN). Depending on the type of underlying storage the volume uses, the migration process is different.

If you are migrating volumes that use SAN-attached MDisks, you must complete the migration steps for “Migrating volumes that use SAN-attached MDisks” on page 254.

Migrating volumes that use SAN-attached MDisks

To migrate volumes that use SAN-attached MDisks and either standard pools or data reduction pools, complete these steps:

1. In the management GUI, select **Pools** > **Pools**.
2. Right-click the non-encrypted parent pool that contains the volumes that you want to migrate and select **Create Child Pool**.
3. On the **Create Child Pool** page, enter the name for the child pool. To create a standard child pool, enter the amount of capacity. Ensure that you select enough capacity to accommodate the migrated volumes. Encryption is selected by default when the system is enabled for encryption. To create a data reduction child pool, continue to the next step. Data reduction child pools have access to the complete free capacity of the parent pool without limit. A data reduction child pool cannot be created if the parent pool does not have free capacity.
4. Click **Create**. After the child pool is created, you can migrate the volumes to the child pool by adding volume copies.
5. In the management GUI, select **Volumes** > **Volumes by Pools**.
6. Select the non-encrypted parent pool to display all the volumes.
7. Right-click the volume and select **Add Volume Copy...**.
8. On the **Add Volume Copy** page, select **Basic** for the type of copy that you are creating. From the list of available pools, select the child pool as the target pool for the copy of the volume.
9. Click **Add**.
10. Repeat these steps to add volume copies to the encrypted child pool for the remaining volumes in the parent pool.
11. After all the copies are synchronized in the encrypted child pool, you can delete the primary copies from the parent pool. The empty parent pool must remain unused to use encrypted volumes in the child pool.

To migrate volumes that use SAN-attached MDisks and either standard pools data reduction pools, complete these steps:

1. In the command line interface, enter the following command to create a standard child pool:

   ```
   mkmdiskgrp -name my_encrypted_child_pool -parentmdiskgrp mypool -encrypt yes
   ```

   where `my_encrypted_child_pool` is the name of the new child pool and `mypool` is the name of the parent pool. Enter the following command to create a data reduction child pool:

   ```
   mkmdiskgrp -name my_encrypted_child_pool -parentmdiskgrp mypool -datareduction yes -noquota
   ```

   where `my_encrypted_child_pool` is the name of the new child pool and `mypool` is the name of the parent pool.
2. Issue the CLI command `addvdiskcopy` to create mirrored copies of the volumes that are in the parent pool, in the new child pool. In the following example, `my_encrypted_child_pool` is the name of the new child pool and `volume1` is the name of the volume that is being copied.

   ```
   addvdiskcopy -autodelete -mdiskgrp my_encrypted_child_pool -vdisk volume1
   ```

   Use the **`-autodelete`** operand to automatically delete the primary copy of the volume after the copy synchronizes.
3. Repeat step 2 until all the volumes from the original parent contain mirrored copies in the new child pool. The empty parent pool must remain unused to use encrypted volumes in the child pool.

Migrating volumes that use internal MDisks

If you have internally-attached enclosures that either do not support encryption or that do not have an encryption license that is activated and enabled and want encrypted volumes, you can still migrate

volumes to encrypted pools. Before you can migrate volumes, you must upgrade your software to a version that supports encryption and obtain an encryption license. After you receive authorization documents for the encryption license, you must activate the license and enable the encryption function on the system before you attempt migration of any volumes. In addition, you must ensure that the unencrypted pool contains enough free extents to complete the migration. The migration fails if enough free extents are not available. You have two options to ensure that the pool contains enough extents:

- If you have spare drives, create a new encrypted array with the spare drives.
- If you do not have spare drives, add unassigned external MDisks to the parent pool and create a new encrypted child pool. If you chose this option and intend to keep the external MDisk in the pool after migration, then use the instructions in “Migrating volumes that use SAN-attached MDisks” on page 254.

These steps assume that these tasks were completed.

To migrate volumes that use internal MDisks in the management GUI, complete these steps:

1. In the management GUI, select **Pools** > **MDisks by Pools**.
2. Select the pool and right-click the MDisk that represents the array that you want to delete and select **Remove**.
3. To add the array back to the pool, select **MDisk by Pools**. The new available internal storage is displayed at the bottom of the pane. Select the icon that represents the drive class of the drives that were formerly in the array that was deleted and select **Assign**. The new array is encrypted by default. After all the arrays are encrypted, the pool is considered encrypted and data on volumes in that pool is also encrypted.
4. To verify that the data on these volumes was migrated to encrypted arrays and are now encrypted, select **Volumes** > **Volume by Pools**.
5. Select the pool and verify that the volumes that are listed are encrypted. A key icon is displayed under the **Encryption** column to indicate that the data on the volume is encrypted.

For migrated volumes that use internal MDisk in the command line interface, complete these steps:

1. To remove an array from pool, enter the following command:

   ```
   rmarray -mdisk mdisk3 pool2
   ```

   In this example, `mdisk3` is the array that is removed from `pool2`.
2. To create a new encrypted array and add the array to the storage pool, enter the following command:

   ```
   mkarray -level raid5 -drive 0:1:2:3 -mdiskgrp pool2
   ```

   In this example, the array is encrypted by default because encryption is enabled on the I/O group for the array. The data that is stored on the volumes that use that array is now encrypted.
3. To verify that the volumes in the pool are encrypted, enter the following command:

   ```
   lsvdisk
   ```

   The encrypt state is **yes** if all the data on the volume is encrypted.

Migration methods for image mode volumes

Several methods can be used to migrate image mode volumes into managed mode volumes.

About this task

To complete any type of migration on an image mode volume, the image mode volume must first be converted into a managed mode volume. The volume is automatically converted into a managed mode volume whenever any kind of migration activity is attempted. After the image-mode-to-managed-mode migration occurs, the volume becomes a managed mode volume and is treated the same way as any other managed mode volume.

If the image mode disk has a partial last extent, this last extent in the image mode volume must be the first to be migrated. This migration is processed as a special case. After this special migration operation occurs, the volume becomes a managed mode volume and is treated in the same way as any other managed mode volume. If the image mode disk does not have a partial last extent, no special processing is completed. The image mode volume is changed into a managed mode volume and is treated the same way as any other managed mode volume.

An image mode disk can also be migrated to another image mode disk. The image mode disk becomes managed while the migration is ongoing, but returns to image mode when the migration is complete.

You can complete the following types of migrations:

- Migrate extents
- Migrate a volume
- Migrate to image mode

Note: Migration commands fail for the following reasons:

- The target or source volume is offline.
- A quorum disk is not defined.
- Defined quorum disks are unavailable.

Correct the offline or quorum disk condition and reissue the command.

Complete the following steps to migrate volumes:

Procedure

1. Dedicate one storage pool to image mode volumes. Ensure that the pool that you select is not a child pool.
2. Dedicate one storage pool to managed mode volumes.
3. Use the migrate volume function to move the volumes.

Migrating volumes between pools by using the CLI

You can migrate volumes between pools by using the command-line interface (CLI).

About this task

You can determine the usage of particular MDisks by gathering input/output (I/O) statistics about nodes, MDisks, and volumes. After you collect this data, you can analyze it to determine which volumes or MDisks are hot. You can then migrate volumes from one storage pool to another.

Complete the following step to gather statistics about MDisks and volumes:

1. Use the Secure Copy (SCP) or SSH File Transfer Protocol (SFTP) to retrieve the dump files for analyzing. For example, issue the following command:

   ```
   scp clusterip:/dumps/iostats/v_*
   ```

   This command copies all the volume statistics files to the AIX host in the current directory.
2. Analyze the memory dumps to determine which volumes are hot. You can determine which MDisks are used heavily to spread the data evenly across all MDisks in the storage pool by migrating the extents.

After you analyze the I/O statistics data, you can determine which volumes are hot. You also need to determine the storage pool that you want to move this volume to. Either create a new storage pool or determine an existing group that is not yet overly used. Check I/O statistics files that you generated and then ensure that the MDisks or volumes in the target storage pool are used less than the MDisks or volumes in the source storage pool.

You can use data migration or volume mirroring to migrate data between storage pools. Data migration uses the command **migratevdisk**. Volume mirroring uses the commands **addvdiskcopy** and **rmvdiskcopy**.

Migrating data by using migratevdisk

You can use the **migratevdisk** command to migrate data between two storage pools. When you issue the **migratevdisk** command, a check is made to ensure that the destination of the migration has enough free extents to satisfy the command. If it does, the command proceeds. The command takes some time to complete.

Notes:

- You cannot use the data migration function to move a volume between storage pools that have different extent sizes.
- Migration commands fail if the target or source volume is offline, there is no defined quorum disk, or the defined quorum disks are unavailable. Correct the offline or quorum disk condition and reissue the command.
- The system supports migrating volumes between child pools within the same parent pool or migrating a volume in a child pool to its parent pool. Migration of volumes fails if source and target child pools have different parent pools. However, you can use **addvdiskcopy** and **rmvdiskcopy** commands to migrate volumes between child pools in different parent pools.

When you use data migration, it is possible for the free destination extents to be consumed by another process. For example, if a new volume is created in the destination parent pool or if more migration commands are started. In this scenario, after all the destination extents are allocated, the migration commands suspend and an error is logged (error ID 020005). To recover from this situation, use either of the following methods:

- Add more MDisks to the target parent pool, which provides more extents in the group and allows the migrations to be restarted. Mark the error as fixed before you reattempt the migration.
- Migrate one or more volumes that are already created from the parent pool to another group. This action frees up extents in the group and allows the original migrations to be restarted.

Complete the following steps to use the **migratevdisk** command to migrate volumes between storage pools:

1. After you determine the volume that you want to migrate and the new storage pool that you want to migrate it to, issue the following CLI command:

```
migratevdisk -vdisk vdisk_name
```

```
 -mdiskgrp
 mdisk_group_name -threads 4
```

2. You can check the progress of the migration by issuing the following CLI command:

```
lsmigrate
```

Migrating data by using volume mirroring

When you use data migration, the volume goes offline if either pool fails. Volume mirroring can be used to minimize the impact to the volume because the volume goes offline only if the source pool fails. You can migrate volumes between child pools or from a child pool to a parent pool by using the **addvdiskcopy** and **rmvdiskcopy** commands instead of using the **migratevdisk** command. Complete the following steps to use volume mirroring to migrate volumes between pools:

1. After you determine the volume that you want to migrate and the new pool that you want to migrate it to, enter the following command:

```
addvdiskcopy -mdiskgrp mdisk_group_name -autodelete vdisk_name
```

where *mdisk_group_name* is the name of the new storage pool and *vdisk_name* is the name of the volume that is being copied. Specify **-autodelete** to automatically delete the original copy of the volume after the copies are synchronized.

2. The copy ID of the new copy is returned. The copies now synchronize such that the data is stored in both storage pools. You can check the progress of the synchronization by issuing the following command:

```
lsvdisksyncprogress
```

Moving a volume between I/O groups

To move volumes between I/O groups non-disruptively, ensure that the hosts that are mapped to the volume support non-disruptive volume move. AIX, Windows, Red Hat, SLES, and VMware hosts are supported. Volumes in a data reduction pool cannot be moved using this method. The cached data that is held within the system must first be written to the system disk before the allocation of the volume can be changed.

About this task

Modifying the I/O group that services the volume can be done concurrently with I/O operations if the host supports non- disruptive volume move. It also requires a rescan at the host level to ensure that the multipathing driver is notified that the allocation of the preferred node has changed and the ports by which the volume is accessed has changed. This can be done in the situation where one pair of nodes becomes over used.

If there are any host mappings for the volume, the hosts must be members of the target I/O group or the migration fails.

Verify that you created paths to I/O groups on the host system. After the system successfully adds the new I/O group to the volume's access set and you moved selected volumes to another I/O group, detect the new paths to the volumes on the host. The commands and actions on the host vary depending on the type of host and the connection method used. These steps must be completed on all hosts to which the selected volumes are currently mapped.

The volume between I/O groups can be moved by using the CLI.

To move a volume between I/O groups by using the CLI, complete the following steps:

Procedure

1. Issue the following command: **addvdiskaccess -iogrp <iogrp id> <name volume> <id name>**
2. Issue the following command: **movevdisk -iogrp destination iogrp -node new preferred node volume id/name**

   The system disables moving a volume if the selected volume is currently performing quick initialization. After the quick initialization completes, you can move the volume to another I/O group.
3. Issue the appropriate commands on the hosts that are mapped to the volume to detect the new paths to the volume in the destination I/O group.
4. After confirming that the new paths are online, remove access from the old I/O group: **rmvdiskaccess -iogrp iogrp id/name volume id/name**
5. Issue the appropriate commands on the hosts that are mapped to the volume to remove the paths to the old I/O group.

Chapter 13. Stretched systems

A stretched system is an extended high availability (HA) method that is supported by SAN Volume Controller systems to enable I/O operations to continue after the loss of half of the system.

A stretched system is also sometimes referred to as a *split system* because one half of the system and I/O group is usually in a geographically distant location from the other, often 10 kilometers or more. A third site is required to host a storage system that provides a quorum disk.

Enhanced stretched systems

The system supports an enhanced stretched system configuration. This configuration can be used regardless of whether the stretched system is configured with or without ISLs between nodes.

Enhanced stretched systems provide the following primary benefits:

- In addition to the automatic failover that occurs when a site fails in a standard stretched system configuration, an enhanced stretched system provides a manual override that can be used to choose which of two sites continues operation.
- Enhanced stretched systems intelligently route I/O traffic between nodes and controllers to reduce the amount of I/O traffic between sites, and to minimize the impact to host application I/O latency.
- Enhanced stretched systems include an implementation of extra policing rules to ensure the correct configuration of a standard stretched system.

Note: In an enhanced stretched system, set the site attribute in the node and controller object.

Stretched system configuration details

You can create an enhanced stretched system configuration where each node on the system is physically on a different site.

When used with mirroring technologies, such as volume mirroring or Copy Services, these configurations can be used to maintain access to data on the system in the event of power failures or site-wide outages.

The enhanced stretched system configuration, with the topology attribute of the system set to stretched, is detailed here. It is possible to non-disruptively move to the current enhanced stretched system configuration by following the final configuration steps that are presented here so that you get better availability and disaster recovery. Contact IBM Remote Technical support center for guidance on changing the topology of an existing system.

In a stretched system configuration, each site is defined as an independent failure domain. If one site experiences a failure, the other site can continue to operate without disruption. You must also configure a third site to host a quorum device that provides an automatic tie-break in the event of a potential link failure between the two main sites. The main site can be in the same room or across rooms in the data center, buildings on the same campus, or buildings in different cities. Different kinds of sites protect against different types of failures.

Sites are within a single location
If each site is a different power phase within a single location or data center, the system can survive the failure of any single power domain. For example, one node can be placed in one rack installation and the other node can be in another rack. Each rack is considered a separate site with its own power phase. In this case, if power was lost to one of the racks, the partner node in the other rack could be configured to process requests and effectively provide availability to data even when the other node is offline due to a power disruption.

Each site is at separate locations
If each site is a different physical location, the system can survive the failure of any single location. These sites can span shorter distances, for example two sites in the same city, or they can be spread

© Copyright IBM Corp. 2025

farther geographically, such as two sites in separate cities. If one site experiences a site-wide disaster, the remaining site can remain available to process requests.

If configured properly, the system continues to operate after the loss of one site. The key prerequisite is that each site contains only one node from each pair of nodes. Simply placing a pair of nodes from the same system in different sites for a stretched system configuration does not provide high availability. You must also configure the appropriate mirroring technology and ensure that all configuration requirements for those technologies are properly configured.

The system supports both Fibre-Channel and Ethernet connections between nodes with a stretched system. Ethernet connections are only supported through RDMA-capable adapters that must be installed and configured on each node within the stretched system. The system supports Fibre Channel and Ethernet (iSCSI) connections to hosts in stretched system environments. However, NVMe-based connections to hosts are not supported in stretched system configurations.

Stretched SVC clusters consisting of four nodes, where two nodes exist in one geographical site and two in another and where volume mirroring is not being used, are susceptible to a race condition relating to rediscovery of backend storage. This may result in the standard upgrade process becoming stalled and possibly needing to be aborted. IBM recommends upgrading SVC clusters in this configuration using the manual software upgrade procedure as documented in the online IBM Documentation.

Stretched system

A stretched system is designed to continue operation after the loss of one failure domain.

The stretched system cannot guarantee continuous operation after the failure of two failure domains. If the enhanced stretched system function is configured, you can enable a manual override for this situation.

The two partner systems can not be in the same production site. However, they can be collocated with the storage system that provides the active quorum disk for the stretched system.

Configuration steps

These additional configuration steps can be done by using the command-line interface (CLI) or the management GUI.

- Each node in the system must be assigned to a site. Use the **chnode** CLI command. If additional nodes are cabled to the system, you can specify these nodes as hot-spare nodes. Hot-spare nodes can nondisruptively take over host I/O operations if any node on the site becomes unavailable. For more information, see the topic about adding hot-spare nodes.
- Each back-end storage system must be assigned to a site. Use the **chcontroller** CLI command.
- Each host must be assigned to a site. Use the **chhost** CLI command
- After all nodes, hosts, and storage systems are assigned to a site, the enhanced mode must be enabled by changing the system topology to `stretched`.
- **Note:** IBM SAN Volume Controller systems support multiple I/O groups on long-term support releases only.

Planning for a stretched topology system

You can plan to meet the requirements of a stretched topology system.

Make sure that:

- Each node of an I/O group must be at a different site (location or power phase).
- There must be managed disks at each site.
- There must be managed disks at a third site for quorum disks.
- There must be enough fibre optic connections between the nodes on each site and the managed disks on the third site.

- The site attribute must be set to 1, 2 or 3 for each controller before the system topology can be set to stretched.

Note: Expansion enclosures are not supported with stretched system topology.

Connections:

- Nodes that are in the same I/O group and separated by more than 100 meters (109 yards) must use longwave Fibre Channel connections. A longwave small form-factor pluggable (SFP) transceiver can be purchased as an optional component, and must be one of the longwave SFP transceivers that are listed at the following website: www.ibm.com/support.
- Avoid using inter-switch links (ISLs) in paths between nodes and external storage systems.

Requirements for stretched systems with Fibre Channel connections

If you are configuring a stretched system that uses Fibre Channel connections, ensure that all SAN and Fibre Channel-specific requirements are met.

Notes:

- Stretched systems are not recommended for models that support nodes with internal flash drives.
- Stretched system Fibre Channel configurations with active/passive controllers such as IBM DS5000, IBM DS4000, and IBM DS3000 systems must be configured with sufficient connections such that all sites have direct access to both external storage systems. For iSCSI configurations with two or more active/passive controllers such as Storwize family systems, the systems must be configured with sufficient connections such that all sites have direct access to both external storage systems. Quorum access for stretched system is possible only through the current owner of the MDisk that is being used as the active quorum disk.

Use the following requirements to configure a stretched system with Fibre Channel connections:

- Directly connect each node to two or more SAN fabrics at the primary and secondary sites (2 - 8 fabrics are supported). In iSCSI connections, connect each node to two or more Ethernet fabrics at the primary and secondary sites. Sites are defined as independent failure domains. A failure domain is a part of the system within a boundary such that any failure (such as a power failure, fire, or flood) within that boundary is contained within the boundary and the failure does not propagate or affect parts outside of that boundary. Failure domains can be in the same room or across rooms in the data center, buildings on the same campus, or buildings in different towns. Different kinds of failure domains protect against different types of faults.
- If a storage system is used at the third site, it must support extended quorum disks. More information is available in the interoperability matrixes that are available at the following website: http://www.ibm.com/support.
- Place independent storage systems at the primary and secondary sites, and use volume mirroring to mirror the host data between storage systems at the two sites. Where possible, set the preferred node of each volume to the node in the same site as the host that the volume is mapped to.
- Connections can vary based on fibre type and small form-factor pluggable (SFP) transceiver (longwave and shortwave).
- Nodes that are in the same I/O group and separated by more than 100 meters (109 yards) must use longwave Fibre Channel or Ethernet connections. A longwave small form-factor pluggable (SFP) transceiver can be purchased as an optional component, and must be one of the longwave SFP transceivers listed at the following website: http://www.ibm.com/support.
- Avoid using inter-switch links (ISLs) in paths between nodes and external storage systems. If this is unavoidable, do not oversubscribe the ISLs because of substantial Fibre Channel traffic across the ISLs. For most configurations, trunking is required. Because ISL problems are difficult to diagnose, switch-port error statistics must be collected and regularly monitored to detect failures.
- Using a single switch at the third site can lead to the creation of a single fabric rather than two independent and redundant fabrics. A single fabric is an unsupported configuration.

- Ethernet port 1 on every node must be connected to the same subnet or subnets. Ethernet port 2 (if used) of every node must be connected to the same subnet (this might be a different subnet from port 1). The same principle applies to other Ethernet ports.
- Some service actions require physical access to all nodes in a system. If nodes in a stretched system are separated by more than 100 meters, service actions might require multiple service personnel. Contact your service representative to inquire about multiple site support.
- Use a third, dedicated site to house a quorum disk or an IP quorum application. For Fibre Channel-connected stretched system, quorum disks or IP quorum applications provide redundancy if communication is lost between the primary and secondary site. In addition both contain configuration metadata that is used to recover the system, if necessary. IP quorum applications are used for when the stretched system connects to iSCSI-attached storage systems. iSCSI storage systems cannot be configured on a third site.

A stretched system locates the active quorum disk or an IP quorum application at a third site. If communication is lost between the primary and secondary sites, the site with access to the active quorum disk continues to process transactions. If communication is lost to the active quorum disk, an alternative quorum disk at another site can become the active quorum disk.

Although a system of nodes can be configured to use up to three quorum disks, only one quorum disk can be elected to resolve a situation where the system is partitioned into two sets of nodes of equal size. The purpose of the other quorum disks is to provide redundancy if a quorum disk fails before the system is partitioned.

This figure illustrates an example stretched system configuration. When used with volume mirroring, this configuration provides a high availability solution that is tolerant of a failure at a single site. If either the primary or secondary site fails, the remaining sites can continue doing I/O operations. In this configuration, the connections between the nodes in the system are greater than 100 meters apart, and therefore must be longwave Fibre Channel connections.

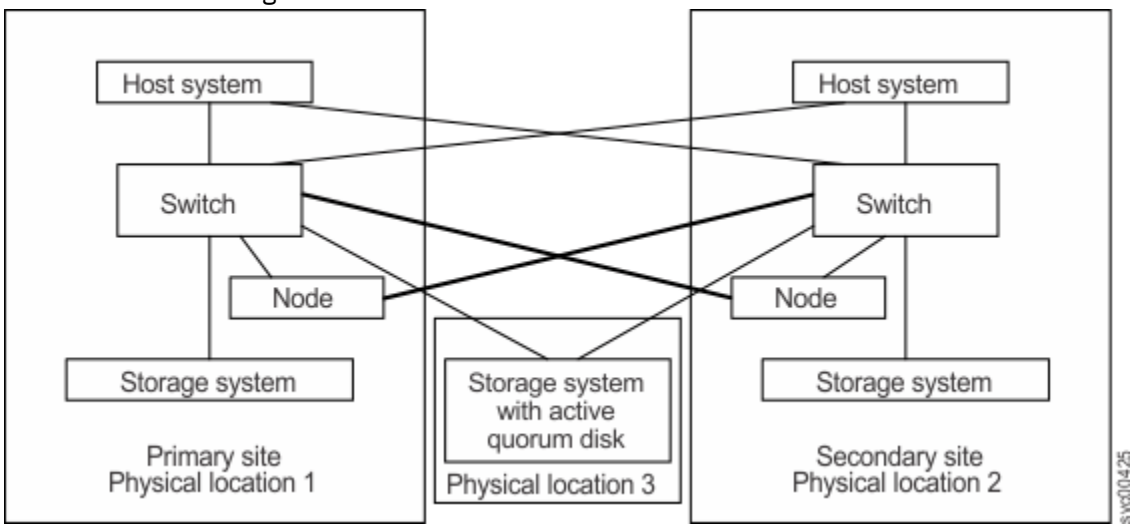


Figure 24. A stretched system with a quorum disk at a third site

In Figure 24 on page 262, the storage system that hosts the third-site quorum disk is attached directly to a switch at both the primary and secondary sites by using longwave Fibre Channel connections. If either the primary site or the secondary site fails, you must ensure that the remaining site retains direct access to the storage system that hosts the quorum disks.

Restriction: Do not connect a storage system in one site directly to a switch fabric in the other site.

An alternative configuration can use an additional Fibre Channel switch at the third site with connections from that switch to the primary site and to the secondary site.

A stretched system configuration is supported only when the storage system that hosts the quorum disks supports extended quorum. Although other types of storage systems can be used to provide quorum disks, access to these quorum disks is always through a single path.

For quorum disk configuration requirements, see the technote *Guidance for Identifying and Changing Managed Disks Assigned as Quorum Disk Candidates*.

Stretched system configuration by using interswitch links

You can use interswitch links (ISLs) in paths between nodes of the same I/O group. If the cable distance between the two production sites exceeds 100 km, potential performance impacts can result.

Using ISLs for node-to-node communication requires configuring two separate SANs, each of them with two separate redundant fabrics:

1. Configure one SAN with two separate fabrics so that it is dedicated for node-to-node communication. This SAN is referred to as a *private* SAN. This private SAN can be used by more than one stretched system if it has enough bandwidth for all these systems. For more information, see "Additional bandwidth requirements" on page 265.
2. Configure one SAN with two separate fabrics so that it is dedicated for host attachment and storage system attachment.

The network hardware that carries the ISLs must maintain the physical separation and independence of the two redundant fabrics. For example, do not connect the two fabrics into a single dark fibre link. If there are two dark fibre links, dedicate one link for each fabric. Do not cross-connect the two fabrics onto each of the two links.

Stretched system configuration rules for stretched system configurations that use ISLs

In a stretched system configuration, a site is defined as an independent failure domain. Different types of sites protect against different types of fault. For example, if configured properly, the system continues to operate after the loss of one failure domain.

However, the system does not guarantee that it can survive the failure of two sites.

- For every storage system, create one zone that contains ports from every node and all storage system ports, unless otherwise stated by the zoning guidelines for that storage system. However, do not connect a storage system in one site directly to a switch fabric in the other site. Instead, connect each storage system only to switched fabrics in the local site. (In stretched system configurations with ISLs in node-to-node paths, these fabrics belong to the public SAN).

  For stretched systems that use the enhanced configuration functions, storage systems that are configured to one of the main sites (1 or 2) need be zoned only to be visible by the nodes in that site. Storage systems in site 3 or storage systems that have no site that is defined must be zoned to all nodes.
- Each node must have direct Fibre Channel connections to at least two SAN fabrics. One is a public fabric and the other is private, in its location. For an example configuration, see Figure 25 on page 264.
- Some service actions require the ability to do actions on the front panel or through the technician port of all nodes in a system within a short-time window. If you use stretched systems, you are required to assist the support engineer and provide communication technology to coordinate these actions between the sites.
- The storage system at the third site must support extended quorum disks. This information is available in the interoperability matrices that are available at the following support website: http://www.ibm.com/support.

Each SAN consists of at least one fabric that spans both production sites. At least one fabric of the public SAN includes also the quorum site. You can use different approaches to configure private and public SANs.

- Use dedicated switches for each SAN.
- Use separate virtual fabrics or virtual SANs for each SAN.

  Note: ISLs must not be shared between private and public virtual fabrics.

To implement private and public SANs with dedicated switches, any combination of supported switches can be used. For the list of supported switches and for supported switch partitioning and virtual fabric options, see the interoperability website: www.ibm.com/support.

Like for every managed disk, all nodes need access to the quorum disk by using the same storage system ports. If a storage system with active/passive controllers (such as IBM DS3000, IBM DS4000, IBM DS5000, or IBM FAStT) is attached to a fabric, the storage system must be connected with both internal controllers to this fabric. This connection is illustrated in Figure 25 on page 264.

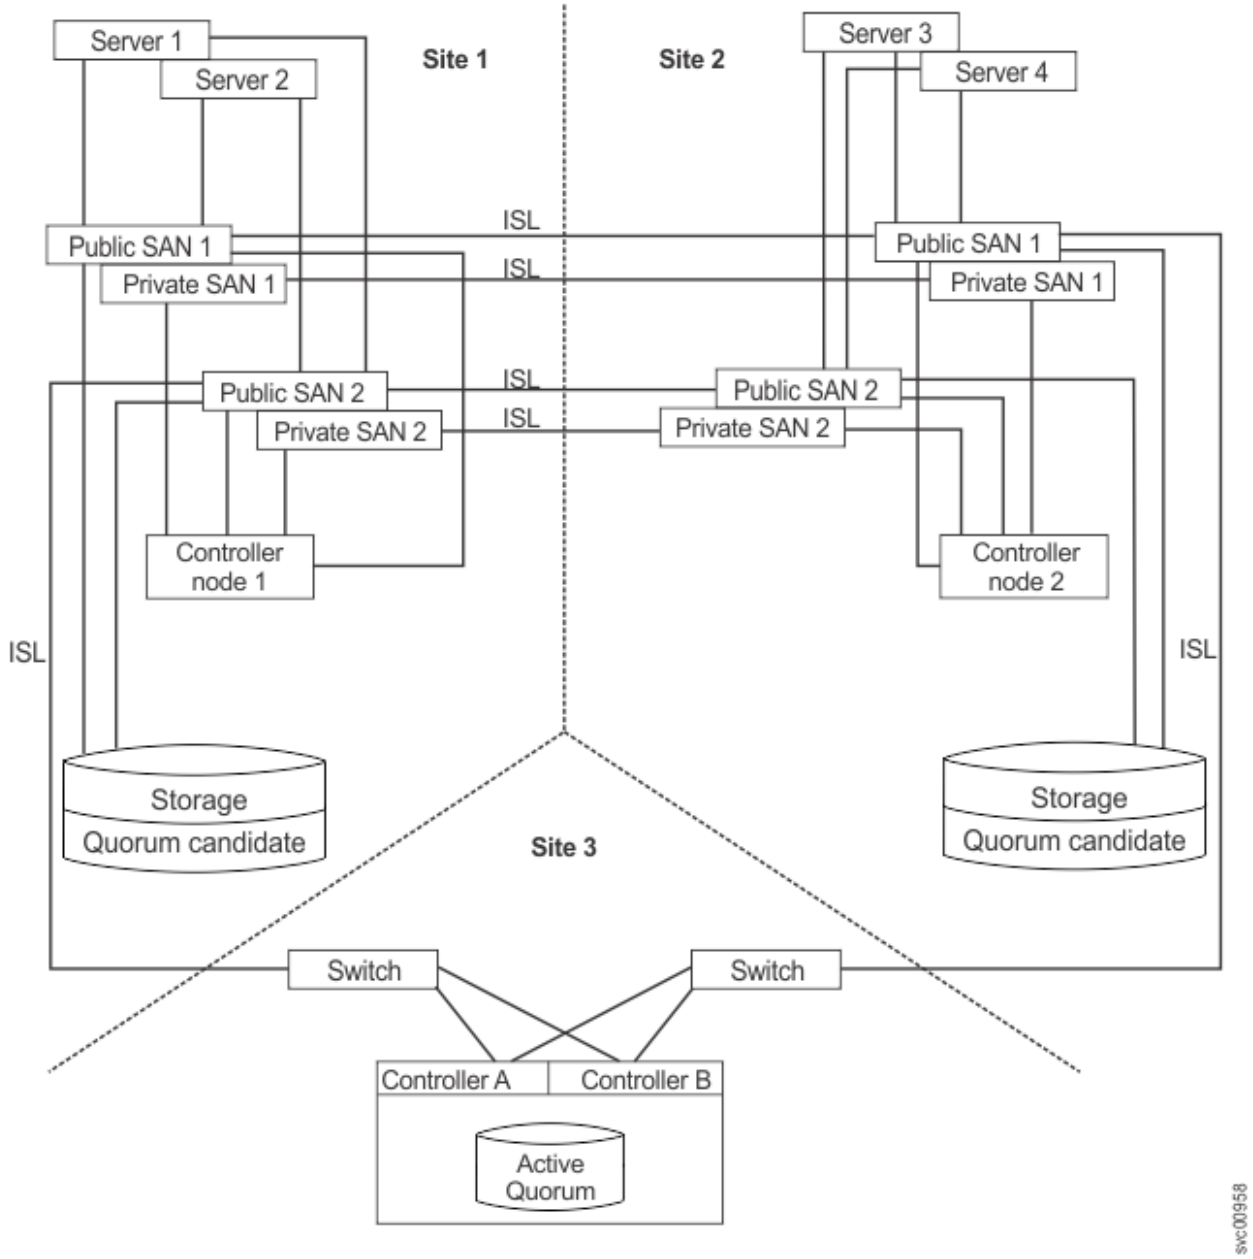


Figure 25. Stretched system nodes with ISLs, DS3000, DS4000, and DS5000 connected to both fabrics

By using FCIP, passive WDM, or active WDM for quorum site connectivity, you can add to the extension. The connections must be reliable. It is strictly required that the links from both production sites to the quorum site are independent and do not share any long-distance equipment. FCIP links are supported also for ISLs between the two production sites in public and private SANs. A private SAN and a public SAN can be routed across the same FCIP link. However, to ensure bandwidth to the private SAN (see also “Additional bandwidth requirements” on page 265), it is typically necessary to configure FCIP tunnels. Similarly, it is permissible to multiplex multiple ISL links across a DWDM link.

Note: It is not required to UPS-protect FCIP routers or active WDM devices that are used only for the node-to-quorum communication.

A stretched system configuration is supported only when the storage system that hosts the quorum disks supports extended quorum. Although the system can use other types of storage systems for providing quorum disks, access to these quorum disks is always through a single path.

For quorum disk configuration requirements, see the *Guidance for Identifying and Changing Managed Disks Assigned as Quorum Disk Candidates* technote at the following website: http://www.ibm.com/support/docview.wss?rs=591&uid=ssg1S1003311.

Additional bandwidth requirements

The connection between each site in the system that is associated with the nodes must ensure a minimum bandwidth of 4 gigabits or two times the peek host-write I/O workload, whichever is higher. The exact usage of this bandwidth depends on whether the site for the system is configured for nodes and storage systems. If the system is not configured as site-aware, half of this bandwidth is required in the public SAN and the other half is required in the private SAN. If the system is configured as site-aware, all of this bandwidth is required in the private SAN. In any case, this bandwidth must be ensured for the traffic between nodes and also between nodes and external storage systems. This bandwidth that is dedicated for system traffic does not include the traffic between nodes and hosts. Hosts that access nodes in the other site require more intersite bandwidth.

Note: To avoid performance bottlenecks, each system node requires two ports that operate at full speed (for example, 2 x 8 gigabits) worth of bandwidth to other nodes within the same system.

For example, a stretched system (no matter how many nodes) with a total of 4 gigabits of stretched system bandwidth on the private SANs is a valid configuration when the peek host-write I/O workload does not exceed 200 megabytes per second.

Operating an eight-node system at its maximum performance requires 64 gigabits of stretched system bandwidth on the private SANs.

Stretched system configuration without using interswitch links

The most simple stretched system configuration is built by attaching each system node directly to the switches in the local and the remote production site. This configuration is done without using interswitch links (ISLs).

Stretched system configuration rules for stretched system configuration without using ISLs

The following rules must be observed when you configure a stretched system without using ISLs:

- For every storage system, create one zone that contains system ports from every node and all storage system ports, unless the zoning guidelines for that storage system state differently. However, do not connect a storage system in one site directly to a switch fabric in the other site. Instead, connect each storage system only to switched fabrics in the local site. .
- Each node must have two direct connections to one or more fabrics at both locations that contain nodes.
- You can include powered components between the system and the switches in a stretched system configuration. For example, if you use Fibre Channel connections, you can use powered dense wavelength division multiplexing (DWDM) Fibre Channel extenders.
- Be prepared to provide and replace longwave SFP transceivers in case they fail.
- Some service actions require the ability to do actions on the front panel or through the technician port of all nodes in a system within a short-time window. If you use stretched systems, you are required to assist the support engineer and provide communication technology to coordinate these actions between the sites.

Configure the stretched system that does not include interswitch links (ISLs) according to the following rules:

- The minimal SAN configuration consists of one switch per production site as two separate fabrics. For highest reliability, use two switches per production site. Single fabric configurations are not supported for stretched I/O group systems.

  Note: Fibre Channel over IP (FCIP) connectivity is not supported between nodes when a stretched system is configured *without* using ISLs. FCIP (or any Fibre Channel extender or SAN router that is supported for intersystem communications) might be used in the connection to the quorum disk at the third site.

  Fibre Channel over IP connectivity is supported (or any Fibre Channel extender or SAN router that is supported for intersystem communications) both between nodes at two different sites and between nodes and the quorum disk at the third site when a stretched system is configured *with* ISLs.
- As with every system, you can use ISLs for host-to-node (with up to three hops) or for node-to-storage (at most one hop) access. If you use Fibre Channel-based connections, configure the SAN zones so that ISLs are not used in paths between nodes. For RDMA-based connections, configure separate physical Ethernet fabrics or use virtual LAN configuration so that ISLs are not used in the paths between nodes.
- In a system with four ports per node, attach half of the ports of each node to the switches in the production site where the node is located.
- In the same system with four ports per node, attach the remaining half of the ports of each node to the switches in the other (remote) production site.
- Connect each storage system at the production sites to switches in the site where the storage system is located.
- Connect the storage system with the active quorum disks to switches in both production sites.
- To avoid fabric topology changes in case of IP errors, it is a good practice to configure Fibre Channel over IP links so that they do not carry ISLs.

  Note: This guideline applies only if Fibre Channel over IP is used to connect to a switch at the third site where the quorum disk is located.

Restriction: Do not connect a storage system in one site directly to a switch fabric in the other site. Connectivity between system ports at one site to storage at the other site that uses ISLs to travel between the locations are not affected by this restriction.

It is strictly required that the links from both production sites to the quorum site are independent and do not share any long-distance equipment.

Configuring an enhanced stretched system

An enhanced stretched system assists with site disaster recovery.

Before you begin

If you are using RDMA-capable Ethernet ports to connect nodes in a stretched system, you must ensure all prerequisites are complete and define all RDMA-capable Ethernet ports before configuring the stretched system. For information, see the topics on Configuration details for using RDMA-capable Ethernet ports to connect nodes and Configuring a system with RDMA-capable Ethernet ports.

About this task

An enhanced stretched topology system can be configured through site allocation and topology change command-line (CLI) commands. The management GUI can also be used to configure a stretched topology with the Modify System Topology wizard. If you are using iSCSI-attached external storage in the stretched system configuration, you also need to add the storage to the system and specify the site for the iSCSI-attached storage system.

You can use either the management GUI or the command-line interface to configure an enhanced stretched system. The management GUI supports a wizard that completes all the steps to create an

enhanced stretched system, including site definitions for nodes, hosts, and external storage systems. The CLI process requires that each of these steps are completed separately.

Using the management GUI

1. Select **Monitoring** > **System**.
2. In the **System Actions** menu, select **Modify System Topology**. Use the wizard to complete the configuration.

Using the command line interface

To configure your clustered system site information in the command-line interface, complete these steps.

Procedure

1. Use the following commands **addnode** or **chnode** and **chcontroller** and **chhost** to assign site information to the nodes, hosts and external storage systems in site 1, 2, and 3.

   Note: Site 3 cannot be assigned to any node and contains only external storage systems that are used for quorum.

   For example, enter the following command to change site definitions for a node:

   ```
   chnode -site site_id node node_id
   ```

   where *site_id* is the identifier for the site and *node_id* is the identifier of the node.

   Enter the following command to see the node configuration details for site 1 on node 1:

   ```
   lsnode 1
   ```

   Enter the following command to assign site information to all the external storage systems by using the following command:

   ```
   chcontroller –site site_id controller_id
   ```

   where

   To view the site for the external storage system for and MDisk information, enter the following commands:

   ```
   lscontroller controller_id
   lsmdisk object_id
   ```

   Issue this command to change to site 1 on host 1:

   ```
   chost -site 1 host_name
   ```

   Issue this command to see the node configuration details for site 1 on node 1:

   ```
   lshost host_name
   ```

2. When site awareness is provided to the entities that are wanted in all the three sites, change the system topology to `stretched`.

   In the CLI, enter the following command:

   ```
   chsystem -topology stretched
   ```

3. You can use the **lssystem** command to verify that the topology is `stretched`.

   ```
    lssystem
   ```

   Note: The quorum disk assignment changes automatically after the system topology is `stretched`. The active quorum disk is at site 3. Sites 1 and 2 each have one inactive quorum disk.

Configuring an external iSCSI storage system

About this task

After you configure the stretched system and all nodes and hosts are assigned to sites, you can add any iSCSI-attached storage and assign it a specific site in the stretched system. The **Modify System Topology** wizard automatically detects any Fibre Channel-attached storage but iSCSI storage must be added manually before the site assignments can be configured. You can use either the management GUI or the command-line interface to add supported iSCSI attached storage to the stretched system.

Using the Management GUI

After you configure the stretched topology, add the iSCSI attached storage by completing the following steps:

1. If you are using iSCSI-attached external storage, you must add the storage to the system and specify the site by using the **Add iSCSI External Storage** wizard. To access the wizard, select **Pools** > **External Storage** > **Add iSCSI External Storage**.
2. Select the type of storage system that you are adding. The system supports iSCSI connections to Storwize family products and to Dell Equallogic that are used as external storage.
3. Enter the information for the source and target ports for the iSCSI connection and specify the site for the iSCSI-attached storage.

Using the command-line interface

After you change the system topology to `stretched` and nodes are assigned sites, complete these steps to add iSCSI attached storage systems to the configuration. If you are using the command-line interface to add iSCSI storage, complete these steps:

1. From the nodes of *site 1*, discover the iSCSI target storage system:

```
detectiscsistorageportcandidate –srcportid 1
 -targetip IP1 -site 1
```

2. Check the status of the discovered targets:

```
lsiscsistorageportcandidate id src_port_id target_ipv4 | target_ipv6
 target_iscsiname iogroup_list configured status
 site_id site_name 0 1 IP1 IQN1 1:1:1:1 no Full 1 site1
```

3. From the nodes of *site 1*, add a session to the discovered targets:

```
addiscsitorageport -site 1 0
```

4. From the nodes of *site 1*, view the established sessions with the target:

```
id src_port_id target_ipv4 | target_ipv6 target_iscsiname
 controller_id iogroup_list status site_id site_name
 0 3 IP1 IQN1 5 1:1:1:1 full 1 site1
```

5. The new storage system can be shown as available by entering the **`lscontroller`** command. Assign a site:

```
chcontroller -site 1 controllerID
```

6. From the system nodes of site 2, repeat steps “1” on page 268 through “5” on page 268 for discovering and establishing sessions to the same iSCSI controller or a different iSCSI controller.

Manual quorum override in a stretched system

A manual quorum override is required when you experience a rolling disaster.

In rare situations, the system is subject to what is referred to as a *rolling disaster*. A rolling disaster occurs when an incident has wide scope, and its effects are felt in multiple steps over an extended time period.

The following example scenario describes a rolling disaster and shows how to recover from that rolling disaster.

An example of a rolling disaster occurs when the following situation is true:

1. The link between the two sites fails, at which point one site uses the automatic quorum feature to continue operation.
2. The system site that has control of the quorum device fails (due to a power outage, for example).

This example leaves the second site as the only site that is potentially capable of continuing data I/O. However, it is unable to do so until it gains control of the quorum device. The MDisks in the second site stop. Nodes at the site display the node error 551, indicating that an insufficient number of nodes are available to form a quorum in a stretched system configuration.

In this scenario, you can run the quorum override command to override the automatic quorum device selection and create a new system that contains the nodes in the second site.

Note:

- To ensure that the system is in the correct state before it is used, the quorum override command can be run only with assistance from support.
- If a fabric disruption occurs while the quorum override command is running, it is possible that a subset of the nodes will update their system ID. The updated nodes display the node error 550. The nodes that were not updated display 551 and the nodes are assigned to two different systems. In this situation, you can run the quorum override command again on one of the nodes that reported the error 551. This command updates all the nodes in the two systems with a new cluster (system) ID. You can then recover data.

Enforcing conditions for a quorum

You must run the **chsystem -topology stretched** command as part of the installation process for the system to make the quorum override command available if a rolling disaster occurs. The quorum override command is not available in systems that do not have the topology set to `stretched`. Before you can use the command, the following prerequisites must be met:

- All I/O groups with two nodes are assigned with one node in site 1 and the other in site 2.
- All storage systems with MDisks must have their site that is defined.

When these prerequisites are met and automatic quorum selection is enabled, the system attempts to assign one quorum device within all three sites. If a site does not have an MDisk suitable to be a quorum device, a quorum device is not assigned to it.

Note: After the **chsystem -topology stretched** command is run, you cannot alter the site assignment of any controller except where that controller is a new controller that has only unmanaged MDisks.

It also does not allow site settings for nodes. This enforcement is required to ensure that the system operates correctly to allow the quorum override command to operate correctly.

When you run the **chsystem -topology standard** command, it again is possible to alter the site settings for nodes and controllers. However, this command disables the override quorum feature. Therefore, aim to run **chsystem -topology stretched** when you complete your changes to re-enable this support.

Converting a volume to a stretched volume

You can use the management GUI or the command-line interface to add copies to volumes to create a stretched volume.

About this task

The **addvolumecopy** command can be used to add a copy to a basic volume, which changes a basic volume into a stretched volume with copies on two separate sites. Before you add a copy to convert a basic volume to a stretched volume, the system must be configured as a stretched system topology. To change a basic volume to a stretched volume in the management GUI, select **Volumes** and right-click the basic volume and select **Add Volume Copies**. On the **Add Volume Copy** page, select the appropriate volume type and specify the pools at each site where the copies are located.

Procedure

To add a copy to basic volume to convert it to a stretched volume, enter the following command, where *site2pool* is the storage pool on the other site.

```
addvolumecopy -pool site2pool volume5
```

Changing the system topology

You can use the command-line interface or the management GUI to change the system topology.

About this task

In a standard topology, the system is configured in a single site. IBM FlashSystem supports only the standard topology. IBM SAN Volume Controller supports the stretched topology to provide high availability where nodes in an I/O group are configured in different sites with volumes configured with a mirrored copy in each site.

Note: The recommended solution for high availability is policy-based high availability.

Use the **-topology** parameter in the **chsystem** command to change the topology of the system. In the management GUI, select **Monitoring** > **System**. In the **System Actions** menu, select **Modify System Topology**.

Procedure

1. To change the system topology from a stretched topology to a standard topology, enter the following command:

   ```
   chsystem -topology standard
   ```

2. To verify that the site definitions have been removed from the enclosure, enter this command:

   ```
   lsnode
   ```

   If either the **site_id** or **site_name** parameters are listed, remove the site definitions from the node by using the **chnode** command.

Chapter 14. Legacy copy services

Starting from 8.7.1, policy-based replication and policy-based high availability replace the remote-copy function. Additionally, volume group snapshot supersedes the FlashCopy function.

See the following pages for content on the legacy copy services:

Related concepts

" Replication" on page 191
Policy-based replication can be used to implement high availability and disaster recovery solutions between systems in different locations.

Related information

"Snapshots" on page 101
A volume group snapshot is a consistent point-in-time copy of a volume group.

FlashCopy function

The FlashCopy function creates a point-in-time copy of data that is stored on a source volume to a target volume. In its basic mode, the FlashCopy function creates copies of content on a source volume to a target volume in a mapping. The function associates a source volume and a target volume in a mapping. If data exists on the target volume, that data is replaced by the copied data.

A FlashCopy mapping defines the relationship that copies data between a source volume and a target volume. A mapping can be created between any two volumes in a system. The volumes do not have to be in the same I/O group or pool. When a FlashCopy operation starts, a checkpoint is made of the source volume. No data is copied at the time a start operation occurs. Instead, the checkpoint creates a bitmap that indicates that no part of the source volume has been copied. Each bit in the bitmap represents one region of the source volume. Each region is called a *grain*. During a read operation to the target volume, the bitmap is used to determine whether the grain was copied. If the grain was copied, the data is read from the target volume. If the grain was not copied, the data is read from the source volume.

After a FlashCopy operation starts, read and write operations to the source volume continue to occur. If new data is written to the source or target volume, the existing data on the source is copied to the target volume before the new data is written to the source or target volume. The bitmap is updated to mark that the grain of the source volume has been copied so that later write operations to the same grain do not recopy the data.

After the copy operation completes, the target volumes contain the contents of the source volumes as they existed at a single point in time, unless target writes are processed. FlashCopy is sometimes described as an instance of a time-zero copy (T 0) or point-in-time copy technology. Although the copy takes some time to complete, the data on the target volume is available immediately. (If needed, data that is still being copied can be accessed from the source.)

Although it is difficult to make a consistent copy of a data set that is constantly updated, point-in-time copy techniques help solve this problem. If a copy of a data set is created that uses a technology that does not provide point-in-time techniques and the data set changes during the copy operation, the resulting copy might contain data that is not consistent. For example, if a reference to an object is copied earlier than the object itself and the object is moved before it is copied, the copy contains the referenced object at its new location, but the copied reference still points to the previous location. You can also assign background copy and cleaning rates to a FlashCopy mapping to control the rate at which updates are propagated to the remote system. FlashCopy mapping copy *rate* values can be from 128 KBps to 2 GBps and can be changed when the FlashCopy mapping is in any state.

The FlashCopy function uses copy-on-write (CoW) to complete write operations to the source volume. The system also supports redirect-on-write (RoW) which is an internal improvement to the existing FlashCopy function. The system automatically activates RoW when the following conditions are met:

- Both source and target volumes are in same data reduction pool with deduplication enabled.

© Copyright IBM Corp. 2025

- Both source and target volumes in the FlashCopy mapping are in the same I/O group.
- Both source and target volumes are not mirrored.

If any of these conditions are not met, the FlashCopy function uses CoW to complete write operations.

More advanced functions allow operations to occur on multiple source and target volumes. Management operations are coordinated to provide a common, single point-in-time for copying target volumes from their respective source volumes. This creates a consistent copy of data that spans multiple volumes. The function also supports multiple target volumes to be copied from each source volume. This can be used to create images from different points in time for each source volume.

FlashCopy mapping guidelines for storage systems

Consider the type of I/O and frequency of update before you create the volumes that you want to use in FlashCopy mappings.

FlashCopy operations complete in direct proportion to the performance of the source and target disks. If you have a fast source disk and slow target disk, the performance of the source disk is reduced because it must wait for the write operation to occur at the target before it can write to the source.

The FlashCopy implementation that is supported on the system copies at least 256 K every time a write is made to the source. This implementation means that any write involves at minimum a read of 256 K from the source, write of the same 256 K at the target, and a write of the original change at the target. Therefore, when an application completes small 4 K writes, this size is converted into 256 K.

Because of this implementation, consider the type of I/O that your application completes during a FlashCopy operation. Ensure that you do not overload the storage. The calculations contain a heavy weighting when the FlashCopy feature is active. The weighting depends on the type of I/O that is completed. Random writes can have a much higher impact than sequential writes. For example, the sequential write would copy the entire 256 K.

You can spread the FlashCopy source volumes and the FlashCopy target volumes between as many storage pools as possible. This configuration limits the potential bottle-necking of a single storage system (assuming that the storage pools contain MDisks from different storage systems). However, this configuration can still result in potential bottlenecks if you want to maintain all your target volumes on a single storage system. You must ensure that you add the appropriate weighting to your calculations.

You can expand volumes for the following reasons:

- To increase the available capacity on a particular volume that is already mapped to a host.
- To increase the size of a volume so that it matches the size of another volume so that it can be used in a FlashCopy mapping.

  Note:

  – For volumes in non-incremental FlashCopy mappings, you can increase the capacity regardless of whether the volume is the source or target volume. For volumes in incremental FlashCopy mappings, both the source and target volumes can be expanded, but the capacity of the target volume must be expanded first.

FlashCopy mapping events

FlashCopy mapping events detail the events that modify the state of a FlashCopy mapping.

Table 39 on page 273 provides a description of each FlashCopy mapping event.

Table 39. FlashCopy mapping events	
Event	**Description**
Create	A new FlashCopy mapping is created between the specified source volume and the specified target volume. The operation fails if any of the following situations is true: • The source volume is already a member of 256 FlashCopy mappings. • The node has insufficient bitmap memory. • The source and target volumes are different sizes.
Prepare	The **prepare** command is directed to either a consistency group for FlashCopy mappings that are members of a normal consistency group or to the mapping name for FlashCopy mappings that are stand-alone mappings. The **prepare** command places the FlashCopy mapping into the preparing state. **Attention:** The **prepare** command can corrupt any data that previously resided on the target volume because cached writes are discarded. Even if the FlashCopy mapping is never started, the data from the target might logically change during the act of preparing to start the FlashCopy mapping.
Flush done	The FlashCopy mapping automatically moves from the preparing state to the prepared state. This is done after all cached data for the source is flushed and all cached data for the target is no longer valid.
Start	When all the FlashCopy mappings in a consistency group are in the prepared state, the FlashCopy mappings can be started. To preserve the cross volume consistency group, the start of all of the FlashCopy mappings in the consistency group must be synchronized correctly regarding I/Os that are directed at the volumes. This alignment is achieved during the **start** command. The following series of events occurs during the **start** command: • New reads and writes to all source volumes in the consistency group are paused in the cache layer until all ongoing reads and writes below the cache layer are completed. • After all FlashCopy mappings in the consistency group are paused, the internal clustered system state is set to allow FlashCopy operations. • After the system state is set for all FlashCopy mappings in the consistency group, read and write operations are unpaused on the source volumes. • The target volumes are brought online. As part of the **start** command, read and write caching is enabled for both the source and target volumes.
Modify	The following FlashCopy mapping properties can be modified: • FlashCopy mapping name • Clean rate • Consistency group • Copy rate (for background copy or stopping copy priority) • Automatic deletion of the mapping when the background copy is complete
Stop	There are two separate mechanisms by which a FlashCopy mapping can be stopped: • You issue a command • An I/O error occurs
Delete	This command requests that the specified FlashCopy mapping is deleted. If the FlashCopy mapping is in the stopped state, the force flag must be used.

Table 39. FlashCopy mapping events (continued)	
Event	**Description**
Flush failed	If the flush of data from the cache cannot be completed, the FlashCopy mapping enters the stopped state.
Copy complete	After all of the source data is copied to the target and there are no dependent mappings, the state is set to copied. If the option to automatically delete the mapping after the background copy completes is specified, the FlashCopy mapping is automatically deleted. If this option is not specified, the FlashCopy mapping is not automatically deleted and can be reactivated by preparing and starting again.
Bitmap Online/ Offline	The node fails.

Grains and the FlashCopy bitmap

When data is copied between volumes, it is copied in units of address space that are known as *grains*.

The grain size can be either 64 KB or 256 KB. The FlashCopy bitmap contains 1 bit for each grain. The bit records whether the associated grain is split by copying the grain from the source to the target. After a FlashCopy mapping is created, the grain size for that FlashCopy mapping cannot be changed. When a FlashCopy mapping is created, if the grain size parameter is not specified and one of the volumes is already part of a FlashCopy mapping, the grain size of that mapping is used. If neither volume in the new mapping is already part of another FlashCopy mapping, and at least one of the volumes in the mapping is a compressed volume, the default grain size is 64 for performance considerations. But other than in this situation, the default grain size is 256.

Write to target volume

A write to the newest target volume must consider the state of the grain for its own mapping and the grain of the next oldest mapping.

- If the grain of the intermediate mapping or the next oldest mapping is not copied, it must be copied before the write is allowed to proceed. This copy is done to preserve the contents of the next oldest mapping. The data that is written to the next oldest mapping can come from a target or source.
- If the grain of the target that is being written is not copied, the grain is copied from the oldest already copied grain in the mappings that are newer than the target (or the source if no targets are already copied). After the copy is complete, the write can be applied to the target.

Read to target volume

If the grain that is being read is split, the read returns data from the target that is being read. If the read is to an uncopied grain on an intermediate target volume, each of the newer mappings is examined to determine whether the grain is split. The read is surfaced from the first split grain that is found. If the newer mappings do not have a split grain, the read is surfaced from the source volume.

Thin-provisioned FlashCopy

You can have a mix of thin-provisioned and standard-provisioned volumes in FlashCopy mappings. One common combination is a standard-provisioned source with a thin-provisioned target, which enables the target to consume a smaller amount of real storage than the source.

For best performance, the grain size of the thin-provisioned volume must match the grain size of the FlashCopy mapping. However, if the grain sizes are different, the mapping still proceeds.

When you create your FlashCopymappings, consider the following information:

- If you are using a standard-provisioned source with a thin-provisioned target, disable background copy and cleaning mode on the FlashCopy map by setting both the background copy rate and cleaning rate

to zero. Otherwise, if these features are enabled, all the source is copied onto the target volume. This action causes the thin-provisioned volume to either go offline or to grow as large as the source.

- If you are using only thin-provisioned source, only the space that is used on the source volume is copied to the target volume. For example, if the source volume has a virtual size of 800 GB and a real size of 100 GB of which 50 GB is used, only the used 50 GB are copied.
- A FlashCopy bitmap contains 1 bit for every grain on a volume. For example, if you have a thin-provisioned volume with 1 TB virtual size (100 MB real capacity), you must have a FlashCopy bitmap to cover the 1 TB virtual size even though only 100 MB of real capacity is allocated.

Legacy Safeguarded Copy based on FlashCopy

The Safeguarded Copy function isolates backup copies from production data, so if a cyberattack occurs, you can quickly recover data from Safeguarded copies.

The system supports Safeguarded snapshots, which use the snapshot function to create point-in-time copies of volume groups that are immutable on the system by using an internal scheduler when you do not have external scheduling applications like IBM Copy Services Manager. For more information, see “Safeguarded snapshots” on page 104.

The Safeguarded Copy function supports the following key characteristics that create cyber-resilient copies of your important data.

Separation of duties
Provides more security capabilities to prevent nonprivileged users from compromising production data. Operations related to Safeguarded backups are restricted to only a subset of users with specific roles on the system.

Administrator
Users with the Administrator role can provision and configure Safeguarded copies and related objects, such as volume groups. However, these users cannot remove or change existing Safeguarded snapshots. For auditing, it is recommended that you create a new Administrator user to configure the Safeguarded snapshots or Safeguarded Copy function. Users with this role are limited in how they can manage and interact with Safeguarded Copy operations.

Security Administrator
Users with the Security administrator role can manage users and security on the entire system and can remove and change Safeguarded backups and Safeguarded backup locations.

Superuser
Users with superuser privileges can configure all objects and complete maintenance tasks on the system. These users can remove and change both Safeguarded backups and Safeguarded policies. For more security, this account can be disabled on the system; however, it can be reenabled for remote support assistance or maintenance tasks.

Restricted Security Administrator
Users with the security administrator role is changed to restricted security administrator when two person integrity (TPI) is enabled on the IBM Storage Virtualize. TPI requires two security administrators to work together to complete critical or risky tasks. For example, a restricted security administrator with an elevated role can remove Safeguarded snapshots.

Protected Copies
Safeguarded copies cannot be mapped directly to hosts to prevent any application from changing these copies.

Automation
The system supports Safeguarded snapshots, which use the snapshot function to create point-in-time copies of volume groups that are immutable can be scheduled with an internal scheduler. The system also supports IBM Storage Copy Data Management and IBM Copy Services Manager as external scheduling applications.

The system supports the following external scheduling applications:

IBM Copy Services Manager

The system supports IBM Copy Services Manager as an external scheduling application. IBM Copy Services Manager coordinates and automates Safeguarded Copy function across multiple systems.

IBM Copy Services Manager controls Copy Services in storage environments. Copy Services are features that are used by storage systems, such as IBM Storage Virtualize system, to configure, manage, and monitor data-copy functions. If you do not have an existing IBM Copy Services Manager license, purchase the IBM Copy Manager for IBM Storage Virtualize license, which includes IBM Copy Services Manager version 6.3.4. This license option is available through iERP/AAS, Passport Advantage®, or your IBM Sales team.If you currently have an existing license for IBM Copy Services Manager, download IBM Copy Services Manager version 6.3.4 or later at Latest Downloads for IBM Copy Services Manager.

Note: For Safeguarded snapshot with IBM Copy Services Manager, use IBM Copy Services Manager version 6.3.4 or later.

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 which eliminates 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 if 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.

Safeguarded policy with IBM Copy Services Manager

To use the Safeguarded policy with IBM Copy Services Manager to create the Safeguarded backups, refer documentation version 8.5.0.x (LTS).

IBM Storage Copy Data Management

IBM Storage Copy Data Management provides management of application-consistent Safeguarded copies that use FlashCopy. IBM Storage Copy Data Management is available to order through IBM standard ordering databases. IBM Storage Copy Data Management license is available on iERP/Advanced Administration System (AAS) as part of the IBM Storage Software Suite. The stand-alone version of IBM Storage Copy Data Management is available on IBM Passport Advantage.

You can also purchase IBM Storage Sentinel, a software bundle (PID 5900APZ) license that combines copy data function and ransomware detection software. The IBM Storage Sentinel package combines the ability to create and manage Safeguarded Copy operations with IBM Storage Copy Data Management software and the ability to automatically detect any compromised data in scheduled backups with the built-in ransomware detection. The IBM Storage Copy Data Management license is available through IBM Passport Advantage.

IBM Storage Sentinel includes the following software licenses:

IBM Storage Copy Data Management version 2.2.16 or later

IBM Storage Copy Data Management provides the following actions:

- Creates Safeguarded policies to schedule backups of IBM Storage Virtualize volumes. Safeguarded backup policies are created and managed by IBM Storage Copy Data Management.

- Creates immutable backups of IBM Storage Virtualize volumes.
- Supports easy to use restore operations so compromised data can be tested and restored back to production volumes.

For the Safeguarded Copy function, Safeguarded source volumes have the following restrictions:

- Safeguarded source volumes cannot be in a Safeguarded backup location.
- Safeguarded source volumes cannot be mirrored volumes.
- Safeguarded source volumes cannot be in an ownership group.
- Safeguarded source volumes cannot be used as cloud backups with the transparent cloud tiering function.

As part of the configuration of the Safeguarded Copy function with IBM Storage Virtualize systems, administrators need access to the IBM Storage Copy Data Management interface and must complete basic installation and configuration.

Security scan software
Supports registering dedicated security scan servers to run ransomware detection software on every Safeguarded backup that is generated as part of the IBM Storage Copy Data Management configuration of the Safeguarded Copy function. To use these features for ransomware detection requires more configuration and setup that is outside the scope of IBM Storage Virtualize documentation.

The two software work together to provide detection of ransomware on Safeguarded backups.

Safeguarded backup

A *Safeguarded backup* is a volume in a Safeguarded backup location. As part of the Safeguarded Copy function, you can add volumes to a volume group and assign a Safeguarded policy to that group. The external scheduling application that is configured with the system applies the policy to all the volumes in the group to create Safeguarded backups. Safeguarded backups are created in the same parent pool as the Safeguarded source volumes. A Safeguarded backup is the target of FlashCopy mapping with Safeguarded source volumes as a source.

The *safeguarded_expiration_time* parameter in the **lsvdisk** command stores expiration time of Safeguarded backups. The *safeguarded_expiration_time* parameter notifies Storage Virtualize software to automatically delete the expired Safeguarded backups without giving the system administrator role the capability of deleting these backups directly. The *retentiondays* parameter in the **startfcconsistgrp** command store the retention period scheduled for Safeguarded backups.

To configure a Safeguarded backup volumes, the **mkvdisk** or **mkvolume** command in the system creates a new target volume in a Safeguarded backup location. The command supports only specific parameters for a volume in the Safeguarded backup location. Ensure that the target volume in the Safeguarded backup location must be created before creating the FlashCopy mapping from the Safeguarded source.

The following actions are restricted on the Safeguarded backups:

- A security administrator role only can delete the Safeguarded backups or Storage Virtualize software can automatically delete the Safeguarded backups.
- Safeguarded backup volume can be the target of FlashCopy mapping if the source volume is a Safeguarded source volume.
- Safeguarded backup volume cannot map to a host.
- Safeguarded backup volume mirror copy cannot be created.
- Safeguarded backup volume cannot be removed or changed by using CLI commands.

Safeguarded backup volume characteristics

The following section explains the characteristics of a Safeguarded backup volume in a Safeguarded backup location.

- The Safeguarded source volumes and Safeguarded backup volumes **-iogrp**, **-preferrednode** matches.
- Safeguarded backup volumes and Safeguarded source volumes size is same.
- Safeguarded backup volumes are *auto expand*.
- The Safeguarded backup volume cache mode is **readwrite**.
- In standard pools, the Safeguarded backup volumes are *thin* and *not compressed*.
- In Data Reduction Pools, the Safeguarded backup volumes *compressed* and *deduplicated*.
- A Safeguarded backup volume is of **striped** type.
- The **easytier** attribute is enabled for Safeguarded backup volumes.
- The **name**, **buffer size**, and **grainsize** parameters are user configurable.

To preserve Safeguarded backups from deleting, the `-keeptarget` parameter in `mkfcmap` command specifies that the target volume and source volume availability must be the same. If the Safeguarded backup volume becomes unavailable, then the Safeguarded source volume is also unavailable. The `-keeptarget` parameter supports FlashCopy mappings that contain Safeguarded source volume or the Safeguarded backup volumes as a source volume.

Testing and recovering Safeguarded backups

If the source volumes are compromised during a cyberattack, then the latest Safeguarded backups may not be useful for recovery. Depending on the external scheduling application that you configure with the system, each application has built-in tools and process that can assist an administrator in testing and recovering compromised data. For example, you can configure the system to use IBM Storage Copy Data Management and as the external scheduling application. IBM Storage Copy Data Management provides management of application-consistent Safeguarded copies that use FlashCopy. IBM Storage Copy Data Management is available to order through IBM standard ordering databases. IBM Storage Copy Data Management license is available on iERP/Advanced Administration System (AAS) as part of the IBM Storage Software Suite. The stand-alone version of IBM Storage Copy Data Management is available on IBM Passport Advantage. You can also purchase IBM Storage Sentinel, a software bundle (PID 5900APZ) license that combines copy data function and ransomware detection software. The IBM Storage Sentinel package combines the ability to create and manage Safeguarded Copy operations with IBM Storage Copy Data Management software and the ability to automatically detect any compromised data in scheduled backups with the built-in ransomware detection. The IBM Storage Copy Data Management license is available through IBM Passport Advantage.

IBM Storage Sentinel includes the following software licenses:

IBM Storage Copy Data Management version 2.2.16 or later

IBM Storage Copy Data Management provides the following actions:

- Creates Safeguarded policies to schedule backups of IBM Storage Virtualize volumes. Safeguarded backup policies are created and managed by IBM Storage Copy Data Management.
- Creates immutable backups of IBM Storage Virtualize volumes.
- Supports easy to use restore operations so compromised data can be tested and restored back to production volumes.

For the Safeguarded Copy function, Safeguarded source volumes have the following restrictions:

- Safeguarded source volumes cannot be in a Safeguarded backup location.
- Safeguarded source volumes cannot be mirrored volumes.
- Safeguarded source volumes cannot be in an ownership group.
- Safeguarded source volumes cannot be used as cloud backups with the transparent cloud tiering function.

As part of the configuration of the Safeguarded Copy function with IBM Storage Virtualize systems, administrators need access to the IBM Storage Copy Data Management interface and must complete basic installation and configuration.

Security scan software

Supports registering dedicated security scan servers to run ransomware detection software on every Safeguarded backup that is generated as part of the IBM Storage Copy Data Management configuration of the Safeguarded Copy function. To use these features for ransomware detection requires more configuration and setup that is outside the scope of IBM Storage Virtualize documentation.

The two software work together to provide detection of ransomware on Safeguarded backups.

If you use IBM Copy Services Manager as an external scheduling application with the system, you can determine the uncompromised Safeguarded backups, IBM Copy Services Manager creates a new target volume in the same parent pool by creating another FlashCopy mapping between the Safeguarded backups as source volumes and the new target volume. After the FlashCopy mapping is created, the new target volume is mapped to a host server or a virtual machine. The user can now test the Safeguarded backups by using the external application that is installed on the server system. The new copy volumes that are created for testing and recovery are created according to the provisioning policy associated with the parent pool. IBM Copy Services Manager uses the pool provisioning policy to create recovery volumes. A user can set recovery copies, such as thin provisioned or deduplication by defining a provisioning policy to their parent pool. For more information, see “Provisioning policies” on page 185.

To recover and restore the data, the user would first test the Safeguarded backups. After testing and determining the uncompromised data, user can use the new FlashCopy target volumes as the source volumes. The new source volumes would be added to a volume group with associated Safeguarded backup policy and the original volume group Safeguarded policy would be removed.

The following image shows the testing and recovering a Safeguarded backups scenario:

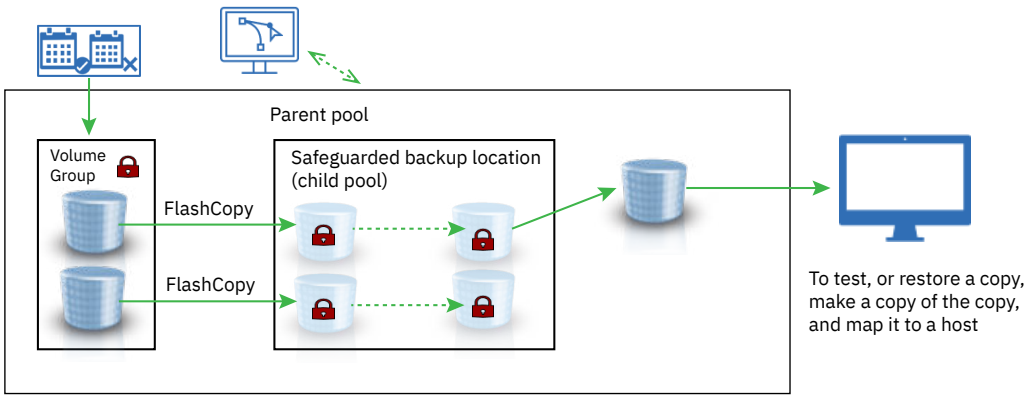


Figure 26. Recovering Safeguarded backups

Suspending Safeguarded Copy operation

You can suspend a Safeguarded copy operation to stop creating 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 Copy operations until the compromised data can be determined. During an attack, Safeguarded backups can contain malicious data or software. By suspending Safeguarded Copy functions, you can examine the existing uncompromised Safeguarded backups. The uncompromised Safeguarded backup can be used to restore or recover the backup of the data using your configured external scheduling application. Safeguarded Copy operation is suspended at the system level. When the Safeguarded Copy function is suspended, your configured external scheduling application does not create new backups in the Safeguarded backup location and prevents the expired backups from being removed. The **safeguarded_copy_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 backups being created which results in a lack of new backups that can be used for future recovery. Before suspending Safeguarded Copy function, contact your support representative to assistance. To resume backups, the Safeguarded Copy function must be turned back on.

Configuring Safeguarded Copy function

You can configure Safeguarded Copy function with the management GUI or the command-line interface (CLI).

The system supports Safeguarded snapshots, which use the snapshot function to create point-in-time copies of volume groups that are immutable on the system by using an internal scheduler when you do not have external scheduling applications like IBM Copy Services Manager.

For more information, see "Snapshots" on page 101.

Configuring Safeguarded snapshots with IBM Copy Services Manager

You can use IBM Copy Services Manager to manage and create Safeguarded snapshots on the system.

Prerequisites

Before you can configure the Safeguarded Copy function on your system, ensure you meet the following prerequisites:

IBM Copy Services Manager Requirements

Ensure that the following requirements are met for IBM Copy Services Manager:

1. If you do not have an existing IBM Copy Services Manager license, purchase the IBM Copy Manager for IBM Storage Virtualize license, which includes IBM Copy Services Manager version 6.3.4. This license option is available through iERP/AAS, Passport Advantage, or your IBM Sales team.
2. If you currently have an existing license for IBM Copy Services Manager, download IBM Copy Services Manager version 6.3.4 or later at Latest Downloads for IBM Copy Services Manager.

   Note: If you are using an existing license, ensure that the licensed capacity is adequate for use of the Safeguarded Copy function. If you need more capacity for Safeguarded Copy function, contact you IBM sales representative to update your licensed capacity for IBM Copy Services Manager.
3. After you download IBM Copy Services Manager, complete the instructions for your installation. IBM Copy Services Manager supports several installation options on different environments. For more information, see Installing IBM Copy Services Manager.
4. During installation, license files can be imported for IBM Copy Services Manager. If the license was not imported during the installation, you need to apply the license to the installation. For more information, see Applying license files after installation.

System Requirements

All systems must be running the 8.5.2 or later release.

For existing systems, ensure that you have completed capacity planning for Safeguarded Copy function.

Creating an Administrator user for IBM Copy Services Manager

Before you can establish the system as a connection endpoint in IBM Copy Services Manager, you need to configure a user with the Administrator role on the IBM Storage Virtualize system. For auditing, it is recommended that you create a new Administrator user to configure the Safeguarded Copy function. Users with this role are limited in how they can manage and interact with Safeguarded Copy operations. The IBM Copy Services Manager uses this role to create Safeguarded snapshots.

Using the management GUI

To create new Administrator user for IBM Copy Services Manager, complete these steps:

1. In the management GUI, select **Access** > **Users by Groups** > **Create User Group**.

2. On the **Create User Group** page, enter a name of the user group, and select **Administrator** for the role.
3. Click **Create**.
4. In the list of user groups, select the user group that you created and select **Create Users**.
5. On the **Create Users** page, enter the name of the user, and select **Local**.
6. Enter and confirm a password.
7. Click **Create**.

After you create the Administrator user, create a connection to the system in IBM Copy Services Manager.

Using the command-line interface (CLI)

To create new Administrator user for IBM Copy Services Manager, complete these steps:

1. To configure a user with an Administrator role, you need to create a user group. Enter the following command.

   ```
   mkusergrp -name group_name -role Administrator
   ```

   where *group_name* specifies the unique user group name and *Administrator* specifies the role for all users that belong to this user group.
2. After creating user group, create a local user to access a system. Enter the following command.

   ```
   mkuser -name user_name -usergrp group_name | group_id -password
   ```

   where *user_name* specifies unique user name. The *group_name* | *group_id* specifies the ID or name of the user group that was created in the previous step. The parameter **-password** specifies the password to be associated with the user.

Creating volume groups and assigning source volumes

A volume group allows for many related volumes to be grouped so that they can be managed together. A volume group also ensures consistency across all volumes in the group when used with snapshot or replication functions.

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.

Using the management GUI

To create a volume group, complete these steps:

1. In the management GUI, select **Volumes** > **Volumes Groups**.
2. Click **Create Volume Group**.
3. On the **Create Volume Group** page, enter the volume group name. If you do not specify a volume group name, the system automatically assigns the name.
4. To assign volumes, select the following options:

   Choose existing volumes
   Select this option if you want to add existing volumes to a volume group by completing these steps:

   a. Select **Choose existing volumes** and click **Next**.
   b. From the list, select the volumes that you want to add to the volume group.
   c. Click **Create Volume Group**.

Using the command-line interface (CLI)

To create and configure a new volume group and assign volumes to that group, complete these steps:

1. To create the new volume group, enter the following command:

   ```
   mkvolumegroup -name volumegroup_name
   ```

   where *volumegroup_name* specifies a volume group name. This command creates the volume group.
2. After creating the volume group, either create new volumes or change existing volumes and assign them to volume group that you created in Step 1.

   If you are assigning a new volume to the volume group, enter the following command:

   ```
   mkvolume -volumegroup <volumegroup_name_or_id> -size <disk_size>
   ```

   where *<volumegroup_name_or_id>* is the name or identifier of the volume group; and *<disk_size>* indicates the capacity that is provisioned for the volume from the parent pool.

   If you are assigning existing volumes to the volume group, enter the following command:

   ```
   chvdisk -volumegroup <volumegroup_name_or_id> <name_or_id>
   ```

   where *<volumegroup_name_or_id>* is the name or identifier of the volume group and *<name_or_id>* of the volume.

Configuring Safeguarded Copy function with IBM Storage Copy Data Management

You can configure Safeguarded Copy function with IBM Storage Copy Data Management.

Prerequisites

Before you can configure the Safeguarded Copy function on your system with IBM Storage Copy Data Management, ensure you meet the following prerequisites:

IBM Storage Copy Data Management Requirements

Purchase a license for IBM Storage Copy Data Management or IBM Storage Sentinel. IBM Storage Copy Data Management provides management of application-consistent Safeguarded copies that use FlashCopy. IBM Storage Copy Data Management is available to order through IBM standard ordering databases. IBM Storage Copy Data Management license is available on iERP/Advanced Administration System (AAS) as part of the IBM Storage Software Suite. The stand-alone version of IBM Storage Copy Data Management is available on IBM Passport Advantage. You can also purchase IBM Storage Sentinel, a software bundle (PID 5900APZ) license that combines copy data function and ransomware detection software. The IBM Storage Sentinel package combines the ability to create and manage Safeguarded Copy operations with IBM Storage Copy Data Management software and the ability to automatically detect any compromised data in scheduled backups with the built-in ransomware detection. The IBM Storage Copy Data Management license is available through IBM Passport Advantage.

IBM Storage Sentinel includes the following software licenses:

IBM Storage Copy Data Management version 2.2.16 or later

IBM Storage Copy Data Management provides the following actions:

- Creates Safeguarded policies to schedule backups of IBM Storage Virtualize volumes. Safeguarded backup policies are created and managed by IBM Storage Copy Data Management.
- Creates immutable backups of IBM Storage Virtualize volumes.
- Supports easy to use restore operations so compromised data can be tested and restored back to production volumes.

For the Safeguarded Copy function, Safeguarded source volumes have the following restrictions:

- Safeguarded source volumes cannot be in a Safeguarded backup location.
- Safeguarded source volumes cannot be mirrored volumes.
- Safeguarded source volumes cannot be in an ownership group.
- Safeguarded source volumes cannot be used as cloud backups with the transparent cloud tiering function.

As part of the configuration of the Safeguarded Copy function with IBM Storage Virtualize systems, administrators need access to the IBM Storage Copy Data Management interface and must complete basic installation and configuration.

Security scan software
Supports registering dedicated security scan servers to run ransomware detection software on every Safeguarded backup that is generated as part of the IBM Storage Copy Data Management configuration of the Safeguarded Copy function. To use these features for ransomware detection requires more configuration and setup that is outside the scope of IBM Storage Virtualize documentation.

The two software work together to provide detection of ransomware on Safeguarded backups.

IBM Storage Virtualize system requirements
All systems must be running the 8.5.1 or later release.

For existing systems, ensure that you complete capacity planning for Safeguarded Copy function. If the Safeguarded Copy function is configured with IBM Copy Services Manager and you want to use IBM Storage Copy Data Management, you unassign all Safeguarded policies that are managed by IBM Copy Services Manager and migrate to IBM Storage Copy Data Management as your scheduling application.

Creating an Administrator user for IBM Storage Copy Data Management

Before you can establish the system as a connection endpoint in IBM Storage Copy Data Management, you need to configure a user with the Administrator role on the IBM Storage Virtualize system. For auditing, it is recommended that you create a new Administrator user to configure the Safeguarded Copy function. Users with this role are limited in how they can manage and interact with Safeguarded Copy operations. The IBM Storage Copy Data Management uses this role to create FlashCopy mappings between the source volumes and the Safeguarded backups on the system.

To create new Administrator user for IBM Storage Copy Data Management, complete these steps:

1. In the management GUI, select **Access** > **Users by Groups** > **Create User Group**.
2. On the **Create User Group** page, enter a name of the user group, and select **Administrator** for the role.
3. Click **Create**.
4. In the list of user groups, select the user group that you created and select **Create Users**.
5. On the **Create Users** page, enter the name of the user, and select **Local**.
6. Enter and confirm a password.
7. Click **Create**.

Using the command-line interface (CLI)

To create new Administrator user for IBM Storage Copy Data Management , complete these steps:

1. To configure a user with an Administrator role, you need to create a user group. Enter the following command.

```
mkusergrp -name group_name -role Administrator
```

where *group_name* specifies the unique user group name and *Administrator* specifies the role for all users that belong to this user group.

2. After creating user group, create a local user to access a system. Enter the following command.

```
mkuser -name user_name -usergrp group_name | group_id -password
```

where *user_name* specifies unique username. The *group_name* | *group_id* specifies the ID or name of the user group that was created in the previous step. The parameter **-password** specifies the password to be associated with the user.

Registering the IBM Storage Virtualize system in IBM Storage Copy Data Management

After you create an Administrator user for IBM Storage Copy Data Management, you must configure and register the system in IBM Storage Copy Data Management interface.

To configure the system in IBM Storage Copy Data Management, complete these steps:

Note: The following instructions are based on another product's interface and can be subject to change. For the most recent version of these steps, see IBM Storage Copy Data Management documentation.

1. In a supported browser, enter the following URL:

   `https://<HOSTNAME>:8443/portal/`

   Where <HOSTNAME> is the IP address of the virtual machine where the IBM Storage Copy Data Management is configured. You now connected to IBM Storage Copy Data Management management GUI.
2. Enter the administrator username and password.
3. After you log in to IBM Storage Copy Data Management GUI, you need to do the following tasks:

   Add the IBM Storage Virtualize system as a site
   Create sites to define a grouping of providers based on their location.

   Registering an Application Server - InterSystems Database
   Establishes a secure communication channel between SCDM and application server.

   Registering an IBM Storage Virtualize provider
   Register the IBM Storage Virtualize system as a storage provider for IBM Storage Copy Data Management.

   Note: As part of registering, you need to add the administrator username that you created in

Specifying IBM Storage Copy Data Management as the external scheduling application

After you configure the system on IBM Storage Copy Data Management, specify IBM Storage Copy Data Management as the external scheduling application.

The management GUI supports two alternatives to specify the external scheduling application. You can set this value for the entire system on the **External Scheduling Applications** page or set the external scheduling application on the **Volume Groups** page.

To specify IBM Storage Copy Data Management as your external scheduling application, complete the following steps:

1. In the management GUI, select **Settings** > **GUI preferences**, and ensure **Display Safeguarded backup policy tile and External Scheduler Application Settings:** is **Visible**.
2. Select **Settings** > **System** > **External Scheduling Applications**.
3. On the **External Scheduling Applications** page, select IBM Storage Copy Data Management.
4. Click **Save**.

To specify IBM Storage Copy Data Management as the external scheduling application on the **Volume Groups** page, complete the following steps:

1. In the management GUI, select **Volumes** > **Volume Groups**.
2. On the **Volume Groups** page, select a volume group and click **Assign Safeguarded Backup Policy**.
3. On the **Assign Safeguarded policy** page, select IBM Storage Copy Data Management.
4. Click **Next** to save changes.

If you are using IBM Copy Services Manager as external scheduling application and you need to migrate to IBM Storage Copy Data Management, see .

Creating Safeguarded backup locations

You can use the GUI and command- line interface for creating Safeguarded backup location.

Using the management GUI

To create a Safeguarded backup location, complete these steps:

1. If the parent pool is a standard pool, enter the amount of capacity that is dedicated to the child pool. If the parent pool is a data reduction pool, the child pool shares capacity with the parent pool.
2. Select **Safeguard** to indicate that the child pool is used as the Safeguarded backup location for immutable backup copies of source volumes.

   Note: Before a Safeguarded policy can be applied, create one Safeguarded backup location for each parent pool. After you create the corresponding Safeguarded backup locations for each parent pool, create a volume group for the all the affected source volumes.
3. Select **Encrypted** to encrypt the child pool. If the parent pool is encrypted, the child pool adopts encryption from the parent.

After you create the Safeguarded backup locations, create volume groups.

Using the command-line interface (CLI)

To creates a new child pool to be used as the Safeguarded backup location for parent pool, enter the following command.

```
mkmdiskgrp -size mdiskgrp_size -parentmdiskgrp mdiskgrp_id | mdiskgrp_name -name pool_name
-safeguarded
```

where *mdiskgrp_size* specifies the child pool capacity. The *mdiskgrp_id* | *mdiskgrp_name* specifies the parent pool from which the volume extents of the child pool are allocated when you create a child pool. The **-safeguarded** parameter specifies whether the child pool is used as a Safeguarded backup location.

Creating volume groups and assigning source volumes

Create a volume group and add source volumes.

Using the management GUI

To create a volume group, complete these steps:

1. In the management GUI, select **Volumes** > **Volumes Groups**.
2. Click **Create Volume Group**.
3. On the **Create Volume Group** page, enter the volume group name. If you do not specify a volume group name, the system automatically assigns the name.
4. To assign volumes, select the following options:

   Choose existing volumes
   Select this option if you want to add existing volumes to a volume group by completing these steps:

   a. Select **Choose existing volumes** and click **Next**.
   b. From the list, select the volumes that you want to add to the volume group.
   c. Click **Create Volume Group**.

Using the command-line interface (CLI)

To create and configure a new volume group and assign volumes to that group, complete these steps:

1. To create the new volume group, enter the following command:

   ```
   mkvolumegroup -name volumegroup_name -external
   ```

   where *volumegroup_name* specifies a volume group name. The *external* defines that the policy is scheduled by an external scheduling application. This command creates the volume group and assigns the policy to the volume group.
2. After creating the volume group, either create new volumes or change existing volumes and assign them to volume group that you created in Step 1.

   If you are assigning a new volume to the volume group, enter the following command:

   ```
   mkvolume -pool <pool_name_or_id> -volumegroup <volumegroup_name_or_id> -size <disk_size>
   ```

   where *<pool_name_or_id>* is the name or identifier of the parent pool that contains the Safeguarded backup location; *<volumegroup_name_or_id>* is the name or identifier of the volume group; and *<disk_size>* indicates the capacity that is provisioned for the volume from the parent pool.

   If you are assigning existing volumes to the volume group, enter the following command:

   ```
   chvdisk -volumegroup <volumegroup_name_or_id> <name_or_id>
   ```

   where *<volumegroup_name_or_id>* is the name or identifier of the volume group and *<name_or_id>* of the volume.

Assigning a Safeguarded policy with IBM Storage Copy Data Management

You can assign one external Safeguarded policy to all the volume groups that are controlled by the IBM Storage Copy Data Management interface.

The system supports one external policy. The external policy is not a predefined policy. An external safeguarded policy cannot be converted into a safeguarded policy after it is created. Only a new created policy can be defined as external at time of its creation.After you create a volume group and assign source volumes in the management GUI, you need to create and assign the Safeguarded policy with IBM Storage Copy Data Management. To create and assign a Safeguarded policy to configured volume groups, access IBM Storage Copy Data Management interface, which is a part of IBM Storage Copy Data Management software bundle.

To create and assign a Safeguarded policy to configured volume groups, complete the following tasks in IBM Storage Copy Data Management:

Note: The following instructions are based on another product's interface and can be subject to change. For the most recent version of these steps, see IBM Storage Copy Data Management documentation.

1. In a supported browser, enter the following URL:

   `https://<HOSTNAME>:8443/portal/`

   Where <HOSTNAME> is the IP address of the virtual machine where the IBM Storage Copy Data Management is configured. This connects you to IBM Storage Copy Data Management management GUI.
2. Enter the administrator username and password.
3. After logging in to IBM Storage Copy Data Management GUI, you need to do the following task:
4. **Creating a Safeguarded Copy SLA policy**
   Create a Safeguarded backup policy in IBM Storage Copy Data Management and assign the policy to detected volume groups on IBM Storage Virtualize system.

Using the command-line interface (CLI)

To assign an external safeguard policy to volume groups, enter the following command:

```
mksafeguardedpolicy -name sgpolicy0 -external
```

Where the *external* parameter defines that the schedule is set by an external software.

Managing the 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 Storage Copy Data Management. For Safeguarded snapshots with internal scheduler, refer to “Snapshots” on page 101.

After you configure Safeguarded Copy function on your system, ensure to regularly test the configuration to ensure that Safeguarded backups are ready in the event of a cyberattack. In addition to testing and recovering, you can also manage objects that are related to the Safeguarded Copy function on the system. These tasks include adding source volumes to Safeguarded volume groups, managing Safeguarded backups after expiration, and other actions. Some administrative actions are completed on the system, while others are completed on the management interface for the external scheduling application, like IBM Storage Copy Data Management or IBM Copy Services Manager.

IBM Copy Services Manager

The system supports IBM Copy Services Manager as an external scheduling application. IBM Copy Services Manager coordinates and automates Safeguarded Copy function across multiple systems. You can manage Safeguarded Copy function with IBM Copy Services Manager using the management GUI and command- line interface.

Adding new source volumes to a Safeguarded volume group

You can add new source volumes to an existing Safeguarded volume group. New volumes use the schedule in the external scheduling software or by using the internal snapshot scheduler.

For releases earlier than 8.5.2, the Safeguarded policy contains copy schedule information for the systems using FlashCopy technology with IBM Copy Services Manager as an external scheduler. For more information, you can refer to the “Configuring Safeguarded snapshots with IBM Copy Services Manager” on page 280.

Using the management GUI

To add source volumes to a volume group, complete these steps:

1. In the management GUI, select **Volumes** > **Volume Groups**.
2. Select the volume group that you want to add volumes to and select **Group Actions** > **Add Volumes**.
3. On the **Add Volume** page, select the volume and select **Add Volume**.

Using the command-line interface (CLI)

To add a volume to a volume group, enter the following command:

```
chvdisk -volumegroup volumegroup_name_or_id volume_name_or_id
```

where *volumegroup_name_or_id* indicates the name or ID of the volume group and *volume_name_or_id* is the name or ID of the source volume that are being added to the volume group.

Recovering a Safeguarded backup

After you configure Safeguarded Copy function on your system, ensure to regularly test the configuration to ensure that Safeguarded backups can be recovered and restored if necessary.

If a cyberattack occurs, a Safeguarded source volume can be compromised for an indefinite amount of time until the breach is detected. In this situation, the most recent Safeguarded backups are not useful for restoring data on the production volume. Effective testing includes identifying versions of Safeguarded backups that can be used to restore the compromised data. Follow the guidelines and direction of your business continuity plan and your recovery point objectives to determine the frequency of testing your configuration.

For recovering the safeguarded backup, refer to the recovery commands for IBM Copy Services Manager.

Migrating to Safeguarded snapshots

If you are currently using IBM Copy Services Manager as your external scheduling application, you can migrate to Safeguarded snapshots. Previous versions of Safeguarded Copy function with IBM Copy Services Manager used FlashCopy mappings to create backups by using an external Safeguarded backup policy. However, IBM Copy Services Manager support Safeguarded snapshots, which reduces the configuration and management complexity of previous versions.

You must remove the Safeguarded policy assigned to the volumes groups in IBM Storage Virtualize for migrating to Safeguard snapshots.

Before you can migrate to Safeguarded snapshots, complete the following prerequisites:

- Update IBM Storage Virtualize to version 8.5.2 or later.
- Update IBM Copy Services Manager to version 6.3.4 or later.

Using management GUI

Configure IBM Copy Services Manager system

To configure IBM Copy Services Manager system, complete these steps:

1. Create an IBM Copy Services Manager snapshot session, and associate it to the volume group in the IBM Copy Services Manager system. For more information, see Creating a Storage Virtualize Snapshot session using a Volume Group.
2. Set up scheduled task in the IBM Copy Services Manager system that runs **`Create Safeguarded Snapshot`** command on the session that is created in step 1. For more information, see Creating scheduled tasks.
3. Validate the creation of Safeguarded snapshot in IBM Copy Services Manager. You can manage and validate Safeguarded snapshots with IBM Copy Services Manager by creating clone or thin-clone. For more information, see Snapshot for Storage Virtualize commands.
4. Repeat these steps for all the Safeguarded backup volume groups as IBM Copy Services Manager sessions have one to one mapping with IBM Storage Virtualize volume groups.

Remove Safeguarded policy in IBM Storage Virtualize

To remove Safeguarded policy in IBM Storage Virtualize, complete these steps:

1. In the management GUI, select **Volumes** > **Volume Groups**.
2. On the **Volume Groups** page, select a volume group and from the overflow menu, select the **View Details**.
3. Select the **Policies** tab, and click **Manage Safeguarded Backup Policy**.
4. On the **Manage Safeguarded Policy** page, select **Remove policy from the volume group**.
5. Click **Apply changes**.
6. Repeat these steps for all the Safeguarded backup volume groups.

7. Verify that the existing Safeguarded backup policies are removed. On the **Volumes Group** page, select a volume group and select the **Policies** tab. The Safeguarded backup policy is not assigned message is displayed. Repeat this step for all the volume groups.

Migrating Safeguarded Copy with FlashCopy to Safeguarded snapshots with internal scheduler

If you are currently using any external scheduling application, you can migrate to Safeguarded snapshot with internal scheduler.

To migrate from Safeguarded Copy policy to Safeguarded snapshot, complete these steps:

Remove Safeguarded policy in IBM Storage Virtualize

To remove Safeguarded policy in IBM Storage Virtualize, complete these steps:

1. In the management GUI, select **Volumes** > **Volume Groups**.
2. On the **Volume Groups** page, select a volume group and from the overflow menu, select the **View Details**.
3. Select the **Policies** tab, and click **Manage Safeguarded Backup Policy**.
4. On the **Manage Safeguarded Policy** page, select **Remove policy from the volume group**.
5. Click **Apply changes**.
6. Repeat these steps for all the Safeguarded backup volume groups.
7. Verify that the existing Safeguarded backup policies are removed. On the **Volumes Group** page, select a volume group and select the **Policies** tab. The Safeguarded backup policy is not assigned message is displayed. Repeat this step for all the volume groups.

Disable external scheduling application in IBM Storage Virtualize

To disable the external scheduling application, complete these steps:

1. In the management GUI, select **Settings** > **GUI preferences** > **GUI Features**, and ensure **Display Safeguarded backup policy tile and External Scheduler Application Settings:** is **Hidden**.

Assign snapshot policy in IBM Storage Virtualize

To assign a snapshot policy to a volume group by selecting **Volume groups**, complete these steps:

1. In the management GUI, select **Volumes** > **Volume groups**.
2. Select the volume group from the list of volume groups. In the overflow menu, select **View Details**.
3. On the **Policies** tab, click **Assign Internal Snapshot Policy**.
4. From the **Assign Internal Snapshot Policy** page, select the snapshot policy from the list of snapshot policies that you want to assign to a volume group. If you select a volume group that is already assigned a snapshot policy, it replaces that policy with this new policy that you selected.
5. If any volumes in the selected volume groups are source volumes in a FlashCopy mapping, the **Ignore existing FlashCopy mappings** is displayed. If you are creating new snapshots or a Safeguarded snapshot, select this option. If this option is not selected, snapshots cannot be created. After you select this option and the snapshot policy is assigned successfully to a volume group, this option no longer displays.

   If volumes in a volume group are target volumes, you cannot create snapshots and **Ignore existing FlashCopy mappings** is not visible. If volumes in a volume group are both source and target volumes, **Ignore existing FlashCopy mappings** is visible, however you cannot create snapshot.
6. Select **Date** and **Choose a time (optional)** for the policy to begin. Based on selected frequency of snapshots, the start time defaults to whenever the policy is assigned to a volume group.
7. Select **Safeguarded** to safeguard the snapshots. If **Safeguarded** is selected, the snapshot policy creates Safeguarded snapshots for all the volumes in the group based on the schedule.
8. Click **Assign Policy** to assign the policy.

Note: If a snapshot policy is assigned to a volume group, the policy's schedule may result in more volume group snapshots being created than can be accommodated within the system's maximum limits. The frequency of snapshots in the snapshot policy should be adjusted to ensure that the limits are not exceeded.

Migrating to IBM Storage Copy Data Management as external scheduling application

If you are currently using IBM Copy Services Manager as your external scheduling application, you can migrate to IBM Storage Copy Data Management.

Using the management GUI

Migration to IBM Storage Copy Data Management from IBM Copy Services Manager requires that all active Safeguarded backup policies are removed from volume groups that use IBM Copy Services Manager as an external scheduling application. After backup policies are removed, remove the URL for the IBM Copy Services Manager interface from the management GUI. Currently, the system only supports one scheduling provider.

To migrate from IBM Copy Services Manager to IBM Storage Copy Data Management, complete these steps:

1. In the management GUI, select **Volumes** > **Volume Groups**.
2. On the **Volume Groups** page, select a volume group, select the **Policies** tab, and click **Manage Safeguarded Backup Policy**.
3. On the **Manage Safeguarded Policy** page, select **Remove policy from the volume group** check box.
4. Click **Apply changes**.
5. Repeat these steps for all the Safeguarded backup volume groups.
6. To verify that the existing Safeguarded backup policies are removed, on the **Volume Groups** page, select a volume group and select the **Policies** tab. Ensure that you see this message: **Safeguarded backup policy is not assigned**. Repeat this step for all the volume groups.
7. In the management GUI, select **Settings** > **GUI preferences**, and ensure **Display Safeguarded backup policy tile and External Scheduler Application Settings:** is **Visible**.
8. Select **Settings** > **System** > **External Scheduling Applications**.
9. On the **External Scheduling Applications** page, select **Remove URL**.
10. Click **Remove** on the verification panel.

Using the command-line interface

Alternatively, you can use following command-line interface steps to migrate to IBM Storage Copy Data Management:

1. Enter the following command:

```
rmsafeguardedpolicy -removefromvolumegroups <safeguarded_policy_id | safeguarded_policy_name>
```

   where *<safeguarded_policy_id | safeguarded_policy_name>* is the ID or name of the Safeguarded policy.
2. Enter the following command:

```
mksafeguardedpolicy -external
```

After removing the URL for IBM Copy Services Manager you can configure IBM Storage Copy Data Management as your external scheduling application.

FlashCopy statistics collection

The system collects FlashCopy specific metrics, which allows support to analyze performance issue.

Introduction

FlashCopy statistics measures the I/O amplification and magnification in FlashCopy layer, and it also measures the latency of the FlashCopy I/O operations. The various cache statistics gives information regarding the various FlashCopy I/O stages, though it is not obvious where the amplification or magnification is originating. Therefore, FlashCopy will track and measure the internal operations that are required to process an I/O in FlashCopy.

Statistics types

There are two types of statistics.

- Aggregate statistics: The aggregate duration in microseconds(XXag) and the maximum duration in microseconds (XXmx).
- Count statistics: A count(XXcn) is just a simple count, this doesn't get reset.

The aggregate statistics(XXag) and count statistics(XXcn) are cumulative. This way if there are two different files (from any time), you can calculate the average response time between the two files. The maximum will be reset after each sample period.

Tables

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

The below table describes the statistics collection for Volume aggregate statistics.

Table 40. Statistics collection for volume aggregate

Statistic name	Description
wlag, wlcn, wlmx	Indicates Write Lock Requests, these will be against the client volume. **Note:** This will include cleaning I/O.
rlag, rlcn, rlmx	Indicates Read Lock Requests, these will be against the client volume.
bcag, bccn, bcmx	Indicates Background Copy Operations, these will be against the target, that is, the volume being copied to.
clag, clcn, clmx	Indicates Cleaning Operations, these will be against the target being cleaned, that is, the volume being copied from.
rwag, rwcn, rwmx	Indicates ROW copy Operations, these will be against the client volume.
twag, twcn, twmx	Indicates Host Write I/O Operations at top of FlashCopy, these will be against the client volume.
trag, trcn, trmx	Indicates Host Read Operation at top of FlashCopy, these will be against the client volume.
bwag, bwcn, bwmx	Indicates Write Operations at bottom of host I/O, these will be against the client volume of the host IO.
brag, brcn, brmx	Indicates Read Operations at bottom of FlashCopy for host I/O, these will be against the client volume of the host IO.

The below table describes the statistics collection for Volume count statistics.

Table 41. Statistics collection for volume count	
Statistic name	**Description**
hwsg	Indicates Host Write Split Grain. The number of grains split due to host I/O for this volume, this will be against the client volume.
hwdg	Indicates Host Write Diff Grain. The number of grains marked as different due to host I/O for this volume, this will be against the client volume.
srcp	Indicates Successful ROW copy. The number of successful ROW copy operations, this will be against the client volume.
twbl	Indicates Host Write Blocks at top of FlashCopy, this will be against the client volume.
trbl	Indicates Host Read blocks at top of FlashCopy, this will be against the client volume.
bwbl	Indicates Write Blocks at bottom of FlashCopy for Host I/O, this will be against the client volume for the host I/O.
brbl	Indicates Read Blocks at bottom of FlashCopy for Host I/O, this will be against the client volume for the host I/O.

Example snippets for volume statistics files.

```
id="mtfc_vdisk0"
ro="431608" wo="134744" wou="16792" rb="256827679" wb="65293344"
rl="1384624" wl="5200676" rlw="305664" wlw="313986" xl="1325749"
wxl="612068" rxl="713681" oro="0" owo="0"
orl="0" owl="0" oiowp="0"
uo="0" ub="0" uou="0" ul="0" ulw="0">
<fc wlag="49953663" wlcn="171517" wlmx="20924" rlag="0" rlcn="0" rlmx="0" bcag="0" bccn="0"
bcmx="0" clag="0" clcn="0" clmx="0" rwag="0" rwcn="0" rwmx="0" twag="174414058"
twcn="13531" twmx="0" trag="0" trcn="0" trmx="0" bwag="1644964" bwcn="11657" bwmx="0"
brag="169298247" brcn="11657" brmx="0" hwsg="0" hwdg="78695" srcp="0" twbl="3018747" trbl="0"
bwbl="5968384" brbl="5968384" />
<ca rh="3509455" d="65314336" ft="128128" wt="0"
fw="65188256" wh="0" v="1920" m="1920"
ri="431608" wi="134789" r="253318224" dav="35731" dcn="36073" sav="2150" scn="30609" teav="0"
tsav="0"  tav="3000"  pp="0"/>
<cl bup="0" bdn="0"/><cpy idx="0">
<ca p="23479360" rh="629009376" ph="23234248" d="67061321"
ft="488576" wt="0" fw="85798816" wh="18677066"
v="7594176" m="553216" pm="1136" ri="1281284"
wi="185325" r="21743776" dav="234688" dcn="51900"
sav="14551" scn="6209" pav="16011" pcn="1365"
teav="0"  tsav="0"  tav="0"  pp="0"/>
```

Chapter 15. Notifications

The system supports the use of Simple Network Management Protocol (SNMP) traps, syslog messages, and Call Home emails to notify you and IBM Support Center when significant events are detected.

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 TCP, UDP, and TLS 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.

Syslog over TLS

The system supports encrypted communication with a syslog server using a TLS connection.

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

The syslog server certificate must be installed in a trust store on the system that has the syslog tag turned on. If the syslog server's certificate is signed by a Certificate Authority (CA), then the CA certificate must be installed in a trust store with the syslog tag that is turned on. See **`Truststore management commands`** for more information about creating and managing trust stores.

© Copyright IBM Corp. 2025

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 syslog server. If the system's certificate is signed by a Certificate Authority (CA), then the CA certificates must be installed on the SYSLOG 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 syslog servers may have requirements about details that must be included in the system certificate. The cluster's fully qualified domain name, and the cluster IP address, should be included in the subject alternative name fields. See **`System certificates`** for more details about generating a new system certificate.

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

Using the management GUI

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

Creating a Syslog server

To configure Syslog notifications, you must create a Syslog server.

1. Go to **Settings** > **Notifications** > **Syslog** > **Add Syslog Server** to add a Syslog server.
2. In the **Add Syslog Server** window, you can define the following protocols:

   UDP
   Use this protocol to reduce system and network overhead.

   TCP
   It is a reliable delivery protocol.

   TLS
   This option uses TLS to send encrypted syslog messages.
3. If you select **TLS**, the **Drag and drop or click here to upload the syslog server certificate** option appears.
   a. If the syslog server's certificate is signed by a CA, upload the root CA certificate.
   b. If the server's certificate is self-signed, upload the self-signed certificate.
   c. If you are creating multiple syslog servers that use the same certificate and the certificate has already been uploaded, select the box **The certificate has already been uploaded**.

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

   When you create a Syslog server that uses TLS, the management GUI creates a new trust store that contains uploaded server certificates. The CLI must be used to modify or remove trust store entries for existing servers.
4. Click **Add**.

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 42. 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 • Security Name
Authentication	Notifications are authenticated, but message content is not encrypted.	• Engine ID • Security Name • Authentication Protocol • Authentication Passphrase
Authentication and Privacy	Notifications are authenticated and message content is encrypted.	• Engine ID • Security Name • Authentication Protocol • Authentication Passphrase • Privacy Protocol • 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 does 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 **Edit properties**.
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 example 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" on page 391 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 **Settings** > **Security** > **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 **Settings** > **Security** > **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 **Settings** > **Security** > **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.

Chapter 16. Network

The system relies on a combination of network infrastructure and connectivity technologies to facilitate data storage, retrieval, and management.

The connectivity in Storage Systems includes Ethernet Connectivity, Fibre Channel (FC) connectivity, Portsets, and Partnerships. You can connect systems to enable migration, data replication, and high-availability solutions.

Ethernet connectivity

IBM FlashSystem storage systems can use Ethernet networks and covers multiple aspects of Ethernet connectivity.

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 system 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 system 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 does not belong to the same subnet as any configured system management IP address, the system forwards the packet to a default gateway. The system selects the default gateway that is associated with the Ethernet port, which has the lowest port number among the ports that are configured with a system management IP address.

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.

Configuring VLAN

Consider these guidelines and tasks before you configure a virtual local area network (VLAN).

Following are the general guidelines before configuring a VLAN:

- Virtual local area network (VLAN) tagging is supported for Ethernet-based traffic.
- VLAN provides network traffic separation at the layer 2 level for transport.
- VLAN tagging by default is disabled for any IP address of a node port. You can use the command-line interface (CLI) to optionally set a VLAN ID for port IP addresses.
- To prevent connectivity issues when a VLAN ID is configured for the IP addresses that are used for host attachment, appropriate VLAN settings on the Ethernet network and servers must also be properly configured.
- A VLAN ID can be configured for IPv4 addresses and IPv6 addresses of Ethernet ports.

© Copyright IBM Corp. 2025

- The same VLAN ID can be configured for all the Ethernet port IP addresses, or a different VLAN ID can be configured on each Ethernet port IP address.
- To ensure reliable connectivity to storage during a node failure, configure the same VLAN ID on the local and failover port on two nodes of an I/O group.
- If different VLAN IDs are set for local and failover port on two nodes of an I/O group, configure both VLAN IDs on the Ethernet switch ports. By doing this, a failover of IP addresses from a failing node to a surviving node succeeds.
- After VLAN is configured, changes to the VLAN settings disrupt traffic to or from those ports.
- Different VLAN configurations on different ports but with the same subnet on the ports is not supported.

  Note: For iSCSI host attach, the corresponding ports on both nodes must be configured with symmetric VLANs. Any asymmetric VLAN configuration between matching ports across nodes is not supported.

Configuring VLAN for hosts

Virtual local area network (VLAN) tagging is supported for Ethernet-based traffic.

- Setting VLAN tags for host attachment sessions that are already configured with a system that is running software before version 7.4.0 is a disruptive process.
- To configure VLAN tags for existing host attachment sessions:
  - Configure VLAN on host ports that are used for Ethernet connectivity.
  - Configure all intervening switches with appropriate VLAN tags according to port connectivity.
  - Configure system ports with appropriate VLAN tags.
- If VLAN tagging needs to be done without disrupting I/Os from hosts, the host should have multiple ports. Multipath drivers on hosts must be correctly configured with multiple Ethernet sessions to each system port. Then, when you configure VLANs, sessions must be brought down one at a time so that I/Os can continue through alternative paths to the system. Depending on switch connectivity, this process might not be possible.

Configuring VLAN ID for iSCSI

You can use management GUI and the command-line interface (CLI) to configure a virtual local area network (VLAN) ID for a node port IP address in multiple ways.

Using the management GUI

Complete these steps to configure VLAN tags on existing host sessions:

1. On the hosts, configure VLAN on the host ports that are used for Ethernet-based connections. Consult the appropriate host documentation to complete this step.
2. On all intervening switches, configure the switches with the VLAN tags per port connectivity. Consult the appropriate switch documentation to complete this step.
3. Verify that the port on which you are configuring VLAN tags is used for host attachment. In the management GUI, select **Settings** > **Network** > **Ethernet Ports**. Ensure that **Yes** is displayed in the host attach column for the selected port. If **No** is displayed, right-click the port and select **Modify iSCSI hosts** and select **Enabled**.
4. Configure ports on the system with appropriate VLAN tags by completing these steps:
   a. In the management GUI, select **Settings** > **Network** > **Ethernet Ports**.
   b. Right-click the port on which to configure VLAN and select **Manage IP Addresses**.
   c. On **Manage IP Addresses** page, select **Add IP Address**. Enter the following information for the IP address that you are adding to the selected port:

      IP address
      Enter the IP address to associate with the selected port.

Type
Select the IP protocol version of the IP address.

Subnet Mask or Prefix
Enter the subnet mask for the IPv4 addresses or enter the prefix for IPv6 addresses.

VLAN
Enter the corresponding VLAN tag that this IP address belongs to.

Portset
Select the name or ID of the portset where the portset type is **Host Attachment**.

d. Click **Save**.

e. If you are updating an existing IP address, select the IP address and select **Edit** from the overflow menu. Enter the new VLAN tag in the **VLAN** field. Click **Save**.

5. After you configure VLAN settings for host attachment, ensure that you set the priority flow control (PFC) setting for the host attachment traffic.

   Note: If you are sharing the same port for different types of traffic, then PFC and Enhanced Transmission Selection (ETS) settings must be configured.

Using the command-line interface

1. To configure a VLAN ID along with a port IPv4 and IPv6 address to a specified Ethernet port, use the **-vlan** parameter in the **mkip** command, as in the following command-line interface (CLI) command:

   ```
   mkip -node node_name | node_id -port port_id -ip ipaddr
   -prefix prefix -vlan vlan_id
   ```

   where *node_name* | *node_id* specifies the name or ID of the node that is being configured, *ipaddr* is the IPv4 or IPv6 address for the Ethernet port, *prefix* is subnet prefix for IPv4 or IPv6, *port_id* specifies the Ethernet port ID. The **-vlan** parameter sets the virtual local area network (VLAN) ID to the value specified by *vlan_id* variable.

2. To remove the VLAN ID and the IP address from a node Ethernet port, enter the following command:

   ```
   rmip ip_address_id
   ```

   where *ip_address_id* specifies the name of the IP address. To list the valid values for the IP address ID, enter the **lsip** command.

System IP addresses

The system IP address provides access to the system management interfaces and access to remote services.

Requirements for storage portsets

The system supports a maximum of one portset of storage type, that is, portset 3. These requirements are specific to storage portsets:

- The maximum number of IP addresses for a storage portset equals the number of Ethernet ports on the node.
- A portset can either contain IPv4 or IPv6 IP or a mix of IPv4 and IPv6 addresses.
- The storage type portset must be already created.

Requirements for System IP addresses

- System IP requirements: The system IP address provides access to the system management interfaces, including the management GUI, CLI, and RestAPI. The system IP address is also used to access remote services like authentication servers, NTP, SNMP, SMTP, and syslog systems, if configured.

- A system address (IPv4 or IPv6) must be configured on Ethernet port 1 (default) or any other Ethernet port.
- Optionally, a second system address can be configured on any Ethernet port except the one where the first system address is configured.
- A maximum of two system addresses, either IPv4, IPv6, or a mix of both, can be configured for system management.
- To ensure system IP failover operations, the Ethernet ports where the system addresses are configured must be connected to the same subnet. The system IP address can fail over to any node in the system.
- System addresses can be configured on any ports.
- IP addresses configured for system management and service access must not be used for Ethernet-based host I/O.
- A management portset (for a storage partition) can include one partition management IP address (IPv4 or IPv6). For more information, see Chapter 11, “Storage partitions,” on page 243.

To assign an IP address to any Ethernet interface I/O port, use the management GUI or the **mkip** command. To change the MTU parameter for specific Ethernet port, use the management GUI or **chportethernet** command.

If the node fails, the address becomes unavailable and the host loses communication with the system via that node. To allow hosts to maintain, access to data, the node-port IP addresses for the failed node are transferred to the partner node in the I/O group. The partner node handles requests for both its own node-port IP addresses and also for node-port IP addresses on the failed node. This process is known as node-port IP failover.

Priority flow control

Priority flow control (PFC) is an Ethernet standard that supports the ability to select the priority of different types of traffic within the network.

With PFC, administrators can reduce network congestion by slowing or pausing certain classes of traffic on ports, thus providing better bandwidth for more important traffic. The Storage Virtualize system supports PFC on various supported Ethernet-based protocols on three types of traffic classes: system, host attachment, and storage traffic.

You can configure a priority tag for each of these traffic classes. The priority tag can be any value 0 - 7. You can set identical or different priority tag values to all these traffic classes. You can also set bandwidth limits to ensure QoS for these traffic classes by using the Enhanced Transmission Selection (ETS) setting on the network. When you plan to configure priority flow control (PFC), follow these guidelines and examples.

To use PFC and ETS, ensure that the following tasks are completed:

1. Ensure that ports support 10 GB or higher bandwidth to use PFC settings.
2. Configure VLAN on the system to use PFC capabilities.
3. Ensure that the same VLAN settings are configured on the all entities, including all switches between the communicating end points.
4.
5. To enable priority flow for host attachment traffic on a port, make sure that the host flag is set to *yes* on the configured IP address on that port.
6. To enable priority flow for storage traffic on a port, make sure that storage flag is set to *yes* on the configured IP address on that port.
7. On the switch, enable Data Center Bridging Exchange (DCBx). DCBx enables switch and adapter ports to exchange parameters that describe traffic classes and PFC capabilities. For these steps, check your switch documentation for details.

8. For each supported traffic class, configure and enable the same priority tag on the switch as COS value in step 4 . For example, if you plan to use priority tag 3 for storage traffic, ensure that it is configured on the IBM Storage Virtualize system through **chsystemethernet -storagecos** command and the priority tag 3 is also enabled on the switch.
9. If you are planning on using the same port for different types of traffic, ensure that ETS settings are configured on the network.

Using the management GUI

To set PFC on the system, complete these steps:

1. In the management GUI, select **Settings** > **Network** > **Priority Flow Control**.
2. For each of following classes of service, select the priority setting for that traffic type:

   System
   Set a value in the range 0 - 7 for the system traffic that includes communication between nodes within the system. The system priority tag is supported on iSCSI connections and on systems that support RDMA over Ethernet connections between nodes. Ensure that you set the same priority tag on the switch to use PFC capabilities.

   Host Attachment
   Set the priority tag in the range 0 - 7 for system to host traffic. The host attachment priority tag is supported on iSCSI connections and on systems that support NVMe over RDMA connections. Ensure that you set the same priority tag on the switch to use PFC capabilities.

   Storage Virtualization
   Set the priority tag in the range 0 - 7 for system to external storage traffic. The storage virtualization priority tag is supported on storage traffic over iSCSI connections. Ensure that you set the same priority tag on the switch to use PFC capabilities.
3. Click **Save**.
4. Make sure that IP is configured with VLAN.

Using the command-line interface

To use PFC on the system by using the command-line interface, follow these steps:

1. To use priority flow control for host attachment and storage traffic, use the **chsystemethernet** command to set the priority tag (COS) values before you set any data IPs by using the **mkip** command. All the host attachment and storage connections can then inherit the COS values that are already set.

   Note: You can also set it after you configure data IPs as existing sessions do not refresh the PCP value immediately. Then, when you try to connect again, the new connection takes the newly configured COS values.
2. Enter the **lssystemethernet** command to determine the current settings for host attachment, system, or storage priority tags. The resulting output displays values in the range 0 - 7 such as the following example:

   ```
   host_attach_cos  0
   system_cos       0
   storage_cos      0
   ```
3. To set the host attachment, system, or storage priority tags, enter the **chsystemethernet** command with values in the range 0 - 7 as in the following example:

   ```
   chsystemethernet -hostattachcos 4 -systemcos 5 -storagecos 6
   ```

   Resulting output:

   ```
   host_attach_cos  4
   system_cos       5
   storage_cos      6
   ```

Where the host attachment priority tag (class of service) is set to 4, the system priority tag is set to 5, and the storage priority tag is set to 6.

4. Make sure that IP is configured with VLAN. If the IP is not configured with VLAN, see "Modifying VLAN for priority flow control" on page 304 to configure the VLAN for host attachment, system, and storage traffic.

Known considerations and limitations for PFC

- PFC is supported on 16 Gbps, 100 Gbps, and 25 Gbps Ethernet adapters.
- PFC is not supported by onboard Ethernet ports.

Considerations for system_cos

  - After an upgrade to 8.5.2 or later releases, you must run **satask stopnode -warmstart** orderly on all the nodes in the system, if the value of *system_cos* is set to a nonzero value.
  - On version 8.5.2 or later releases, if you change *system_cos* values with the **chsystemethernet** command, you must run **satask stopnode -warmstart** orderly on all the nodes in the system.

Warmstart ensures that the *system_cos* value is reflected correctly for system traffic.

Modifying VLAN for priority flow control

You can change Virtual local area network (VLAN) settings to separate network traffic for Ethernet transport.

The system supports VLAN configuration for system, host attachment, storage virtualization, and IP replication traffic. VLAN is a prerequisite for using priority flow control (PFC). VLAN tags msvc_chnodeipust be assigned to IP versions for which you want to use PFC capabilities and VLAN should be configured end to end including all the switches. PFC creates different classes for all types of traffic that improves network performance during times of congestion. To ensure that PFC operates as expected, the network and its components must be configured correctly. For example, if a single port is used for multiple types of traffic (host attachment, and storage) then the performance of these traffic types can be affected. The system supports setting system-wide PFC-setting on traffic for host attachment traffic and storage virtualization traffic. PFC capabilities are not supported on IP replication traffic.

Configuring VLAN for host attachment

When you configure VLAN tagging for host connections, follow these guidelines:

- To configure VLAN tagging without disrupting I/O from hosts, all hosts need multiple ports and multipath drivers. In addition, each host must be configured with multiple sessions to each node on the system.
- Each session needs to be stopped one at a time to ensure that I/O can continue on alternative paths to the system.

If you are running previous versions of the system, configuring VLAN tags for host sessions is a disruptive process.

Complete these steps to configure VLAN tags on existing host sessions:

1. On the hosts, configure VLAN on the host ports that are used for Ethernet-based connections. Consult the appropriate host documentation to complete this step.
2. On all intervening switches, configure the switches with the VLAN tags per port connectivity. Consult the appropriate switch documentation to complete this step.
3. Verify that the port on which you are configuring VLAN tags is used for host attachment. In the management GUI, select **Settings** > **Network** > **Ethernet Ports**. Ensure that **Yes** is displayed in the host attach column for the selected port. If **No** is displayed, right-click the port and select **Modify iSCSI hosts** and select **Enabled**.

4. Configure ports on the system with appropriate VLAN tags by completing these steps:
   a. In the management GUI, select **Settings** > **Network** > **Ethernet Ports**.
   b. Right-click the port on which to configure VLAN and select **Manage IP Addresses**.
   c. On **Manage IP Addresses** page, select **Add IP Address**. Enter the following information for the IP address that you are adding to the selected port:

      IP address
      Enter the IP address to associate with the selected port.

      Type
      Select the IP protocol version of the IP address.

      Subnet Mask or Prefix
      Enter the subnet mask for the IPv4 addresses or enter the prefix for IPv6 addresses.

      VLAN
      Enter the corresponding VLAN tag that this IP address belongs to.

      Portset
      Select the name or ID of the portset where the portset type is **Host Attachment**.

   d. Click **Save**.
   e. If you are updating an existing IP address, select the IP address and select **Edit** from the overflow menu. Enter the new VLAN tag in the **VLAN** field. Click **Save**.
5. After you configure VLAN settings for host attachment, ensure that you set the priority flow control (PFC) setting for the host attachment traffic.

   Note: If you are sharing the same port for different types of traffic, then PFC and Enhanced Transmission Selection (ETS) settings must be configured.

Configuring VLAN for storage virtualization

VLAN tagging of the IP addresses configured for external storage systems is supported. For the systems and external storage systems to access each other, the same VLAN ID must be configured for ports on nodes on the system and the ports on the external storage system. In addition, VLAN settings must be configured on the Ethernet network to prevent any connection issues. When more than two ports are available, configure different ports for iSCSI hosts and storage systems.

Complete these steps to configure VLAN tags on storage systems:

1. On the storage systems, configure VLAN on the ports that are used for Ethernet connections. Consult the appropriate storage documentation to complete this step.
2. On all intervening switches, configure the switches with the VLAN tags per port connectivity. Consult the appropriate switch documentation to complete this step.
3. Verify that the port on which you are configuring VLAN tags is used for storage system attachment. In the management GUI, select **Settings** > **Network** > **Ethernet Ports**. Ensure that **Enabled** is displayed in the Storage Port column for the selected port. If **Disabled** is displayed, right-click the port and select **Modify Storage Ports** to enable the port for iSCSI connections to storage systems.
4. Configure ports on the system with appropriate VLAN tags by completing these steps:
   a. In the management GUI, select **Settings** > **Network** > **Ethernet Ports**.
   b. Right-click the port on which to configure VLAN and select **Manage IP Addresses**.
   c. On **Manage IP Addresses** page, select **Add IP Address**. Enter the following information for the IP address that you are adding to the selected port:

      IP address
      Enter the IP address to associate with the selected port.

      Type
      Select the IP protocol version of the IP address.

Subnet Mask or Prefix
Enter the subnet mask for the IPv4 addresses or enter the prefix for IPv6 addresses.

VLAN
Enter the corresponding VLAN tag that this IP address belongs to.

Portset
Select the name or ID of the portset where the portset type is **Storage**.

Note: Portset3 (ID=3) is created by default and is reserved for storage traffic. Portset3 (ID =3) is the only portset of type Storage.

d. Click **Save**.

e. If you are updating an existing IP address, select the IP address and select **Edit** from the overflow menu. Enter the new VLAN tag in the **VLAN** field. Click **Save**.

5. After you configure VLAN settings for storage virtualization, ensure that you set PFC settings for the system-to-storage traffic.

Note: If you are sharing the same port for different types of traffic, then PFC and Enhanced Transmission Selection (ETS) settings must be configured.

Planning for a direct-attached iSCSI configuration

A direct-attached configuration dedicates entire port bandwidth to be used for a specific connection. Planning must account for the volume of expected traffic before you decide how many ports are used in a direct-attached configuration. Direct attached connections are supported on the same speed at both ends.

The system offers multiple options in deciding how to create a direct-attached configuration. Bandwidth requirements for a direct-attached configuration make it important to determine the volume of expected traffic in deciding the number of ports. Ensure you have accounted for sufficient bandwidth, in particular node to node as the bandwidth between nodes determines write speed for replicated workloads.

For example, if a system to system direct-attached link is configured between nodes in an I/O group, then the link carries cache replication traffic and it might carry inter-node traffic such as FlashCopy data. Hence, enough system to system bandwidth needs to be available so that:

- Cache replication for host writes can occur without any bandwidth bottleneck
- All possible intra-node traffic can be carried without any bandwidth bottleneck.

The following guidelines are provided for direct-attached iSCSI configurations.

Note: One or more Ethernet host adapters must be installed for the system to support these direct attachment configurations.

1. iSCSI direct attachment to host systems.

   The system allows direct attachment connectivity between its iSCSI ports and the host ports. Host systems can be connected to 10 Gbps, 25 Gbps, or 100 Gbps Ethernet ports on the system. No special configuration is needed for host systems that are using this configuration.

2. iSCSI direct attachment between nodes in a system.

   Direct connection of the system Ethernet ports without using an Ethernet switch is supported. Such direct connections between the system nodes might be useful in small configurations where no switch exists. It can also be used to connect nodes in the same I/O group to provide a dedicated connection for mirroring the fast write cache data. Ensure that sufficient bandwidth is provisioned between nodes to accommodate all the intra-node traffic.

For more information about host compatibility, see IBM System Storage Interoperation Center (SSIC).

Note:

- AIX host attach is generally not supported with direct-attached iSCSI configuration and requires SCORE approvals for enabling support.

- iSCSI direct-attached over policy-based high availability configuration is generally not supported and requires SCORE approvals.
- For each connection, the host interfaces must be on different subnets. For instance, two interfaces from the same host must be on different subnets.

Fibre channel connectivity

The Fibre Channel protocol (FCP) is used to establish connectivity over FC infrastructure.

Systems are defined by supported configuration and zoning rules. A system requires single-initiator zoning for all large configurations that contain more than 64 host objects. See zoning for hosts and zoning for external storage system.

Auto zoning

Auto zoning enables inline communication between storage systems and Fibre Channel (FC) switches to automate zoning. The auto zoning feature automatically creates and manages zones on supported FC switches, simplifying host-to-storage connectivity and reducing errors. By automatically forming zones between host ports and storage target ports, auto zoning reduces manual administrative effort, configuration errors, and deployment time.

IBM FlashSystem auto zoning feature provides a complete solution for Fibre Channel zone management that integrates seamlessly with the Brocade Target Driven Zoning (TDZ) technology. This feature is designed to address a wide range of challenges typically faced by SAN administrators, including:

- Creating, modifying, and deleting zones within the Fibre Channel fabric
- Identifying optimal zone definitions to achieve optimal multi-path configurations for host servers
- Managing zoning in accordance with SAN best practices
- Backing up and restoring zone configurations to ensure reliability and recoverability
- Automatically adapting zone definitions by tracking changes in host or storage connectivity

With auto zoning feature, storage target ports manage zoning operations on the FC fabric. Auto zoning creates peer zones, which provide an optimized zoning model:

Peer zone contents

- One principal device is typically a storage array target port.
- Multiple non-principal devices are typically host initiator ports.

Peer zone communication rules

- The principal device can communicate with all nonprincipal devices, and non-principal devices can communicate with the principal device.
- Nonprincipal devices cannot communicate with each other.

Peer zones replace traditional single-initiator zoning and maintain consistent and predictable zoning patterns.

Supported Switch Platforms

The auto zoning feature is supported by Brocade FC switches that meet the following conditions:

- Fabric OS (FOS) version 9.1.1 or later
- The TDZ feature must be enabled on the switch ports where the storage target ports are connected. The feature can be enabled through the switch CLI. Refer to the switch documentation for detailed steps to enable this capability.
- The TDZ feature does not need to be enabled on host ports.

The TDZ feature must be explicitly enabled by the fabric administrator for each relevant port. The auto zoning feature initiates automatic zoning actions only after the TDZ is enabled.

Fabric Connectivity Requirements

For auto zoning to operate correctly and predictably, the fabric connectivity of the storage target ports and connected host ports must meet specific requirements. These considerations can ensure consistent zone formation, proper path distribution, and stable host-storage communication across fabrics.

- Auto zoning is supported only for user-defined FC portsets.
  - All FC I/O port IDs included in an auto zoning enabled portset must demonstrate symmetric fabric connectivity. The same FC I/O port ID from every node in the cluster must be connected to the same fabric. For example, FC IO Port ID 1 from all nodes must connect to Fabric A.
  - Ideally, a portset should contain an equal number of target FC I/O ports from each redundant fabric.
  - Host initiator ports must be connected to the same fabric as the target ports defined within the portset that you want to associate with host object.

Auto zoning creates peer zones between host initiator ports and the target ports that share fabric.

Auto zoning Workflow

To start using auto zoning when adding new hosts in an existing configuration:

1. Ensure that fabric requirements are met as described above.
2. Create an auto zoning enabled Fibre Channel portset.
3. Create the host by using the FDMI hostname or by specifying a list of host port WWPNs.
4. Associate host with the auto zoning enabled portset.

Zoning policies

Auto zoning supports configurable zoning policies that determine how initiator WWPNs are matched with target port WWPNs when forming peer zones.

Zoning policies are applied at the portset level, and different portsets can use different policies based on deployment requirements.

Available zoning policies are:

one_to_one connectivity

This policy enforces a single path between every initiator port and one target port on each canister within the same fabric.

- Each initiator WWPN is zoned with only one FC IO Port ID from the portset. Each FC IO Port ID represents a single target port from each canister.
- Zoning is performed between the initiator port and the portset target port, provided both ports are within the same fabric.
- Initiator ports are evenly distributed across all target ports from the portset to achieve better workload distribution.
- After zoning, that initiator WWPN is not zoned with any additional target ports from the portset, even if multiple suitable target ports exist on that fabric.

This policy is optimal for environments requiring strict path control or host configurations where minimal connectivity is preferred.

one_to_one_all_fabrics

This policy also creates one-to-one zoning, but it is designed for host environments where initiator WWPNs can not be visible on all fabrics during initial discovery.

- Each initiator WWPN is zoned with one target port WWPN from each fabric but will have a single host path from fabric where Host port is connected .
- This policy is useful where initiator WWPNs do not appear upfront in the fabric
  - SAN-boot scenarios

– NPIV (N_Port ID Virtualization) WWPN configurations

This policy ensures reliable boot-time zoning even when initiator WWPN visibility is limited or dynamic.

Zone Naming Convention

Zones are created that uses the following format, `IBM_SV_<Target Port WWPN>_00`, for example, IBM_SV_50050768101655c4_00.

This predictable naming scheme allows easy identification of IBM Storage managed zones in the switch fabric.

Zone Name Prefix Configuration

Storage administrator has an option to specify a zone name prefix for all automatically created zones. This prefix can be configured by using the chsystem command.

- The zone name prefix must be configured before the first auto zoning enabled portset is created on the system.
- Once the prefix is set on a system that contains any auto-zone-enabled portset, the prefix cannot be modified.
- If no prefix is provided, the default naming scheme is used, for example, `IBM_SV_<target-WWPN>_00`.

This mechanism ensures naming consistency across all auto-generated zones in the fabric.

If the user configures the prefix FC_Infinity, it is automatically prepended to the zone name. For example, if **`Prefix`** is FC_Infinity and **`Target Port WWPN`** is 50050768101655c4, then the resulting auto-generated zone name becomes: FC_Infinity_IBM_SV_50050768101655c4_00.

This example demonstrates how user-defined prefixes are applied to the default zone naming format.

Configuring auto zoning in the management GUI

In the management GUI, if auto zoning was not configured automatically during system setup or created manually on an auto zoning enabled portset. Select **Settings** > **Network** > **Fibre Channel Portsets**, and click **Auto configure ports and portsets**.

Note: The auto configure ports and portsets does not appear when auto zoning-enabled ports exist in the system.

To verify auto zoning capability, click **Create portset**. If **Auto zoning** is **Enabled**, the portset supports auto zoning.

When adding a host in **Hosts** > **Add Host**, portsets with auto zoning enabled are preselected.

Related information

Configuration limits

`mkportset`
`mkhost`
`lshostzone`
`lsportfc`
`lsfabricport`
`lsportset`
`addhostport`
`addfcportsetmember`
`chhost`
`rmhost`
`rmfcportsetmember`
`chsystem`

Fabric Device Management Interface

The Fabric Device Management Interface (FDMI) enables any storage endpoint to register itself to the Fibre Channel (FC) fabric and query the HBA and port details of the entire fabric.

The FDMI functionality is divided into these two categories:

FDMI registration
The IBM Storage Virtualize system registers with the FC fabric and provides the system information to the FDMI database, such as node names, port speed, and software version.

FDMI discovery
The IBM Storage Virtualize system queries the FC fabric and collects the information about all the other devices on the fabric. Information is also collected about the devices that are not zoned to connect to the IBM Storage Virtualize system. Host server details appear in FDMI discovery only if the host operating system registers FDMI information with the FC fabric. Legacy operating systems that do not register FDMI information are not discovered through FDMI discovery.

To view the FDMI information, run the **lsfabricport** command that provides HBA and port attributes such as fdmi_host_name, fabric_name, os, and speed. FDMI discovery runs automatically every 6 hours and collects information about the fabric. However, to get the updated/latest fdmi data, use **detectfabricport** command that triggers FDMI discovery on all connected fabrics.

The same details are captured as a part of Call Home inventory data. To get the list of unconfigured host candidate ports, run the **lsfabricport -hostcandidate** command.

Fibre Channel port numbers and worldwide port names

Fibre Channel (FC) ports are identified by their physical port number and by a worldwide port name (WWPN).

The physical port numbers identify FC adapters and cable connections when you run service tasks. Worldwide port names (WWPNs), which uniquely identify the devices on the SAN, are used for tasks such as FC switch configuration. The WWPNs are derived from the worldwide node name of the node in which the ports are installed.

The vital product data (VPD) for each controller enclosure contains a unique WWNN value for each node canister. Each WWNN, in turn, defines the WWPN that each node canister allocates to the ports on an FC host interface adapter. All systems can support an optional quad-port 16 Gbps Fibre Channel (FC) host interface adapter.

Table 43 on page 310 shows the WWPN values for each port when a quad-port FC host interface adapter is installed in a control enclosure.

Table 43. Examples of WWPN values for the system

Host interface adapter	Port number	Canister 1 WWPN	Canister 2 WWPN
Quad-port 16 Gbps	1	5005 0768 0D00 0000	5005 0768 0D00 0001
	2	5005 0768 0D04 0000	5005 0768 0D04 0001
	3	5005 0768 0D08 0000	5005 0768 0D08 0001
	4	5005 0768 0D0C 0000	5005 0768 0D0C 0001

N-Port ID Virtualization

The system supports N_Port ID virtualization (NPIV) technology for Fibre Channel (FC) connections. NPIV is an industry standard technology that provides the capability to share one physical FC adapter as multiple virtual ports. One physical port on an FC adapter can use several world wide port names (WWPNs) and world wide node names (WWNNs) to register with the SAN fabric.

NPIV is a method for virtualizing a physical Fibre Channel port that is used for host I/O. When NPIV is enabled, ports do not come up until they are ready to service I/O, which improves host behavior around node unpends. In addition, path failures due to an offline node are masked from host multipathing.

NPIV functionality allows different WWPNs to be used for host I/O and clustering on the same physical Fibre Channel port. When NPIV is enabled, host I/O WWPNs do not come online until they are ready to service I/O, which can improve host behavior when a node is coming online. In addition, path failures due to an offline node are masked from hosts and their multipathing drivers do not need to do any path recovery because the I/O group partner can take the WWPNs on the offline node.

By default, NPIV is enabled on the system. However, you must accurately zone your system for NPIV configuration. For more information, see “Zoning requirements for N_Port ID virtualization” on page 311.

Zoning requirements for N_Port ID virtualization

All of the N_Port IDs for the two nodes in an I/O group must be associated with the same fabric.

The port ID of a virtual port remains constant regardless of the node that it is active on. Ports of the two nodes in an I/O group must be attached to the same fabric, while ports of different I/O groups can be on different fabrics. Figure 27 on page 311 shows a configuration that is not valid when using NPIV.

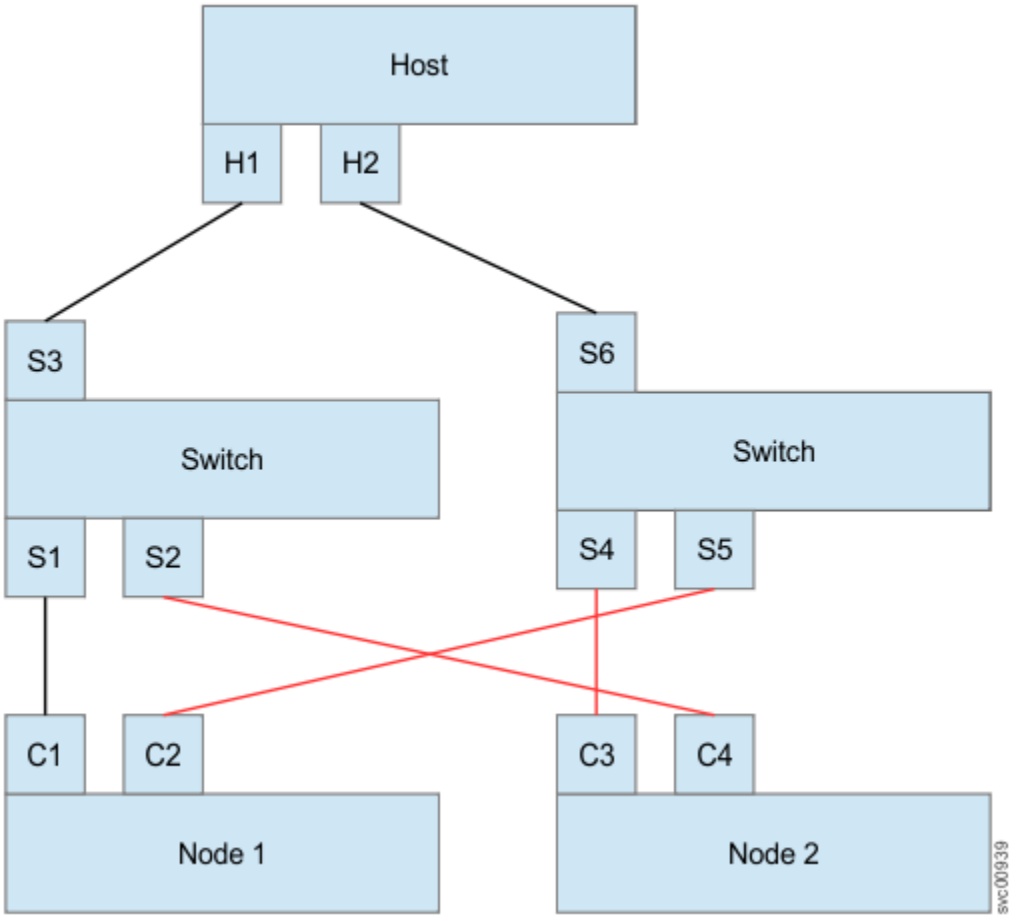


Figure 27. A configuration that is not valid for use with NPIV

The connections to switch ports S2, S4, and S5 must be exchanged, leading to the following correct configuration that is shown in Figure 28 on page 312.

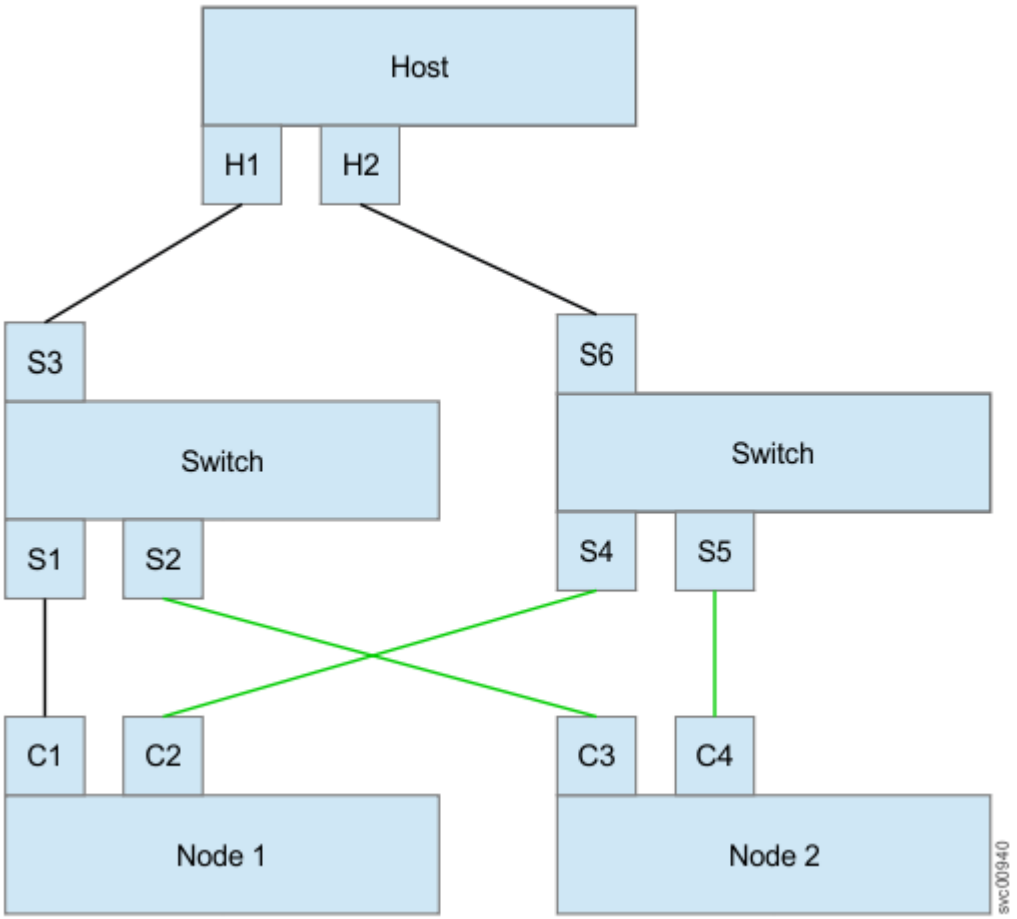


Figure 28. A configuration that is valid for use with NPIV and (WWPN-based) zoning

The same situation can occur when using a single switch fabric with hard (port based) zoning, as shown in Figure 29 on page 313.

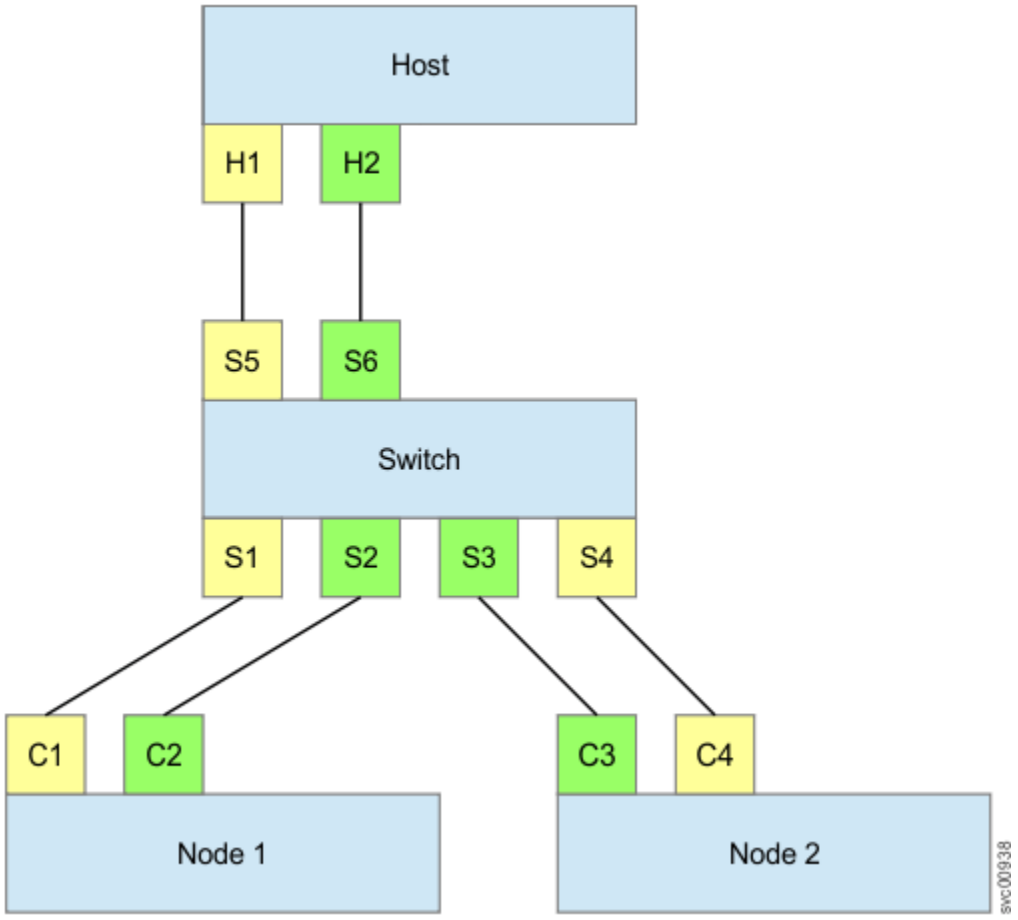


Figure 29. A single switch configuration that is not valid for use with NPIV and hard (port-based) zoning

Two nodes are connected to one host by using one switch. However, port-based zoning is used to ensure that each port of the host can detect only one port of each node. S1, S4, and S5 are in one zone; and S2, S3, and S6 are in the other zone.

During an IP failover, NPIV ports move from C1 to C3 and C2 to C4 (or the other way around). The configuration is not valid with NPIV as the visibility of ports move between H1 and H2. To use port-based zoning in this situation, two zones are required: One zone contains S1, S3, S5, C1 and C3. Another zone contains S2, S4, S6, C2 and C4. This configuration is shown in Figure 30 on page 314.

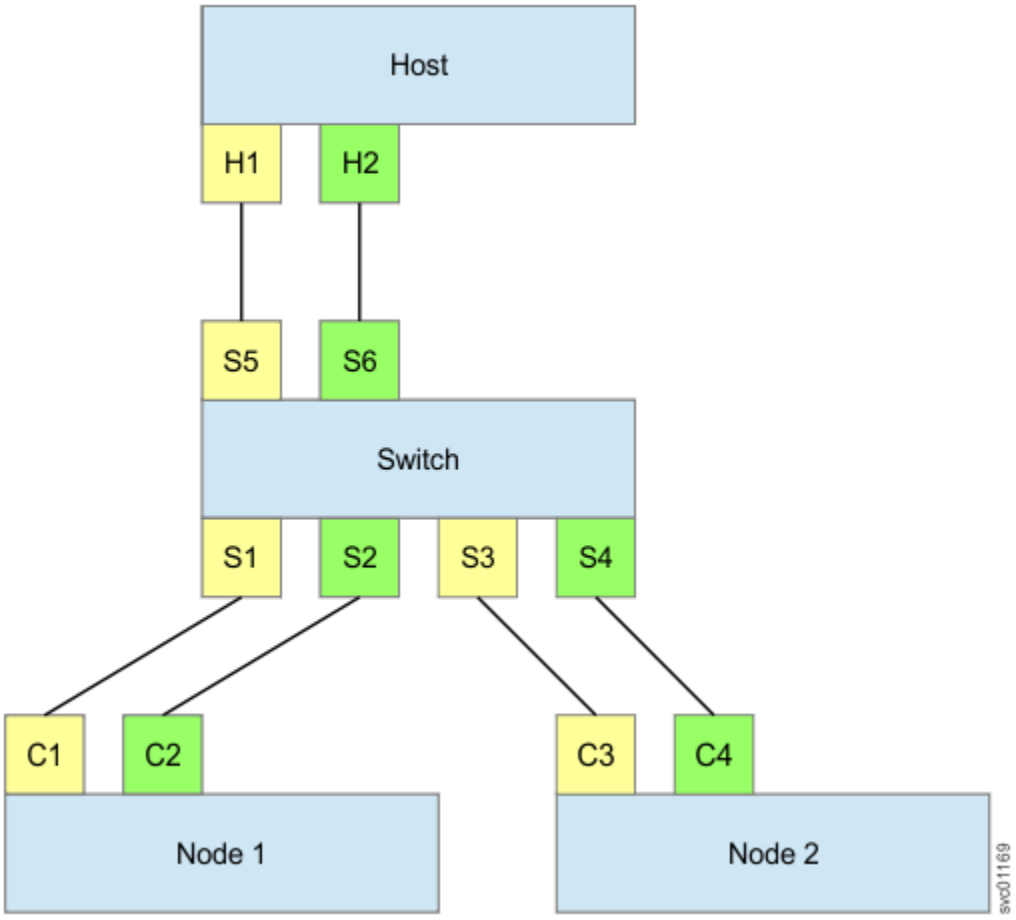


Figure 30. A single switch configuration that is valid for use with NPIV and hard (port-based) zoning

This situation does not arise when using a single switch fabric with soft (WWPN based) zoning. However, care needs to be taken in a multiswitch fabric where the node-to-host route passes through different switches for the two nodes in an I/O group.

If fabric changes are required, make them before you switch each I/O group to transitional state. I/O groups are switched to transitional state to help move them to NPIV. In transitional state, an I/O group contains both physical and virtual Fibre Channel ports that are currently being used. After you switch to transitional state, an attempt is made by the system software to warn about mismatched fabric issues. It does not warn about issues that are caused by port-based zoning.

Failure to adhere to the fabric zoning rules might cause undesirable behavior on the host during NPIV port failure. Examples include a 2-minute pause in I/O or reduced path redundancy.

Planning for connectivity

Plan to provide the network infrastructure and the storage network infrastructure that your system requires.

For more information about host compatibility, see IBM System Storage Interoperation Center (SSIC).

Planning for a direct-attached configuration

A direct-attached configuration dedicates entire port bandwidth to be used for a specific connection. Planning must account for the volume of expected traffic before you decide how many ports are used in

a direct-attached configuration. Direct attached connections are supported on the same speed at both ends.

The system offers multiple options in deciding how to create a direct-attached configuration. Bandwidth requirements for a direct-attached configuration make it important to determine the volume of expected traffic in deciding the number of ports. Ensure you have accounted for sufficient bandwidth, in particular node to node as the bandwidth between nodes determines write speed for replicated workloads.

For example, if a system to system direct-attached link is configured between nodes in an I/O group, then the link carries cache replication traffic and it might carry inter-node traffic such as FlashCopy data. Hence, enough system to system bandwidth needs to be available so that:

- Cache replication for host writes can occur without any bandwidth bottleneck
- All possible intra-node traffic can be carried without any bandwidth bottleneck.

The following guidelines are provided for direct-attached configurations.

1. Fibre Channel direct attachment to host systems.

   The system allows direct attachment connectivity between its Fibre Channel ports and host ports. Host systems can be connected to 16 Gbps to 64 Gbps Fibre Channel ports on the system. No special configuration is needed for host systems that are using this configuration.

   For more information about host compatibility, see IBM System Storage Interoperation Center (SSIC).
2. Fibre Channel direct attachment between nodes in a system.

   Direct connection of the system Fibre Channel ports without using a Fibre Channel switch is supported. Such direct connections between the system nodes might be useful in small configurations where no Fibre Channel switch exists. It can also be used to connect nodes in the same I/O group to provide a dedicated connection for mirroring the fast write cache data. Ensure that sufficient bandwidth is provisioned between nodes to accommodate all the intra-node traffic.

Note: Direct attached connections are supported on the same speed at both ends.

Planning for more than four fabric ports per node canister

You can use more than four fabric ports per node canister to improve peak load I/O performance, but careful planning is needed.

A *fabric port* is a Fibre Channel port. If you use more than four fabric ports per node, you must either use the **`localfcportmask`** and **`partnerfcportmask`** commands or be careful with your fabric zoning configuration.

Careful zoning improves resilience and prevents lease expires due to port errors, low fabric buffer credits, shared adapter resources, or push-back from remote links. Table 44 on page 315 displays a port allocation scheme.

Table 44. Port allocation scheme

Adapter / Port	4 ports	8 ports	12 ports	16 ports	SAN Fabric
Adapter 1 Port 1	Host and Storage	Host and Storage	Host and Storage	Host and Storage	A
Adapter 1 Port 2	Host and Storage	Host and Storage	Host and Storage	Host and Storage	B
Adapter 1 Port 3	Intracluster and Replication	Intracluster	Intracluster	Intracluster	A
Adapter 1 Port 4	Intracluster and Replication	Intracluster	Intracluster	Intracluster	B
Adapter 2 Port 1	-	Host and Storage	Host and Storage	Host and Storage	A

Table 44. Port allocation scheme (continued)					
Adapter / Port	**4 ports**	**8 ports**	**12 ports**	**16 ports**	**SAN Fabric**
Adapter 2 Port 2	-	Host and Storage	Host and Storage	Host and Storage	B
Adapter 2 Port 3	-	Intracluster or Replication	Replication or Host and Storage	Replication or Host and Storage	A
Adapter 2 Port 4	-	Intracluster or Replication	Replication or Host and Storage	Replication or Host and Storage	B
Adapter 3 Port 1	-	-	Host and Storage	Host and Storage	A
Adapter 3 Port 2	-	-	Host and Storage	Host and Storage	B
Adapter 3 Port 3	-	-	Intracluster	Intracluster	A
Adapter 3 Port 4	-	-	Intracluster	Intracluster	B
Adapter 4 Port 1	-	-	-	Host and Storage	A
Adapter 4 Port 2	-	-	-	Host and Storage	B
Adapter 4 Port 3	-	-	-	Replication or Host and Storage	A
Adapter 4 Port 4	-	-	-	Replication or Host and Storage	B
localfcportmask	1100	11001100 or 00001100	110000001100	0000110000001100	-
remotefcportmask	1100	00000000 or 11000000	000011000000	1100000011000000	-

- Host refers to host objects defined in the system.
- Replication refers to nodes that are part of a different cluster.
- Storage refers to controller objects defined in the system if external storage is being used.
- Intracluster refers to nodes within the same cluster.
- The word "and" indicates that both types are used.
- The word "or" indicates that one of the options must be selected. If using replication, preference should be given to replication.

If a node canister receives more than 16 logins from another node, then it causes node error 860.

You can ensure that no more than 16 logins are received by following these guidelines:

- Zone the SAN fabric to reduce the number of paths between the nodes.

- Apply a local Fibre Channel port mask (if the nodes are within the same system) or partner Fibre Channel port mask (if the nodes are in different systems) to reduce the number of ports that are being used for node-to-node communication.
- Provide a combination of the two (zoning and port masks).

To avoid receiving node error 860 and to maximize performance on your system, follow these guidelines:

- For redundancy, use a minimum of two ports. If ports are being dedicated for different types of traffic, assign a minimum of two ports for each type of traffic.
- Within a system, up to 50% of the overall data that is transferred across the SAN is transmitted between the nodes. However, for read-intensive workloads, the figure is much less. Therefore, if ports are being dedicated for different types of traffic, assign between 1/4 and 1/2 of the overall ports for node-to-node communication within a system.
- For replication between systems, the connection between the systems is usually the bottleneck. Other than for redundancy, there is no point in having more SAN connections than there is bandwidth. For example, if two sites are connected with a 10 Gbps link, two 8 Gbps ports are sufficient. Systems usually are configured with two ports per node for replication traffic primarily for redundancy. For systems with larger numbers of nodes, it might be appropriate to have connections only from a subset of the nodes for replication to a remote system. The system automatically forwards replication traffic between local nodes so that all nodes can still participate in replication.

In the examples that follow, consider this configuration that is shown in Figure 31 on page 317:

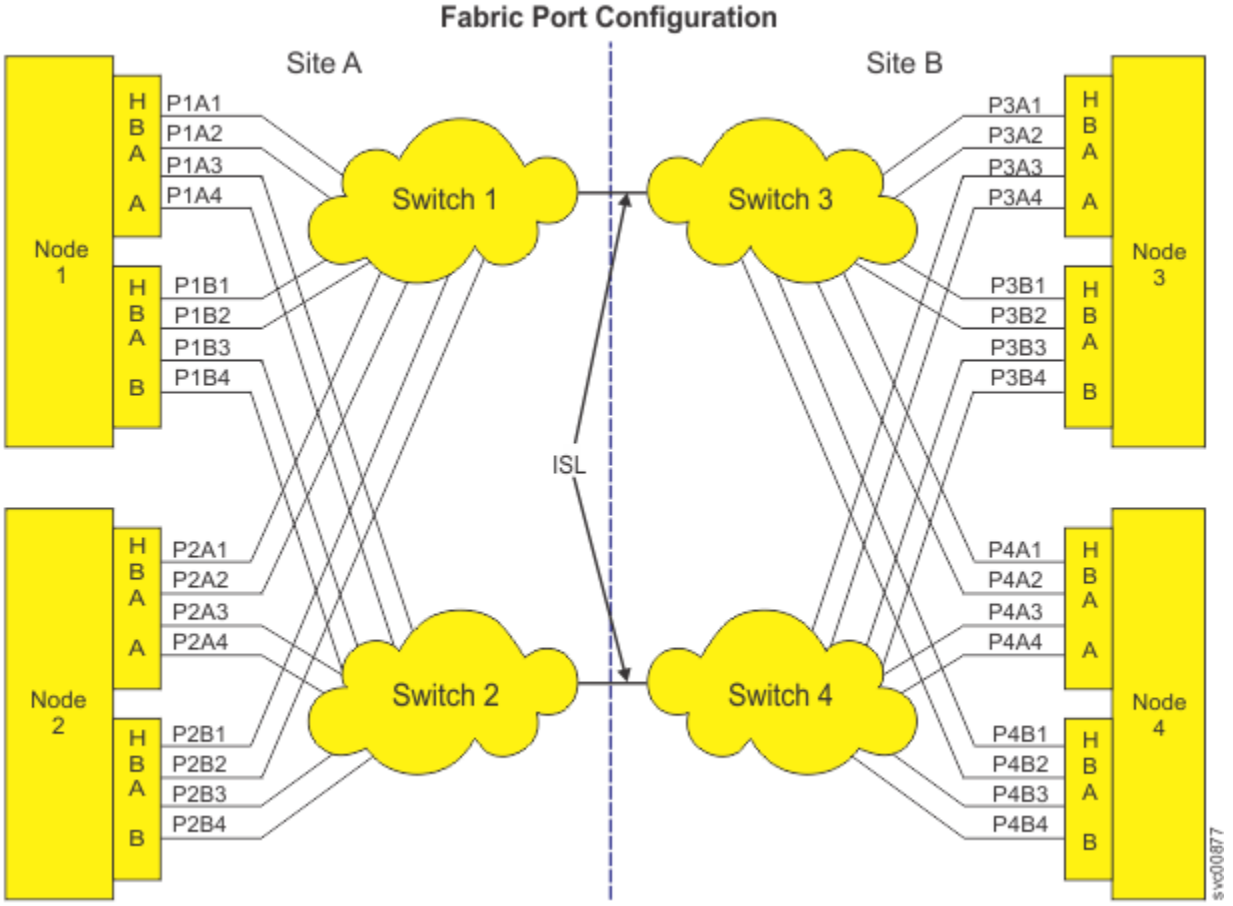


Figure 31. Fabric port configuration

- Four nodes, where nodes 1 and 2 are in a system on site A. Nodes 3 and 4 are in a system on site B.
- Each site has 2 switches (switch 1 and 2 on site A; switch 3 and 4 on site B).
- Each node has 2 adapters (A and B).

- Each adapter has 4 ports where the port is named P *nodeid adapterid portnumber*. For example, port 3 on adapter B of node 4 is named P4B3.

Example 1: Zoning only

Each node in the system is configured in the following way:

- Four ports are used for host, controller, and local node connections.
- Four ports are used for replication.

The following zones are required for two sites, A and B:

Zone sets for site A:

- Zone for storage controllers and local node connections (Fibre Channel adapter HBA)
  - Zone 1 - [Node ports P1A1, P1A2, P2A1, P2A2 plus all storage controller ports in site A on switch 1]
  - Zone 2 - [Node ports P1A3, P1A4, P2A3, P2A4 plus all storage controller ports in site A on switch 2]
- Zones for host attachment (Fibre Channel adapter HBA A)
  - The same ports on the first adapter are zoned to be visible to hosts. It is recommended that each host is in a separate zone, and a host with two ports that are connected to two fabrics have two zones, one per fabric. The actual number of zones might be two multiplied by the number of hosts. A host with four ports that are connected to two fabrics should have four zones. In this case, the actual number of zones might be four multiplied by the number of hosts. In the four-port host case, the zones for one host might be:
    - Zone 3 - [host port, P1A1 and P2A1 for multipath fail over] on switch 1
    - Zone 4 - [host port, P1A2 and P2A2 for multipath fail over] on switch 1
    - Zone 5 - [host port, P1A3 and P2A3 for multipath fail over] on switch 2
    - Zone 6 - [host port, P1A4 and P2A4 for multipath fail over] on switch 2

    If the host only had two ports, it might have Zone 3 and Zone 5, or Zone 4 and Zone 6.
- Zones for replication (Fibre Channel adapter HBA B only):

  To ensure replication use the second Fibre Channel adapter exclusively, it is required that each local or remote port pair has its own zone:

  Zone 7 - [P1B1,P3B1]
  Zone 8 - [P1B1,P3B2]
  Zone 9 - [P1B1,P4B1]
  Zone 10 - [P1B1,P4B2]
  Zone 11 - [P1B2,P3B1]
  Zone 12 - [P1B2,P3B2]
  Zone 13 - [P1B2,P4B1]
  Zone 14 - [P1B2,P4B2]
  Zone 15 - [P2B1,P3B1]
  Zone 16 - [P2B1,P3B2]
  Zone 17 - [P2B1,P4B1]
  Zone 18 - [P2B1,P4B2]
  Zone 18 - [P2B1,P4B2]
  Zone 19 - [P2B1,P4B2]
  Zone 20 - [P2B2,P3B2]
  Zone 21 - [P2B2,P4B1]
  Zone 22 - [P2B2,P4B2]

Zone sets for site B:

This zone set is essentially a mirror image of the zoning for site A.

- Zone for storage controllers and local node connections (Fibre Channel adapter HBA A):
  - Zone 23 - [Node ports P3A1, P3A2, P4A1, P4A2 plus all storage controller ports in site B on switch 3]
  - Zone 24 - [Node ports P3A3, P3A4, P4A3, P4A4 plus all storage controller ports in site B on switch 4]
- Zones for host attachment (Fibre Channel adapter HBA A)
  - The same ports on the first adapter are zoned to be visible to hosts. Only zones for one host are listed. It is recommended that each host is in a separate zone. A host with two ports that are connected to two fabrics has two zones, one per fabric. The actual number of zones might be two multiplied by the number of hosts. A host with four ports that are connected to two fabrics should have four zones. In this case, the actual number of zones might be four multiplied by the number of hosts. In the four-port host case, the zones for one host might be:
    - Zone 25 - [host ports, P3A1 and P4A1 for multipath fail over] on switch 3
    - Zone 26 - [host ports, P3A2 and P4A2 for multipath fail over] on switch 3
    - Zone 27 - [host ports, P3A3 and P4A3 for multipath fail over] on switch 4
    - Zone 28 - [host ports, P3A4 and P4A4 for multipath fail over] on switch 4

    If the host had only two ports, it might have Zone 25 and Zone 27, or Zone 26 and Zone 28.
- Zones for replication (Fibre Channel adapter HBA B only):

  To ensure replication use the second Fibre Channel adapter exclusively, it is required that each local or remote port pair has its own zone:

  Zone 29 - [P1B3,P3B3]
  Zone 30 - [P1B3,P3B4]
  Zone 31 - [P1B3,P4B3]
  Zone 32 - [P1B3,P4B4]
  Zone 33 - [P1B4,P3B3]
  Zone 34 - [P1B4,P3B4]
  Zone 35 - [P1B4,P4B3]
  Zone 36 - [P1B4,P4B4]
  Zone 37 - [P2B3,P3B3]
  Zone 38 - [P2B3,P3B4]
  Zone 39 - [P2B3,P4B3]
  Zone 40 - [P2B3,P4B4]
  Zone 41 - [P2B4,P3B3]
  Zone 42 - [P2B4,P3B4]
  Zone 43 - [P2B4,P4B3]
  Zone 44 - [P2B4,P4B4]

Example 2: Port masking only

Assuming the same configuration as in example 1, you can achieve the same result by using port masking rather than zoning. The positions in the mask represent the Fibre Channel I/O port IDs with ID 1 in the rightmost position. These IDs are shown by the **lsportfc** command. In this example ports A1, A2, A3, A4, B1, B2, B3, and B4 correspond to FC I/O port IDs 1, 2, 3, 4, 5, 6, 7 and 8.

For remote partnership communication, apply a mask that limits communication to ports B1, B2, B3, and B4 by applying a port mask of 11110000 to both systems. Use the command `chsystem -partnerfcportmask 11110000`.

For local node-to-node communication, apply a mask that limits communication to ports A1, A2, A3, and A4 by applying a port mask of 00001111 to both systems. Use the command `chsystem -localfcportmask 00001111`.

It is then possible to have a single zone that contains all the node ports, as the port masking ensures that the number of logins is no more than 16.

With this setup, zoning for host connectivity and storage connectivity does not need to consider the rules about more than 16 logins.

Zoning examples

The following zoning examples describe ways for zoning a switch. In the examples, a list of port names that are inside brackets ([]) represent a single zone whose zone members are the list of ports shown.

Example 1

Consider the SAN environment in the following example:

- Two nodes (nodes A and B).
- Nodes A and B each have four ports.
  - Node A has ports A0, A1, A2, and A3.
  - Node B has ports B0, B1, B2, and B3.
- Two hosts that are called P and Q.
- Each of the two hosts has two ports, as described in the following table.

Table 45. Two hosts and their ports

P	**Q**
P0	Q0
P1	Q1

- Two switches that are called X and Y.
- Two storage systems I and J.
- Each of the two storage systems has ports as described in the following table.

Table 46. Two storage systems and their ports

I	**J**
I0	J0
I1	J1
I2	
I3	

The following tasks comprise an example configuration:

1. Attach half the host and node ports 1 (A0, A1, B0, B1, P0, Q0) to switch X.
2. Attach half the host and node ports 3 (A2, A3, B2, B3, P1, Q1) to switch Y.
3. Attach half the storage system ports (I0, I1, J0) to switch X.
4. Attach half the storage system ports (I2, I3, J1) to switch Y.
5. Create one zone per host port (one port per node) on switch X:

   [A0, B0, P0]
   [A1, B1, Q0]

Figure 32 on page 321 illustrates this example.

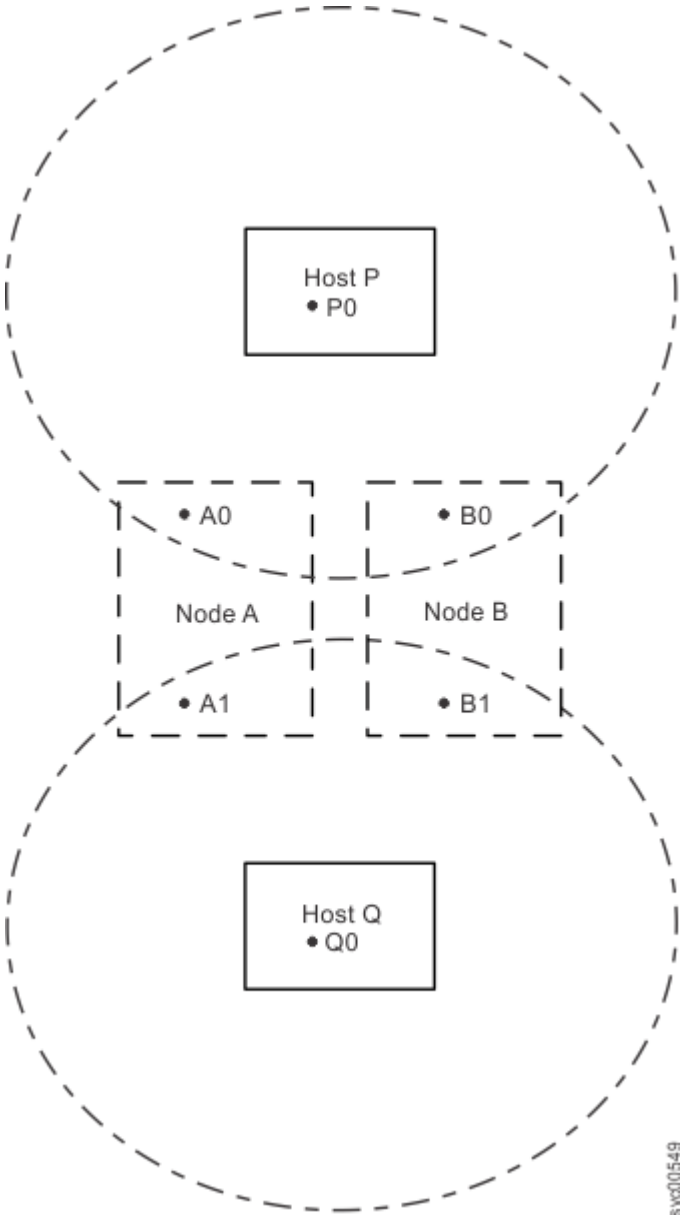


Figure 32. An example of a host zone

6. Create one storage zone per storage system on switch X:

   [A0, A1, B0, B1, I0, I1]
   [A0, A1, B0, B1, J0]

   Figure 33 on page 322 illustrates this example.

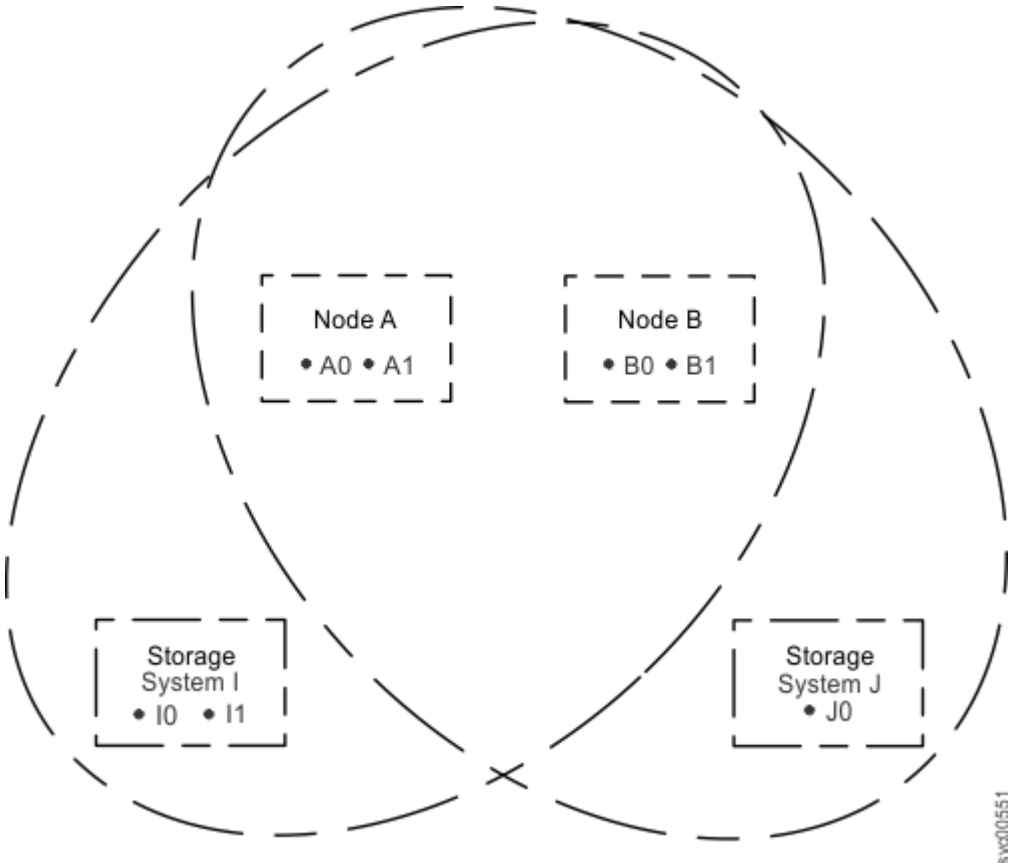


Figure 33. An example of a storage system zone

7. Create one internode zone on switch X:

   [A0, A1, B0, B1]

   Figure 34 on page 322 illustrates this example.

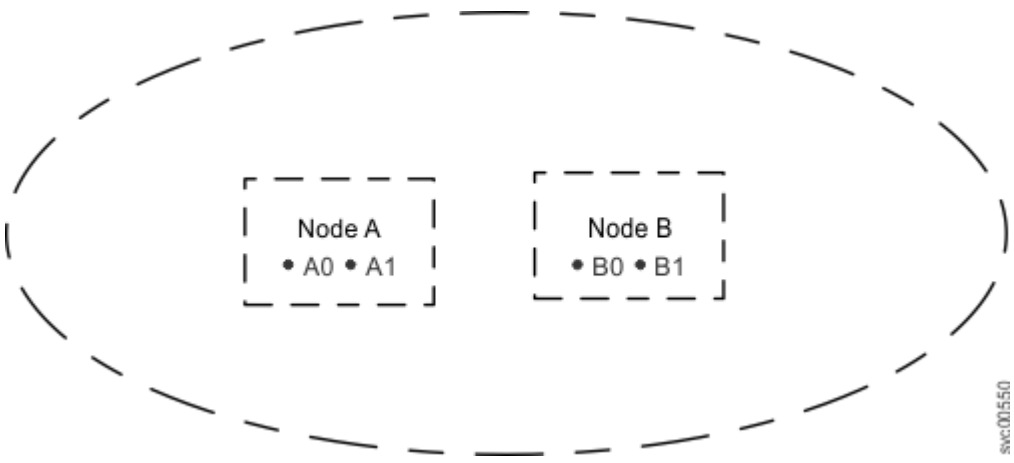


Figure 34. An example of a system zone

8. Follow the same steps "5" on page 320 through "7" on page 322 to create the following list of zones for switch Y:

   One zone per host port:

   [A2, B2, P1]
   [A3, B3, Q1]

   Storage zone:

[A2, A3, B2, B3, I2, I3]
[A2, A3, B2, B3, J1]

One internode zone:

[A2, A3, B2, B3]

Example 2

The following example describes a SAN environment that is like the previous example except for the addition of four hosts that have two ports each.

- Two nodes that are called A and B.
- Nodes A and B have four ports each.
  - Node A has ports A0, A1, A2, and A3.
  - Node B has ports B0, B1, B2, and B3.
- Six hosts that are called P, Q, R, S, T, and U.
- Four hosts have four ports each and the other two hosts have two ports each as described in the following table.

Table 47. Six hosts and their ports

P	**Q**	**R**	**S**	**T**	**U**
P0	Q0	R0	S0	T0	U0
P1	Q1	R1	S1	T1	U1
P2	Q2	R2	S2		
P3	Q3	R3	S3		

- Two switches that are called X and Y.
- Three storage systems
- Each storage system has ports as described in the following table.

Table 48. Three storage systems and their ports

I	**J**	**K**
I0	J0	K0
I1	J1	K1
I2		K2
I3		K3
		K4
		K5
		K6
		K7

The following tasks comprise a different example configuration:

1. Attach half of the host and node ports 1 (A0, A1, B0, B1, P0, P1, Q0, Q1, R0, R1, S0, S1, T0, U0) to switch X.
2. Attach half of the host and node ports 1 (A2, A3, B2, B3, P2, P3, Q2, Q3, R2, R3, S2, S3, T1, U1) to switch Y.
3. Attach half of the storage system ports (I0, I1, J0, K0, K1, K2, K3) to switch X.

4. Attach half of the storage system ports (I2, I3, J1, K4, K5, K6, K7) to switch Y.
5. Create one zone per host port (one port per node) on switch X:

   [A0, B0, P0]
   [A1, B1, P1]
   [A0, B0, Q0]
   [A1, B1, Q1]
   [A0, B0, R0]
   [A1, B1, R1]
   [A0, B0, S0]
   [A1, B1, S1]
   [A0, B0, T0]
   [A1, B1, U0]

   Attention: Hosts T and U (T0 and U0) and (T1 and U1) are zoned to different ports so that each port is zoned to the same number of host ports.
6. Create one storage zone per storage system on switch X:

   [A0, A1, B0, B1, I0, I1]
   [A0, A1, B0, B1, J0]
   [A0, A1, B0, B1, K0, K1, K2, K3]
7. Create one internode zone on switch X:

   [A0, A1, B0, B1]
8. Follow the same steps “5” on page 324 through “7” on page 324 to create the following list of zones for switch Y:

   One zone per host port:

   [A2, B2, P2]
   [A3, B3, P3]
   [A2, B2, Q2]
   [A3, B3, Q3]
   [A2, B2, R2]
   [A3, B3, R3]
   [A2, B2, S2]
   [A3, B3, S3]
   [A2, B2, T1]
   [A3, B3, U1]

   Storage zone:

   [A2, A3, B2, B3, I2, I3]
   [A2, A3, B2, B3, J1]
   [A2, A3, B2, B3, K4, K5, K6, K7]

   One internode zone:

   [A2, A3, B2, B3]

Fibre Channel port masking

With Fibre Channel port masking, you control the use of Fibre Channel ports. You can control whether the ports are used to communicate to other nodes within the same local system, and if they are used to communicate to nodes in partnered systems. Fibre Channel port masking does not affect host or storage traffic. It gets applied only to node-to-node communications within a system and replication between systems.

Usefulness of Fibre Channel port masking

- The setup of Fibre Channel port masks is useful when you have more than four Fibre Channel ports on any node in the system, as it saves setting up many SAN zones.
- Fibre Channel I/O ports are logical ports, which can exist on Fibre Channel platform ports.
- The system has two Fibre Channel port masks. The local port mask controls connectivity to other nodes in the same system, and the partner port mask control connectivity to nodes in remote, partnered systems. By default, all ports are enabled for both local and partner connectivity.
- The port masks apply to all nodes on a system; a different port mask cannot be set on nodes in the same system. You do not have to have the same port mask on partnered systems.

Fibre Channel port masking

Port numbers refer to the Fibre Channel I/O port IDs that are shown by the **lsportfc** command.

A port mask is a string of zeros and ones. The last digit in the string represents port one. The previous digits represent ports two, three, and so on. If the digit for a port is "1", the port is enabled and the system attempts to send and receive traffic on that port. If it is "0", the system does not send or receive traffic on the port. If there are not sufficient digits in the string to specifically set a port number, the port is disabled for traffic.

For example, if the local port mask is set to "00001100" on a node with eight Fibre Channel ports, ports 3 and 4 are able to connect with other nodes in the system. The other Fibre Channel ports are not able to be used for node-to-node communication.

Note: The **lsfabric** CLI command shows all of the paths that are possible on the system (as defined by zoning) independent of their usage. Therefore, the command output includes paths that will not be used because of port masking.

The port masks are set by using the **chsystem** CLI command.

Port masking example

The following table shows recommended port masking for systems with 4, 8, and 12 Fibre Channel ports. The number of 0s and 1s in the port mask examples match the number of ports.

Table 49. Port masking configuration

Card / Port	4 ports	8 ports	12 ports
Card 1 Port 1	Host/Storage/Inter-node	Host/Storage	Host/Storage
Card 1 Port 2	Host/Storage/Inter-node	Host/Storage	Host/Storage
Card 1 Port 3	Host/Storage/ Replication*	Inter-node	Inter-node
Card 1 Port 4	Host/Storage/ Replication*	Inter-node	Inter-node
Card 2 Port 1		Host/Storage	Host/Storage
Card 2 Port 2		Host/Storage	Host/Storage
Card 2 Port 3		Host/Storage/ Replication*	Host/Storage/ Replication*
Card 2 Port 4		Host/Storage/ Replication*	Host/Storage/ Replication*
Card 3 Port 1			Host/Storage

Table 49. Port masking configuration (continued)			
Card / Port	**4 ports**	**8 ports**	**12 ports**
Card 3 Port 2			Host/Storage
Card 3 Port 3			Host/Storage
Card 3 Port 4			Host/Storage
localfcportmask	0011	00001100	000000001100
partnerfcportmask	1100	11000000	000011000000
* Use for host or storage when no replication is in place.			

Fibre Channel SAN configuration details

Apply the following configuration details for Fibre Channel, Fibre Channel switches (FC Gateway, FCF) switches to ensure that you have a valid configuration.

Configuring your SAN with at least two independent switches, or networks of switches, ensures a redundant fabric with no single point of failure. If one of the two SAN fabrics fails, the configuration is in a degraded mode, but is still valid. Maintain separate fabrics for FC. If you attempt to combine these fabrics, you might risk adding more paths to the volumes. Supported configurations allow a maximum of eight paths. A SAN with only one fabric is a valid configuration but risks loss of access to data if the fabric fails. SANs with only one fabric are exposed to a single point of failure.

For Fibre Channel connections, the nodes must be connected to either SAN switches or directly connected to a host port. The system requires that a minimum of two Fibre Channel ports from each node to be connected to the SAN. This configuration provides connections to each of the counterpart SANS that are in the redundant fabric. When iSCSI hosts are attached to nodes, Ethernet switches must be used.

All back-end storage systems must always be connected to SAN switches only. Multiple connections are permitted from redundant storage systems to improve data bandwidth performance. A connection between each redundant storage system and each counterpart SAN is not required. For example, in an IBM DS4000 configuration that contains two redundant storage systems, only two storage system minihubs are usually used. Storage system A is connected to counterpart SAN A, and storage system B is connected to counterpart SAN B. Any configuration that uses a direct physical connection between the node and the storage system is not supported.

When you attach a node to a SAN fabric that contains core directors and edge switches, connect the node ports to the core directors. Then, connect the host ports to the edge switches. In this type of fabric, the next priority for connection to the core directors is the storage systems, leaving the host ports connected to the edge switches.

A SAN must follow all switch manufacturer configuration rules, which might place restrictions on the configuration. Any configuration that does not follow switch manufacturer configuration rules is not supported.

Fibre Channel switches and interswitch links

The system supports distance-extender technology, including dense wavelength division multiplexing (DWDM) and Fibre Channel over IP (FCIP) extenders, to increase the overall distance between local and remote systems. If this extender technology involves a protocol conversion, the local and remote fabrics are regarded as independent fabrics, limited to three ISL hops each.

With ISLs between nodes in the same system, the inter-switch links (ISLs) are considered a single point of failure. Figure 35 on page 327 illustrates this example.

- If Link 1 or Link 2 fails, the system communication does not fail.
- If Link 3 or Link 4 fails, the system communication does not fail.

- If ISL 1 or ISL 2 fails, the communication between Node A and Node B fails for some amount of time. The node is not recognized, even though the nodes still have a connection.

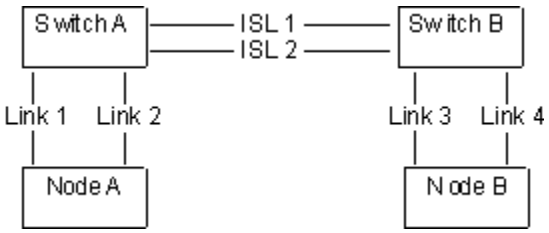


Figure 35. Fabric with ISL between nodes in a system

To ensure that a Fibre Channel link failure does not cause nodes to fail when ISLs are between nodes, you must use a redundant configuration. This configuration is illustrated in Figure 36 on page 327. With a redundant configuration, if any one of the links fails, communication on the system does not fail.

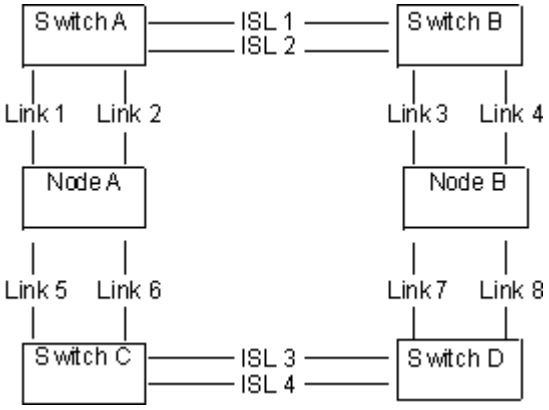


Figure 36. Fabric with ISL in a redundant configuration

Fibre Channel servers and system connections to the existing Fibre Channel SAN

The following examples show the various supported configurations.

Figure 37 on page 328 shows a system that is connected to a Fibre Channel forwarder switch along with any FC hosts and FC storage systems. The connections are 10 Gbps Ethernet. The Fibre Channel forwarder is linked to the existing Fibre Channel SAN by using Fibre Channel ISLs. Any Fibre Channel hosts or storage systems remain on the existing Fibre Channel SAN. The connection to the system can be via the SAN (if the system is connected via Fibre Channel).

1.

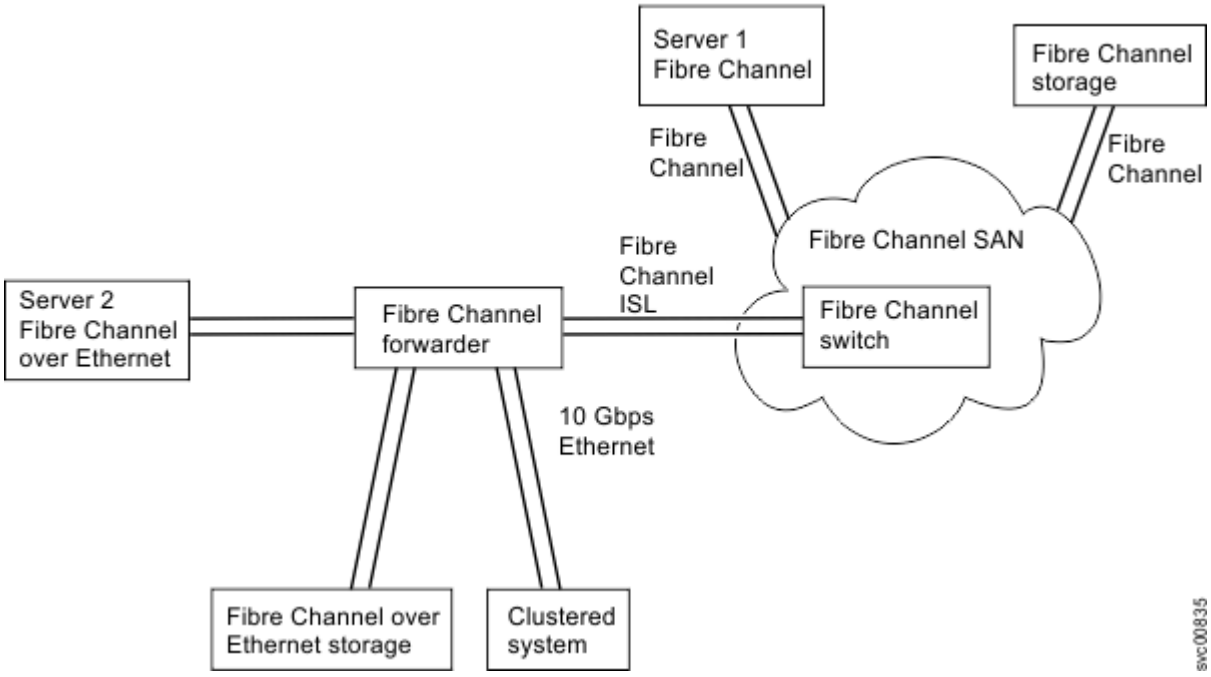


Figure 37. Fibre Channel forwarder linked to existing Fibre Channel SAN

The second example, Figure 38 on page 328, is almost the same as the first example but without an existing Fibre Channel SAN. It shows a system that is connected to a Fibre Channel forwarder switch along with any FC hosts and FC storage systems. The connections are 10 Gbps Ethernet.

2.

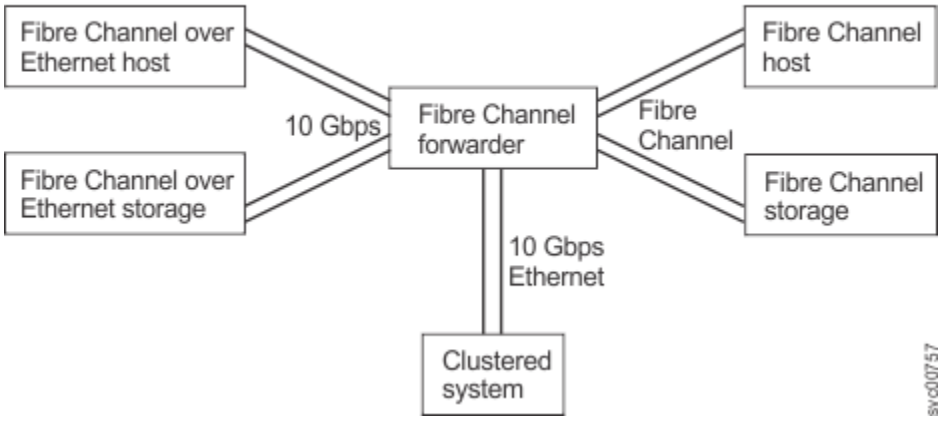


Figure 38. Fibre Channel forwarder linked to hosts and storage systems without an existing Fibre Channel SAN

In the third example, Figure 39 on page 329, a Fibre Channel host connects into the Fibre Channel ports on the Fibre Channel forwarder. The system is connected to a Fibre Channel forwarder switch along with any FC storage systems. The connections are 10 Gbps Ethernet. The Fibre Channel forwarder is linked to the existing Fibre Channel SAN by using Fibre Channel ISLs. Any Fibre Channel hosts or storage systems remain on the existing Fibre Channel SAN. The FC host connects to a 10 Gbps Ethernet switch (transit switch) that is connected to the Fibre Channel forwarder.

3.

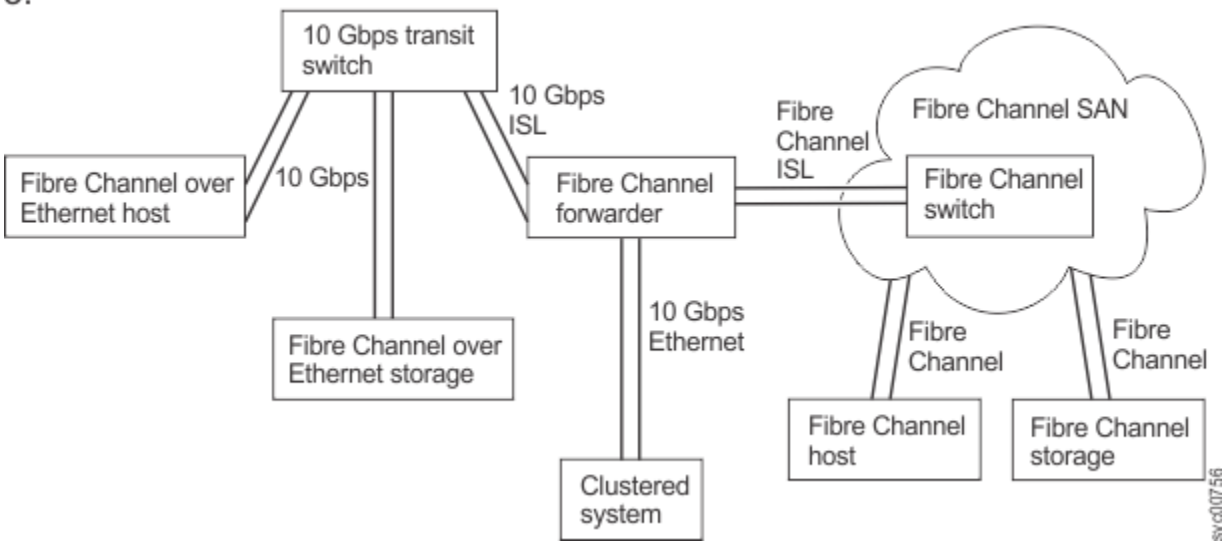


Figure 39. Fibre Channel host connects into the Fibre Channel ports on the Fibre Channel forwarder

The fourth example, Figure 40 on page 329, is about the same as the previous example but without an existing Fibre Channel SAN. The Fibre Channel hosts connect to Fibre Channel ports on the Fibre Channel forwarder.

4.

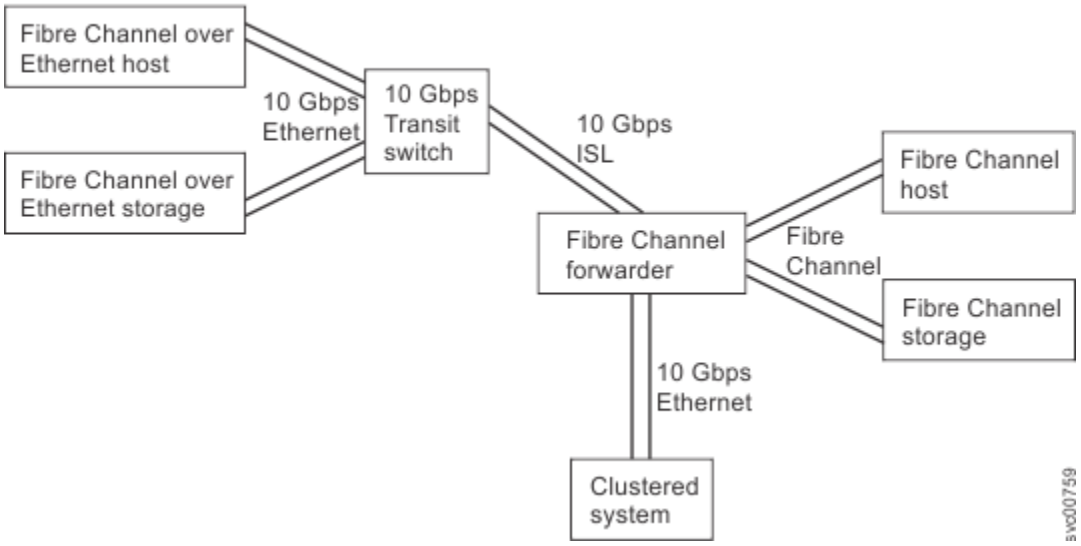


Figure 40. Fibre Channel host connects into the Fibre Channel ports on the Fibre Channel forwarder without an existing Fibre Channel SAN.

ISL oversubscription

Complete a thorough SAN design analysis to avoid ISL congestion. Do not configure the SAN to use system to system traffic or a system to storage system traffic across ISLs that are oversubscribed. For host to system traffic, do not use an ISL oversubscription ratio that is greater than 7 to 1. Congestion on the ISLs can result in severe performance degradation and I/O errors on the host.

Example SAN configurations

Examples of typical ways to configure your system and an associated Fibre Channel network are provided to guide you in setting up your SAN configuration.

SAN fabric function is available only if a Fibre Channel option is installed. Otherwise, this information does not apply.

Note: The illustrations in these examples do not show expansion enclosures as they have no relationship to SAN connectivity.

Figure 41 on page 330 illustrates a small SAN configuration. Two Fibre Channel switches are used to provide redundancy. The control enclosures and each host system are connected to both Fibre Channel switches. If you are using external storage systems with the system, connect as shown in this illustration.

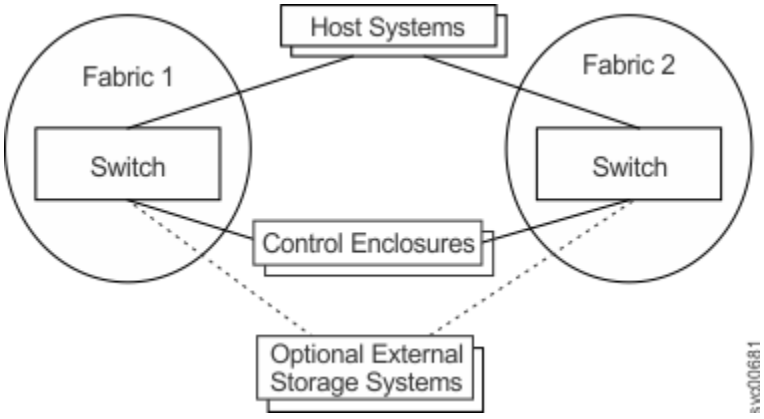


Figure 41. Simple SAN configuration

Figure 42 on page 330 illustrates a medium-sized configuration with external storage systems. The Fibre Channel SAN fabric consists of switches that are interconnected with interswitch links (ISLs). For redundancy, connect each control enclosure and external storage system to two fabrics. The example fabric attaches the control enclosures and the storage systems to the core switch. No ISL hops exist between node canister ports or between the control enclosures and the external storage systems. The example fabric attaches the control enclosure and the storage systems to the core switch. No ISL hops exist between node canister ports or between the control enclosure and the external storage systems.

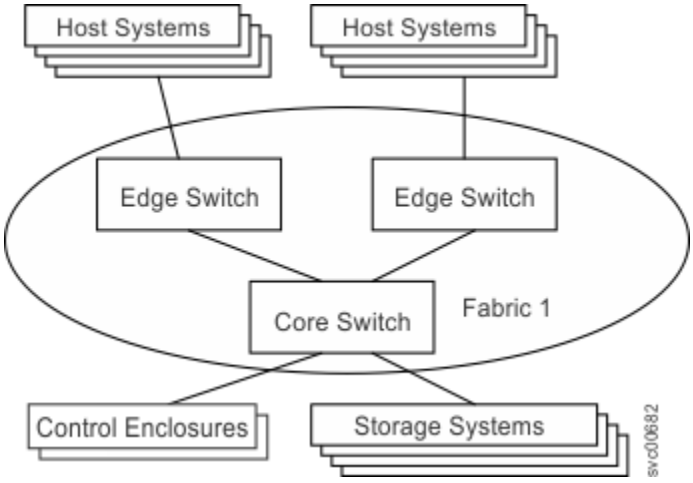


Figure 42. SAN configuration with a medium-sized fabric

Figure 43 on page 331 illustrates a large fabric that consists of two core Fibre Channel switches and edge switches that are interconnected with ISLs. For redundancy, use two fabrics with each host system, node, and storage system that is being connected. Both fabrics attach the nodes to both core fabrics and distribute the storage systems between the two core switches. This configuration ensures that no ISL hops exist between nodes or between nodes and the storage systems.

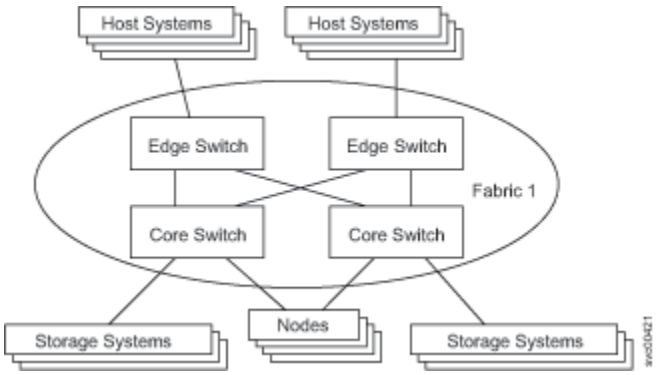


Figure 43. SAN configuration with a large fabric

Figure 44 on page 331 illustrates a fabric where the host systems are at two different sites. A long-wave optical link is used to interconnect switches at the different sites. For redundancy, use two fabrics and at least two separate long-distance links. If many host systems are at the remote site, use ISL trunking to increase the available bandwidth between the two sites.

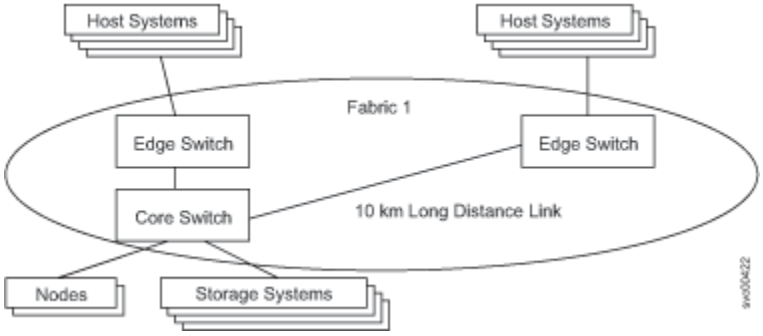


Figure 44. SAN configuration across two sites

System configuration by using SAN fabrics with long-distance fiber connections

Each system that uses SAN fabric switches can use shortwave or longwave optical fiber connections to connect to application hosts, storage systems, or remote systems.

The maximum distance between the system and host or the system and the storage system depends on the Fibre Channel speed and cable type. For all speeds and cable types, the maximum distance is 300 m for shortwave optical connections and 10 km for longwave optical connections.

When you use longwave optical fiber connections, follow these guidelines:

- For disaster recovery, each system must be regarded as a single entity, which includes the storage system that provides the quorum disks for the system. Therefore, the system and quorum disks must be co-located.
- The maximum distance between nodes within a system is 100 meters. There can be a large fiber-cable distance between the same nodes in the system. However, the nodes must be physically co-located for effective service and maintenance.

  For example, the nodes can be 300 meters away for 2 Gbs connections. For 4 Gbs connections, nodes can be 150 meters away from the SAN fabric. In both configurations, the maximum cable distance between the nodes is 600 meters. However, the nodes must be physically within 100 meters of each other.
- All nodes in a system must be on the same IP subnet so that the nodes can assume the same system or service IP address.

Determining and changing the Fibre Channel network speed

The speed of the Fibre Channel (FC) ports is auto-negotiated with the FC switch or host to which the ports are connected.

Determining the FC port speed

The ports operate at the highest speed at which a reliable connection can be established. Different ports and nodes in a system can operate at different speeds.

You can use several methods to determine the port speed of a FC port on a node.

Determining the FC port speed with the management GUI

If the node is in a system that is operational, you can determine the FC port speed with the management GUI.

To view the FC port speed in the management GUI, complete these steps:

1. Select **Monitoring** > **System**.
2. On the **System -- Overview** page, select the directional arrow next to the node that contains the FC port.
3. On the **Node Details** page, under the rear view of the node, hover over the FC port to display details about the port including the port speed. You can also select **Fibre Channel Port** to display port information for all configured ports in the **Component Details** tile.

Determining the FC port speed with the service assistant GUI

You can determine the FC port speed with the service assistant GUI.

To view the FC port speed in the service assistant GUI, complete these steps:

1. Point your web browser to `serviceaddress/service` for the node that you want to work on.
2. Log on to the service assistant using the superuser password.
3. Select the **Port** tab under **Node Details** to view the speed of the FC port.

Determining the FC port speed from the CLI

If the node is in a system that is operational, you can determine the FC port speed using the CLI commands.

To view the FC port speed of a system, issue the following command:

```
svcinfo lsportfc 0
```

The following output is an example of what you might see. The FC port speed is shown in bold.

```
id 0
fc_io_port_id 1
port_id 1
type fc
port_speed 16Gb
node_id 1
node_name node1
WWPN 500507680C212C72
```

HTTP proxy server

As part of Call Home with cloud services configuration, you can choose to use your existing HTTP proxy server within a network to manage connections between the system and the support center. A HTTP or

internal proxy server manages connections between your internal network and any entity outside of your network.

Call Home with cloud services or remote support assistance requires connections to the support center for quick and efficient problem resolution for your system. Instead of opening ports for these external connections through your firewall, you can specify an internal proxy server for better security. In addition to these features, an internal proxy server can be defined to connect to cloud-based services, such as an authentication service provider. To use an internal proxy server, a DNS server is also required to resolve hostnames to IP addresses. To create a DNS server, select **Settings** > **Network** > **DNS**.

To define an internal proxy server, use the **Settings** > **Network** > **Internal Proxy Server** panel in the management GUI.

To define an internal proxy server using the command-line interface, see **mkdnsserver** and **mkproxy** commands.

Portsets

Portsets manage traffic on the Ethernet and Fibre Channel (FC) ports in the system for connectivity to hosts, storage, partnered systems (replication), management applications, and IP quorum applications. A portset is a system-wide object that restricts a host or remote device to communicate only through the ports and logical addresses that are contained in that portset.

Fibre Channel portsets

A Fibre Channel portset contains:

- A collection of physical FC port locations (FC I/O port IDs) on the nodes.
- One or more host objects (FCP-SCSI or FC-NVMe).

When a host is mapped to an FC portset, the system communicates with the host only through the FC ports that are members of that portset.

Auto zoning can be enabled on a Fibre Channel portset. For more information, see .

Ethernet portsets

An Ethernet portset contains:

- A collection of node IP addresses, each associated with a physical Ethernet port location (and VLAN or prefix as applicable).
- And either:
  - One or more Ethernet host objects (host portset), or
  - One or more Storage controller objects (storage portset), or
  - A collection of an IP partnership link object (replication portset), or
  - One or more management IP address (System IP and partition management portset)
  - An quorum IP address (in a quorum portset)

The system communicates to the remote devices through the local IP addresses contained within the portset.

Default and reserved portsets

By default, a system initializes the following portsets (IDs are reserved and used by the system):

Ethernet Host portset (ID 0)

- This portset contains all IP addresses configured on the node ports.

- This portset defines the connectivity for any Ethernet (iSCSI, NVMe-TCP, NVMe-RDMA) hosts that are not included in user-defined portsets.

Fibre Channel Host portset (ID 64)

- Initially, this portset contains all Fibre Channel ports in the system.
- This portset defines the connectivity for any Fibre Channel hosts that are not included in user-defined portsets. FC-NVMe hosts must be associated with a user-defined portset, and thus this portset does not apply.
- During IP partnership configuration, new FC ports are added to 64 by default. For more information, see "IP partnership requirements" on page 339.

Ethernet management portset (ID 72)

- The management portset provides system management IP addresses for interfaces such as the GUI, CLI, and REST API.

Ethernet quorum portset (ID 105)
Used for IP quorum communication. This is not created by default.

Ethernet replication portsets (IDs 1 and 2)
Used for IP partnerships.

Ethernet storage portset (ID 3)
Used to define connectivity to iSCSI storage.

Additional fibre-channel host portsets are created during system setup. Host and quorum portsets can be assigned to an ownership group to control management of the portset. For more information, see "Ownership groups" on page 211.

Configuring Ethernet portsets

Consider the following requirements when configuring ethernet portsets.

Each IP address in a portset must be configured on a separate Ethernet port. A portset can contain a mix of IPv4 and IPv6 addresses unless otherwise stated in the following concepts pages.

- Chapter 5, "Hosts ," on page 107, for "Host portsets" on page 110
- "IP partnership requirements" on page 339, for Replication portset
- "Quorum" on page 502, for Quorum portset
- "System IP addresses" on page 301, for Management portset

You can use the management GUI or the command-line interface (CLI) to create a portset.

Creating Ethernet portsets by using the management GUI

To create an Ethernet portset in the management GUI, complete the following steps:

- To create or manage an Ethernet portset in the management GUI, select **Settings** > **Network** > **Ethernet Portsets**.
- To create or assign IP addresses to the portset, select **Settings** > **Network** > **Ethernet Ports**.
- Verify that each port displays the **Configured** state after it has been added to the portset.
  - Right-click the port and select either **Modify Replication**, **Modify iSCSI hosts**, or **Modify Storage Ports**. The traffic type for the port must match the traffic for the portset that you created.

Creating portsets by using the command-line interface

To configure Ethernet portsets and assign IP addresses, use the following commands:

- mkportset
- mkip

Step 1: Create a portset

Run the following command:

```
mkportset -name portset_name -type portset_type -ownershipgroup owner_id | owner_name
-replicationportsetlinkuid link_uid
```

Step 2: Assign IP addresses to the portset

After you create the portset for the traffic type, run:

```
mkip -node node_name -port port_id -portset portset_id | portset_name -ip <x.x.x.x> -gw
gateway -vlan vlan_id -prefix subnet_prefix
```

Configuring Fibre Channel portsets

A **Fibre Channel portset** defines a group of Fibre Channel I/O ports that a host can use for connectivity. Fibre Channel portsets are used only for host attachment. By assigning hosts to specific portsets, you can control which Fibre Channel ports they use to access volumes.

You can create multiple Fibre Channel portsets to isolate traffic for different hosts or workloads. Multiple hosts can be mapped to the same portset; however, a host can only be mapped to one portset.

Each portset can include up to four Fibre Channel I/O ports, which are defined using their port ID. The port ID applies across all nodes in the system. A Fibre Channel I/O port can belong to multiple portsets. A portset must include at least one port before a host can be assigned to the portset.

If a Fibre Channel host is not explicitly mapped to a portset, it is automatically mapped to the default portset to maintain connectivity. FC-NVMe hosts can only be mapped to user-defined portsets.

You can enable auto zoning to simplify zoning configuration. For more information, see .

You can use the management GUI or the command-line interface (CLI) to create a portset.

Creating portsets by using the management GUI

To create Fibre Channel portsets, complete the following tasks:

1. **Create a Fibre Channel portset**
   You can create a new Fibre Channel portset by navigating to **Settings** > **Networks** > **Fibre Channel Portsets** in the management GUI. Select **Create Portset** and enter the required configuration details. After the portset is created, it is available for host assignment and connectivity configuration.
2. **Assign Fibre Channel ports to a portset**
   After creating a portset, assign ports from the **Fibre Channel Ports**. Use the **Assign Portset** option from the overflow menu to associate the selected ports with the portset. The system then updates the configuration so that these ports participate in the portset's zoning and connectivity.
3. **Create a host and assign it to a portset**
   After configuring portsets and assigning ports, you can assign hosts or host clusters to the appropriate portset.

Creating portsets by using the command-line interface

To create a Fibre Channel portset, complete the following steps in the command-line interface:

1. To create a Fibre Channel portset, enter the following command:

   ```
   mkportset <name> -porttype fc
   ```

   Where **<name>** is the name of the portset.
2. To associate a Fibre Channel IO port ID with portset, enter the following command:

   ```
   addfcportsetmember -portset <portset_name | portset_id> -fcioportid <fc_io_port_id>
   ```

where *<portset_name | portset_id>* is the name or ID of the portset and *<fc_io_port_id>* is the ID for the Fibre Channel IO port that you want to assign the portset to.

3. To verify that the mapping between the Fibre Channel IO port ID and the associated portset, enter the following command:

```
lsfcportsetmember
```

In the results that display, verify that the portset is assigned to the correct Fibre Channel IO port ID. In the following example, the Fibre Channel IO port ID 1 was assigned to portset4:

```
id   fc_io_port_id portset_id portset_name owner_id owner_name
 4    1             4          portset4     0        ownershipgroup0
```

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.

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 “Partnerships using IP Connectivity” on page 339.

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

Replication requires REST API connectivity between the partnered systems. Replication setup using the management GUI requires management GUI connectivity between the systems. Ensure that firewalls between the systems allow traffic to ports 6443, 7443, and 443 on the system management IP addresses.

Note: The communication on ports 6443, 7443, and 443 between the two partnered systems must occur on the default management IP address of each system.

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 “Configuring system certificates” on page 393.

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 replication layer by default, but the system can be configured to be in the storage 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 .

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.

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.

-
-
-

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

Connect to the GUI on either system. Select **Settings** > **System** > **Partnerships** and select **Create Partnership**.

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 that are configured by 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.

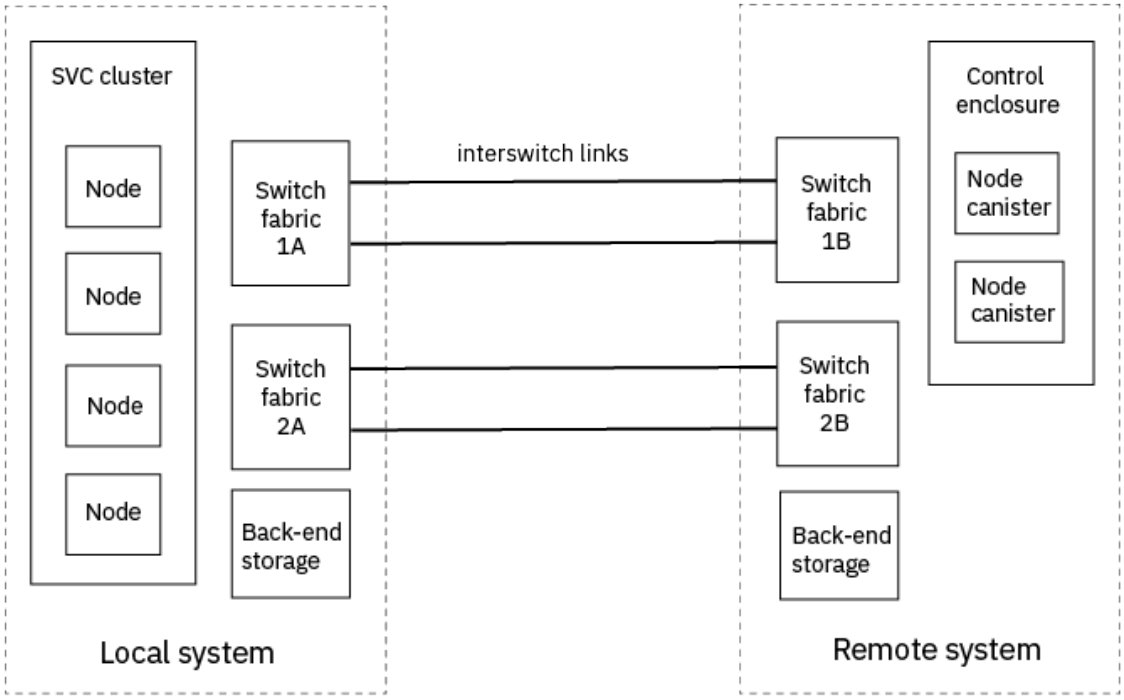


Figure 45. Redundant fabrics

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 or Private ports

Best practice is to ensure that cluster ports performing node to node traffic are separated 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.

Note: Policy-based high availability environments require isolation between public and private traffic to protect performance. Dedicated ISLs or switch-based QoS can be used to prioritize private traffic and maintain isolation.

Note: IBM FlashSystem does not support systems with more than one I/O group. IBM SAN Volume Controller systems continue to support multiple I/O groups on long-term support releases only.

Policy-based replication partnership 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.

Host zoning

For policy-based high availability, hosts need to be zoned so they can see the host ports from both systems. If hosts are created with auto zoning enabled Fibre Channel portset on primary site, configuring policy-based high availability links each host to its partner system and automatically creates the required zones.

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.

Partnership over native IP brings with it changes to some fundamental assumptions in that no name server is in an Ethernet fabric to detect and discover replication ports. There are specific configuration settings such that the local and remote systems can discover each other's replication capable IP addresses. TCP/IP sessions can then be set up across those IP addresses to establish a partnership.

Partnerships over IP Connectivity are provided as follows:

- “Short-distance partnerships using RDMA” on page 348.
- Long-distance partnerships using TCP.
- “Long-distance Ethernet partnerships using TCP” on page 345

IP partnership requirements

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

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).
- Ethernet portsets group Ethernet ports to control port usage, assign IP addresses, and define host connectivity policies.

- 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.
- IBM SAN Volume Controller systems continue to support multiple I/O groups on long-term support releases only.
- If your system supports Ethernet ports of several speeds, **only one** of those speeds can be used to configure all of the replication 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 replication 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 replication port group for that link.
- No more than two intersite links are supported.
- 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:

- Use a different port for IP partnership traffic, which is not used for iSCSI host, iSCSI virtualization, iSCSI backend, and RDMA clustering. Also, use a different VLAN ID for iSCSI host I/O and IP partnership traffic.

Requirements for replication portsets

Replication portsets support IP partnerships between systems. These requirements are specific to replication portsets:

- Replication portset can have a maximum of one IP address per node.
- All IP addresses in the same replication portsets must be either IPv4 or IPv6 addresses. You cannot mix IP protocol versions on replication portsets.
- Each IP partnership can be mapped to two portsets, one for each link between systems.

  Single link partnership

  For a partnership with a single link, a single portset can be defined in the **portset1** field on the **Create Partnership** in the management GUI.

  Dual link partnership

  For a partnership with dual links, a second portset must be defined in the **portset2** field in the management GUI.

  For a partnership with a single link, a single portset can be defined in the **portset1** field on the **Create Partnership** page in GUI.

- Portsets replace the requirement for creating replication groups for IP partnerships. During updates of the software, any IP addresses that are assigned to groups with an existing IP partnership are automatically moved to a corresponding portset. For example, if group 1 is defined on the system before the update then IP addresses from that replication group are mapped to portset 1 after the update. Similarly, the IP address in group 2 is mapped to portset 2.
- A system can be part of up to three IP partnerships.
- You can create IP partnerships between systems that are in the same layer. In other words, both systems in the 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 343.
- 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.
- On systems that support IP compression:
  - 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.
- Review the appropriate hardware guide to understand which ports can be used for IP partnerships.
- Using an Ethernet switch to convert a 25 Gbps to a 1 Gbps IP partnership, or a 10 Gbps to a 1 Gbps partnership is not supported.
- The IP infrastructure on both partnership sites must match.
- When an IP partnership is configured between two systems with one system on an older build and system management IP is configured on the port other than the default port in a newer build. The **lssystemip** command on the older build shows stale port_id value as 1 and 2 for remote system. Until the other system is upgraded with latest build, it continues to show stale port_ids. To get actual port_id, **lsip** can be used in the remote system.
- For more information, see "Portsets" on page 333.

Requirements for short-distance partnerships using RDMA

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

- Up to two short-distance partnership using RDMA can be configured per system.
- 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 Gbps or 25 Gbps or 100 Gbps Ethernet ports with RDMA (iWARP) capabilities to establish a short-distance partnership using RDMA between systems. Support may depend on the specific hardware model. For model-specific requirements and restrictions, see supported hardware documentation for your product.
- 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 does 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.
- 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.

Asymmetric Ethernet port configurations (100 Gbps – 25 Gbps) considerations and risks

Asymmetric high-speed Ethernet (100 Gbps – 25 Gbps) configurations are supported in the following scenarios:

- The host write bandwidth requirement aligns with the effective I/O bandwidth.
- Partnership bandwidth is limited by the minimum port speed or link bandwidth, which works as the limiting factor for total bandwidth consumption. For details, see the *link_bandwidth_mbits* attribute of the **lspartnership** CLI command.
- Quality of service (QoS) mechanisms, such as priority flow control, are used to regulate bandwidth demand.

Asymmetric 100 Gbps – 25 Gbps configurations are suitable for short-term or transitional use. Long-term deployments are not recommended. Because the uplink and downlink speeds differ, the configuration can lead to bandwidth oversubscription and increase the risk of congestion during sustained or burst workloads.

In an oversubscribed environment, effective I/O bandwidth is limited by the lowest port-to-port speed in the data path. As a result, overall throughput is driven by the slowest link. If host I/O generation exceeds this limit, congestion occurs. Without appropriate host-side I/O throttling or flow-control mechanisms,

traffic continues to exceed available network capacity, which can worsen congestion. Oversubscription leads to packet drops and retransmissions, which results in higher latency, reduced throughput, and degraded performance. In certain conditions, such as during failure or recovery events, these effects can influence high-availability behavior and service continuity.

Although asymmetric high-speed Ethernet configurations can support temporary transitions, careful alignment of bandwidth, I/O rate limits, and congestion-management settings is required to reduce performance and availability risks.

Related tasks

"Planning secured IP partnerships" on page 343
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.

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.

Before you begin

To enable a secured IP partnerships, install an encryption license on both of the partnered systems. See "Licensing encryption" on page 457 for more information, or contact support for assistance.

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 "Configuring system certificates" on page 393.

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 or creating an internally signed certificate" on page 394.
- 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" on page 397.

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.

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.
- 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.

Ensure that the following network requirements are met to configure IP replication successfully:

- These TCP ports must be open for inbound and outbound access:

  3260
  Required for system management communication.

  3265
  Required for data path connections.
- For secured IP partnerships, ensure that the following ports are also unblocked

  500
  Used by the internet key exchange (IKE) protocol to manage encryption keys.

  4500
  Essential for secure VPN communications.

  IPSEC-ESP
  Required for application-based firewalls.

Long-distance Ethernet partnerships using TCP

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 “Licensing encryption” on page 457.

Portsets replace the requirement for creating groups for IP partnerships. Dedicated portsets can be created for replication 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 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.

Note:

IBM Storage FlashSystem does not support systems with more than one I/O group. IBM SAN Volume Controller systems continue to support multiple I/O groups on long-term support releases only. Chapter 10, “FlashSystem grid,” on page 235 can be used to scale out a storage system.

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

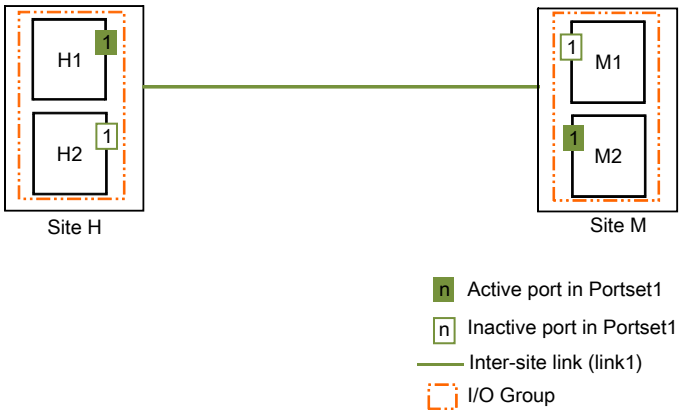


Figure 46. One intersite link, one I/O group per system

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. Replication 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, 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.

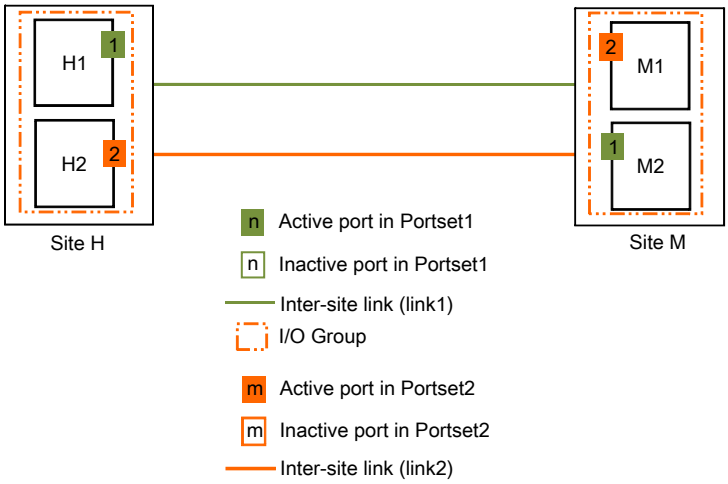


Figure 47. Two intersite links, one I/O group per system

Intersite link planning

Specific intersite link requirements must be met when you are planning to use IP partnership for Policy-based replication or Policy-based HA.

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

- On systems that support IP compression, 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 "Creating IP partnership" on page 349.
- 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 48 on page 347 describes the scaling of host I/O.

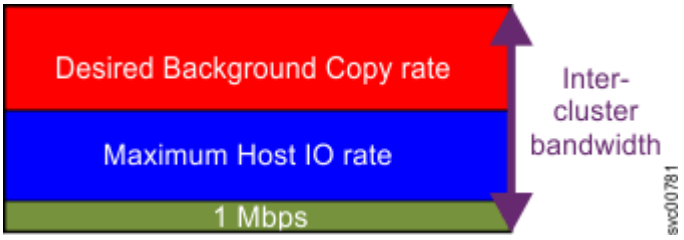


Figure 48. Scaling of host I/O

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.

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

Figure 49 on page 348 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.

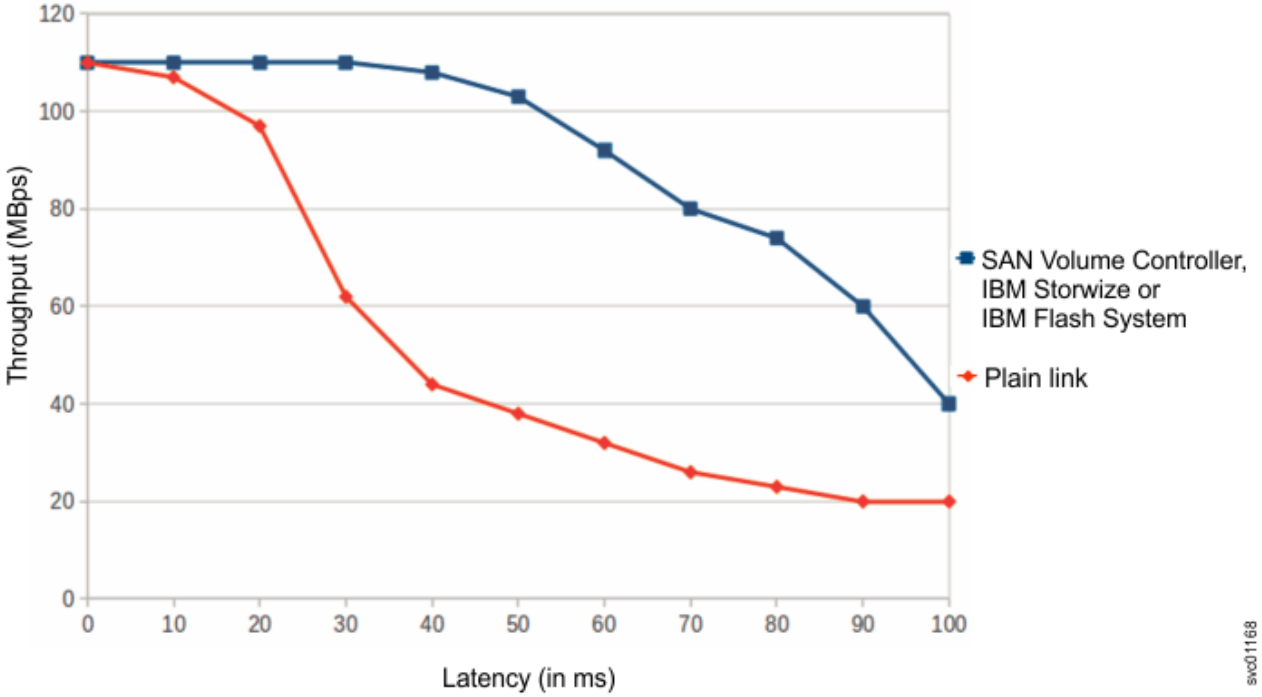


Figure 49. Product transfer rate comparison for systems built with IBM Storage Virtualize

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 339 for additional considerations.

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

Short-distance partnerships using RDMA

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.

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 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.

Network design requirements
Partnerships using RDMA require a network design that provides near-lossless Ethernet characteristics and avoids oversubscription on the replication path.

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.

Partnerships using RDMA on high-speed Ethernet networks require a network design that provides near-lossless Ethernet characteristics and avoids oversubscription on the replication path.

For short-distance partnerships that uses 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 the IP addresses within a *highspeedreplication* type portset must be of the identical type that is, either IPv4 or IPv6.

Note: When a node is added to a system configured for a short-distance partnership using RDMA, any clustering or node IP addresses configured on the node are removed.

Creating IP partnership

After creating portset and assigning IP addresses to the portset, you have to create IP partnership for replication.

Using the management GUI

To create IP partnership using the management GUI, complete these steps:

1. Select **Settings** > **System** > **Partnerships** and select **Create Partnership**.
2. On the **Create Partnership** page, select either IP (long distances using TCP) or IP (short distances using RDMA).
3. Enter the following information to configure the partnership:

   Partner IP/FQDN
   Select the IP address or the fully qualified domain name (FQDN) of the partner system. 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**.

   Secured IP partnerships

   Enter the following information to obtain the certificate:

   a. Select **Secured IP partnerships** to secure the connection between systems.

b. Click **Test Connection** to retrieve and validate the authorized-signed certificate of the remote system.

c. Select **View certificate** to display the certificate from the remote system. Verify that the certificate is signed by a trusted certificate authority, and that the subject and issuer names are the expected values.

d. If the retrieved certificate is signed by an unrecognized authority, select **Upload File** to upload the root certificate or certificate chain from the certificate authority that signed the partner system's certificate.

Secured IP partnerships support Both internally signed certificates and externally signed certificates.

Link Bandwidth
Enter the amount of bandwidth used for replication between systems in the partnership.

Partner System's CHAP Secret (Optional):
Specify the CHAP secret for the partner system if you plan to use Challenge Handshake Authentication Protocol (CHAP) to authenticate connections between the systems in the partnership.

Compression enabled
On systems that support IP compression, select this option if data on the local system should be compressed prior to transfer to the partner system. To fully enable compression in an IP partnership, each system must support compression.

Portset Link 1
In single link configurations between partnered systems, select the portset that provides IP addresses for replication traffic.

Portset Link 2 (optional)
In dual link configurations between partnered systems, select the portset that provides IP addresses for replication traffic. Dual links provide redundancy for connections between systems.

Using the CLI

For creating secured IP partnerships, the authorities and certificate installation is a prerequisite.

You can export the certificate from the remote system and import it on the local system in order to allow configuration changes to be coordinated across systems without manual authentication. Use the **chsystemcert** command to create and manage certificates on the system.

If the certificate is signed by the system's root certificate authority, or it is a self-signed certificate, then use the following command to export the entire certificate chain on the remote system. If the certificate is signed by a trusted third-party certificate authority, then the root certificate must be retrieved from the third-party certificate authority.

```
chsystemcert –export
```

The certificate is exported to the `/dumps/certificate.pem` directory on the configuration node.

After you complete all these steps, run **lspartnership** command on each system in the configuration and verify the following items:

- Portsets are configured for all links between the systems.
- All systems are in the Fully Configured state.

For more information, see **lspartnership** command.

Creating portset

You can create portset for IP partnership by using both the management GUI and the CLI.

Create and assign IP addresses to the Ethernet portset

To create an Ethernet portset in the management GUI, complete the following steps:

- To create or manage an Ethernet portset in the management GUI, select **Settings** > **Network** > **Ethernet Portsets**.
- To create or assign IP addresses to the portset, select **Settings** > **Network** > **Ethernet Ports**.
- Verify that each port displays the **Configured** state after it has been added to the portset.
  - Right-click the port and select either **Modify Replication**, **Modify iSCSI hosts**, or **Modify Storage Ports**. The traffic type for the port must match the traffic for the portset that you created.

Using the CLI

To create portset using the CLI, enter the following command:

```
mkportset -name portset_name -type portset_type -ownershipgroup owner_name
```

where *portset_name* is the name of the portset and *portset_type* is **replication**. The value *owner_name* indicates the name of the ownership group to which the portset belongs. It is an optional value.

Configuring VLAN for IP partnerships

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

- 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.

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

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.

Migrating partnerships

It is possible to migrate existing partnerships from IPv4 to IPv6 deployments, and from one Ethernet speed link to a different Ethernet speed link. It is not possible to convert a partnership between Fibre Channel, IP or RDMA.

Migrate an existing partnership from IPv4 to IPv6 deployments

To migrate from IPv4 to IPv6, the following requirements must be met:

- System IPs have IPv6 addresses configured.
- Data path IPs (IPs configured by using **mkip**) have IPv6 addresses.

Note: You can assign IPv6 addresses to ports while IP partnerships are active. However, you cannot add them to portsets.

Complete the following procedure:

1. Stop the IP partnership on both systems.
   a. On system C1: #svctask **chpartnership -stop** <**systemid C2**>
      - The command **lspartnership** reports as `fully_configured_stopped` on system C1.
      - The command **lspartnership** reports as `fully_configured_remote_stopped` on system C2.
   b. On system C2: #svctask **chpartnership -stop** <**systemid C1**>
      - The command **lspartnership** reports as `fully_configured_stopped` on system C1.
      - The command **lspartnership** also reports as `fully_configured_stopped` on system C2.
2. Add the IPv6 IP addresses, which are configured on data path IP Ports (by using **mkip**), in respective portsets.
   - On system C1: **mkip -node** <**id | name**> **-port** <**id**> **-portset** <**id | name**> **-ip** <**ipv6 address**> **-prefix** <**subnet prefix**> **-gw** <**ipv6 gateway address**> **-vlan** <**vlan_id**>
   - On system C2: **mkip -node** <**id | name**> **-port** <**id**> **-portset** <**id | name**> **-ip** <**ipv6 address**>**-prefix** <**subnet prefix**> **-gw** <**ipv6 gateway address**> **-vlan** <**vlan_id**>
3. Modify the IP partnerships to do discovery over IPv6 addresses.
   - On system C1: #svctask **chpartnership -type ipv6 -clusterip** <**ipv6_ipaddr_of_cluster_c2**> <**systemid C2**>
   - On system C2: #svctask **chpartnership -type ipv6 -clusterip** <**ipv6_ipaddr_of_cluster_c1**> <**systemid C1**>
   - On system C1: #svctask **chpartnership -start** <**systemid C2**>
   - On system C2: #svctask **chpartnership -start** <**systemid C1**>

     Note: Since new data path over IPv6 addresses are now available, partnership will first change to `not_present`, and then to `fully_configured`. If it remains in `not_present`, monitor the node error/DMP if it is getting triggered and check the appropriate DMP in the troubleshooting section.

This same procedure can be applied to migrate from IPv6 to IPv4 by applying suitable substitutes for IPv4 instead of IPv6.

Migrate an existing partnership to new IP addresses

When you change one or both system and data path IP addresses of the system, replication needs to be temporarily stopped during the procedure. However, host I/O can continue. After the procedure is complete, the v need to resynchronize only the host write operations that are submitted while the partnership was stopped.

Complete the following procedure:

1. Stop the IP partnership on both systems.
   a. On system C1: #svctask **chpartnership -stop** <**systemid C2**>
      - The command **lspartnership** reports as `fully_configured_stopped` on system C1.
      - The command **lspartnership** reports as `fully_configured_remote_stopped` on system C2.
   b. On system C2: #svctask **chpartnership -stop** <**systemid C1**>
      - The command **lspartnership** reports as `fully_configured_stopped` on system C1.
      - The command **lspartnership** also reports as `fully_configured_stopped` on system C2.
2. Reconfigure the IP addresses on one or both systems by using the following commands:

chsystemip -clusterip <**new ip**> **-port** <**ethernet port number**>

mkip -node <**id | name**> **-port** <**id**> **-portset** <**id | name**> **-ip** <**ipv4 | ipv6 address**> **-prefix** <**subnet prefix**> **-gw** <**ipv4 | ipv6 gateway address**> **-vlan** <**vlan_id**>

3. For each reconfigured system IP, reconfigure the partnership on the remote system:

   If system IP changed on system C1, then on system C2: #svctask **chpartnership -clusterip** <**system IP of system C1**> <**systemid C1**>

   If system IP changed on system C2, then on system C1: #svctask **chpartnership -clusterip** <**system IP of system C2**> <**systemid C2**>

4. Restart the partnerships:

   On system C1: #svctask **chpartnership -start** <**systemid C2**>

   On system C2: #svctask **chpartnership -start** <**systemid C1**>

5. Restart any stopped relationships.

Migrate an existing partnership to different links

Before you attempt to migrate existing deployments from to a different link, refer to the limitations and considerations requirements in the IP partnership configuration topic.

Note: You cannot mix link speeds. For example, if there are two links, both links should either be 10 Gbps links or both should be 1 Gbps links.

This procedure applies to two systems that are in IP partnership with each other over two 1 Gbps links.

Complete the following procedure:

1. Stop the IP partnership on both systems.
   a. On system C1: #svctask **chpartnership -stop** <**systemid C2**>
      - The command **lspartnership** reports as `fully_configured_stopped` on system C1.
      - The command **lspartnership** reports as `fully_configured_remote_stopped` on system C2.
   b. On system C2: #svctask **chpartnership -stop** <**systemid C1**>
      - The command **lspartnership** reports as `fully_configured_stopped` on system C1.
      - The command **lspartnership** also reports as `fully_configured_stopped` on system C2.
2. Add the IP addresses, which are configured on data path IP Ports (by using **mkip**), on 10 Gbps links in respective portsets.
   - On system C1: #svctask **mkip -node** <**id | name**> **-port** <**id**> **-portset** <**id | name**> **-ip** <**ipv4 | ipv6 address**>**-prefix** <**subnet prefix**> **-gw** <**ipv4 | ipv6 gateway address**> **-vlan** <**vlan_id**>
   - Remove the existing the 1 Gbps ports from the portset by using the following command:

     #svctask **rmip ip_ address_id**
   - On system C2: #svctask **mkip -node** <**id | name**> **-port** <**id**> **-portset** <**id | name**> **-ip** <**ipv4 | ipv6 address**>**-prefix** <**subnet prefix**> **-gw** <**ipv4 | ipv6 gateway address**> **-vlan** <**vlan_id**>
   - Remove the existing the 1 Gbps ports from the portset by using the following command:

     #svctask **rmip ip_ address_id**
3. Start the IP partnerships.
   - On system C1: #svctask **chpartnership -start** <**systemid C2**>
   - On system C2: #svctask **chpartnership -start** <**systemid C1**>

Note: Since new data path over IPv6 addresses are now available, partnership will first change to `not_present`, and then to `fully_configured`. If it remains in `not_present`, monitor the node error/DMP if it is getting triggered and check the appropriate DMP in the troubleshooting section.

This procedure completes the migration from 1 Gbps links to 10 Gbps links. This same procedure can be applied to migrate from 10 Gbps to 1 Gbps or 25 Gbps by applying suitable substitutes for 1 Gbps or 25 Gbps ports instead of 10 Gbps ports.

Intersystem link for host traffic

If you use the intersystem link for host traffic, ensure that you have sufficient bandwidth to support all sources of load.

Scenario: The hosts in a local clustered system can read and write to the volumes in a remote system

In this scenario, the hosts in the local system also exchange heartbeats with the hosts that are in the remote system. Because the intersystem link is being used for multiple purposes, you must have sufficient bandwidth to support the following sources of load:

- Local host to remote volume I/O traffic or remote host to local volume I/O traffic.
- Local-host-to-remote-host heartbeat traffic.

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 50. 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 51 on page 354 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 51. 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

- 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.

Chapter 17. Cyber resilience

Ransomware threat detection

Ransomware threat detection is a process that identifies and responds to security threats before they can damage data or systems.

Ransomware threat detection uses information theory concepts such as Shannon Entropy and other statistics to characterize the current workload. This type of analysis requires byte-by-byte processing and is computationally expensive.

FlashCore Module (FCM) drives provide computational storage offload to compute and summarize these statistics, which the FlashSystem aggregates and analyzes to notify Storage Insights of a potential threat. With drives that run FCM 4.2 firmware or later, statistics about data are processed within the FCM drives, and presented to the FlashSystem on a per volume group basis. This is collected on a regular basis, and aggregated across all participating drives, resulting in a set of statistics per volume group. These statistics are then processed on the system to determine its threat potential.

© Copyright IBM Corp. 2025

Chapter 18. 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 “User authentication” on page 426.

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 “Remote Authentication” on page 428.

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 “Users” on page 221.

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” on page 222.

© Copyright IBM Corp. 2025

Object-based access control
An ownership group defines a subset of users and objects within the system. For more information on configuring an ownership group, see “Configuring ownership groups” on page 214.

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 ” on page 224.

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 487.

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” on page 222.

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 “Login message” on page 234.

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 405.

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 431.

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 “Partnerships using IP Connectivity” on page 339.

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.

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 "Security protocol levels" on page 361.

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 "Configuring system certificates" on page 393.

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.

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.

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" on page 486.

Volume protection

Volume protection prevents active volumes or host mappings from being deleted inadvertently if the system detects recent I/O activity. For more information, see "Volume protection " on page 84.

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 Storage Copy Data Management.

For more information, refer to the "Configuring Safeguarded Copy function" on page 280.

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. Ethernet portsets group Ethernet ports to control port usage, assign IP addresses, and define host connectivity policies. For more information, refer to "Portsets" on page 333.

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 "Encryption" on page 455.

Managing security

The system supports several security-related features that can help protect system data and resources from unauthorized access.

User security and 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.

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.

Encryption is a technology that uses cryptography to help 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.

SSL/TLS security controls

The system supports a choice of security levels (2 to 7) to enforce a minimum level of SSL (Secure Sockets Layer) or TLS (Transport Layer Security) that can be used to access the system. The security levels 5, 6, and 7 support TLS 1.3, which provides more secure cipher suites along with a faster handshake process. Only clients that support the minimum SSL or TLS level that is enforced by the system are able to establish secure connections.

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 52. Supported SSL/TLS security levels

Security level	Description	Minimum security allowed	Supported Java SSL ciphers	Supported OpenSSL ciphers
1	Sets the system to disallow SSL version 3.0.	TLS 1.1	• 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	• ECDHE-RSA-AES256-GCM-SHA384 • ECDHE-ECDSA-AES256-GCM-SHA384 • ECDHE-RSA-AES256-SHA384 • ECDHE-ECDSA-AES256-SHA384 • ECDHE-RSA-AES256-SHA • ECDHE-ECDSA-AES256-SHA • DHE-DSS-AES256-GCM-SHA384 • DHE-RSA-AES256-GCM-SHA384 • DHE-RSA-AES256-SHA256 • DHE-DSS-AES256-SHA256 • DHE-RSA-AES256-SHA • DHE-DSS-AES256-SHA • DHE-RSA-CAMELLIA256-SHA • DHE-DSS-CAMELLIA256-SHA • ECDHE-ECDSA-DES-CBC3-SHA • DHE-RSA-SEED-SHA • DHE-DSS-SEED-SHA • DHE-RSA-CAMELLIA128-SHA • ECDH-RSA-AES256-GCM-SHA384 E • ECDH-ECDSA-AES256-GCM-SHA384 • ECDH-RSA-AES256-SHA384

Table 52. Supported SSL/TLS security levels (continued)

Security level	Description	Minimum security allowed	Supported Java SSL ciphers	Supported OpenSSL ciphers
	Sets the system to disallow SSL version 3.0.	TLS 1.1	• 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	• ECDH-ECDSA-AES256-SHA384 • ECDH-RSA-AES256-SHA • ECDH-ECDSA-AES256-SHA • AES256-GCM-SHA384 • AES256-SHA256 • AES256-SHA • CAMELLIA256-SHA • PSK-AES256-CBC-SHA • ECDHE-RSA-AES128-GCM-SHA256 • ECDHE-ECDSA-AES128-GCM-SHA256 • ECDHE-RSA-AES128-SHA256 • ECDHE-ECDSA-AES128-SHA256 • ECDHE-RSA-AES128-SHA • ECDHE-ECDSA-AES128-SHA • DHE-DSS-AES128-GCM-SHA256 • DHE-RSA-AES128-GCM-SHA256 • DHE-RSA-AES128-SHA256 • DHE-DSS-AES128-SHA256 • DHE-RSA-AES128-SHA • DHE-DSS-AES128-SHA • ECDHE-RSA-DES-CBC3-SHA

Table 52. Supported SSL/TLS security levels (continued)

Security level	Description	Minimum security allowed	Supported Java SSL ciphers	Supported OpenSSL ciphers
	Sets the system to disallow SSL version 3.0.	TLS 1.1	• SSL_ECDHE_ECDSA_WITH_AES_256_CBC_SHA • SSL_ECDHE_RSA_WITH_AES_256_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 • SSL_ECDHE_ECDSA_WITH_3DES_EDE_CBC_SHA • SSL_ECDHE_RSA_WITH_3DES_EDE_CBC_SHA • SSL_RSA_FIPS_WITH_3DES_EDE_CBC_SHA • SSL_RSA_WITH_3DES_EDE_CBC_SHA • SSL_ECDH_ECDSA_WITH_3DES_EDE_CBC_SHA • SSL_ECDH_RSA_WITH_3DES_EDE_CBC_SHA • SSL_DHE_DSS_WITH_3DES_EDE_CBC_SHA • SSL_DHE_RSA_WITH_3DES_EDE_CBC_SHA	• DHE-DSS-CAMELLIA128-SHA • EDH-RSA-DES-CBC3-SHA • EDH-DSS-DES-CBC3-SHA • ECDH-RSA-AES128-GCM-SHA256 • ECDH-ECDSA-AES128-GCM-SHA256 • ECDH-RSA-AES128-SHA256 • ECDH-ECDSA-AES128-SHA256 • ECDH-RSA-AES128-SHA • ECDH-ECDSA-AES128-SHA • ECDH-RSA-DES-CBC3-SHA • ECDH-ECDSA-DES-CBC3-SHA • AES128-GCM-SHA256 • AES128-SHA • SEED-SHA • CAMELLIA128-SHA • DES-CBC3-SHA • PSK-AES128-CBC-SHA • PSK-3DES-EDE-CBC-SHA • KRB5-DES-CBC3-SHA

Table 52. Supported SSL/TLS security levels (continued)

Security level	Description	Minimum security allowed	Supported Java SSL ciphers	Supported OpenSSL 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 • 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	• ECDHE-RSA-AES256-GCM-SHA384 • ECDHE-ECDSA-AES256-GCM-SHA384 • ECDHE-RSA-AES256-SHA384 • ECDHE-ECDSA-AES256-SHA384 • DHE-DSS-AES256-GCM-SHA384 • DHE-RSA-AES256-GCM-SHA384 • DHE-RSA-AES256-SHA256 • ECDH-RSA-AES256-GCM-SHA384 E • ECDH-ECDSA-AES256-GCM-SHA384 • ECDH-RSA-AES256-SHA384 • ECDH-ECDSA-AES256-SHA384 • AES256-GCM-SHA384 • AES256-SHA256 • AES256-SHA • ECDHE-RSA-AES128-GCM-SHA256 • ECDHE-ECDSA-AES128-GCM-SHA256

Table 52. Supported SSL/TLS security levels (continued)

Security level	Description	Minimum security allowed	Supported Java SSL ciphers	Supported OpenSSL ciphers
	Sets the system to disallow SSL version 3.0, TLS version 1.0, and TLS version 1.1.	TLS 1.2	• 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	• ECDHE-RSA-AES128-SHA256 • ECDHE-ECDSA-AES128-SHA256 • DHE-DSS-AES128-GCM-SHA256 • DHE-RSA-AES128-GCM-SHA256 • DHE-RSA-AES128-SHA256 • DHE-DSS-AES128-SHA256 • ECDH-RSA-AES128-GCM-SHA256 • ECDH-ECDSA-AES128-GCM-SHA256 • ECDH-RSA-AES128-SHA256 • ECDH-ECDSA-AES128-SHA256 • AES128-GCM-SHA256 • AES128-SHA256 • AES128-SHA • DES-CBC3-SHA • ECDHE-RSA-AES128-GCM-SHA256 • ECDHE-ECDSA-AES128-GCM-SHA256

Table 52. Supported SSL/TLS security levels (continued)

Security level	Description	Minimum security allowed	Supported Java SSL ciphers	Supported OpenSSL ciphers
	Sets the system to disallow SSL version 3.0, TLS version 1.0, and TLS version 1.1.	TLS 1.2	• 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	• ECDHE-ECDSA-AES128-GCM-SHA256 • ECDHE-RSA-AES128-SHA256 • ECDHE-ECDSA-AES128-SHA256 • DHE-DSS-AES128-GCM-SHA256 • DHE-RSA-AES128-GCM-SHA256 • DHE-RSA-AES128-SHA256 • DHE-DSS-AES128-SHA256 • ECDH-RSA-AES128-GCM-SHA256 • ECDH-ECDSA-AES128-GCM-SHA256 • ECDH-RSA-AES128-SHA256 • ECDH-ECDSA-AES128-SHA256 • AES128-GCM-SHA256 • AES128-SHA256 • AES128-SHA • DES-CBC3-SHA

Table 52. Supported SSL/TLS security levels (continued)

Security level	Description	Minimum security allowed	Supported Java SSL ciphers	Supported OpenSSL 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 • 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	• ECDHE-RSA-AES256-GCM-SHA384 • ECDHE-ECDSA-AES256-GCM-SHA384 • ECDHE-RSA-AES256-SHA384 • ECDHE-ECDSA-AES256-SHA384 • DHE-DSS-AES256-GCM-SHA384 • DHE-RSA-AES256-GCM-SHA384 • DHE-RSA-AES256-SHA256 • ECDH-RSA-AES256-GCM-SHA384 E • ECDH-ECDSA-AES256-GCM-SHA384 • ECDH-RSA-AES256-SHA384 • ECDH-ECDSA-AES256-SHA384 • AES256-GCM-SHA384 • AES256-SHA256

Table 52. Supported SSL/TLS security levels (continued)

Security level	Description	Minimum security allowed	Supported Java SSL ciphers	Supported OpenSSL ciphers
	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_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	• ECDHE-RSA-AES128-GCM-SHA256 • ECDHE-ECDSA-AES128-GCM-SHA256 • ECDHE-RSA-AES128-SHA256 • ECDHE-ECDSA-AES128-SHA256 • DHE-DSS-AES128-GCM-SHA256 • DHE-RSA-AES128-GCM-SHA256 • DHE-RSA-AES128-SHA256 • DHE-DSS-AES128-SHA256 • ECDH-RSA-AES128-GCM-SHA256 • ECDH-ECDSA-AES128-GCM-SHA256 • ECDH-RSA-AES128-SHA256 • ECDH-ECDSA-AES128-SHA256 • AES128-GCM-SHA256 • AES128-SHA256

Table 52. Supported SSL/TLS security levels (continued)

Security level	Description	Minimum security allowed	Supported Java SSL ciphers	Supported OpenSSL 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 • SSL_DHE_DSS_WITH_AES_256_GCM_SHA384 • SSL_ECDHE_ECDSA_WITH_AES_128_GCM_SHA256 • SSL_DHE_DSS_WITH_AES_128_GCM_SHA256	• ECDHE-ECDSA-AES256-GCM-SHA384 • DHE-DSS-AES256-GCM-SHA384 • ECDHE-ECDSA-AES128-GCM-SHA256 • DHE-DSS-AES128-GCM-SHA256
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** • 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 1.3** • AES-256-GCM-SHA384 • CHACHA20-POLY1305-SHA256 • AES-128-GCM-SHA256 • AES-128-CCM-8-SHA256 • AES-128-CCM-SHA256

Table 52. Supported SSL/TLS security levels (continued)

Security level	Description	Minimum security allowed	Supported Java SSL ciphers	Supported OpenSSL ciphers
	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 1.2** 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: • 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	**For TLS 1.2** The security level 5 supports all the OpenSSL ciphers supported at the security level 3. A few additional TLS 1.2 specific OpenSSL ciphers supported at level 5 are: • ECDHE-RSA-AES256-GCM-SHA384 • ECDHE-ECDSA-AES256-GCM-SHA384 • ECDHE-RSA-AES256-SHA384 • ECDHE-ECDSA-AES256-SHA384 • DHE-RSA-AES256-GCM-SHA384 • DHE-RSA-AES256-SHA256 • AES256-GCM-SHA384 • AES256-SHA256 • ECDHE-RSA-AES128-GCM-SHA256 • ECDHE-ECDSA-AES128-GCM-SHA256

Table 52. Supported SSL/TLS security levels (continued)

Security level	Description	Minimum security allowed	Supported Java SSL ciphers	Supported OpenSSL ciphers
	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	• 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	• ECDHE-RSA-AES128-SHA256 • ECDHE-ECDSA-AES128-SHA256 • DHE-RSA-AES128-GCM-SHA256 • DHE-RSA-AES128-SHA256 • AES128-GCM-SHA256 • AES128-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 • TLS_CHACHA20_POLY1305_SHA256 • TLS_AES_128_GCM_SHA256 • TLS_AES_128_CCM_8_SHA256 • TLS_AES_128_CCM_SHA256	• AES-256-GCM-SHA384 • CHACHA20-POLY1305-SHA256 • AES-128-GCM-SHA256 • AES-128-CCM-8-SHA256 • AES-128-CCM-SHA256
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	AES-256-GCM-SHA384

The following table describes the SSH security levels supported by the system:

Table 53. SSH algorithms supported at each security level

Security level	Key Exchange	Cipher Suite	MAC Algorithm	Host Key Algorithms
1	• curve25519-sha256 • curve25519-sha256@libssh.org • ecdh-sha2-nistp256 • ecdh-sha2-nistp384 • ecdh-sha2-nistp521 • diffie-hellman-group-exchange-sha256 • diffie-hellman-group16-sha512 • diffie-hellman-group18-sha512 • diffie-hellman-group14-sha256 • diffie-hellman-group14-sha1 • diffie-hellman-group1-sha1 • diffie-hellman-group-exchange-sha1	aes256-ctr aes192-ctr aes128-ctr chacha20-poly1305@openssh.com aes256-gcm@openssh.com aes128-gcm@openssh.com aes256-cbc aes192-cbc aes128-cbc	hmac-sha2-256 hmac-sha2-512 hmac-sha2-256-etm@openssh.com hmac-sha2-512-etm@openssh.com hmac-sha1	rsa-sha2-256 rsa-sha2-512 ecdsa-sha2-nistp521 ecdsa-sha2-nistp521-cert-v01@openssh.com ssh-rsa ssh-rsa-cert-v01@openssh.com

Table 53. SSH algorithms supported at each security level (continued)

Security level	Key Exchange	Cipher Suite	MAC Algorithm	Host Key Algorithms
2	• curve25519-sha256 • curve25519-sha256@libssh.org • ecdh-sha2-nistp256 • ecdh-sha2-nistp384 • ecdh-sha2-nistp521 • diffie-hellman-group-exchange-sha256 • diffie-hellman-group16-sha512 • diffie-hellman-group18-sha512 • diffie-hellman-group14-sha256 • diffie-hellman-group14-sha1	aes256-ctr aes192-ctr aes128-ctr chacha20-poly1305@openssh.com aes256-gcm@openssh.com aes128-gcm@openssh.com	hmac-sha2-256 hmac-sha2-512 hmac-sha2-256-etm@openssh.com hmac-sha2-512-etm@openssh.com hmac-sha1	rsa-sha2-256 rsa-sha2-512 ecdsa-sha2-nistp521 ecdsa-sha2-nistp521-cert-v01@openssh.com ssh-rsa ssh-rsa-cert-v01@openssh.com
3	• curve25519-sha256 • curve25519-sha256@libssh.org • ecdh-sha2-nistp256 • ecdh-sha2-nistp384 • ecdh-sha2-nistp521 • diffie-hellman-group-exchange-sha256 • diffie-hellman-group16-sha512 • diffie-hellman-group18-sha512 • diffie-hellman-group14-sha256	aes256-ctr aes192-ctr aes128-ctr chacha20-poly1305@openssh.com aes256-gcm@openssh.com aes128-gcm@openssh.com	hmac-sha2-256 hmac-sha2-512 hmac-sha2-256-etm@openssh.com hmac-sha2-512-etm@openssh.com	rsa-sha2-256 rsa-sha2-512 ecdsa-sha2-nistp521 ecdsa-sha2-nistp521-cert-v01@openssh.com ssh-rsa ssh-rsa-cert-v01@openssh.com

Table 53. SSH algorithms supported at each security level (continued)				
Security level	**Key Exchange**	**Cipher Suite**	**MAC Algorithm**	**Host Key Algorithms**
4	• curve25519-sha256 • curve25519-sha256@libssh.org • ecdh-sha2-nistp256 • ecdh-sha2-nistp384 • ecdh-sha2-nistp521	aes256-ctr aes192-ctr aes128-ctr aes256-gcm@openssh.com aes128-gcm@openssh.com	hmac-sha2-256 hmac-sha2-512 hmac-sha2-256-etm@openssh.com hmac-sha2-512-etm@openssh.com	rsa-sha2-256 rsa-sha2-512 ecdsa-sha2-nistp521 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.

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 "System Certificates" on page 391.

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 54 on page 377 displays the protocols that are supported at security level 7.

Table 54. 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 55 on page 377 displays Java SSL ciphers that are supported at security level 7.

Table 55. Java SSL ciphers supported at security level 7

Java SSL ciphers
TLS_AES_256_GCM_SHA384

Table 56 on page 378 displays OpenSSL security ciphers that are supported by security level 7.

Table 56. OpenSSL ciphers supported at level 7 (`chsecurity -sslprotocol 7`)		
Cipher	**Bulk encryption algorithm**	**Hashing algorithm**
AES-256-GCM-SHA384	AES-256-GCM	SHA384

Table 57 on page 378 displays the protocols that are supported at security level 6.

Table 57. 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 58 on page 378 displays Java SSL ciphers that are supported at security level 6.

Table 58. 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 59 on page 378 displays OpenSSL security ciphers that are supported by security level 6.

Table 59. 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 60 on page 378 displays the protocols that are supported at security level 5.

Table 60. 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

Table 60. Protocols supported at level 5 (continued)	
Protocol level	**Is it supported?**
SSL 3 and earlier	No

Table 61 on page 379 displays Java SSL ciphers that are supported at security level 5.

Table 61. 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 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 62 on page 379 and Table 63 on page 380 displays OpenSSL security ciphers that are supported by security level 5.

Table 62. 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

Table 62. OpenSSL ciphers supported at level 5 for TLS 1.3 (`chsecurity -sslprotocol 5`) (continued)		
Cipher	**Bulk encryption algorithm**	**Hashing algorithm**
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 63. 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
AES128-SHA256	RSA	RSA	AES(128)	SHA256

Table 64 on page 381 displays the protocols that are supported at security level 4.

Table 64. 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 65 on page 381 displays Java SSL ciphers that are supported at security level 4.

Table 65. 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 66 on page 381 displays OpenSSL security ciphers that are supported by security level 4.

Table 66. 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 67 on page 381 displays the protocols that are supported at security level 3.

Table 67. 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 68 on page 381 displays Java SSL ciphers that are supported at security level 3.

Table 68. Java SSL ciphers supported at security level 3
Java SSL ciphers
SSL_ECDHE_ECDSA_WITH_AES_256_CBC_SHA384

Table 68. Java SSL ciphers supported at security level 3 (continued)
Java SSL ciphers
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
SSL_DHE_DSS_WITH_AES_128_GCM_SHA256

Table 69 on page 383 displays OpenSSL security ciphers that are supported by security level 3.

Table 69. 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
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

Table 69. OpenSSL ciphers supported at level 3 (`chsecurity -sslprotocol 3`) (continued)				
Cipher	**Kx**	**Au**	**Enc**	**Mac**
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 70 on page 384 displays the protocols that are supported at security level 2.

Table 70. 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 71 on page 384 displays Java SSL ciphers that are supported at security level 2.

Table 71. 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
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

Table 71. Java SSL ciphers supported at level 2 (continued)
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 72 on page 386 displays OpenSSL security ciphers that are supported by security level 2.

Table 72. 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
DHE-RSA-AES128-GCM-SHA256	DH	RSA	AESGCM(128)	AEAD

Table 72. OpenSSL ciphers supported at level 2 (`chsecurity -sslprotocol 2`) (continued)				
Cipher	**Kx**	**Au**	**Enc**	**Mac**
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 73 on page 387 lists all supported ports and describes how they can be used.

Table 73. 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 UDP	6514 (TCP) 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
HTTPS redirect for GUI access	Inbound	TCP	443	Mandatory

Table 73. TCP and UDP ports that are supported (continued)

Service	Traffic direction	Protocol	Port	Service type
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
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 74. SSH algorithms supported at each SSH security level (chsecurity -sshprotocol)

Security level	Key Exchange	Cipher Suite	MAC Algorithm	Host Key Algorithms
1	• curve25519-sha256 • curve25519-sha256@libssh.org • ecdh-sha2-nistp256 • ecdh-sha2-nistp384 • ecdh-sha2-nistp521 • diffie-hellman-group-exchange-sha256 • diffie-hellman-group16-sha512 • diffie-hellman-group18-sha512 • diffie-hellman-group14-sha256 • diffie-hellman-group14-sha1 • diffie-hellman-group1-sha1 • diffie-hellman-group-exchange-sha1	aes256-ctr aes192-ctr aes128-ctr chacha20-poly1305@openssh.com aes256-gcm@openssh.com aes128-gcm@openssh.com aes256-cbc aes192-cbc aes128-cbc	hmac-sha2-256 hmac-sha2-512 hmac-sha2-256-etm@openssh.com hmac-sha2-512-etm@openssh.com hmac-sha1	rsa-sha2-256 rsa-sha2-512 ecdsa-sha2-nistp521 ecdsa-sha2-nistp521-cert-v01@openssh.com ssh-rsa ssh-rsa-cert-v01@openssh.com

Table 74. SSH algorithms supported at each SSH security level (chsecurity -sshprotocol) (continued)

Security level	Key Exchange	Cipher Suite	MAC Algorithm	Host Key Algorithms
2	• curve25519-sha256 • curve25519-sha256@libssh.org • ecdh-sha2-nistp256 • ecdh-sha2-nistp384 • ecdh-sha2-nistp521 • diffie-hellman-group-exchange-sha256 • diffie-hellman-group16-sha512 • diffie-hellman-group18-sha512 • diffie-hellman-group14-sha256 • diffie-hellman-group14-sha1	aes256-ctr aes192-ctr aes128-ctr chacha20-poly1305@openssh.com aes256-gcm@openssh.com aes128-gcm@openssh.com	hmac-sha2-256 hmac-sha2-512 hmac-sha2-256-etm@openssh.com hmac-sha2-512-etm@openssh.com hmac-sha1	rsa-sha2-256 rsa-sha2-512 ecdsa-sha2-nistp521 ecdsa-sha2-nistp521-cert-v01@openssh.com ssh-rsa ssh-rsa-cert-v01@openssh.com
3	• curve25519-sha256 • curve25519-sha256@libssh.org • ecdh-sha2-nistp256 • ecdh-sha2-nistp384 • ecdh-sha2-nistp521 • diffie-hellman-group-exchange-sha256 • diffie-hellman-group16-sha512 • diffie-hellman-group18-sha512 • diffie-hellman-group14-sha256	aes256-ctr aes192-ctr aes128-ctr chacha20-poly1305@openssh.com aes256-gcm@openssh.com aes128-gcm@openssh.com	hmac-sha2-256 hmac-sha2-512 hmac-sha2-256-etm@openssh.com hmac-sha2-512-etm@openssh.com	rsa-sha2-256 rsa-sha2-512 ecdsa-sha2-nistp521 ecdsa-sha2-nistp521-cert-v01@openssh.com ssh-rsa ssh-rsa-cert-v01@openssh.com

Table 74. SSH algorithms supported at each SSH security level (chsecurity -sshprotocol) (continued)				
Security level	**Key Exchange**	**Cipher Suite**	**MAC Algorithm**	**Host Key Algorithms**
4	• curve25519-sha256 • curve25519-sha256@libssh.org • ecdh-sha2-nistp256 • ecdh-sha2-nistp384 • ecdh-sha2-nistp521	aes256-ctr aes192-ctr aes128-ctr aes256-gcm@openssh.com aes128-gcm@openssh.com	hmac-sha2-256 hmac-sha2-512 hmac-sha2-256-etm@openssh.com hmac-sha2-512-etm@openssh.com	rsa-sha2-256 rsa-sha2-512 ecdsa-sha2-nistp521 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.

SSH rules

Use the **SSH Rules** page to customize settings related to secure shell (SSH) access. SSH is used to authenticate users to the command-line interface.

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.

   SSH listening port
   Indicates the SSH port number on which SSH server is listening. The default value is 22.
3. Click **Save**.

System Certificates

To ensure secure communication between systems, applications, and browsers, IBM Storage FlashSystem allows you to manage system certificates using the management GUI.

System certificates authenticate the FlashSystem to other devices and applications, enabling encrypted and trusted communication. Depending on the security needs, you can choose between:

- Internally signed certificates (simpler, auto-renewable)
- Externally signed certificates (trusted by public CAs, manually managed)

Certificate signing options

Internally signed certificate

- **Use case**: Ideal for internal environments or where ease of management is key.
- **How it works**: FlashSystem includes a built-in root Certificate Authority (CA) from version 8.5.3 onward.
- **Benefits**:
  - Automatically created during system setup.
  - Can be auto-renewed, preventing expiry-related disruptions.
  - Root certificate can be exported and added to truststore for browsers, key servers (Generic KMIP servers) and other devices.
- **Limitations**: Only trusted within environments where the root certificate is manually added.

Note: FlashSystem uses a built-in root CA to sign multiple system certificates for different purposes (web access, key server, internal communication). The root certificate is the trust anchor, while system certificates are scoped for specific functions.

Externally signed certificate

- **Use case**: Best for public-facing systems or when compliance requires third-party trust.
- **How it works**: You generate a Certificate Signing Request (CSR) and submit it to a public or enterprise CA.
- **Benefits**: Trusted by most browsers and external systems.
- **Limitations**:
  - Manual renewal required before expiry.
  - System alerts you 30 days before expiration.

Understanding the certificate types

Root certificate

- Generated during first boot (version 8.5.3 and above).
- Shared across all nodes in a system.
- Cannot be modified by users.
- Valid for 20 years, using a 4096-bit RSA key.
- Distinguished Name (DN) format:
  - **O**: International Business Machines Corporation
  - **OU**: IBM Storage Virtualize Root CA
  - **CN**: Serial number of the configuration node

System certificate

- **Signed by**:
  - Internal root CA
  - External CA
- **Presented to**:
  - Web browsers
  - Key servers
  - Internal communication modules
- **Supported key types**:
  - RSA 2048 / RSA 4096

  - ECDSA 384 / ECDSA 521
- **SSL protocol level 4**:
  - Does not support RSA key exchange.
  - Use ECDSA if SSL level 4 is required.

Certificate scope values

Each system certificate has a scope ID that defines its purpose:

Table 75. System certificate scope values

Scope ID	Use case	Description
0	Default	Used for web browser access, key servers, and general system authentication.
1	Key server	Specifically for communication with KMIP-compatible key servers.
3	Internal communication	Used for internal communication purpose, such as certificate for and with remote system.

Best practices

- Use internally signed certificates for ease of management in closed environments.
- Use externally signed certificates when interoperability or compliance with external systems is required.

Configuring system certificates

Secure communications between systems requires that certificates are configured on the system. To ensure that the certificate and public ID of the system are valid and secure, each system has a digital Secure Sockets Layer (SSL) certificate.

About this task

Before you create a request for a specific certificate, ensure that your current browser does not have restrictions on the type of keys that are used for certificates. Some browsers limit the use of specific types of keys for security and compatibility issues.

If a specific certificate is signed by an intermediate certificate authority (CA), then the full chain of certificates must be installed. To install the signed certificate and certificate authority (CA) certificates, create a single file that contains the full chain of certificates. The file should include the signed certificate and the intermediate CA certificates. The root CA certificate can be included but is optional.

To manage the certificate that is installed on a system, use either the **`svctask chsystemcert`** (only for default), **`svctask mksystemcertstore`**, **`svctask chsystemcertstore`**, **`svctask rmsystemcertstore`** commands in the command-line interface (CLI) or click **Settings** > **Security** > **Certificates** in the management GUI to do the following tasks:

- Upload signed certificate
- Update certificate
- Cancel outstanding signing request
- Export system certificate
- Export root certificate

- Remove system certificate

For specific directions on completing these tasks, see “System Certificates” on page 391.

Updating or creating an internally signed certificate

You can create an internally signed certificate for the system by using the management GUI or command-line interface (CLI).

Note: Before updating the certificates, keep the following factors into consideration:

- Updating the certificate logs you out of the current management GUI session, requiring a fresh login. For features that support chain of trust checking, the secure connection is not interrupted when updating the internally signed certificate.
- IBM Security Guardium® Key Lifecycle Manager key servers do not currently support chain of trust checking with IBM Storage Virtualize. The new internally signed certificate must be exported to the key servers to re-establish a secure connection.
- If you are using IBM Security Verify multifactor authentication, users who are required to use multifactor authentication cannot log in to the management GUI after the certificate is renewed. The new internally signed certificate must be exported using the CLI and added as a new signer certificate to IBM Security Verify for successful authentication.

Using the management GUI

To configure an internally signed certificate, complete these steps:

1. If the root certificate has not already been exported, export the root certificate to other systems, web browsers, and devices that require secure communications with the system. For more information, see .
2. In the management GUI, select **Settings** > **Security** > **Certificates**.
3. On the Certificates page, the certificate details are displayed.

   Automatic renewal of the system certificate

   The system certificate can be renewed automatically if it is signed by the system root CA. Turn automatic renewal on by going to **Settings** > **Security** > **Certificates** > **Go to the particular certificate** and set automatic renewal to **On**. The default validity period of the system certificate is one year. If automatic renewal is **On**, the system certificate is renewed thirty days before the expiry date. If the validity period of the system certificate is fewer than thirty days, then the certificate will be renewed after eight hours.

   The renewed certificate contains all of the same field values, key type and validity period details as the previous certificate.
4. Go to specific certificate, click the overflow menu and select **Replace** or click **Create** and then select the scope of the certificate from the drop-down menu.
5. Select **Internally Signed Certificate** for the certificate type.
6. If you are already using certificates, the **Certificate Details** are automatically populated. You can update any of the following details:

   Key type
   Select the cryptographic key type that is used to generate the certificate.

   Validity days
   Enter the number of days the certificate is valid for. The maximum number of days that are allowed is 9000.

   Country
   Enter the two-letter country code or location, for example, 01 for US.

   State
   Enter the name of the state where the system requesting the certificate is located.

City
Enter the name of the city where the system is located.

Organization name
Enter the name of the organization.

Organizational unit
Enter the name of organizational unit.

Common name
Enter the common name for the certificate.

Subject alternative name
Subject alternative name is the hostname of the system.

Web browsers, and other features that use certificate authentication, require a **Subject Alternative Name**, which is an extension to the Internet standard for public key certificates. The **Subject Alternative Name** extension is used to match the domain name and site certificate and can be an email address, an IP address, a URI, or a DNS name. A certificate can contain a collection of these values so that the certificate can be used on multiple sites.

The **Subject Alternative Name** field can include the management IP addresses for the cluster or DNS names, the service IP addresses for each node in the cluster or DNS names, and any IP addresses configured for IP replication.

For example, if the system has a management DNS name of `cluster.company.com`, and service DNS names of `node1.company.com` and `node2.company.com`, enter these values in the **Subject Alternate Name** field. For multiple values, list each value on a separate line within the box of the **Subject Alternate Name** field:

```
DNS:cluster.company.com
DNS:node1.company.com
DNS:node2.company.com
IP:196.192.0.20
```

Email Address
Enter the email address.

7. Click **Create and install**. The certificate is updated in the main panel.
8. If using multifactor authentication with IBM Security Verify which uses the main system certificate as a signer certificate, export the new main system certificate and install it as a signer certificate in IBM Security Verify. The management GUI is unavailable until the new certificate is added, so the CLI must be used to export the new certificate.

Using the command-line interface (CLI)

Use the following steps to generate an internally signed certificate in the command-line interface.

To generate an internally signed system certificate that uses RSA 2048 key type and expires in one year, enter the following command:

- Using **`mksystemcertstore`** command

```
svctask mksystemcertstore -systemsigned -scope default -commonname virtualize -country GB
-locality Manchester -org IBM -orgunit Systems -email certificates@support.ibm.com -keytype
rsa2048 -validity 365 -subjectalternativename "DNS:test.ibm.com"
```

```
svctask mksystemcertstore -systemsigned -scope internal_communication -commonname
virtualize -country GB -locality Manchester -org IBM -orgunit Systems -email
certificates@support.ibm.com -keytype rsa2048 -validity 365 -subjectalternativename
"DNS:test.ibm.com"
```

- Using **chsystemcert** command (only for default)

```
chsystemcert -mksystemsigned -commonname virtualize -country GB -locality Manchester -org
IBM -orgunit Systems -email certificates@support.ibm.com -keytype rsa2048 -validity 365
-subjectalternativename "DNS:test.ibm.com"
```

After the internally signed system certificate is created, it is automatically installed on the system.

Note: You must use the **-subjectalternativename** parameter to include the management hostname or IP address.

For example, to add a DNS name to the Subject Alternative Name extension, include the following parameter in the **chsystemcert** or **mksystemcertstore** CLI command: **-subjectalternativename "DNS:dns.mysystem.com"** For multiple values, use a recommended delimiter to separate each entry for the **-subjectalternativename** parameter.
Delimiters can be mixed:

Table 76. Recommended delimiters

Delimiter Name	Symbol	Example
Space	(space)	**-subjectalternativename "DNS:dns.myco.com IP:1.2.3.20 URI:http:\ \www.myco.com email:support@myco.com"**
Comma	(,)	**-subjectalternativename "DNS:dns.myco.com,IP:1.2. 3.20,URI:http:\ \www.myco.com,email:suppo rt@myco.com"**
Semi-colon	(;)	**-subjectalternativename "DNS:dns.myco.com;IP:1.2. 3.20;URI:http:\ \www.myco.com;email:suppo rt@myco.com"**
Newline (for Linux or UNIX operating systems)	(\n)	**-subjectalternativename "DNS:dns.myco.com\nIP:1.2 .3.20\nURI:http:\ \www.myco.com\nemail:supp ort@myco.com"**
Tab (for Linux or UNIX operating systems)	(\t)	**-subjectalternativename "DNS:dns.myco.com\tIP:1.2 .3.20\tURI:http:\ \www.myco.com\temail:supp ort@myco.com"**
Carriage return (for Windows operating systems)	(\r)	**-subjectalternativename "DNS:dns.myco.com\rIP:1.2 .3.20\rURI:http:\ \www.myco.com\remail:supp ort@myco.com"**
Carriage return with newline (for Windows operating systems)	(\r\n)	**-subjectalternativename "DNS:dns.myco.com\r\nIP:1 .2.3.20\r\nURI:http:\ \www.myco.com\r\nemail:su pport@myco.com"**

For more information, see the **chsystemcert** or **mksystemcertstore** CLI command.

Requesting and installing an externally signed certificate

If your current externally signed certificate expires or is about to expire, you can request a new signed certificate from a third-party certificate authority. The externally signed certificate can be installed by using the management GUI or command line interface (CLI).

Note: Before requesting and installing the certificates, take the following factors into consideration:

- The certificate signing request must be generated on the system. IBM Storage Virtualize does not support certificate signing requests that are generated outside of the system.
- The signed certificate must contain the following X509v3 Key Usage extensions: Digital Signature. It must also contain the following X509v3 Extended Key Usage Extensions: TLS Web Server Authentication and TLS Web Client Authentication. Ensure that the Certificate Authority (CA) used to sign the certificate includes these extensions.
- Updating the certificate logs you out of the current management GUI session, requiring a fresh login. For features that allow chain of trust checking, the secure connection is not interrupted when updating the externally signed certificate.
- By default, IBM Security Guardium Key Lifecycle Manager key servers do not currently enable chain of trust checking with IBM Storage Virtualize. From version 4.2 onwards, this setting can be enabled. Consider enabling this setting if signed certificate chains are used in order to avoid the export of refreshed certificates to the key server.
- If you are using multifactor authentication with IBM Security Verify, the management GUI are unavailable when you update the certificate. The new externally signed certificate must be exported using the CLI and added as a new signer certificate to IBM Security Verify for successful authentication.

In the management GUI, select **Settings** > **Security** > **Certificates** and select **Replace** from the overflow menu.

In the Certificates page, select externally signed certificate and complete the form to create a request for a signed certificate for particular use case. The resulting certificate signing request (CSR) should be exported to the third-party certificate authority so that it can be signed. After you receive the externally signed certificate from the third-party certificate authority, use the management GUI or command line interface (CLI) to upload and install the signed certificate on the system.

Using the management GUI

To configure an externally signed certificate, complete these steps:

1. In the management GUI, select **Settings** > **Security** > **Certificates**.
2. Click **Create** and then select the scope of the certificate from the drop-down menu.
3. Select **Externally Signed Certificate** for the certificate type.
4. If you are already using certificates, the **Certificate Details** are automatically populated. You can update the details. For more information on the certificate details, see step 6 in .
5. Click **Generate CSR**. This action presents a dialogue window to download the file. Select a location on your local machine to save the file.
6. If the signing request is generated, you have the option to cancel the outstanding signing request in case of an error. To cancel any outstanding signing request, complete these steps:
   a. Select **Settings** > **Security** > **Certificates**.
   b. Go to specific certificate, select **Cancel CSR** and click **Yes**.
7. Share the generated CSR file to the trusted third-party CA. If the CA is a public CA, then it can take some time for the CA to verify your identity before issuing the signed certificate. When it is ready, download the signed certificate file from the CA. You can also download any intermediate CA certificates that were used to sign the request. The files must all be in PEM format.

8. If intermediate CAs are used to sign the certificate request, create a single file that contains the contents of the signed certificate and the contents of each intermediate CA certificate, concatenated together. As of version 9.1.0, the full certificate chain, including the root CA certificate, is mandatory and must be included when installing a certificate for the default use case.
9. Go to specific certificate, select **Install certificate**.
10. On the **Upload Signed Certificate** page, select the new signed certificate that you downloaded to your device or laptop.
11. Click **Install**. This action uploads the certificate onto the system. For more information, see .
12. If you are using IBM Security Guardium Key Lifecycle Manager key server, which does not currently support chain of trust checking, export the new system certificate and install it on the key server.
13. If you are using multifactor authentication with IBM Security Verify, which uses the system certificate as a signer certificate, export the new system certificate and install it as a signer certificate in IBM Security Verify. The management GUI is unavailable until the new certificate is added as a signer certificate to IBM Security Verify.

Using the command-line interface (CLI)

In the command-line interface, enter the following command to create a new certificate request and upload the new certificate:

1. • Enter the **chsystemcert** command to create a new certificate request: (only for default)

   ```
   chsystemcert -mkrequest -keytype ecdsa521 -country GB -state Hampshire -locality
   Hursley -org MYCO -orgunit Storage -commonname svcsystem1.myco.com -email admin@myco.com
   -subjectalternativename "DNS:test.ibm.com"
   ```

   • Enter the **mksystemcertstore** command to create a new main system certificate request:

   ```
   svctask mksystemcertstore -externalsigned -scope default -keytype ecdsa521 -country
   GB -state Hampshire -locality Hursley -org MYCO -orgunit Storage -commonname
   svcsystem1.myco.com -email admin@myco.com -subjectalternativename "DNS:test.ibm.com"
   ```

   The default certificate request is automatically written to `/dumps/certificate.csr`. For more information on delimiters, see “Updating or creating an internally signed certificate” on page 394.
2. Use the Secure Copy (SCP) or SSH File Transfer Protocol (SFTP) to copy the file `/dumps/certificate.csr` from the system to your local machine. Share the generated CSR file to the trusted third-party CA. If the CA is a public CA, then it may take some time for the CA to verify your identity before issuing the signed certificate. When it is ready, download the signed certificate file from the CA. You should also download any intermediate CA certificates that were used to sign the request. The files must all be in PEM format.
3. If the intermediate CA's are used to sign the certificate request, create a single certificate chain file that contains the contents of the signed certificate and the contents of each intermediate CA certificate, concatenated together. As of version 9.1.0, the full certificate chain, including the root CA certificate, is mandatory and must be included when installing a certificate for the default use case.
4. Use the SCP or SFTP to copy the certificate back onto the system in the file `/dumps/certificate.pem`, where `certificate.pem` is the name of the certificate.
5. After you copy the signed CA certificate to the system, enter the following command:

   • Using **chsystemcert** command: (only for default)

   ```
   chsystemcert -install -file /dumps/certificate.pem
   ```

   • Using **chsystemcertstore** command:

   ```
   chsystemcertstore -install -scope default -file /dumps/certificate.pem
   ```

   where **/dumps/certificate.pem** is the absolute path name of the signed certificate or certificate chain file.

For more information, see the **chsystemcert** or **chsystemcertstore** CLI command.

Export system certificate

You can export the current root certificate and system certificate by using the management GUI or command line interface (CLI).

You must add the system's root certificates and system certificates of the system to the trust stores of other systems, web browsers, key servers, or any other devices to authenticate the identity and secure the communication channel. If the browser, device, or application supports chain of trust checking, then only the root certificate is needed.

Using the management GUI

To export the system certificate, complete these steps:

1. In the management GUI, select **Settings** > **Security** > **Certificates**.
2. Go to the specific certificate, click the overflow menu and select **Export**. Select End point Certificate or Root CA Certificate.
3. In the **Export Certificate** page, you can either download the certificate file by clicking **Export** or copy the contents of the file to clipboard by clicking **Copy to clipboard**.

Using the command line interface (CLI)

In the command-line interface, enter the following command to export the certificate:

1. • Enter the following command to export the certificates to the file `/dumps/certificate.pem` (only for default)

   ```
   chsystemcert -export
   ```

   Note: The **chsystemcert -export** command exports the full chain of trust. Alternatively, to export the system's root certificate before the system initialization, use the **satask exportrootcertificate** command. Run this command on the node that is used to do the system initialization. You can also use the **sainfo lsrootcertificate** command to view the details of the system's root certificate.

   • Enter the following command to export the certificates to the file `/dumps/certificate.pem` (only for default) or `/dumps/system_certificate_slot_<usecaseid>.pem`

   ```
   svctask chsystemcertstore -export -scope default
   ```

   Note: The **chsystemcertstore -export -scope default** command exports the full chain of trust. The certificate is exported to the `/dumps/system_certificate_slot_<usecaseid>.pem` directory on the configuration node. For default slot, the certificate is exported to `/dumps/certificate.pem`.
2. • Enter the following command to export the Root CA certificate to file `/dumps/rootcacertificate.pem` (only for default)

   ```
   chsystemcert -exportrootcacert
   ```

   For more information, see **chsystemcert** command.

   • Enter the following command to export the Root CA certificate to file `/dumps/rootcacertificate.pem` (only for default) or `/dumps/system_rootcacertificate_slot_<usecaseid>.pem` (for non default scope)

   ```
   svctask chsystemcertstore -exportrootca -scope default
   ```

   For more information, see **chsystemcertstore** command.

Note: The certificate is exported to the `/dumps/system_rootcacertificate_slot_<usecaseid>.pem` directory on the configuration node. For default slot, the certificate is exported to `/dumps/rootcacertificate.pem`.

3. After the certificate is exported, use the Secure Copy (SCP) or SSH File Transfer Protocol (SFTP) to copy the certificate from the system.

Removing system certificate

You can remove certificate by using the management GUI or command-line interface (CLI).

Using the management GUI

To remove the system certificate, complete these steps:

1. In the management GUI, select **Settings** > **Security** > **Certificates**.
2. Go to specific certificate, click the overflow menu and select **Remove**.

Using the command-line interface (CLI)

Enter the following command to remove the specific certificate:

```
rmsystemcertstore keyserver
```

Note: You can not remove the certificate for scope value 'default' (ID 0).

For more information, see the **rmsystemcertstore** CLI command.

Partition Certificates

You can manage the partition certificates to enable secure access to partitions from external application or clients such as plugins or a REST API client using management GUI or command-line interface (CLI).

About this task

Before you create partition certificates, determine how certificates are signed based on your current security requirements. The system supports the following ways to sign partition certificates:

Internally signed certificate

Externally signed certificate
For more information on internally and externally signed certificates, see .

Consider the following additional information about certificates:

Root certificate
For more information on root certificates, see .

Partition certificate
The system certificate is signed by the system's root CA or a trusted third-party CA. This certificate is presented to other devices such as external applications or clients such as plugins or a REST client, to authenticate the partition and create a secure connection.

The system certificate supports the following key types:

- RSA 4096
- RSA 2048
- ECDSA 521
- ECDSA 384

SSL protocol level 4 restricts the use of RSA key exchange cipher suites. Therefore, you cannot enable SSL protocol level 4 when using a certificate with an RSA key. Likewise, you cannot generate a certificate with an RSA key if SSL protocol level 4 is set.

Configuring partition certificates

Secure connections enable external applications or clients, such as plug-ins or REST clients, to access partitions. To ensure that the certificate and public ID of the partition are valid and secure, each partition has a digital Secure Sockets Layer (SSL) certificate.

About this task

To manage the certificate that is installed on a system, use either the **svctask mkpartitioncertstore**, **svctask chpartitioncertstore**, **svctask rmpartitioncertstore** commands in the command-line interface (CLI) or click **Settings** > **Security** > **Certificates** in the management GUI to do the following tasks:

- Upload signed certificate
- Update certificate
- Cancel outstanding signing request
- Export partition certificate
- Export root certificate
- Remove partition certificate

For specific directions on completing these tasks, see "Partition Certificates" on page 400.

Updating or creating an internally signed partition certificate

You can create an internally signed partition certificate by using the management GUI or command-line interface (CLI).

Using the management GUI

To configure an internally signed partition certificate, complete these steps:

1. If the root certificate has not already been exported, export the root certificate to other systems, web browsers, and devices that require secure communications with the system. For more information, see .
2. In the management GUI, select **Settings** > **Security** > **Certificates**.
3. On the Certificates page, the certificate details are displayed.

   Automatic renewal of the system certificate

   The partition certificate can be renewed automatically if it is signed by the system root CA. Turn automatic renewal on by going to **Settings** > **Security** > **Certificates** > **Go to the particular certificate** and set automatic renewal to **On**. The default validity period of the system certificate is one year. If automatic renewal is **On**, the system certificate is renewed thirty days before the expiry date. If the validity period of the system certificate is fewer than thirty days, the system attempts renewal during its next scheduled check (performed every eight hours).

   The renewed certificate contains all of the same field values, key type and validity period details as the previous certificate.
4. Go to any partition certificate, click the overflow menu and select **Replace** or click **Create** and then select the scope of the certificate from the drop-down menu.
5. Select **Internally Signed Certificate** for the certificate type.
6. If you are already using certificates, the **Certificate Details** are automatically populated. You can update any of the following details:

   Key type
   Select the cryptographic key type that is used to generate the certificate.

Validity days
Enter the number of days the certificate is valid for. The maximum number of days that are allowed is 9000.

Country
Enter the two-letter country code or location, for example, 01 for US.

State
Enter the name of the state where the system requesting the certificate is located.

City
Enter the name of the city where the system is located.

Organization name
Enter the name of the organization.

Organizational unit
Enter the name of organizational unit.

Common name
Common name is UUID of current partition.

Subject alternative name
Subject alternative name is the IP address of the current partition.

Subject Alternate Name field:

```
IP:196.192.0.20
```

Email Address
Enter the email address.

7. Click **Create and install**. The certificate is updated in the main panel.

Using the command-line interface (CLI)

To generate an internally signed partition certificate that uses RSA 2048 key type and expires in one year, enter the following command:

```
svctask mkpartitioncertstore -systemsigned -partition partition0 -commonname BC30C5EB-
FAB9-59FE-9880-5F7F479753B8 -country GB -locality Manchester -org IBM -orgunit Systems
-email certificates@support.ibm.com -keytype rsa2048 -validity 365 -subjectalternativename
"IP:192.165.10.15"
```

For more information, see the **mkpartitioncertstore** CLI command.

Requesting and installing an externally signed partition certificate

You can request a new signed partition certificate from a third-party certificate authority. The externally signed partition certificate can be installed by using the management GUI or command-line interface (CLI).

Note: Before requesting and installing the partition certificates, take the following factors into consideration:

- The certificate signing request must be generated on the system. IBM Storage Virtualize does not support certificate signing requests that are generated outside of the system.
- The signed certificate must contain the following X509v3 Key Usage extensions: Digital Signature. It must also contain the following X509v3 Extended Key Usage Extensions: TLS Web Server Authentication and TLS Web Client Authentication. Ensure that the Certificate Authority (CA) used to sign the certificate includes these extensions.

In the management GUI, select **Settings** > **Security** > **Certificates** and select **Replace** from the overflow menu.

In the Certificates page, select externally signed certificate and complete the form to create a request for a signed certificate for specific partition. The resulting certificate signing request (CSR) should be exported to the third-party certificate authority so that it can be signed. After you receive the externally signed certificate from the third-party certificate authority, use the management GUI or command line interface (CLI) to upload and install the signed certificate on the system.

Using the management GUI

To configure an externally signed partition certificate, complete these steps:

1. In the management GUI, select **Settings** > **Security** > **Certificates**.
2. Click **Create** and then select the scope of the certificate as **Partition** from the drop-down menu.
3. Select **Externally Signed Certificate** for the certificate type.
4. If you are already using certificates, the **Certificate Details** are automatically populated. You can update the details. For more information on the certificate details, see step 6 in .
5. Click **Generate CSR**. This action presents a dialogue window to download the file. Select a location on your local machine to save the file.
6. If the signing request is generated, you have the option to cancel the outstanding signing request in case of an error. To cancel any outstanding signing request, complete these steps:
   a. Select **Settings** > **Security** > **Certificates**.
   b. Go to a specific certificate, select **Cancel CSR** and click **Yes**.
7. Share the generated CSR file to the trusted third-party CA. If the CA is a public CA, then it can take some time for the CA to verify your identity before issuing the signed certificate. When it is ready, download the signed certificate file from the CA. You can also download any intermediate CA certificates that were used to sign the request. The files must all be in PEM format.
8. If intermediate CAs are used to sign the certificate request, create a single file that contains the contents of the signed certificate and the contents of each intermediate CA certificate, concatenated together. The root CA certificate is not mandatory, but can optionally be included.
9. Go to specific certificate, select **Install certificate**.
10. On the **Upload Signed Certificate** page, select the new signed certificate that you downloaded to your device or laptop.
11. Click **Install**. This action uploads the certificate onto the system. For more information, see .

Using the command-line interface (CLI)

In the command-line interface, enter the following command to create a new partition certificate request and upload the new partition certificate:

1. Enter the following command to create a new partition certificate request:

```
svctask mkpartitioncertstore -externalsigned -partition partition0 -commonname BC30C5EB-
FAB9-59FE-9880-5F7F479753B8 -keytype ecdsa521 -country GB -state Hampshire -locality Hursley
-org MYCO -orgunit Storage -email admin@myco.com -subjectalternativename "IP:192.165.10.15"
```

   The certificate request is automatically written to `/dumps/partition_certificate_slot_<partitionid>.csr`.
2. Use the Secure Copy (SCP) to copy the file `/dumps/partition_certificate_slot_<partitionid>.csr` from the system to your local machine. Share the generated CSR file to the trusted third-party CA. If the CA is a public CA, then it may take some time for the CA to verify your identity before issuing the signed certificate. When it is ready, download the signed certificate file from the CA. You should also download any intermediate CA certificates that were used to sign the request. The files must all be in PEM format.
3. If the intermediate CA's are used to sign the certificate request, create a single certificate chain file that contains the contents of the signed certificate and the contents of each intermediate CA

certificate, concatenated together. The root CA certificate is not mandatory, but can optionally be included.

4. Use the SCP to copy the certificate back onto the system in the file `/dumps/certificate.pem`, where `certificate.pem` is the name of the certificate.
5. After you copy the signed CA certificate to the system, enter the following command:

```
svctask chpartitioncertstore -install -partition partition0 -file /dumps/certificate.pem
```

   where **/dumps/certificate.pem** is the absolute path name of the signed certificate or certificate chain file.

For more information, see the **chpartitioncertstore** CLI command.

Export partition certificate

You can export the current root certificate and partition certificate by using the management GUI or command line interface (CLI).

Using the management GUI

To export the partition certificate, complete these steps:

1. In the management GUI, select **Settings** > **Security** > **Certificates**.
2. Go to the partition certificate, click the overflow menu and select **Download**. Select End point Certificate or Root CA Certificate.
3. In the **Export Certificate** page, you can either download the certificate file by clicking **Export** or copy the contents of the file to clipboard by clicking **Copy to clipboard**.

Using the command line interface (CLI)

In the command-line interface, enter the following command to export the certificate:

1. Enter the following command to export the certificates to the file `/dumps/partition_certificate_slot_<partitionid>.pem`

```
svctask chpartitioncertstore -export -partition partition0
```

   Note: The **svctask chpartitioncertstore -export -partition partition0** command exports the full chain of trust. The certificate is exported to the `/dumps/partition_certificate_slot_<partitionid>.pem` directory on the configuration node.
2. Enter the following command to export the Root CA certificate to file `/dumps/partition_certificate_slot_<partitionid>.pem`

```
svctask chpartitioncertstore -exportrootca -partition partition0
```

   Note: This command exports only the root certificate for the specific partition. The certificate is exported to the `/dumps/partition_rootcacertificate_slot_<partitionid>.pem` directory on the configuration node.

   For more information, see **chpartitioncertstore** command.
3. After the certificate is exported, use the Secure Copy (SCP) to copy the certificate from the system.

Removing partition certificate

You can remove partition certificate by using the management GUI or command-line interface (CLI).

Using the management GUI

To remove the partition certificate, complete these steps:

1. In the management GUI, select **Settings** > **Security** > **Certificates**.
2. Go to the partition certificate, click the overflow menu and select **Remove**.

Using the command-line interface (CLI)

Enter the following command to remove the partition certificate:

```
svctask rmpartitioncertstore partition1
```

For more information, see **rmpartitioncertstore** command.

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.

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.

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.

Related information

"Configuring multifactor authentication with IBM Security Verify" on page 406
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 411
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 421
The system integrates with Duo Security to provide multifactor authentication for system users.

"Configuring user groups for multifactor authentication with Duo Security" on page 422
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 77. 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 78. 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.
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: https://hostname.com/mfa https://hostname/service/mfa

Table 78. Sign-on tab (continued)		
Field	**Action**	**Details**
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.	
Access policies	Complete these steps: a. Deselect **Use default policy**. b. Click the **Edit** icon. c. Select **Always require 2FA in all devices**. 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 79. 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.

Related information

“Multifactor authentication” on page 405
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.

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** and 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.

2. 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** and 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.

   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** and 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.

2. 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** and 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.

   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.

Changing MFA credentials of system

You can change multifactor authentication credentials for the system with the management GUI or command-line interface (CLI).

If credentials are updated on the multifactor authentication provider, these new credentials must be changed on the system. If the credentials are not updated on the system, the system cannot communicate with the authentication providers and user cannot be authenticated to access the system.

Using the management GUI

To update multifactor authentication credentials on the system, complete these steps:

1. Select **Settings** > **Security** > **Multifactor Authentication**.
2. Select the **Authentication provider** to change the credentials.
3. In **Authentication server** enter the correct host name and port number.
4. For changing credentials for IBM Security Verify on Storage Virtualize GUI, enter the new Client ID and Client secret under the **OpenID Credentials** heading.
5. For changing credentials for IBM Security Verify on Storage Virtualize CLI, enter the new Client ID and Client secret under the **API Client Credentials** heading.
6. For changing credentials for Duo Security on Storage Virtualize GUI, enter the new Client ID and Client secret under the **OpenID Credentials** heading.
7. For changing credentials for Duo Security on Storage Virtualize CLI, enter the new Integration key and Secret key under the **API Client Credentials** heading.
8. Click **Save**.

Using the command-line interface

- Use the command **chauthmultifactorverify** to change the credentials for IBM Security Verify. By running the following command you can change the credentials.

```
chauthmultifactorverify -openidclientid openid_client_id
                        -openidclientsecret openid_client_secret
                        -cliclientid cli_client_id
                        -cliclientsecret cli_client_secret
```

  where *openid_client_id* specifies the IBM Security Verify OpenID Connect client ID for the system. It is required to enable the multifactor authentication (MFA) for login to the Storage Virtualize GUI. The *openid_client_secret* specifies the IBM Security Verify OpenID Connect client secret for the system. The *cli_client_id* specifies the IBM Security Verify CLI client ID for the system. It is required to enable the

multifactor authentication (MFA) for login to the Storage Virtualize CLI. The *cli_client_secret* specifies the IBM Security Verify CLI client secret for the system.

- Use the command **chauthmultifactorduo** to change the credentials for Duo Security. By running the following command you can change the credentials.

```
chauthmultifactorduo -openidclientid openid_client_id
                     -openidclientsecret openid_client_secret
                     -integrationkey integration_key
                     -secretkey secret_key
```

where *openid_client_id* specifies the Duo Security OpenID Connect client ID for the system. It is required to enable the multifactor authentication (MFA) for login to the Storage Virtualize GUI. The *openid_client_secret* specifies the Duo Security OpenID Connect client secret for the system. The *integration_key* specifies the Duo Security CLI integration key for the system. It is required to enable the multifactor authentication (MFA) for login to the Storage Virtualize CLI. The *secret_key* specifies the Duo Security CLI secret key for the system.

Changing the failmode parameter for multifactor authentication

The failmode parameter determines how the system responds if a user logs into the system when multifactor authentication is enabled, but the authentication server is unavailable. Failmode setting can be managed by command-line interface (CLI).

Use the failmode parameter in the **chauthmultifactorverify** command to set one of the following values:

secure
When the authentication server is unavailable, specify this value to prevent any user with multifactor authentication enabled from accessing the system. When a user logs into the system, the login attempt fails.

insecure
When the authentication server is unavailable, specify this value to allow any user with multifactor authentication enabled to access the system. When a user logs into the system, the login attempt succeeds.

Enter the following command to prevent access to the system for users with multifactor authentication enabled:

```
chauthmultifactorverify -failmode secure
```

Enter the following command to allow access to the system for user with multifactor authentication enabled:

```
chauthmultifactorverify -failmode insecure
```

User authentication

You can configure authentication and authorization for users of the system. The system supports multifactor authentication, single sign-on, local user authentication and LDAP authentication for remote users.

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.

In addition to these first factor authentication methods for local users, you can enable the system to require multifactor authentication for all local users defined in a user group. To use multifactor authentication on the system, you need to configure a supported authentication service and enable multifactor authentication on the system and user groups. For example, the system integrates with IBM Security Verify that provides second factor authentication. Local users must be also manually added to the supported authentication service and set up their second factors.

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.

The system also supports requiring multifactor authentication for remote users. As with local users, you must configure a supported authentication service and enable the feature on the system and remote user groups. However, remote users can be automatically managed with IBM Security Verify Bridge for Directory Sync. 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 information on how to configure directory synchronization for remote users with Duo Security, see Duo Directory Synchronization.

In addition to multifactor authentication, the system also supports single sign-on for remote users only. Remote users are authenticated to all applications through a single set of credentials. Single sign-on requires that the feature is enabled at the system level and on user groups.

The system supports several authentication methods. However, it's important to understand the key differences and limitations of each authentication method. The following table summarizes the differences and limitations of each supported authentication method:

Table 80. Authentication methods comparison

	Multifactor Authentication	**Single Sign-on**	**LDAP authentication**
First factor handled by	IBM Storage Virtualize	Identity Provider (IdP)	LDAP server
Second factor handled by	Multifactor authentication service (for example, IBM Security Verify)	Identity Provider (IdP)	NA
CLI login support	Yes	No	Yes
GUI login support	Yes	Yes	Yes
REST API support	No	No	Yes
Local user support	Yes	No	No
Remote user support	Yes	Yes	Yes

Table 80. Authentication methods comparison (continued)			
	Multifactor Authentication	**Single Sign-on**	**LDAP authentication**
Supported services	• IBM Security Verify • Duo Security	• IBM Security Verify • Microsoft Active Directory Federation Services • Microsoft Azure Active Directory • Okta	• Microsoft Active Directory • IBM Security Directory Server • Others (for example, OpenLDAP)
Protocol	OpenID Connect (OIDC)	OpenID Connect (OIDC)	LDAP

Related information

"Managing users and user groups" on page 224
You can use the management GUI or the command-line interface to manage user and user groups on the system.

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** and then select the type of LDAP server that is used for authentication.
3. 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.
4. 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.

5. 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.

6. 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** and 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.

   b. 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** and 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.

      Repeat these steps for all remote users.

7. 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 X.XX.XX.XXX -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.

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 405. 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 81. Authentication methods 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
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.

	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 • IBM Security Directory Server • 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” on page 453.

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 432.

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 82. 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 PingFederate

PingFederate can be configured as the authentication provider for the system.

Note:

- To configure single sign-on with PingFederate, you must upgrade the system to 8.7.0. or later release.
- 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 432.

Prerequisites

The following prerequisite steps on PingFederate must be completed before you can configure single sign-on on the system:

1. Install PingFederate. For more information, see Installing PingFederate.
2. Configure your DataStore where you will create your users and user groups. For example, Configure LDAP/ Active Directory/ Ping Directory.
3. Once your PingFederate server is up and running, you can launch it using:

   ```
   https://<pingfederateserver>:9999/pingfederate/app#/
   ```

   Login using Administrator credentials to setup PingFederate server.
4. Go to **SYSTEM** > **OAuth Settings**.
   a. Under Scope Management, add common scope “openid”.
   b. Under Authorization Server Settings, configure Token Endpoint Base URL to `https://<pingfederateserver>:9031`. Also, update “Password Credential Validator”, “Access Token Manager” once you configure it in Step 6 and Step 11.

      For “Required Scope” mention “openid”.
5. Configure Data Store that is mentioned in Step 2. Go to **SYSTEM** > **Data Stores**. For more information, see Managing datastores.
6. Configure Password Credential Validator, under **SYSTEM** > **Password Credential Validators**. For more information, see Password Credential Validators.

   Note: Select Type as “LDAP Username Password Credential Validator”
7. Configure IdP Adapters. Go to **AUTHENTICATION** > **IdP Adapters**.
   - For type, select "HTML Form IdP Adapter".
8. Create Policy Contract. Go to **AUTHENTICATION** > **Policies** > **Policy Contracts**.
9. Create Policy. Go to **AUTHENTICATION** > **Policies** > **Policies**. Select checkbox for **IDP AUTHENTICATION POLICIES**. Create New Policy by clicking “Add Policy”. This policy determines what happens on login failure as well as login success. Here, under POLICY, select IdP Adapter created in Step 7. For FAIL, “Done” can be selected. And for SUCCESS, select Policy Contract created in Step 8.
10. Go to **AUTHENTICATION** > **Policy Contract Grant Mapping**. Create Mapping of Policy Contract.
11. Configure fields in Access Token Management. Go to **Applications** > **OAuth** > **Token Management**. Configure settings for type “JSON Web Tokens”.
12. Add FlashSystem system details into PingFederate now. Go to **Applications** > **OAuth** > **Clients**.
    a. Make a note of value that you provide for CLIENT ID. This value will need to be provided in FlashSystem later for “client id”.
    b. For CLIENT AUTHENTICATION, select “CLIENT SECRET”.
    c. Make a note of value that will get auto generated in CLIENT SECRET. This value will need to be provided in FlashSystem later for “client secret”.
    d. For redirection URLs, enter `https://<flashsystem-cluster>/sso`.
    e. For ALLOWED GRANT TYPES, select ‘Authorization Code’ and ‘Implicit’.
    f. Select checkbox for “REQUIRE PROOF KEY FOR CODE EXCHANGE”.

    For rest of the fields, default values can be kept.
13. Configure Access Token Mappings. Go to **Applications** > **OAuth** > **Access Token Mappings**. Add Mapping between policy contract and access token manager created in step before.
14. Configure fields that will get returned in ID Token. Go to **Applications** > **OAuth** > **OpenID Connect Policy Management**. Create new policy.

a. Under ACCESS TOKEN MANAGER, select the access token manager that is created in step before.

b. In the Extended contract, add attributes “group”, “preferred_username”.

c. In Attribute Scopes, map above attributes under openid scope.

d. For Attribute Sources & User Lookup, add your DataStore details as Attribute source. Fill details like LDAP Directory Search & LDAP Filter, which will match to your schema.

e. Under Contract Fulfilment, fill appropriate source for all the attributes.

   i) For ***sub***, select source as Access Token and Value as sub.

   ii) For ***preferred_username***, select source as Access Token and Value as sub.

   iii) For ***group***, ensure that the PingFederate returns group information in such a way that group names will match the group names created on FlashSystem.

Note: If you want to create group names only with CN, then in PingFederate, select source as “Expression” and in value give OGNL Expression in such a way that it will return group names in comma separated array to client (for example, FlashSystem GUI).

Using the management GUI

To configure single sign-on with PingFederate, complete these steps:

1. Select **Settings** > **Security** > **Single Sign-on**.
2. Enter the OpenID Configuration Endpoint URL of the authentication server as:

   ```
   https://<pingfederateserver>:9031/.well-known/openid-configuration
   ```

3. For the **OpenID Credentials**, add the **Client ID** and **Client Secret** that you copied earlier in prerequisites in Step 12
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://<pingfederateserver>:9031/.well-known/
openid-configuration -clientid xxxxxxxx-xxxx-xxxx-xxxx-xxxxxxxxxxxx -clientsecret xxxxxxxx
-userclaim preferred_username -groupclaim group -enable
```

The values for the **-clientid** and the **-clientsecret** are the Client ID and Client Secret that you copied earlier in prerequisites in Step 12. 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 432.

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 **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 83. 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 84. 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.

Table 84. 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 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, 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.
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: a. Deselect **Use default policy**. b. Click the **Edit** icon. c. Select **Always require 2FA in all devices**. 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.

Table 84. Sign-on tab (continued)		
Field	**Action**	**Details**
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 432.

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 85. 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 86. 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.	

Table 86. Server application (continued)		
Field	**Action**	**Details**
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: **Hostname** `https://hostname/sso` **Fully-qualified domain name** `https://hostname.com/sso` **IP address** `https://1.2.3.4/sso` **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 87. Configure Application Credentials		
Field	**Action**	**Details**
Generate a shared secret	a. Select **Generate a shared secret**. b. After the shared secret is automatically generated, click **Copy to clipboard**. **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 88. 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. b. Click **Add**	

12. Click **Next**.
13. The following table shows the required fields and actions for the **Choose Access Control Policy** step.

Table 89. Choose Access Control Policy		
Field	**Action**	**Details**
Choose an access control policy	From the options, select one of the following supported access policies: **Permit everyone** Select this option if you are only configuring single sign-on without any additional factors to authenticate users. **Permit everyone and require MFA** 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 90. Configure Application Permissions		
Field	**Action**	**Details**
Permitted scopes	Select the following scopes: • Select **allatclaims**. • 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 Entra ID

Microsoft Entra ID can be configured as the authentication provider for the system.

Prerequisites

The following prerequisite steps on Microsoft Entra ID must be completed before you can configure single sign-on on the system:

1. Create a subscription for Microsoft Entra ID.

2. Make sure Microsoft Entra ID tenant is reachable from FlashSystem.
   a. Eighter configure proxy server to reach to Microsoft Entra ID tenant from FlashSystem.
   b. Modify firewall rules to allow traffic between Microsoft Entra ID and FlashSystem.
3. For more information, see the prerequisites in “Configuring single sign-on” on page 432.

Configuration on Microsoft Entra ID

To configure single sign-on with Microsoft Entra ID, complete these steps:

1. In the Microsoft Entra ID interface, from the menu select **Identity** > **Application** > **App registrations** > **New registration**.

Table 91. Register an application page

Field	Action
Name	Enter a name to identify FlashSystem on Microsoft Entra ID.
Supported account types	Select **Accounts** in this organizational directory only.
Redirect URI	Select **Web**, `https://cluster_hostname/sso` `https://cluster_ip/sso`

2. Click **Register**. After the system is saved as a new app registration, the application reloads with the **Overview** page selected.

Copy the important information from the application overview page

To find the OpenID Endpoint configuration, select **Overview** > **Endpoints** and copy the OpenID Connect metadata document.

On the **Overview** page, copy the **Application (client) ID**.

Go to **Manage** > **Certificates & secrets** and click **New client secret**. Add appropriate description and copy the value displayed in "Value".

Create group and user claims

For group claims:

- Go to **Token configuration** and click **Add groups claim**. For setting **Select group types to include in Access, ID, and SAML tokens**, select checkbox **Security groups**.
- For **Customize token properties by type**, select **sAMAccountName** for ID, Access, and SAML.

For user claims:

- Click **Add optional claim**. Select token type as ID and select the checkbox for “preferred_username”.
- Click **Add optional claim**. Select token type as Access and select the checkbox for “preferred_username”.

Enable ID token for the application

Navigate to **Authentication** tab. For **Implicit grant and hybrid flows**, ensure **ID tokens** (used for implicit and hybrid flows) is selected.

Configure users and group in Microsoft Entra ID

Create users in Microsoft Entra ID and add users in the group. Make sure when you navigate to desired user, it reflects correct group memberships.

Create group and copy the “Object ID” of that group.

Assign the desired users to your application in Microsoft Entra ID

- Navigate to **Identity** > **Application** > **Enterprise applications** > **Your application**.
- Navigate to **Users and groups** and click **Add user/group**.
- Select the appropriate users.

Configuration on FlashSystem

To configure single sign-on with FlashSystem, complete these steps:

Single sign-on settings

- Go to **Settings** > **Security** > **Single Sign-on**.

  Fill up information for OpenID configuration endpoint, Client ID, and Client Secret from the information copied in Microsoft Entra ID steps.

  For User claim, give string **preferred_username** and for group claims, give string **groups**.
- 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, select the checkbox and ensure a proxy server is defined on the system.
- Click **Save**. On the confirmation page, click **Confirm** to enable single sign-on for the system.

Configure groups on FlashSystem

- Create user group with the name that matches the Object ID of the group that is created on Microsoft Entra ID.
- Select the role you wish to assign and then select the checkbox for **Remote Authentication**.

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 432.

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 92. 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 93. 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 432.

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 94. 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 95. 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.
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, `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 96. OIDC Response section		
Field	**Action**	**Details**
Scopes	Select **openid** and **profile**.	After selecting profile, a list of IdP Attribute and Claim mappings will appear. 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. 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 97. 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:

```
chauthsinglesignon -oidcconfigurationendpoint https://sso-tenant.sso.duosecurity.com/oidc/
clientid/.well-known/openid-configuration -clientid xxxxxxxx-xxxx-xxxx-xxxx-xxxxxxxxxxxx
-clientsecret xxxxxxxx -userclaim username -groupclaim groups -enable
```

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
```

Supported authentication providers

The system uses the OpenID Connect (OIDC) protocol to communicate with supported authentication providers. The system does not currently support the Security Assertion Markup Language (SAML) protocol.

Multifactor authentication
The system supports the following MFA providers:

- IBM Security Verify
- Duo Security

Single Sign-on
The system supports the following SSO providers:

- IBM Security Verify
- Microsoft Active Directory Federation Services (AD FS)
- Microsoft Entra ID
- Okta
- Duo Security
- PingOne
- PingFederate

Single sign-on providers are often used to provide multifactor authentication in addition to single sign-on. For example, Microsoft AD FS supports a range of third-party authentication providers for additional multifactor authentication capability such as Duo Security, RSA and Okta. For more information, see Configure Additional Authentication Methods for AD FS.

Note: Connections to authentication providers might break if you change the SSL security level. This happens if the authentication providers do not support any of the ciphers and protocols allowed at the set level, particularly TLS 1.3. In this case, you can revert to a lower level of security to restore connectivity.

Managing proxy server for SSO login

When using single sign-on, the system must be able to contact the single sign-on service.

Consider your network settings and whether the single sign-on service is accessible to the system. If the single sign-on service is not accessible, then firewall rules need to be added to allow communication, or alternatively a proxy server can be used. In this case, a proxy server can be used to permit traffic between the system and the single sign-on service.

Note: If the firewall rules or proxy server is not configured correctly, the system is not able to access the single sign-on service and users are not able to authenticate to the system by using single sign-on.

Using the management GUI

Complete the following steps:

1. Select **Settings** > **Security** > **Single-Sign on**.
2. In **Proxy server**, check the box **Use proxy server** so single sign-on can use proxy server.
3. Click **Save**.

Using the command-line interface

Use the command **chauthsinglesignon** to specify whether the proxy server should be used.

```
chauthsinglesignon -proxy yes/no
```

If proxy value is set to *yes* then the traffic is routed through the proxy defined on the system. If proxy value is set to *no* then the traffic is not routed through the proxy server. The proxy server can be created through **mkproxy**, for more information refer to **mkproxy command**.

Changing SSO credentials of system

Changing credentials for the Storage Virtualize application can be completed through management GUI or command-line interface (CLI).

If the security administrator changes the single sign-on credentials used for the application, the new credentials must be updated on the system. If the credentials are not updated, the system will not be able to communicate with the authentication service and users will not be able to authenticate successfully to the system.

Using the management GUI

Complete the following steps:

1. Select **Settings** > **Security** > **Single Sign-on**.
2. In **Authentication server** enter the OpenID configuration URL for the SSO provider.
3. Select **OpenID Credentials** and enter the new client ID and client secret.
4. Click **Save**.

Using the command-line interface

Use the command **chauthsinglesignon** to change the credentials. By running the following command you can change the credentials.

```
chauthsinglesignon -clientid client_id
                   -clientsecret client_secret
```

where *clientid client_id* specifies the OpenID Connect client ID for the system, required to authenticate with the single sign-on provider. The *clientsecret client_secret* specifies the OpenID Connect client secret for the system, required to authenticate with the single sign-on provider.

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 help 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 methods.
- 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 unauthorized proxy 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 helps ensure 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 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.

Key management methods

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.

Before you activate and enable encryption, you must determine the method of accessing key information during times when the system requires an encryption key to be present. The system requires an encryption key to be present during the following operations:

- System power-on
- System restart
- User initiated rekey operations
- System recovery

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 help 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 using key servers or internal key management method to manage encryption keys instead.

It is possible to simultaneously configure USB flash drives and key servers to help 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.

You can migrate from the existing external methods (USB flash drives and key servers) to internal key management method without affecting the encrypted objects.

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.
- It is not recommended to configure the external methods along with the internal key management method.
- It is recommended to configure encryption recovery key while enabling the encryption using external methods or internal key management method.

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).

Encryption on cloud accounts

If you are using encryption to protect data that is copied to cloud storage, the cloud account is always synchronized with the system encryption settings. If both USB flash drives and key servers are configured, the cloud account that is created supports both of these methods. If just one encryption method is configured and the other is disabled, the cloud account supports encryption with the remaining configured encryption method. To ensure that the cloud account supports encryption, one or both methods must be configured with active keys when the cloud account is created.

Note: The internal key management method does not support encrypted cloud accounts.

If a cloud account is created with one encryption method, you can configure the second method later, but the cloud account must be online while the configuration occurs. After the second method is configured, the cloud account will support both key providers.

Licensing encryption

Before you can configure encryption on the system, you must have the encryption license. If you intend to use encryption of data-in-flight to secure IP connections between partnered systems, you also require an encryption license.

On current-generation systems running IBM Storage Virtualize 9.1.2 or later, the encryption license is installed at manufacturing and no additional installation steps are required.

On older-generation systems, the encryption license must be installed and activated by using the procedures described in this document before encryption can be configured. If you have not already purchased a license, contact a customer representative.

Before you begin

To obtain license keys, you need the machine type and model (MTM), serial number (S/N), and machine signature to manually activate the keys.

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**.
2. 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 is displayed on the Properties page.

Using the management GUI

Within the management GUI, you can activate an encryption license on the system by two ways. 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 the encryption will be activated, and select **Actions** > **Activate License Manually**.
   b. On the **Activate License Manually** page, you must retrieve the license keys by completing the form at https://www.ibm.com/servers/eserver/ess/landing/index.html . 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 the Secure Copy (SCP) or SSH File Transfer Protocol (SFTP) 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 internal key management

Internal key management uses the internal boot drive to manage the main encryption key for the system.

Internal key management method provides user an efficient way to configure and manage the encryption without needing the external methods such as USB flash drives and key servers. It avoids the additional overhead of maintaining the external methods.

The internal key management method simplifies the encryption configuration by providing a one-step interface to manage the encryption. This method performs the rekey, for example, the regeneration of key operation every 24 hours internally and without requiring any user intervention. This helps in adhering to the security policies of the organization of regenerating the encryption key in regular intervals.

The internal key management method stores the main encryption key in an encrypted form on the internal boot drive of every node of the system. This method uses the onboard Trusted Platform Module (TPM) chip to encrypt the main encryption key. The system helps ensure that the key is securely distributed across all the nodes of the system, during enable process, rekey process, and recovery process.

Note: It is recommended to configure the encryption recovery key along with the internal key management method.

If you have configured the encryption by using the external methods such as USB flash drives or key servers, then you can switch to the internal key management method without affecting the encrypted objects configured on your system.

Note: It is not recommended to configure both internal key management and external methods.

Enabling encryption with internal key management

The system supports enabling encryption that uses the internal boot drive to store encryption keys. Encryption with internal key management can be enabled by using the management GUI.

Using management GUI

To enable the encryption with internal key management that uses the management GUI, complete these steps:

1. In the management GUI, select **Settings** > **Security** > **Encryption** and click **Enable Encryption**.
2. On the Welcome page, select **Internal Key Management**, then click **Next**.
3. After all the validation checks are successfully performed, the encryption key is generated and securely stored on the internal boot drive.
4. It is recommended to configure the Encryption Recovery Key along with the internal key management. For more information, see .

Using the command-line interface

To enable the encryption with internal key management, enter the following command:

```
chencryption -internal enable
```

The new key has been created successfully when the **`internal_key_status`** field shows enabled and **`internal_key_last_rekey_time`** field shows the timestamp of key generation upon executing **`lsencryption`** command. For more information, see **`chencryption`** and **`lsencryption`** commands.

Rekeying a system with internal key management

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.

The internal key management method performs rekey operation every 24 hours without user intervention that means the new key is generated every 24 hours.

The command-line interface or the management GUI option is unavailable to initiate rekey operation, as it is performed internally by the system.

After successful completion of the rekey operation,

- The last generated field in the **Encryption** page of management GUI shows the updated rekey time.
- The **`lsencryption`** command output shows the updated rekey time in **`internal_key_last_rekey_time`** field.

The rekey operation fails if any of the following conditions occur in the system,

- All the nodes in the system are not online.
- The system upgrade is in progress.
- A system is performing another encryption task using **`chencryption`** command.
- A rekey operation has timed out.
- A TPM module on one or more nodes have failed.

In such case, a system performs retry operation after a specified time interval. If after finite retry attempts for example, 6, the system fails to perform rekey operation then a fix procedure notify user that the rekey operation has failed. The system attempts the rekey operation after 24 hours. For more information, you can refer to the fix procedures.

Disabling encryption with internal key management

Encryption with internal key management can be disabled using the management GUI.

Note: For security, encryption methods (including the internal key management) 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, internal key management is automatically disabled in the process. See for instructions on disabling encryption using the management GUI.

To disable only the internal key management, refer to the instructions described in Using the command-line interface.

Using the command-line interface

Follow these steps to disable the internal key management:

1. Identify the configuration node of the system. For more information, see Configuration node.
2. Connect your computer to the technician port of the configuration node. For more information, see the Node canisters page in the relevant Hardware Guide.
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 internal key management, enter the following command:

```
chencryption -internal disable
```

For more information, see **chencryption** command.

5. The internal key management has been disabled successfully when the `internal_key_status` field is licensed. For more information, see **lsencryption** command.

Switching to and from the internal key management method

You can switch to internal key management from the existing methods (key servers or USB flash drives) or switch from internal key management to external methods without impacting the encrypted objects and data.

Ensure that the encryption recovery key is available for validation if configured on the system. Without validating the encryption recovery key, the migration between key management methods cannot be performed.

Note: Switching to internal key management method disables the currently enabled external method. This does not disable or regenerate the encryption recovery key.

Before you initiate migration (switch) to internal key management method, make sure that the internal key management is accessible.

To switch to internal key management from the existing methods (key servers or USB flash drives) or switch from the internal key management to external methods, go to **Settings** > **Security** > **Encryption** panel in the management GUI and select **Edit encryption method**.

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 the following key servers to handle key management on the system.

- IBM Security Guardium Key Lifecycle Manager key server
- Generic KMIP key servers
  - Thales CipherTrust Manager key server
  - Gemalto SafeNet KeySecure key server
  - Fortanix DSM key server
  - Utimaco ESKM key server
  - Entrust CSP Vault key server
  - HashiCorp Vault key server

Note: For more information on the supported key servers, see IBM Storage Virtualize Supported Key Servers.

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. Also, 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

When planning for key server encryption, the following items are important to consider.

SSL certificates

- IBM Storage Virtualize and key servers use SSL certificates to authenticate each other and create a secure connection.
- The system supports certificates signed by both the internal root certificate authority and a trusted thrid-party CA.
- The correct certificates and their certificate authority certificates must be installed in IBM Storage Virtualize and on the key servers to establish a secure connection.
- The connection between IBM Storage Virtualize and key servers becomes interrupted if certificates are renewed or expire. Installing the certificate authority certificates prevents this interruption by allowing the endpoint certificates to be renewed without disrupting the connection.
- IBM Security Guardium Key Lifecycle Manager key servers do not currently support chain of trust checking with IBM Storage Virtualize. The new system certificates generated must be exported to the key servers to re-establish a secure connection.
- If you are using multifactor authentication with IBM Security Verify, the management GUI is unavailable when you update the certificate. The new certificate generated must be exported using the CLI and added as a new signer certificate to IBM Security Verify for successful authentication.

System requirements

Only one type of key server is supported at this time. The system implements the KMIP that is sent over an SSL connection between the client and the server. Support is provided for internal-signed and external CA-signed certificates. The system validates the server's SSL certificate and conforms to the KMIP standard. Existing SAS hardware needs access to at least one master key to unlock and needs to be able to respond to key server master keys. Enabling key servers for the first time is a simple procedure. Once key server encryption is enabled, the type can be configured and enabled, server end-points can be created, and then the keys can be prepared and committed.

Security requirements

Security of all key server communications is governed by TLS 1.2 and TLS 1.3 protocols. Encryption keys are distributed between nodes in the system using TLS 1.2 and TLS1.3. The system uses AES-256 encryption that uses OpenSSL library interfaces. To establish a connection between the key server and the system, the key server or services must support the configured TLS version.

IP addresses and ports

All nodes that want to communicate with key servers must have their service IP address configured. A node must have its full service IP stack configured (address, gateway, mask) in order for that node to be a candidate for attempting to contact the key server. Key servers are typically set up on a private LAN, and this requires enforcement of service IP addresses. If only a subset of nodes have service IP addresses set, then those nodes without a service IP address log an error. The IP address that the user supplies must be the one that the system uses to communicate with the key server.

Each key server has a TCP port associated with its access. Since a key server serves multiple clients, the system allows the user to use a different port for each server and enables access for this port

when required. KMIP server conformance mandates that TCP port 5696 is supported, so this is the default port for the server end point.

Key generation policy and key database

If key server encryption is enabled, then the key server generates and manages the master keys. The node generates all other keys.

The key database can be clustered or unclustered depending on the type of key server that is used. For unclustered key servers, the user needs to consider backup and replication of the key database. IBM Security Guardium Key Lifecycle Manager is an example of a key server product where replication must be configured for encryption keys to be shared automatically between IBM Security Guardium Key Lifecycle Manager instances. Without replication configured, manual backup and restore operations must be used. Other products might self-replicate, so other key server instances automatically have any new keys created. For IBM Security Guardium Key Lifecycle Manager, complete backups and restores by following the IBM Security Guardium Key Lifecycle Manager user guide. The SafeNet KeySecure server and Thales CipherTrust Manager support both clustered and unclustered key servers configuration.

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 configuring IBM Security Guardium Key Lifecycle Manager 5.0 or later, you need to create a catalog for the specified device group in FlashSystem (for example, **SPECTRUM_VIRT**) under **Endpoint Management** > **Catalog**. This catalog must be based on the GPFS catalog.

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 Generic KMIP key servers

Encryption key servers create and manage encryption keys that are used by the system. In environments with many systems, key servers distribute keys remotely without requiring physical access to the systems.

Generic KMIP key servers 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 Gemalto SafeNet 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 username 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. For more information on migrating key servers, see .
- If you are using Generic KMIP key server to create and manage keys, determine whether the system needs a username and password to authenticate to the Generic KMIP 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.
- If you are using HashiCorp Vault as a KMIP server, note that Vault uses an internal KMIP Certificate Authority (CA) to generate certificates for clients that authenticate by using the KMIP protocol. You cannot import external KMIP certificate authorities. All KMIP client authentication must use certificates that are signed by the internal KMIP CA. Ensure that KMIP clients trust the internal CA before you configure them for Vault KMIP integration.

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** 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 a system signed certificate or an external CA signed certificate, see **satask exportrootcertificate** command and **chsystemcertstore** 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 Generic KMIP key servers

For generic KMIP 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 Generic KMIP key server is installed on the system and in use. Either add the server certificate for each Generic KMIP 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.

   Note: IBM Storage FlashSystem supports secret size (password) up to 32 bytes and does not support secret size of 64 bytes.

   - 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 Key Management Interoperability Protocol (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. The IBM Storage FlashSystem certificate must be trusted by the Generic KMIP key servers. It is recommended to use a certificate signed by a certificate authority (CA) for all Generic KMIP key servers. If the system's root Certificate Authority (CA) is used to sign the certificate, then the system's root certificate must be updated in the Generic KMIP key server 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 updated in Generic KMIP key server that can be used for KMIP.

   Note: For various KMIP key servers, certificate requirements are different and as mentioned:

   - For Fortanix DSM key server, the certificate contains "UUID" in the CN field.
   - For Thales CipherTrust Manager and Utimaco key servers, the certificate contains "username" in the CN field.
6. 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 Generic KMIP key servers 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 **Generic KMIP server** 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 only use a certificate from a trusted third party, or a system signed 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 Generic KMIP key server, 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 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 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 Generic KMIP key servers 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 a system signed certificate or an external CA signed certificate, see **satask exportrootcertificate** command and **chsystemcertstore** 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.

Note: For more information on the supported Generic KMIP key servers, see IBM Storage Virtualize Supported Key Servers.

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 499.
2. Connect your computer to the technician port of the configuration node. For more information, see the Node canisters page in the relevant Hardware Guide.
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.

Changing the primary key server

Using the management GUI

To change which key server is the primary key server, complete these steps:

1. In the management GUI, select **Settings** > **Security** > **Encryption**.
2. Expand **Key Servers** and right-click a key server that is not marked as the primary and select **Make Primary**.
3. The selected key server is now the primary key server.

Using the command-line interface

Migrating between key management methods

The migration between different key server types cannot be performed directly from one type of key server to another. USB flash drives encryption must be used to facilitate this process. If you want to migrate from one type of key server to another, you first must migrate from your current key servers to USB encryption, and then, migrate from USB to the other type of key servers.

To migrate from key servers to USB flash drives, use the command-line interface only. During migration, the system supports simultaneous configuration of both key management methods. After the migration completes, you can disable the old key management method.

Using the management GUI

During migration, the system does not disable the currently configured key management method until the new method is configured completely. Therefore, encrypted data can still be accessed with the current key until the migration is completed. For example, if you were migrating from USB flash drives to key servers, the old keys on the USB flash drive are still available until the key server encryption is configured. However, at least one of the USB flash drives with the current encryption key must be inserted into the system before you can migrate to the key server. After the key servers are configured, the old keys on the USB flash drive can no longer decrypt data on the system. Dispose of any old USB flash drives according to your recommended procedures for disposal of sensitive information. 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. After keys are migrated to the key servers, you can use the command-line interface to disable USB ports.

Note: The management GUI supports migration from USB flash drives to a key server encryption method only. To migrate from key servers to USB flash drives, you must use the command-line interface.

Before migrating to key server-based encryption, ensure that at least one USB flash drive with the current encryption key is inserted into the system. To migrate encryption from a USB flash drive to key servers, complete these steps:

1. In the management GUI, select **Settings** > **Security** > **Encryption**.
2. On the **Encryption** page, verify the following information:
   a. Expand **USB Flash Drives** and verify that USB flash drives are configured and detected in the system.
   b. Expand **Key Servers** and verify that key servers are not configured on the system.
3. Under **Key Servers**, click **Configure**.

   Note: To enable key servers to manage encryption keys on versions 8.1.0.0 or later, a USB rekey operation must be performed.
4. Select the type of key server that you are using to manage encryption keys. The system supports IBM Security Guardium Key Lifecycle Manager and Generic KMIP server key servers to handle key management on the system. If you currently use Gemalto SafeNet Key Secure key servers, you can migrate to Thales CipherTrust Manager.

   Note: If you want to migrate to a Generic KMIP key server, select **Generic KMIP server**. For more information on the supported Generic KMIP key servers, see IBM Storage Virtualize Supported Key Servers.
5. Enter the name, IP address or fully qualified 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. If you selected **IBM SGKLM (with KMIP)**, the **Key Server Options** page displays. 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. If you selected **Generic KMIP server**, the **Key Server Credentials** page displays. If you enabled authentication to the key server that uses a username and password, enter that username and password. These credentials are used to authenticate the system each time it connects to the key server. The username and password must match the credentials that are configured on the key server.
8. On the **Key Server Certificate** page, you must upload all the necessary key server certificates to the system.

   Note: It is recommended to use a certificate signed by a certificate authority (CA) for all Generic KMIP key servers.
9. If using IBM Security Guardium Key Lifecycle Manager, export the system certificate and install that on the key servers.
10. If you are using Generic KMIP servers, then the IBM Storage Virtualize certificate must be signed by a CA.
    - If IBM Storage Virtualize certificate is signed by the system's root CA, complete these steps:
      a. Select **Settings** > **Security** > **Certificates**.
      b. Click **Export Root Certificate**.
      c. Install the root certificate as an external CA on the key servers and add it to the list of external CAs that can be used for Key Management Interoperability Protocol (KMIP).
    - If IBM 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 key servers and add it to the list of external CAs that can be used for KMIP.
11. Return to the **System Encryption Certificate** page and select **The system's public key certificate has been transferred to each configured key server**.
12. On the **Disable USB Encryption** page, select **Yes** and click **Next**.
13. On the **Summary** page, verify the configuration for the key servers and click **Finish**. After the configuration completes, keys that were stored on the USB flash drives are not valid. Ensure that you dispose of all USB flash drives securely.

Using the CLI

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 to Thales

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 Key Management Interoperability Protocol (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 by Thales 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

Migrating between different key server types

The migration between different key server types cannot be performed directly from one type of key server to another. USB flash drives encryption must be used to facilitate this process. If you want to migrate from one type of key server to another, you first must migrate from your current key servers to USB flash drives encryption, and then, migrate from USB flash drives to the other type of key servers. For example, migrating an IBM Storage Virtualize system that is servers that are configured with IBM Security Guardium Key Lifecycle Manager key servers to SafeNet KeySecure key servers.

Certificates that are used for encryption key servers

When you enable encryption with key servers, two types of certificates are required to ensure secure communication between the system and the encryption key server.

Certificates are the primary method that is used by the key servers to authenticate the system and for the system to authenticate to the key servers. The exchange of these certificates verifies that access to the encryption keys that are stored on the key servers is allowed. The authentication of the system ensures that the key servers do not give access to keys to an untrusted party. The authentication of the key servers ensures that the system does not ask for sensitive keys to be stored by an untrusted party. Security of the system relies on two factors. First, the public certificates of the key servers and the system must be exchanged securely so that each device can trust the other. Second, the key servers and the system must keep their private key, which is associated with the certificate, secure.

The key server certificates, which are used by the system to verify the key servers, must be transferred to the system and installed. When you are adding key servers to the system, either a certificate that is signed by a trusted third party, or a self-signed certificate can be used. A combination of these certificates can be used if you are adding more than one key server. If a combination of self-signed and CA-signed certificates are installed, then the self-signed key server certificates that are installed on the system take priority over any CA-signed key server certificates that are installed on the system.

Note: It is recommended to use a certificate signed by a certificate authority (CA) for all Generic KMIP key servers.

If all key servers are using a certificate that is signed by the same certificate authority, it is recommended that the CA certificate is installed in the key server type's settings (for example, **`chkeyserverisklm -sslcert`** ...) rather than the key server endpoints (for example, **`mkkeyserver -sslcert`** ...). The CA certificate for the key server type's settings is used only when you connect to a key server if the key server endpoint does not have a certificate that is installed. If you are using individual self-signed certificates for each key server, the certificates are installed against each key server endpoint. If you are using key server certificates that are part of a chain of trust, see "Using certificate chains for key servers" on page 473.

If the key servers are part of the same system and are configured to replicate their keys and certificates with each other, then the key server certificates are copied to all of the key servers. The system can connect to all key servers with the same server certificate. In addition, the system's certificate must be installed on each of the configured key servers. The key server administrator accepts the system certificate to grant access to the key servers. To configure the system certificate for secure communications, select **Settings** > **Security** > **Certificates**.

Note: For more information on the supported Generic KMIP key servers, see IBM Storage Virtualize Supported Key Servers.

Using certificate chains for key servers

In order for the system to authenticate by using SSL with a key server that uses signed certificate chains, the system requires multiple certificates to be installed. All certificates in the certificate chain (excluding the server certificate itself, that is, the leaf certificate) must be installed. For example, if a key server X's certificate is signed by certificate authority Y, where certificate authority Y's certificate is signed by Root certificate authority Z, then the system requires the certificates Y and Z to be installed on the system to successfully communicate with key server X. To install this certificate, create a single certificate file that

contains the full chain of certificates, excluding the key server's certificate. The first certificate in the file is the CA certificate, which was used to sign the leaf of the certificate chain (the key server certificate). The last certificate in the file is the root CA certificate. In between the first certificate in the file and the root certificate, all of the intermediate CA certificates are specified in order. If multiple intermediate certificates exist, include these certificates in order from leaf to root.

The resulting certificates file must contain the first intermediate certificate at the top of the file, followed by subsequent intermediate CA certificates in order, one-by-one. Finally, the root CA certificate is last at the bottom of the file. The resulting file contains only the PEM data for each certificate with anything else, such as comments, removed from the file. Each certificate in the file is separated by a new line and each certificate begins with the header "-----BEGIN CERTIFICATE-----" and end with "-----END CERTIFICATE-----".

An example certificate file, for a certificate chain of depth 3, might look like the following example. The example includes two certificates. The included certificates are the intermediate CA certificate Y first, followed by the root CA certificate Z. The key server certificate X is not included. If a bigger certificate chain is used, then more intermediate CA certificates would be included in the certificate file.

```
-----BEGIN CERTIFICATE-----
MIIF4DCCA8igAwIBAgICEAAwDQYJKoZIhvcNAQELBQAwgYUxCzAJBgNVBAYTAkdC
MRAwDgYDVQQIDAdFbmdsYW5kMRMwEQYDVQQHDApNYW5jaGVzdGVyMQwwCgYDVQQK
DANJQk0xDDAKBgNVBAsMA1NURzEPMA0GA1UEAwwGcm9vdENBMSIwIAYJKoZIhvcN
AQkBFhNqYW1pZXByeUB1ay5pYm0uY29tMB4XDTE4MTEwNzIxNTg1OVoXDTI4MTEw
NDIxNTg1OVoweDELMAkGA1UEBhMCR0IxEDAOBgNVBAgMB0VuZ2xhbmQxDDAKBgNV
BAoMA0lCTTEMMAoGA1UECwwDU1RHMRcwFQYDVQQDDA5pbnRlcm1lZGlhdGVDQTEi
MCAGCSqGSIb3DQEJARYTamFtaWVwcnlAdWsuaWJtLmNvbTCCAiIwDQYJKoZIhvcN
AQEBBQADggIPADCCAgoCggIBAKbE3NRBJHGcOLBdRRe9InwQqw2kIQ7dDMyAOFbh
/EcuZ03ex0q8HYgJmzGaSwjFZlFBOOZ3eZGWxMgFzKljX8Q7LuQBMi/t7Vxh9Qec
TDcjJ4jtClgFFzlYHFNk0dFJ9sryrhuNMeel64nafJGQQ+dy+vONbqBZfgaMXSbZ
DgQUmmtVqALn+PgFKOjRIMSTpm4hddxn0yg+Av1rAYlRkbpE0B69ktLRZMtnb0s/
dXE73LTMjA7+62P4/DHlU6cBs3ljURsV59OpYrGnTDO1mFIo8Ot89aHQotW/SWCm
RYIsxZ7Hfh6P/fdhScSdJXyyz8Nfs8qr2Ec3Vz2dss0k3MI0E4L2lyKcYaPgUmxF
DlO39ZEcffLrUsfI2GN9RuvE2umqk6FgSCbDP3mxjMxUs3rB8oItU2g5sJ1MVGqb
gYUtS/19hPpjDB/yAjI5Dlv0uUlefBTMPASpLEzF8AKkkVg98Vmkkx5eYjuu3Qpr
rHXN8CVGvhw9TUYzA9Pt9POmq6icMoBHpzujGcDrMflM8aE3sI+EmsFnbZnESeje
CWpc6Q8EZaxnnnHixfMBzX7Al8IIhnzeNUL5x46oUaFzpthYN3FlwWN/W8UDMPPr
h9mh82/CfZqol7suOh5WFWBbA2YGqRQacezMwxteAmvoCvxXhjducU1AzPLkwyPd
xXUBAgMBAAGjZjBkMB0GA1UdDgQWBBSXo1kphy8JAvB25abXPxEDUgHccTAfBgNV
HSMEGDAWgBQ6vl82Zxr3V7UJHz0StpCWmfakQzASBgNVHRMBAf8ECDAGAQH/AgEA
MA4GA1UdDwEB/wQEAwIBhjANBgkqhkiG9w0BAQsFAAOCAgEAIPyFwnJrVyEuC+Hi
8tj0L/aIJabuT8QNmZcqcSW4dfvDDupp7VUiheexCr9btuYIUGrL59clrrJwPFRw
iZTDq3sOQQQz8keYEwXGsjTFafFcdkFIKg1rD+JZojrCvMvpGn1Zt0eqCkla7PsB
pbaXRIbh/ybyafxzPhBYeFDdl9BQQf9vsa1DkNABsx80hy/lHjohaLlG5C+aG+5r
2bUXRNhgl3lSToQkD/JgdvBbgycIcriiL4u6mTJbnpiCB2JbBdyuAzTkdNAZStCv
kB/N+E+7agjx32IxHwjr1FjQZuyqRbHcsvz3f5VB4pGX/6+xeFQn6DlUewvB6CBg
OqsUddu+FKgFw4f3l59UVC1KicQ4Hokbst9WJGkSMgdvQOXyxwapHrsGiTTbzekG
dD7XUdUiLGcT1CDB3NrcHsJ1a2PCUgMQpeEsWDyUspBrvo40w3jfSoKtFmwp3o6a
vx9esOWzevLzAbv7YGZD/iXTQbUk2iNJjt5lRIee20GxZK+/BV36378U0CNKwdUt
mKAaOHY/DsGzOghGdvCz8/g5zhPN+bw9ZXm116cAygIJt2xytVy1cLvtbY2Vzgd9
zE1hCdm6oOfijOQ6F1JAzE/vewTxxOuNzp/T4LLGmrXtWwqhkrFrhRUiwsRb5zjY
XIh94QnyIz09Oj/W0qxllTPIGxk=
-----END CERTIFICATE-----
-----BEGIN CERTIFICATE-----
MIIF8jCCA9qgAwIBAgIJAMmFTDgLwpj8MA0GCSqGSIb3DQEBCwUAMIGFMQswCQYD
VQQGEwJHQjEQMA4GA1UECAwHRW5nbGFuZDETMBEGA1UEBwwKTWFuY2hlc3RlcjEM
MAoGA1UECgwDSUJNMQwwCgYDVQQLDANTVEcxDzANBgNVBAMMBnJvb3RDQTEiMCAG
CSqGSIb3DQEJARYTamFtaWVwcnlAdWsuaWJtLmNvbTAeFw0xODExMDcyMTEwMzda
Fw0zODExMDIyMTEwMzdaMIGFMQswCQYDVQQGEwJHQjEQMA4GA1UECAwHRW5nbGFu
ZDETMBEGA1UEBwwKTWFuY2hlc3RlcjEMMAoGA1UECgwDSUJNMQwwCgYDVQQLDANT
VEcxDzANBgNVBAMMBnJvb3RDQTEiMCAGCSqGSIb3DQEJARYTamFtaWVwcnlAdWsu
aWJtLmNvbTCCAiIwDQYJKoZIhvcNAQEBBQADggIPADCCAgoCggIBAM1wapd+YEWk
1/INfxjA+0mwwmz4fFgbyIjLtPjh6Z9KMRx8x368PXZ7Y1br/Zt9rbMx0cwxEO6b
U6XG/h91Fk4mCPKsjFRImlrFEHWDZ+/mv38SgF1yqTJ4Cb1ZjZS6bY0VwLo0gRiu
hUjwMUvoIksUc9+dauuQU8TtllqmXfOCVdsONiixlTCQy+sqbceYucEZItXkFWCH
6VAKrbyQS2QZP8KAoCw9j1skiUAqJ9hVwryQ602chqqzUGoAJbSBL2Xn8JQM2LEn
2KlVoIQFASTmnyKuaqeBHcIB5KCSDCQcy8TtybO/X38YmpYzfHC+L7Tr3gcgfEZf
vAitoPTFjqqBlXbNR5Vl8t8aMNq++WP42iI5sv40NWrSqsURXMxFFkMV++YYaeq9
kra5WCfeZMSYOzolM5fNBLrkA1sfFWqRH6IMGFSxlg+GRlbXS47RwpgncI/GSStN
RR9mfYiFU05VbnZSlLMhYZZv3KoZars+V5KNGfi95zBcMocinq/SVFXCRtIZfp3A
M6/OKL7mziA4NgjaaybmNbQOZGD4/eZDRSljSoMpL4xgs9BMEv4QyfKbOJXMQHqi
to7jNJRkpz3yxlwkLAGkigyCKmQFAzj0ykB3V+NwmLTJezKBolOvdjda+2rkP2/b
EvCBHIqgAP9t9EOLMDke8v0dyryN+ESJAgMBAAGjYzBhMB0GA1UdDgQWBBQ6vl82
Zxr3V7UJHz0StpCWmfakQzAfBgNVHSMEGDAWgBQ6vl82Zxr3V7UJHz0StpCWmfak
QzAPBgNVHRMBAf8EBTADAQH/MA4GA1UdDwEB/wQEAwIBhjANBgkqhkiG9w0BAQsF
AAOCAgEAiGo8x8nGlOVXQhkVzLzwoXMHNDHTpn13u8SGkX68zVaulvwuXOXl4d6b
QMXrMLQ3KonrKSpBCgTbyjDHp+6/NiFUs0t+O4gZczQHx8CgyjkRYmhrtcNW/RGb
```

```
66jYTkZpcTbJsrxOyzJY04Z0cr5Qxi9lgj+MHFPcSj3WJxWidurkQCh/AOSERuYg
LhNxNihomfKplKGUBKAhwTjBW7abIN4u/15iNaM/67X+/MOF1IKjB9U6ikgIjY5D
VetEEJhnYLjuzHpCmNeFZ1hbySMaiKP3kUNKSidEh5Oi+UP2UujgZi8VMPZApStW
AmRVCWtpiwcBD2WqsUuL+XcpqWBZ7GcNa89maMcybTjPAl0y2phi27ZXm78LmpCS
uK/UbRfpUpAsxiIPK1tnFkKcqdWBYNuOt1GIW9ADsYytDjo14X+i7axJYPtiqW+H
UxY/6b16c2/YOeCwIq+1g4882VtAoQfGmODC2AhhFqHrFhfOQ3Ag86x9scUjQl+J
4uYQkKsR1xW/y9Tx9dpqZZrZEu3B91SrtvA/6GtGFFs7jyeRCPdhb8+KF2NBonNl
FZC+p8qEeMSaCbBdI+ptL7vKQzzwKT/72Ts5BsABn4ZHXsD/39x0HPqSfWcEKR1e
s8hisZjl8mRRb5+sHRsmNoh9ki1AjIUTSjn/O+94Xahc0LvMWJk=
-----END CERTIFICATE-----
```

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.
2. Connect your computer to the technician port of the configuration node. For more information, see the Node canisters page in the relevant Hardware Guide.
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, key servers, or internal key management). 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, key server encryption, or internal key management encryption) and should not be used to 'stand-alone'.

The encryption recovery key is recommended supplement encryption method while enabling the encryption method (such as USB flash drives encryption, key server encryption or internal key management encryption) on the management GUI.

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, key server encryption, or internal key management 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.

Note: A rekey operation needs to be performed after the upgrade from 8.7.0 to a later software version so that rekey time can be updated in GUI or CLI.

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, network-attached key servers, or internal boot drive attached internal key management. 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 unlockencryption -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.
2. Connect your computer to the technician port of the configuration node. For more information, see the Node canisters page in the relevant Hardware Guide.
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.

Adding nodes or enclosures to an encryption-enabled system

You can add nodes or enclosures to the IBM Storage Virtualize system to increase its capacity. If encryption is enabled on the system, you must manually copy the current encryption key.

The system automatically shares the current encryption key in memory with nodes or enclosure when they are added, but does not automatically write the key file to any new USB flash drives. The current encryption key file is automatically written to new USB flash drives only when you enable encryption for the first time or rekey the system. To copy the current encryption key, you can use either of the following methods:

- Manually copy the current encryption key file from an existing USB flash drive to any new USB flash drives that were received with the new node or enclosure.
- Rekey the system.

Copying the current encryption key by using a USB flash drive

If encryption is configured with USB flash drives, you can copy the existing key file from the USB flash drive that contains the current encryption key for the system. You can use either the management GUI or the command-line interface to identify the USB flash drive.

Identifying the USB flash drive with the current encryption key

If you have multiple USB flash drives, you must first identify the USB flash drive that has a valid copy of the current encryption key for the system.

Using the management GUI

To identify the USB flash drive with the current encryption key by using the management GUI, complete these steps:

1. Insert a USB flash drive into an available USB port.
2. In the management GUI, select **Settings** > **Security** > **Encryption**
3. On the **Encryption** page, expand **USB flash drives**.
4. Check the value for **State**. If the key on the USB flash drive is valid, **validated** is shown. If a different value is shown, the USB flash drive does not have the right key.

Using command-line interface

To identify the USB flash drive with current encryption key from command-line interface, enter the following command:

```
lsportusb
```

The output shows the details of the USB port where the USB flash drive is inserted and the state of the USB flash drive. If the key on the USB flash drive is valid, **validated** is shown for **state**. If a different value is shown, the USB flash drive does not have the right key. For more information, see **lsportusb**command.

Note: Insert only one USB flash drive at a time and repeat the procedure to copy the key for each USB flash drive separately.

Copying the encryption key to a new USB flash drive

1. To identify the file name of the current encryption key for the system, enter the following command:

   ```
   lsencryption
   ```

   The key file name is displayed in the **usb_key_filename** field:

   ```
   encryptionkey_000000E03BA000F0_000701DD00000001_my-storage-system-name
   ```

2. Insert the USB flash drive with the current encryption key into another system, such as your personal workstation or a server in your environment. Ensure that the drive is mounted.

3. Navigate to the file system of the USB flash drive with the current encryption key, locate the file with file name identified using **lsencryption** command on the device, and copy the located file.
4. Insert the new USB flash drive into the same system where you inserted the USB flash drive with the current encryption key. Ensure that the drive is mounted.
5. Navigate to the file system of the new USB flash drive, copy the encryption key, and safely unmount and remove the drive from the system.
6. Insert the new USB flash drive into the encryption-enabled system and run the **lsportusb** command to ensure that the copy of the key is valid and copied correctly.

Rekeying the system

You can rekey the system and manage the encryption keys when you use key server, USB flash drives, or both. Rekeying writes a new encryption key to all the USB flash drives currently inserted into the system. If encryption is configured with USB flash drives, it is possible to create new keys and store them on USB flash drives. If you configured key servers to manage encryption keys, you can generate new keys with the encryption key servers. If you use both a key server and USB flash drives, you must rekey each of them individually considering one at a time.

Tip: Before you rekey the system, ensure that the new nodes or enclosures are added to the system.

Key server
If you use a key server for encryption, see "Rekeying a system with key servers" on page 468.

USB flash drive
If you use USB flash drives for encryption, see "Rekeying a system with USB flash drives" on page 477.

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 the Node canisters page in the relevant Hardware Guide.

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** command.
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** command.
4. To disable encryption on a system that uses key servers to manage encryption keys, enter the following command:

   ```
   chencryption -keyserver disable
   ```

   The encryption key that uses the key server has been disabled successfully. For more information, see **chencryption** command.
5. To verify whether encryption 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** command.
6. To disable encryption on a system that uses USB flash drives to manage encryption keys, enter the following command:

   ```
   chencryption -usb disable
   ```

   The encryption key that uses the USB flash drives has been disabled successfully. For more information, see **chencryption** command.
7. To verify whether encryption key that uses the USB flash drives 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** command.
8. To disable encryption on a system that uses internal key management to manage encryption keys, enter the following command:

   ```
   chencryption -internal disable
   ```

   The encryption key that uses the internal key management has been disabled successfully. For more information, see **chencryption** command.
9. To verify whether encryption key that uses the internal key management has been disabled successfully, run **lsencryption** command and check that the `internal_key_status` is licensed. For more information, see **lsencryption** command.

Secure data deletion

The system provides methods to securely erase data from a drive or boot drive when a node or node canister is decommissioned.

Secure data deletion effectively erases or overwrites all traces of existing data from a data storage device. The original data on that device becomes inaccessible and cannot be reconstructed. You can securely delete data on individual drives and on a boot drive of a control enclosure. The methods and commands that are used to securely delete data enable the system to be used in compliance with European Regulation EU2019/424.

Deleting data from SAS drives

You can use the **chdrive -task erase** *drive_id* command to securely delete data from a SAS solid-state drive (SSD) or SAS hard disk drive (HDD) drive. The secure erase data functionality supports most devices. However, it is not supported on a few older drive models. When the command is used against the unsupported drive models, the `CMMVC6626E: The task was not initiated because a command was rejected by the drive that you have specified` error appears. Also, the command fails if the drive is an active member of an array, the drive is in auto manage mode, or a firmware upgrade is in progress on the drive or the enclosure. If the command process did not complete or fails, you cannot access the data on the drive; you can only enter the **chdrive -task erase** command again.

To see the status of the data deletion task and the estimated completion time, you can use the **lsdriveprogress** command. During the data deletion process, the drive activity LED indicator flashes green. When the process completes, the drive activity LED indicator becomes solid green.

Note: Do not power off the enclosure or remove the drive until the data deletion process completes. Otherwise, you might not see any progress from the **lsdriveprogress** command and the drive remains offline. If the drive is offline, wait until the activity LED indicator is a solid green before you try to use or remove the drive again.

The below table compares the methods that the system uses to securely delete data from the drives. The method varies according to the commands that each type of drive can support. The completion time for the erase procedure also varies, depending on the amount of data and the method that is used to delete the data. In each case, when the operation completes, the result is that the data on the drive effectively becomes impossible to access.

Table 98. Comparison of methods to securely delete data from drives

Priority	Deletion type	Method	Completion time
1	Cryptographic erase	Changes the encryption key and makes the data inaccessible.	Instant
2	Block erase	Quickly raises and lowers the voltage level of the storage element. Physical blocks are altered with a vendor-specific value.	Fast
3	Data overwrite	Replaces the existing data with random data.	Slow

The methods that are used to securely delete data vary according to manufacture, drive type, and drive firmware. For more information, see the documentation that is provided by the drive manufacturer.

If a drive supports more than one data deletion method, the system uses the highest-priority method. Consider the following examples:

- If a SAS solid-state drive (SSD) supports all of the methods, cryptographic erase is used to delete the data.
- If a SAS SSD supports block erase and data overwrite, the system uses the block erase method to delete the data.

- If a SAS hard disk drive (HDD) supports cryptographic erase and data overwrite, the system uses cryptographic erase to delete the data.

For more information about the drives that are supported by the system, see the Control enclosure replaceable units page in the relevant Hardware Guide.

Deleting data from NVMe drives

You can use the **chdrive -task erase** command to delete data from an NVM Express (NVMe) drive. The drive is marked as `offline` and the deletion process starts. Use the **lsdriveprogress** command to monitor the status of the data deletion process and the estimated completion time.

You can also securely delete data from NVMe drives by using the **chdrive -task format** *drive_id* command. When you issue **chdrive -task format**, the system changes the internal key in the NVMe drive and then erases the user data. All previous data in the drive becomes unreadable.

Before an NVMe drive is added to an encryption pool, the drive is encrypted with a new key. Because that encryption key is removed during the formatting process, the data on the NVMe drive can no longer be read.

Deleting data from boot drives

You can also securely delete the data on a boot drive of a control enclosure. To do so, use the **satask rescuenode -secureerase -force** command. The system then uses the block erase method to delete most of the vital product data (VPD) from the boot drive. Only the serial number, FRU part number, and MTM information remain in the VPD.

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.

Content

The 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 and numbers.
- 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 the next login under any of these conditions:
  - The password expired.
  - The administrator created new accounts with temporary passwords.

Expiration and lockout

The administrator can create the following rules for password expiration:

- Set a password expiration limit.
- Set a password to expire immediately.
- Set a number of failed login attempts before the account is locked.
- Set a time period before locked accounts can be used.
- Automatic log out for inactivity.
- Locking the 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. By default, user account locking is enabled for all users except the superuser. If a user fails to meet the password policy requirements, their account is automatically locked for 24 hours. You can configure user locking settings to alter the response in case user fails password policy.

Important: When upgrading from a previous release, the existing user locking settings are retained. If user locking was disabled in the previous release, it will remain disabled after the upgrade. An event will be raised notifying administrator to recommend appropriate action to enable user locking.

- Expire all passwords and require you to create a new password on the next login.

Two person integrity

Use two person integrity (TPI) to prohibit critical and risky tasks in the system from being executed by a single security administrator and by requiring the involvement of two security administrators.

TPI requires two security administrators to work together to complete certain tasks. Protecting data is an important part of IBM Storage Virtualize, and TPI helps mitigate the chance of data loss, prevent inadvertent mistakes on operations, and enhance security.

Requirements for enabling TPI:

- Ensure to have two users with the security administrator role.
- The two users can be local, remote, or a combination of both.
- If using remote users, a remote user group of security administrator role must be defined on the system and the remote authentication service must be enabled.

Requirements for disabling TPI:

After TPI is enabled, a user with an approved TPI request can disable TPI.

When you enable TPI, the users that belong to user groups of security administrator role are assigned the restricted security administrator role instead. However, their user groups retain their security administrator role.
Once TPI is enabled, a role elevation request and approval process is required to perform certain sensitive tasks:

- The restricted security administrator can issue a role elevation request on its own behalf to complete certain tasks in the system.
- Another restricted security administrator or a security administrator must approve the role elevation request.
- For example, this role elevation request and approval process is required to remove a Safeguarded snapshot.
- The restricted security administrators or security administrators can approve or deny role elevation requests, cancel role elevation requests, or revoke a role elevation request that was approved.

Available actions for a restricted security administrator that has an approved role elevation request.

- Create, change, or remove security administrator user groups.
- Create, change, or remove locally defined remote users.
- Change the non security administrator user group attribute on an existing local user to a security administrator user group.
- Modify attributes on existing local users that are members of the security administrator user groups.
- Change the role of existing non security administrator user groups to the security administrator role.
- Change the security administrator role of an existing user group to a non security administrator role.
- Remove and change Safeguarded backups and Safeguarded backup locations.
- Delete Safeguarded snapshots.

- Use a provisioning policy to define a set of rules that are applied when volumes are created within a storage pool or child pool.
- Change the single sign-on credentials that are used for the system.
- Remove the Safeguarded snapshot policy association from a volume group.

Configuring two person integrity (TPI)

Enable TPI when you want to prohibit critical and risky tasks in the system from being executed by a single security administrator and by requiring the involvement of two security administrators.

When TPI is enabled, the users that belong to user groups of security administrator role are assigned the restricted security administrator role. However, their user groups retain their security administrator role.

The restricted security administrator can submit a TPI request on its own behalf to complete certain tasks in the system. Another restricted security administrator or a security administrator must approve the TPI request. For example, this request and approval process is required to remove a Safeguarded snapshot.

Prerequisites

Before you enable TPI, ensure that you have at least two users in the system with the security administrator role. The users can be local, remote, or a combination of both.

Using the management GUI

To enable TPI, complete the following steps:

1. In the management GUI, select **Settings** > **Security** > **User Access** > **Two person integrity**.
2. Select **Enabled**.
3. Click **Save**.
4. To apply the TPI changes, click **Logout**.
5. After you sign in, the page displays that the current role is updated to restricted security administrator, and **Manage Role Elevation Requests** and **Request Elevated role** are displayed in the user menu list

Click **Snooze for 5 minutes** to stay on the page for 5 more minutes to review the settings.

Note: Snooze for 5 minutes is available only if user role is changed from security administrator to restricted security administrator.

To disable TPI, complete the following steps:

1. In the management GUI, select **Settings** > **Security** > **User Access** > **Two person integrity**.
2. Toggle the **Enabled** switch to the off position.
3. Click **Save**.

Using command-line interface

To enable TPI, enter the following command. Any security administrator can enable TPI.

```
chsecurity -twopersonintegrity yes
```

To disable TPI, enter the following command. After TPI is enabled, a user with an approved TPI request can disable TPI.

```
chsecurity -twopersonintegrity no
```

Managing two person integrity (TPI)

Use two person integrity (TPI) to prohibit critical and risky tasks in the system from being executed by a single security administrator and by requiring the involvement of two security administrators.

TPI requires two security administrators to work together to complete certain tasks. Protecting data is an important part of IBM Storage Virtualize, and TPI helps mitigate the chance of data loss, prevent inadvertent mistakes on operations, and enhance security.

When you enable TPI, the users that belong to user groups of security administrator role are assigned the restricted security administrator role instead. However, their user groups retain their security administrator role.
Once TPI is enabled, a role elevation request and approval process is required to perform certain sensitive tasks:

- The restricted security administrator can issue a role elevation request on its own behalf to complete certain tasks in the system.
- Another restricted security administrator or a security administrator must approve the role elevation request.
- For example, this role elevation request and approval process is required to remove a Safeguarded snapshot.
- The restricted security administrators or security administrators can approve or deny role elevation requests, cancel role elevation requests, or revoke a role elevation request that was approved.

Viewing elevation requests

To access the Manage Role Elevation Requests page, select **Access** > **Manage Role Elevation Requests**.

- You must have the security administrator role or restricted security administrator role to access the page.
- After your request is approved, you must log out of the management GUI. Then, log in again for the setting to take effect.
- The maximum number of active role elevation requests is four.
- You can't approve your own requests.
- Approved requests are displayed until their duration expires.
- Denied, revoked, and canceled requests are automatically removed from the page.
- If no requests are pending, information about how to request an elevated role is displayed.

Viewing the audit Log

To view a list of completed actions, select **Access** > **Audit Log**. If TPI is enabled, the Elevated Role column is displayed; if TPI is disabled, the Elevated Role column is hidden.

Host authentication

Host authentication allows user to configure host authentication mode for iSCSI hosts.

The iSCSI host authentication mode determines which hashing algorithms are used during iSCSI CHAP authentication.

Two authentication modes are supported:

Disabled (Default)
The default iSCSI host authentication mode imposes no restrictions on the hashing algorithms used. It allows MD5, SHA1, SHA256, and SHA3-256.

Enabled
When enabled, this mode restricts the use of weaker hashing algorithms. Only stronger algorithms, such as SHA256 and SHA3-256, are permitted, while MD5 and SHA1 are disallowed.

Note: If you are using the system as iSCSI backend, then enabling FIPs mode is not a supported configuration on iSCSI backend system.

Prerequisites

Ensure that the following prerequisite tasks are completed on the system before you configure host authentication:

1. For enabling this mode, ensure that all active iSCSI sessions are logged out. For details on active sessions, use the **`svcinfo lshostiplogin`** command.
2. Once enabled, all previously logged-out sessions that use weaker algorithms must update their hashing algorithm and undergo rediscovery.

Using the management GUI

To configure host authentication, complete these steps:

1. In the management GUI, select **Settings** > **Security** > **Host Authentication** > **iSCSI host authentication mode**.
2. Select **Enabled** and then click **Save**.

Ransomware threat detection

The system can analyze data patterns for its ransomware threat potential. With IBM Storage Insights, if the data patterns are found to be anomalous, an alert is raised.

When an alert is raised, affected volumes within a volume group automatically have a snapshot triggered. This snapshot is marked with an anomaly status of **Detected threat triggered**. This snapshot may contain some of the anomalous workload as it was triggered during that workload.

It is good practice to configure a schedule for creating regular snapshots of these volume groups. If an earlier volume group snapshot exists when an anomalous workload is detected, the latest snapshot will be marked as **Retention extended**, and its retention time extended to avoid early expiration.

Volume group snapshots created while an alert exists are marked as **Potential threat present**.

To configure ransomware threat detection, use the **Settings** > **Security** > **Ransomware Threat Detection** panel in the management GUI.

Chapter 19. 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.

Active - active capability

IBM Storage FlashSystem is an active-active dual controller system. Active-active is a term that describes the I/O processing mechanisms in a storage controller.

Assuming a system is made of a pair of controller nodes. In this case, an active-active controller pair would process I/O for a specific volume through either node. The opposite mechanism would be an active-passive controller system. In active-passive, only one node processes I/O for a specific volume. If a passive node receives I/O, it forwards the I/O to the active node to process. IBM Storage FlashSystem use an active-active processing mechanism.

In addition to the processing mechanism, a path model provides logical connections between a host system and the storage controller nodes. A path can be classified as **optimized** (preferred) or **non-optimized** (non-preferred). The path model is independent of the I/O processing mechanism.

As shown in Figure 50 on page 493, the system can be installed and configured to present optimized paths from all nodes, as defined by the standard SCSI multipath terminology. Configuring all paths as **optimized** means that either node can process the I/O operation in a symmetric manner. For read operations, I/O is processed entirely on the node that receives the I/O. For write operations, the I/O is coordinated by the node that receives the I/O with any required cache mirroring processing. To ensure data consistency, only one node handles any subsequent destage operations.

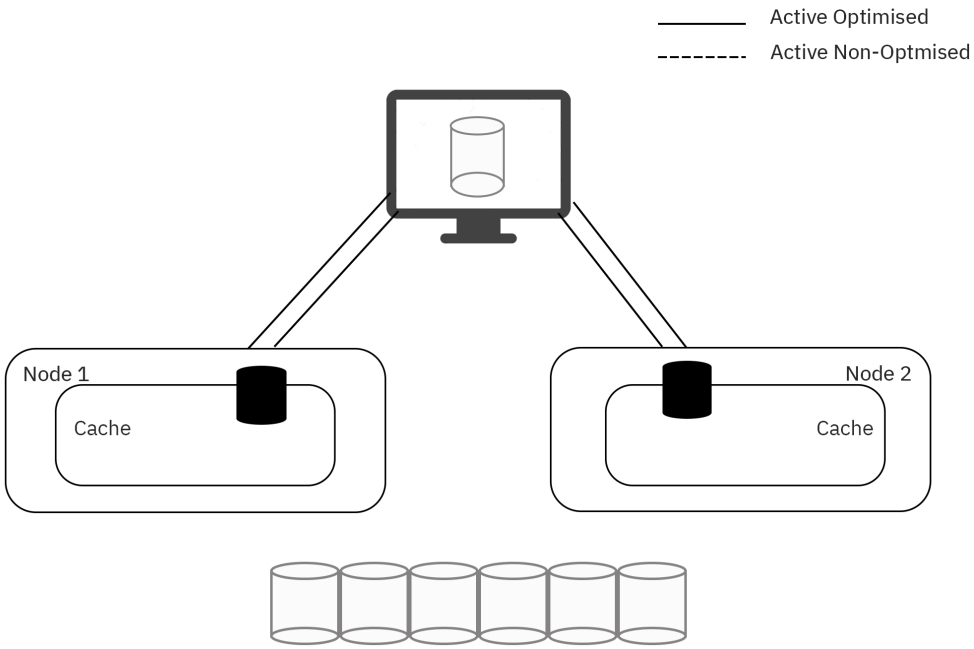


Figure 50. Active-active with all paths optimized

© Copyright IBM Corp. 2025

Enhancements available with FlashSystem

When installing a new IBM Storage FlashSystem node, the default configuration uses both **optimized** and **non-optimized** paths as shown in Figure 51 on page 494.

The default optimized and non-optimized path configuration provide the following performance advantages.

- An increased probability of a read cache hit by ensuring maximum I/O read requests for a specific volume are sent that uses the optimized paths, and therefore the same node receives. Read cache is not mirrored and is only stored on the node that receives the read request.
- The read-ahead or cache prefetch algorithm is more likely to detect a sequential read stream if the same node receives all read I/O for a specific volume.
- Each volume is allocated to a preferred node by using round-robin algorithm. The preferred node presents optimized paths to that volume and vice versa. This ensures an even balancing of workload across nodes.

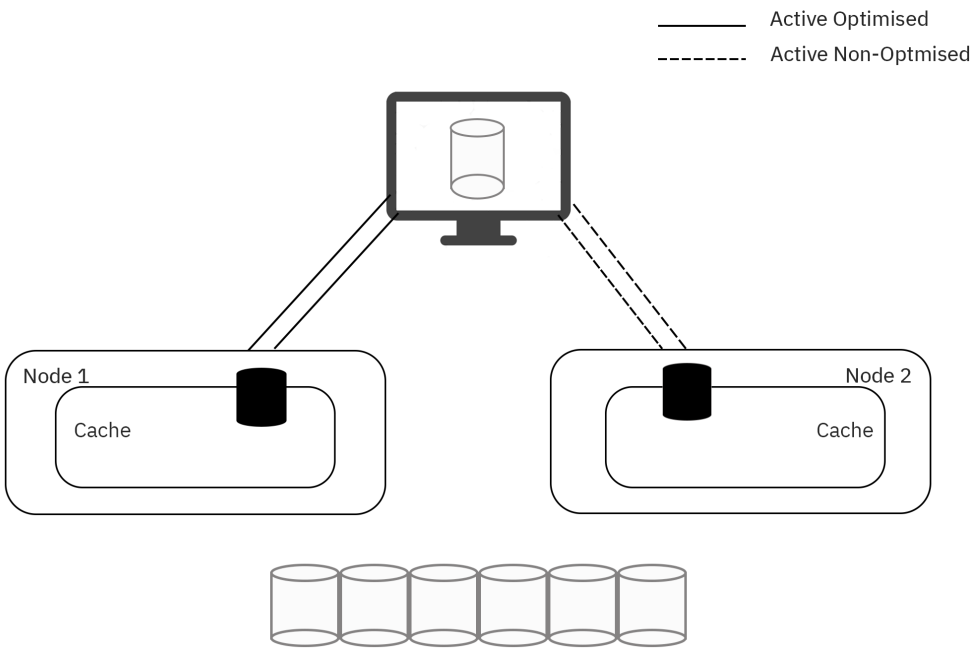


Figure 51. Active-active with default optimized and non-optimized paths

Active-active I/O flow

To further demonstrate the nature of active-active behavior, consider Figures 3 - 6 that shows the I/O flow for read and write operations through a FlashSystem.

Consider a read is issued to a volume. In this case, no matter if the system is using optimized paths or not, if the read is sent through a path to Node 1, the entire read operation is completed by Node 1 alone.

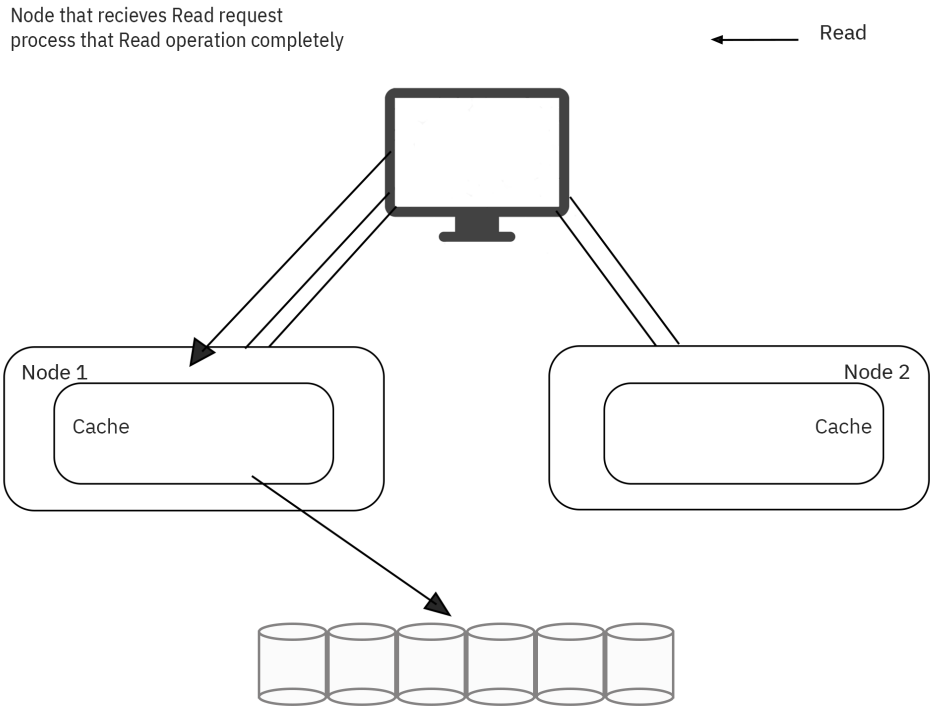


Figure 52. Active-active Reads through Node 1

Similarly, as shown in Figure 53 on page 495, the same read is issued to the same volume, but this time a path is provided by Node 2. The entire read operation is completed by Node 2 alone. This function is the standard definition of an Active-Active controller.

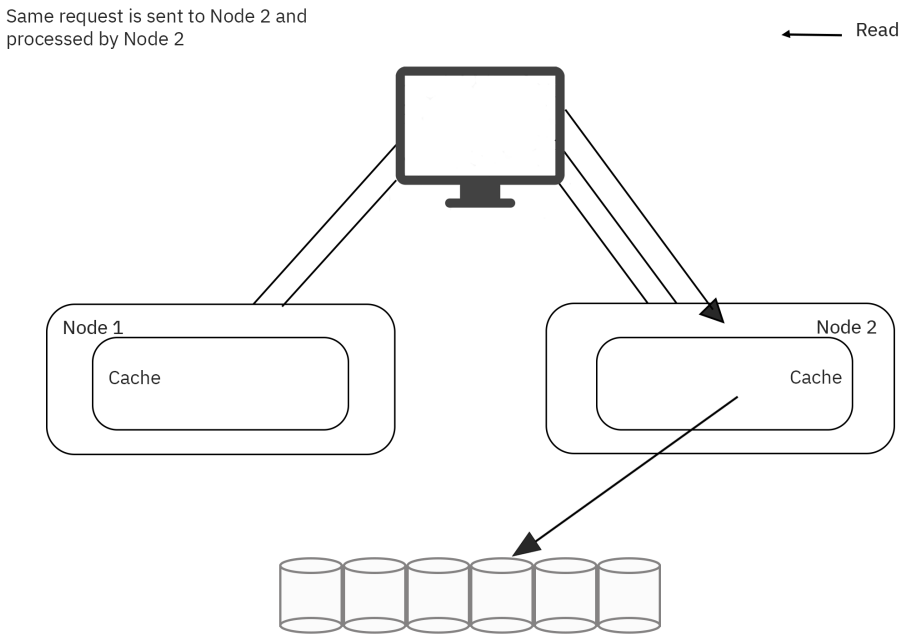


Figure 53. Active-active read through Node 2

As shown in Figure 54 on page 496, the write I/O flow is shown. In this case, the volume is written through a path to Node1. For each volume, the system automatically designates one node as the **owner**

node for the volume. In this example, Node 1 is set as the owner node that defines which node owns the eventual destage operation and happens to match the optimized or non-optimized assignment used.

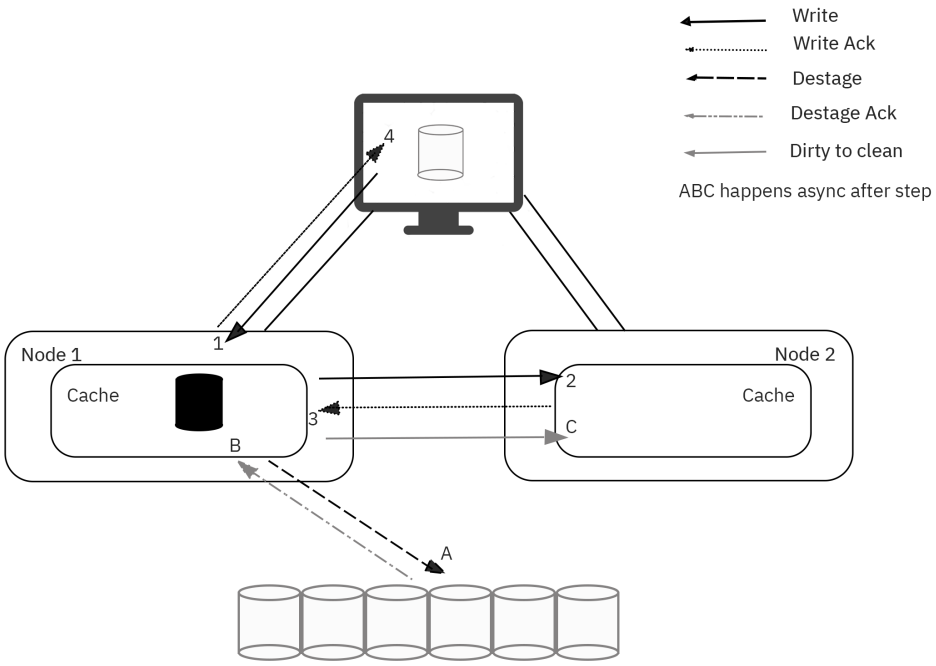


Figure 54. Active-active - Write to Node 1

In Figure 54 on page 496, steps 1 - 4 shows steps to complete a write operation. The initial write, the cache mirror operation, and the acknowledgments back to the host. Also, step A to C shows the required steps to asynchronously complete a write destage.

Note: A storage controller must have a mechanism as stated in the preceding example, where only one node conducts the destage operation. If such mechanism is not implemented, it can lead to data corruption.

In Figure 55 on page 497, it shows the same write I/O flow for the same I/O when the operation is submitted through Node 2. Node 1 is still the owner of the volume. Therefore, Node 1 completes the destage operation. However, all other processing is managed by Node 2.

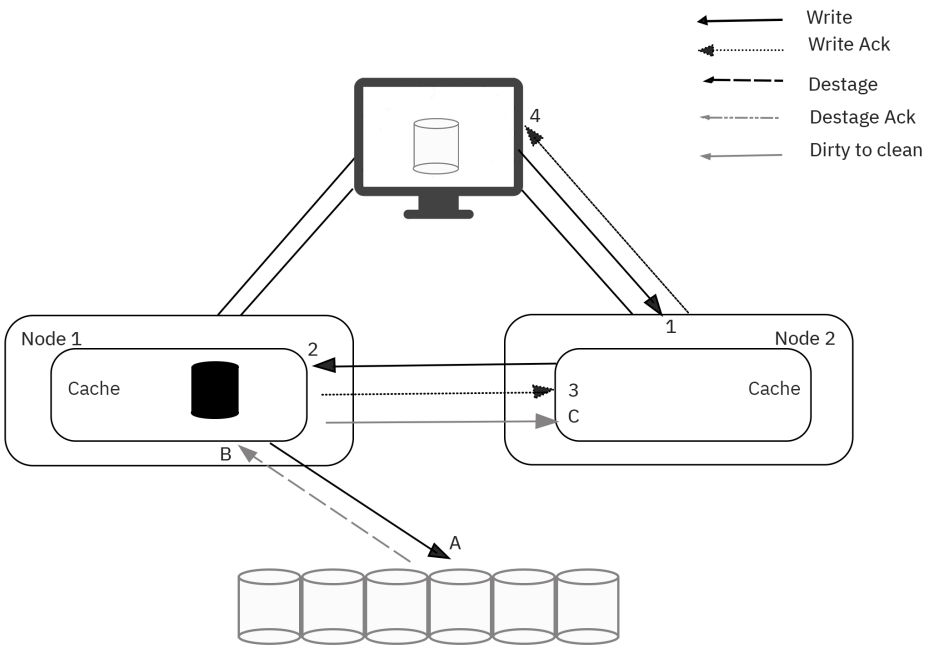


Figure 55. Active-active - Write to Node 2

In both cases, the number and nature of the required internal operations are identical. The same write data ends up in cache on both nodes, and one node completes destage of the write data to disk.

As explained in the preceding sections, active-active processing is always present on the IBM Storage FlashSystem and in addition it can be configured to use optimized and non-optimized paths. These two concepts do not interfere with each other. Therefore, I/O mechanism of the path model should not be confused with the active-active processing.

I/O group

The pair of nodes within a single enclosure is known as an *input/output (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. 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 **Settings** > **System** > **I/O Groups**. In the **I/O Groups** page, you can view the configured I/O groups on the system. The **I/O Groups** 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. Although a system can have multiple nodes within it, the nodes manage I/O in independent pairs.

Figure 56 on page 498 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).

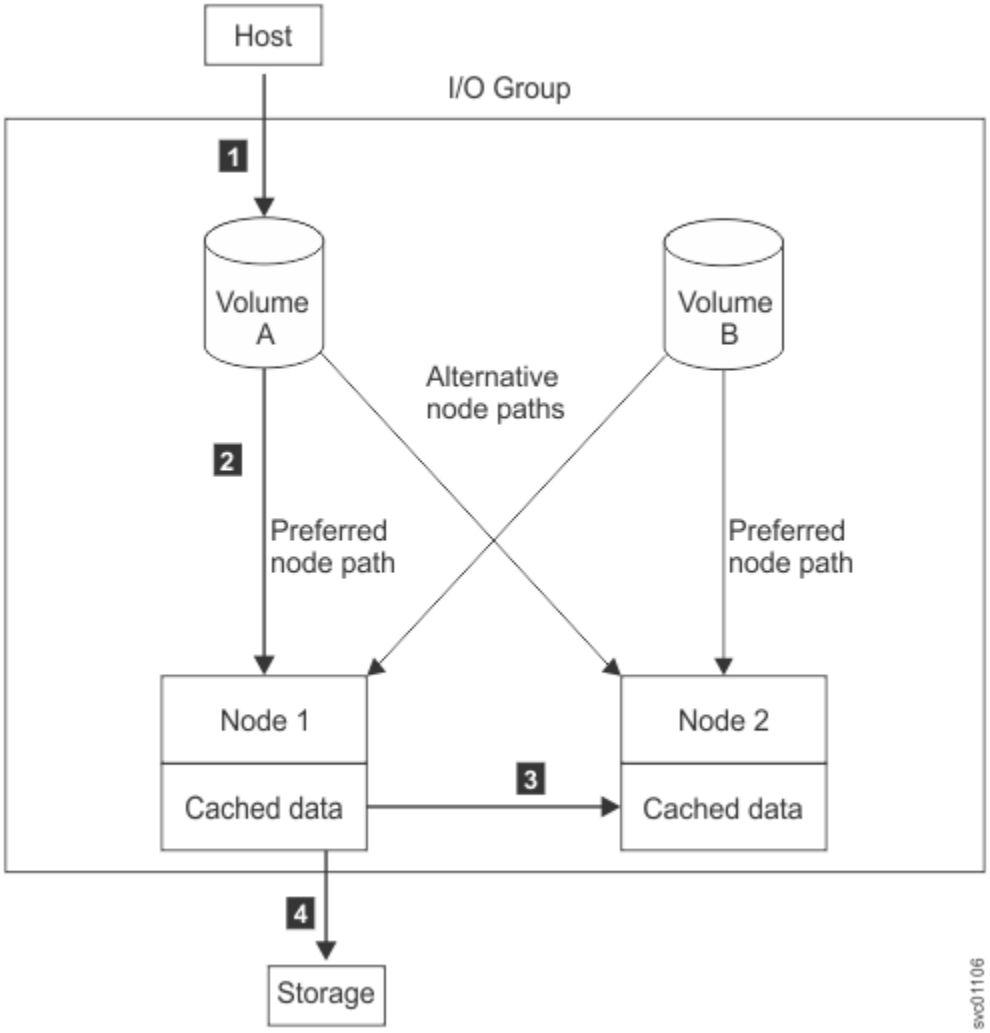


Figure 56. Example of a write operation in an I/O group

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.

A system that is configured on version 8.6.3 or later does not permit multiple I/O groups by default because of the additional features that are supported only on systems with a single I/O group. IBM SAN Volume Controller systems continue to support multiple I/O groups on long-term support releases only. FlashSystem grid can be used to scale out a storage system.

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.

Note: IBM FlashSystem does not support systems with more than one I/O group. IBM SAN Volume Controller systems continue to support multiple I/O groups on long-term support releases only.

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. For more details, see **Expansion enclosures** topic in the relevant hardware guide.

Nodes use Fibre Channel to communicate with hosts, external storage, other nodes (outside the enclosure) and Ethernet ports to communicate with hosts and external storage 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.

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 99 on page 500 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 99. 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. **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.

Table 99. Node status (continued)	
Status	**Description**
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.

Hot-spare node

When you add nodes, you can specify up to four of them as *hot-spare nodes*. A hot-spare node can become online (handling I/O operations) automatically if needed.

For example, if a node fails, an available hot-spare node that matches the failed node is activated automatically and moves to the **Online Spare** state. The hot-spare node handles I/O operations for the failed node until it comes back online. After the node returns to the system, the hot-spare node returns to the **Spare** state, which indicates it can be automatically swapped for other failed nodes on the system.

The loss of a node, either for unplanned reasons, such as hardware failure, or planned outages, such as upgrades, can result in loss of redundancy or degraded system performance. To reduce this possibility, a hot-spare node is kept powered on and visible on the system. A hot-spare node has active system ports, but no host I/O ports, and is not part of any I/O group. If a node fails or is upgraded, this spare node joins the system and assumes the place of the failed node, restoring redundancy. Only host connection on Fibre Channel ports that support node port virtualization (NPIV) can be used for hot-spare nodes. The hot-spare node uses the same N_Port ID Virtualization (NPIV) worldwide port names (WWPNs) for its Fibre Channel ports as the failed node, so host operations are not disrupted. The hot-spare node retains its node identifier when it was the spare. During an upgrade, the spare node is added to the system when a node is removed. The hot-spare node replaces each node that shuts down for upgrade in a system.

Note: IBM SAN Volume Controller systems support multiple I/O groups on long-term support releases only.

You can assign node pairs to specific I/O groups and then assign the extra nodes as hot-spare nodes. When a hot-spare node is added to the system, it is in **Spare** state, which indicates that it is not part of an I/O group. If a node in an I/O group fails, a hot-spare node automatically replaces that node and becomes a part of the I/O group. While the hot-spare node is in the I/O group, it is in the **Online Spare** state and returns to the **Spare** state when the original node rejoins the I/O group. A system can contain up to four spares at any time, which includes any hot-spare nodes that are online as spare nodes. Ensure that all cabling is correct to ensure that the system detects the nodes. If a node is not detected, review the installation information that was included with the system.

The following restrictions apply to using hot-spare node on the system:

-
- Hot-spare nodes can be used with Fibre Channel-attached external storage only.
- Hot-spare nodes cannot be used with SAS-attached storage.
- Hot-spare nodes cannot be used with iSCSI-attached storage.
- Hot-spare nodes cannot be used with storage that is directly attached to the system.

Criteria that requires an exact match

- Memory capacity
- Fibre Channel port ID
- Compression support
- Site

Criteria that is recommended to match, but can be different

- Hardware type
- CPU count
- Number of Fibre Channel ports

Note:

IBM SAN Volume Controller systems continue to support multiple I/O groups on long-term support releases. Chapter 10, “FlashSystem grid,” on page 235 can be used to scale out a storage system.

If the nodes on the system support and are licensed to use encryption, the hot-spare node must also support and be licensed to use encryption.

If an adapter PCI slot location on the spare node does not match the active nodes, an active node cannot be replaced by a spare node by using the command. If the user encounters the error , it means that the command failed because the specified node does not have a status of *candidate*. It is recommended to have adapters in the same slot for spare nodes and the active node for replace command to work.

When the online spare node is put into the Service state, it is immediately removed back to spare and 5 minutes later rejoin as online spare, in the cluster. Instead of putting the online spare into service, wait for the original to come back or remove the online spare and then perform their maintenance.

Adding hot-spare nodes by using the command-line interface

To add a spare node to the system, enter the following command, where *panel_name* is the name of the node that is displayed in the service assistant or in the output of command.

```
addnode -panelname panel_name -spare
```

For enhanced stretched configurations, use the following command to specify the name and site location of the spare node.

```
addnode -panelname panel_name -spare -site site_id
```

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 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.

Stretched 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 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 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.

There are two types of IP quorum applications:

- A **system** IP quorum application is only used for single-system tie-break if half of the nodes are unavailable.
- A **partnership** IP quorum application is only used to arbitrate between systems in a policy-based high availability partnership. 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 enabling 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.

For more information on the support article for the IP quorum application requirements, see IBM Storage Virtualize IP quorum application requirements.

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 quorum application on the storage system by using the GUI.

For partnership IP quorum applications:

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

For system IP quorum applications:

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

For more information on configuring quorum IP addresses, see .

Deploying the IP quorum application on the server

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

1. Create a dedicated directory for each IP quorum application instance. Multiple instances cannot be started from the same directory.
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 or quorum 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. -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 .-_/]
5. -rotation (optional) default = 5. Specify the maximum number of log files. Allowed values:
[1-10]
6. -size (optional) default = 5120. Specify the log file size in kb. Allowed values:
[1024-10240]
7. -version / -v (optional) display the vrmf of the system when the app was generated,
timestamp and generation id.
```

Reviewing and monitoring the IP quorum application

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

Using the management GUI

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

For more information on configuring quorum IP addresses, see .

Quorum IP addresses

Quorum IP addresses refer to the IP address configured in each node for connecting to the IP quorum application. Quorum IP addresses can optionally be used to make a more secure connection to the IP quorum application.

Quorum IP address requirements

- All Quorum IP addresses can be either on the same VLAN or not assigned to any VLAN.
- The quorum IP address and the application communicate through port 1260. Ensure that this port is open on the system's firewall and on the host to allow quorum traffic.

Quorum portset

A portset of type quorum can be configured to control which IP addresses on the node are used to communicate with the IP quorum application.

Quorum portset requirements

- One quorum IP address per node can be configured by using quorum portset.
- All quorum IP addresses must be in the same VLAN or without VLAN.

Configuring Quorum IP addresses

You can configure quorum IP address on the storage system by using the GUI.

Using the management GUI

In the management GUI, select **Settings** > **Network** > **Quorum IP addresses**.

To use quorum IP addresses:

1. Configure quorum IP addresses on each node.

2. Generate a new Quorum application through the management GUI. For more information, see .
3. Deploy the newly created JAR file.
4. Go to **Settings** > **System** > **Quorum applications** to verify the quorum application connectivity to storage system nodes.

System topology

The system topology can be set up in following ways.

- *Standard* topology, where all nodes in the system are at the same site.

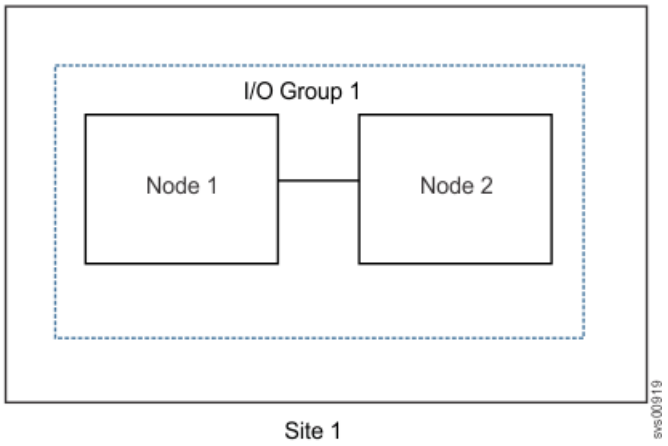


Figure 57. Example of a standard system topology

- *Stretched* topology, where each node of an I/O group is at a different site. When one site is not available, access to a volume can continue but with reduced performance. This topology is only available on SAN Volume Controller.

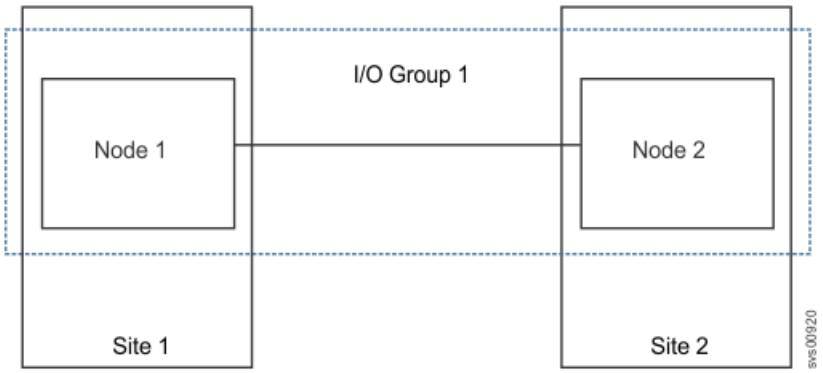


Figure 58. Example of a stretched system topology

Summary of system topology and volumes

Table 100 on page 506 summarizes the types of volumes that can be associated with each system topology.

Table 100. System topology and volume summary

Topology	Volume Type		
	Basic	**Stretched**	**Custom**
Standard	X		X
Stretched	X	X	X

System layers

The *system layer* affects how the system interacts with another system and other external systems that run IBM Storage Virtualize software.

Systems that run IBM Storage Virtualize software can either be in the *replication layer* or the *storage layer*.

Differences between the storage layer and the replication layer

This table compares the characteristics and requirements of storage layer and replication layer systems.

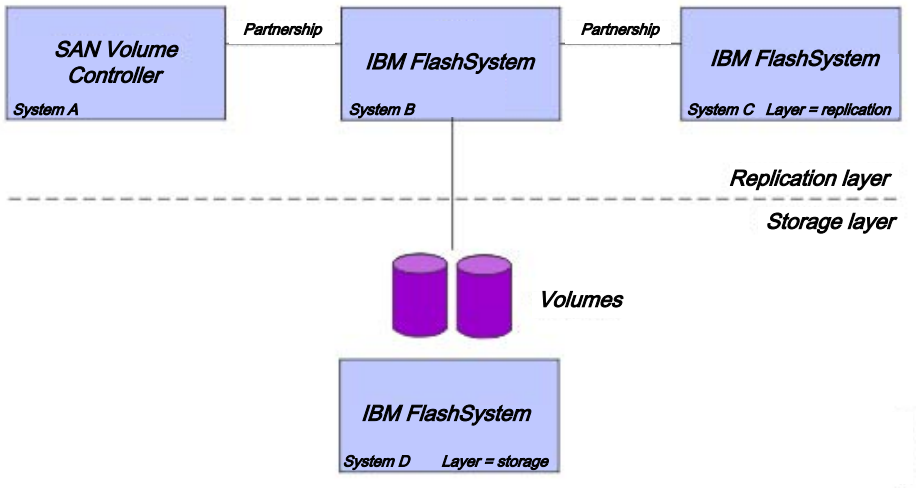
Table 101. Comparison of storage layer and replication layer systems

Storage layer	Replication layer
• The system can complete replicate with other storage-layer systems. • The system can provide external storage for replication-layer systems. • The system cannot use another Storwize family or FlashSystem family system that is configured with storage-layer as external storage.	• The system can complete replicate with other replication-layer systems. • The system cannot provide external storage for other IBM Storage Virtualize systems. • The system can use another Storwize family or FlashSystem family system that is configured with storage-layer as external storage.

A FlashSystem family system is in the replication layer by default, but the layer can be changed. For example, you might want to change a system to the storage layer to replicate with another system in the storage layer. Systems in a FlashSystem grid must be in the replication layer.

To create a partnership between systems, both systems must be in the same layer.

Figure 59. Example configuration with systems in different layers



Changing the layer of a FlashSystem family system

To change the storage layer of an existing system with internal storage before you add a second system into the SAN zone, you do not need to stop I/O operations.

However, if the system has Fibre Channel (FC) connections to another FlashSystem family, or system in the SAN fabric, I/O operations must be stopped temporarily. In this scenario, the FC ports must be disabled (for example, by unplugging all the FC ports, changing zoning, disabling switch ports) before you change the system layer. Then, you must re-enable the FC ports.

Before you change the system layer, the following conditions must be met:

- No host object can be configured with worldwide port names (WWPNs) from other IBM Storage Virtualize systems.
- No system partnerships can be defined.
- No IBM Storage Virtualize systems can be visible on the SAN fabric.

To change the layer of a system, follow these general steps:

1. Use the **lssystem** command to check the current system layer.
2. Configure SAN zoning so that the system whose layer you are changing is not visible to any other system.
3. Ensure that all other configuration requirements are met.
4. Run the **detectmdisk** command to ensure that the system detects any fabric changes that were done in step "3" on page 508.
5. Enter the one of the following commands to switch the system to the appropriate layer.
   - To switch to the replication layer, enter the **chsystem -layer replication** command.
   - To switch to the storage layer, enter the **chsystem -layer storage** command.
6. If the SAN configuration was changed in step "2" on page 508, restore the normal SAN configuration.
7. Run the **detectmdisk** command to update the list of MDisks.

Licensed functions

The base license that is provided with your system includes the use of its basic functions. However, there are also extra licenses that can be purchased to expand the capabilities of your system. Administrators are responsible for purchasing extra licenses and configuring the systems within the license agreement, which includes configuring the settings of each licensed function on the system.

The system supports differential, capacity, and key-based licensing. For example, with external virtualization differential licensing charges different rates for different types of storage, which provides cost-effective management of capacity across multiple tiers of storage. Licensing for these functions is based on the number of storage capacity units (SCUs) purchased. With other functions, like remote mirroring and FlashCopy, the license grants a specific number of terabytes of capacity for that function. Key-based licenses require an authorization code and key to be downloaded to the system before the function can be used.

Differential licensing is the term that is used to refer to the measurement of external virtualization and compression licenses, where different types of storage are charged at different rates. Differential licensing uses a storage capacity unit (SCU)-based license scheme. For each SCU, a client is entitled to manage different amounts of usable capacity based on a storage category classification. Table 102 on page 508 identifies the storage categories and SCU ratios for each category.

Table 102. Supported SCU licenses, drive classes, and ratios

License	Drive Classes	SCU Ratio
SCM	Storage Class Memory (SCM) devices	1 SCU equates to 1.00 TiB usable of Category 1 storage
Flash	All flash devices, other than SCM drives	1 SCU equates to 1.18 TiB usable of Category 1 storage
Enterprise	10K or 15K RPM drives	1 SCU equates to 2.00 TiB usable of Category 2 storage
Nearline	Nearline Serial ATA (SATA) drives	1 SCU equates to 4.00 TiB usable of Category 3 storage

Table 103 on page 509 provides an example of how SCU licensing might work based on category types.

Table 103. Example of SCU usage			
Storage category	**TiB usable**	**SCU ratio**	**Total SCUs**
1	42	1.18	35
2	400	2.00	200
3	800	4	200
Totals	1242		435

The system supports the following licensed functions:

External Virtualization
Enter the total number of storage capacity units (SCUs) you are licensed to virtualize across tiers of storage on your system. The system supports virtualization licensing for different tiers of storage. A license can be purchased for a specific quantity of storage capacity units (SCU) which can be used to virtualize a designated number of tebibytes (TiB) of storage. The used capacity for each tier can be monitored to learn how the virtualization license is distributed across the following tiers of storage:

- Storage class memory (SCM)
- Tier 0 Flash
- Tier 1 Flash
- Enterprise
- Nearline

FlashCopy
A FlashCopy base license provides you with an amount of capacities equal to the internal storage. Beyond this capacity, FlashCopy requires additional licensed capacity, which is charged per terabyte. The FlashCopy function copies the contents of a source volume to a target volume. The FlashCopy function is also used to create cloud snapshots of volumes in systems that have transparent cloud tiering enabled. The used capacity for FlashCopy mappings is the sum of all of the volumes that are the source volumes of a FlashCopy mapping and volumes with cloud snapshots. The Used TiB value indicates the amount of TiB that this system has used of the total amount of licensed TiBs for all FlashCopy operations on the system.

Note: For the 4983-AH8 model, FlashCopy license is applicable to external storage only and is charged per terabyte. FlashCopy operations that are used on internal storage are part of the base license.

Encryption
In addition to these licensed functions, you can also purchase a separate license to use software-based encryption. Software-based encryption protects data at rest and protects other keys that are used on self-encrypting drives. For each enclosure in the system, you must have an encryption license.

Bitmap space configuration

As part of Copy Services, the system supports policy-based replication. If you are currently using relationships to replicate data between systems, you may need to reallocate bitmap space to use policy-based replication.

Before you begin

The total I/O group memory that can be dedicated to these functions is not defined by the physical memory in the system. The I/O group memory is constrained by the software functions that use the memory.

To use policy-based replication on the system requires a set amount of allocated I/O group memory.

If you plan to use policy-based replication with other functions, use the following table to determine the required allocated memory for other features:

Table 104. Product comparison of required I/O group memory for policy-based replication

Product	Total I/O group memory	Required I/O group memory for policy-based replication	Available memory for other features (RAID and volume mirroring)
IBM Storage FlashSystem 9500	5920 MiB	4609 MiB	1311 MiB
IBM Storage FlashSystem 9200	3628 MiB	2561 MiB	1067 MiB
IBM Storage FlashSystem 7300	3628 MiB	2561 MiB	1067 MiB
IBM Storage FlashSystem 7200	3628 MiB	2561 MiB	1067 MiB
IBM SAN Volume Controller	3628 MiB	2561 MiB	1067 MiB
IBM Storage FlashSystem 5045 (requires a minimum of 32 GiB memory in each node canister)	896 MiB	713 MiB	183 MiB
IBM Storage FlashSystem 5200	2088 MiB	1537 MiB	551 MiB
IBM Storage Virtualize for Public Cloud	2088 MiB	1537 MiB	551 MiB

For example, an IBM Storage FlashSystem 9500 system has a maximum of 5920 MiB memory that is used by replication, volume mirroring, RAID, and policy-based replication. If policy-based replication is enabled, it uses 4609 MiB, the remaining 1311 MiB can be allocated between replication, volume mirroring, and RAID with maximum limits as 1024 MiB, 512 MiB and 800 MiB each in order. If policy-based replication is disabled, a maximum of 1824 MiB can be allocated between replication, volume mirroring, and RAID with maximum limits as 1024 MiB, 512 MiB and 800 MiB each in order.

In planning the installation for a system, consider the future requirements for the advanced functions.

This table describes the configuration of the bitmap space in a system.

Table 105. Bitmap space default configuration

Copy Service	Minimum allocated bitmap space	Default allocated bitmap space	Maximum allocated bitmap space	Minimum[1] functionality when using the default values
Volume mirroring	0	20 MiB	512 MiB	40 TiB of mirrored volumes, or 40 TiB of non-mirrored formatting standard-provisioned volumes, or 20 TiB of mirrored formatting standard-provisioned volumes
RAID	0	40 MiB	800 MiB	80 TiB writable capacity using RAID 0, 1, or 10. Slightly less than 120 TiB writable capacity in five-disk RAID 6 array.

[1] The sum of all bitmap memory allocation for all functions must not exceed 1824 MiB. The actual amount of functions might increase based on settings such as grain size and strip size. RAID is subject to a 15% margin or error.

The following tables describe the amount of bitmap space necessary to configure the various Copy Services functions and RAID:

This table provides an example of the amount of memory that is required for remote mirroring functions and volume mirroring.

Table 106. Examples of memory required

Function	Grain size	1 MiB of memory provides the following provisioned capacity for the specified I/O group
Volume mirroring or format	256 KiB	2 TiB of mirrored provisioned capacity

Notes:

- For volume mirroring, the 512 MiB of memory space enables 1 PiB of total provisioned capacity.
- Provisioned capacity and formatting of standard-provisioned volumes both use the mirroring bitmap space.

Before you specify the configuration changes, consider the following factors:

- The smallest possible bitmap is 4 KiB; therefore, a 512-byte volume requires 4 KiB of bitmap space.

On existing systems, also consider these factors:

This table shows the distributed RAID requirements for bitmap memory.

Table 107. Distributed RAID requirements	
RAID level	**Approximate required bitmap memory (MiB)**
RAID 1	Drive capacity * (count + 2 * stripe) / (stripe * 256 * 8)
RAID 6	Drive capacity * (count + 3 * stripe) / (stripe * 256 * 8)
Note: Drive capacity is in GiB, count = drive count, Stripe = stripe width, strip = strip length (in KiB). • For FlashCore Modules, the effective capacity should be used and not the physical capacity.	

You can also use the **lsiogrp** and **chiogrp** command-line interface (CLI) commands to modify the settings.

FlashCopy and volume group snapshot capacity configuration

FlashCopy and volume group snapshot capacity limits are determined by various factors.

- The physical memory of the system.
- The platform type in the system.
- The system's grain size. It can be 256 KiB or 64 KiB.

For volume group snapshots there are two limits, Production Capacity Allowance (PCA) and Snapshot Capacity Allowance (SCA).

- PCA is the provisioned capacity that is directly accessible from a host.
- SCA is the provisioned capacity held in snapshots.

The following table shows the maximum values for PCA and SCA.

Table 108. Volume group snapshot capacity limits

Product	Node canister memory	Grain size	Max PCA	Max SCA
IBM Storage FlashSystem 9600	768 GiB	256 KiB	4 PiB	400 PiB
	1536 GiB			800 PiB
	768 GiB	64 KiB	1 PiB	100 PiB
	1536 GiB			200 PiB
IBM Storage FlashSystem 9500	512 GiB	256 KiB	40 PiB	200 PiB
	1024 GiB			400 PiB
	1536 GiB			800 PiB
	512 GiB	64 KiB	10 PiB	50 PiB
	1024 GiB			100 PiB
	1536 GiB			200 PiB
IBM Storage FlashSystem 9200	128 GiB	256 KiB	4 PiB	4 PiB
	384 GiB			200 PiB
	768 GiB			300 PiB
	128 GiB	64 KiB	1 PiB	1 PiB
	384 GiB			50 PiB
	768 GiB			75 PiB

Table 108. Volume group snapshot capacity limits (continued)

Product	Node canister memory	Grain size	Max PCA	Max SCA
IBM Storage FlashSystem 7600	384 GiB	256 KiB	4 PiB	200 PiB
	768 GiB			400 PiB
	384 GiB	64 KiB	1 PiB	50 PiB
	768 GiB			100 PiB
IBM Storage FlashSystem 7300	128 GiB	256 KiB	20 PiB	20 PiB
	384 GiB			200 PiB
	768 GiB			300 PiB
	128 GiB	64 KiB	5 PiB	5 PiB
	384 GiB			50 PiB
	768 GiB			75 PiB
IBM Storage FlashSystem 7200	128 GiB	256 KiB	4 PiB	4 PiB
	384 GiB			200 PiB
	768 GiB			300 PiB
	128 GiB	64 KiB	1 PiB	1 PiB
	384 GiB			50 PiB
	768 GiB			75 PiB
IBM Storage FlashSystem 5600	128 GiB	256 KiB	4 PiB	0 PiB
	256 GiB			200 PiB
	128 GiB	64 KiB	1 PiB	0 PiB
	256 GiB			50 PiB
IBM Storage FlashSystem 5300	32 GiB	256 KiB	4 PiB	4 PiB
	128 GiB			4 PiB
	256 GiB			200 PiB
	32 GiB	64 KiB	1 PiB	1 PiB
	128 GiB			1 PiB
	256 GiB			50 PiB
IBM Storage FlashSystem 5200	32 GiB	256 KiB	4 PiB	4 PiB
	128 GiB			4 PiB
	256 GiB			200 PiB
	32 GiB	64 KiB	1 PiB	1 PiB
	128 GiB			1 PiB
	256 GiB			50 PiB

Table 108. Volume group snapshot capacity limits (continued)				
Product	**Node canister memory**	**Grain size**	**Max PCA**	**Max SCA**
IBM SAN Volume Controller SV3	512 GiB	256 KiB	40 PiB	200 PiB
	1024 GiB			400 PiB
	1536 GiB			800 PiB
	512 GiB	64 KiB	10 PiB	50 PiB
	1024 GiB			100 PiB
	1536 GiB			200 PiB

Any volume that has snapshots consume from the PCA an amount equal to the provisioned capacity of the volume. When a new snapshot is added additional PCA is consumed temporarily, once the snapshot has been added successfully this additional PCA is freed.

For example, when the first snapshot is added to a volume group, of 100 TiB, 100 TiB of PCA will be consumed and 100 TiB of SCA will be consumed. When a second snapshot is added another 100 TiB of PCA is used along with 100 TiB of SCA. Once the snapshot has been added successfully the additional PCA will be freed, this means these two snapshots consume 100 TiB of PCA and 200 TiB of SCA. Adding more snapshots increase SCA but the PCA remain at 100 TiB provided there is buffer 100 TiB of PCA free.

Volumes of type clone and thin-clone will also consume PCA equal to the volume's provisioned capacity. For a clone, once the clone has completed its background population process the PCA is freed . Similarly, an in-place restore from a snapshot consumes PCA until the restore process has completed.

Note: PCA and SCA are consumed in chunks:

- If the grain size is 256 KiB, then the chunks are 8 GiB, so a volume of 21 GiB would consume 24 GiB of PCA and SCA.
- If the grain size is 64 KiB, then the chunks are 2 GiB, so a volume of 21 GiB would consume 22 GiB of PCA and SCA.

If FlashCopy is used it consumes from the PCA. The consumption of the allowance is based on the provisioned capacity of the FlashCopy mapping's target volume.

The following table shows the upper capacity limits that can be achieved for FlashCopy target volumes.

Table 109. FlashCopy capacity limits		
Product	**Grain size**	**Max Target Capacity Allowance (TCA)**
IBM Storage FlashSystem 9600	256 KiB	4 PiB
	64 KiB	1 PiB
IBM Storage FlashSystem 9500	256 KiB	8 PiB
	64 KiB	2 PiB
IBM Storage FlashSystem 9200	256 KiB	4 PiB
	64 KiB	1 PiB
IBM Storage FlashSystem 7600	256 KiB	4 PiB
	64 KiB	1 PiB
IBM Storage FlashSystem 7300	256 KiB	4 PiB
	64 KiB	1 PiB

Table 109. FlashCopy capacity limits (continued)		
Product	**Grain size**	**Max Target Capacity Allowance (TCA)**
IBM Storage FlashSystem 7200	256 KiB	4 PiB
	64 KiB	1 PiB
IBM Storage FlashSystem 5600	256 KiB	4 PiB
	64 KiB	1 PiB
IBM Storage FlashSystem 5300	256 KiB	4 PiB
	64 KiB	1 PiB
IBM Storage FlashSystem 5200	256 KiB	4 PiB
	64 KiB	1 PiB
IBM SAN Volume Controller SV3	256 KiB	8 PiB
	64 KiB	2 PiB

Note:

1. The TCA is consumed from the PCA, meaning that using FlashCopy reduces the amount of PCA for volume group snapshot.
2. An incremental FlashCopy mapping consumes twice the amount of TCA than a regular FlashCopy mapping.
3. If a volume is highly available or in a volume group with policy-based replication, then it consumes PCA equal to the provisioned capacity of the volume.
4. On the IBM SAN Volume Controller SV3, the PCA is shared with volume mirroring.

Bitmap Space

The PCA, SCA, and TCA are related to the bitmap space available to volume group snapshot and FlashCopy. While it is better to plan volume group snapshot and FlashCopy capacity in terms of PCA, SCA, and TCA, the limits and consumption values comes from the different bitmap limits. The bitmap space is consumed in chunks of 4 KiB pages and each page corresponds to:

- 8 GiB of capacity if the grain size is 256 KiB.
- 2 GiB of capacity if the grain size is 64 KiB.

The values for a particular system can be found using the **lsiogrp** command view. This view has values:

Table 110. FlashCopy capacity limits	
Value	**Description**
flash_copy_total_memory	The amount of bitmap space in use for PCA.
flash_copy_free_memory	The amount of bitmap space reserved but not in use for PCA.
flash_copy_maximum_memory	The maximum amount of bitmap space available for PCA.
flash_copy_max_legacy_memory	The maximum amount of bitmap space available for TCA.
flash_copy_used_legacy_memory	The amount of bitmap space in use for TCA.
snapshot_volatile_total_memory	The amount of bitmap space in use for SCA.

Table 110. FlashCopy capacity limits (continued)	
Value	**Description**
snapshot_volatile_free_memory	The amount of bitmap space reserved but not in use for SCA.
snapshot_volatile_max_memory	The maximum amount of bitmap space available for SCA.

Note: If **snapshot_volatile_max_memory** is zero then the SCA comes out of the PCA.

The bitmap space is normally managed automatically and objects are created or deleted. However, once space is reserved it is not automatically freed. If the space needs to be freed then **chiogrp -feature flash** can be used to adjust the amount of bitmap space reserved for PCA(including TCA) and **chiogrp -feature snapshot** can be used to adjust the amount of bitmap space reserved for SCA.

Throttles

Throttles are a mechanism to control the amount of resources that are used when the system is processing I/Os on supported objects. The system supports throttles on hosts, host clusters, volumes, copy offload operations, and storage pools. If a throttle limit is defined, the system either processes the I/O for that object, or delays the processing of the I/O to free resources for more critical I/O operations.

When you configure throttles on the system, keep in mind the following guidelines:

- The throttle limit is a per node limit. For example, if a throttle limit is set for a volume at 100 IOPS, each node on the system that has access to the volume allows 100 IOPS for that volume. Any I/O operation that exceeds the throttle limit are queued at the receiving nodes.
- The multipath policies on the host determine how many nodes receive I/O operations and the effective throttle limit.
- If more than one throttle applies to an I/O operation, the lowest and most stringent throttle is used.

For each type of throttle, you can create either a IOPS or bandwidth limit or both. The throttle limit controls whether the processing of the host I/O continues or if it is delayed for later processing. Several reasons can exist to delay processing for specific hosts and host clusters. For example, if a host or host cluster has I/O intensive workloads, such as data mining systems, you can create a throttle limit that automatically delays the processing of I/O operations that exceed the configured throttle. Each host in the system can have a throttle defined or you can define throttle limits for a host cluster. In a host cluster throttle, all the hosts in the host cluster share the throttle limit.

Note: If you are creating a host cluster, any hosts that are added cannot have throttles configured. Throttles can be applied to the entire host cluster.

The system also supports throttles to delay processing of I/O operations for volumes. If storage systems provide storage to a wide variety of applications, then production volumes with more critical I/O can be competing with volumes that have lower priority operations. For example, volumes that are used for backup or archive operations can have I/O intensive workloads, potentially taking bandwidth from production volumes. Volume throttle can be used to limit I/Os for these types volumes so that I/O operations for production volume are not affected.

You can also create throttles for systems that have copy offload features enabled, such as offloaded data transfer (ODX) on Microsoft Windows Server 2012 or for XCOPY/WRITESAME features on VMware hosts. Copy offload frees up hosts and speeds the copy process by offloading the operations of certain types of hosts to a storage system. For systems with these features enabled, administrators can define throttles to delay processing for copy offloads to free bandwidth for other more critical operations. When a throttle for copy offload is defined, the throttle applies across the entire system. Like other throttles on the system, you can set IOPS throttles, bandwidth throttles, or both; however, bandwidth throttles are more effective for copy offload operations.

Throttles can be defined for storage pools to control I/O operations on back-end storage systems. Storage pool throttles can be used to avoid overwhelming the back-end storage and be used with virtual volumes.

You can set throttles for both parent and child pools. Throttles work in a hierarchy when they are applied to both the parent and child pool. Throttles that are defined in the parent pool applies to volumes in both parent and child pools. However, throttles that are defined on the child pool only apply to the volumes in that child pool. For mirrored volumes that are in different pools, only the throttling that applies to the primary copy is considered, regardless of which copy the I/O is directed to. In this case, any throttles defined on the secondary copy of the volume are ignored. Only throttles on the primary copy are used, whether the pool is parent or child pool. When throttles are defined in both parent and child pools, the more restrictive throttle applies to the volume in the child pool, but the I/O is counted against both pools. Consider these examples:

- A parent pool is configured with a throttling limit of 5K IOPS and its single child pool had a limit at 3K IOPS. When I/O operations to a volume in the child pool reaches the throttling limit, the volume cannot receive any more I/O during that time. However, volumes in the parent pool can still use 2K of IOPS to process during that time.
- A parent pool is configured with a throttling limit of 4K IOPS . The parent pool also has two child pool and each child pool has a throttling limit of 3K of IOPS.
- A parent pool is configured with a throttling limit of 5K IOPS . The parent pool has also has two child pool without a throttle defined. All in all, the volumes in each pool cannot handle more than 4K IOPS but each of these pools can could serve its IOPS limit if the volumes in the other pools consume at a lower rate of IOPS In this example, the volumes in the parent and both child pools can reach the 5K IOPS together, regardless of which pool the volumes are located.

Throttles for volumes, hosts, host clusters, copy offload, and storage pools can be configured in the management GUI and by using the command-line interface (CLI).

Chapter 20. Support

The system provides several options to help you to identify and resolve problems that may require assistance from service support personnel.

Call Home

Call Home connects the system to support personnel who can monitor and respond to system events to ensure that your system remains up and running.

If a software or hardware error occurs, Call Home notifies the support center, and automatically opens a service request. Call Home sends service-related information to inform support personnel so they can quickly act to resolve the problem.

Cloud services

Prerequisites

If you are using Call Home with cloud services to send notifications to the IBM Support, you need to ensure that the system can connect to the IBM Support.

Before you configure Call Home with cloud services, ensure that the following prerequisites are configured on the system:

- Ensure that all the nodes on the system have internet access.
- Ensure that a valid service IP address is configured on each node on the system.

Call Home with cloud services support the following configurations that have extra network requirements.

Table 111. Supported network configurations for Call Home with cloud services

Supported configuration	DNS configuration	Connection requirements
Call Home with cloud services with an internal proxy server	Required	Configure a proxy to allow outbound traffic on port 443 to `esupport.ibm.com`
Call Home with cloud services with a DNS server	Required	Configure a firewall to allow outbound traffic on port 443 to `esupport.ibm.com`.

Call Home with cloud services sends notifications directly to a centralized file repository that contains troubleshooting information that is gathered from your system. Support personnel can access this repository and be assigned issues automatically as problem reports. This method of transmitting notifications from the system to support removes the need for you to create problem reports manually.

As part of Call Home with cloud services configuration, you can define an internal proxy server within your network to manage connections between the system and the IBM Support. If an SSL inspection proxy is in use, you must define the proxy server and add its SSL certificate to enable Call Home to work correctly. For more information, see .

To configure Call Home with cloud services, use the **Edit settings** option on the **Settings** > **Support** > **Call Home** panel in the management GUI.

Email notifications

For redundancy, email notifications can also be configured. Call Home with email notifications sends notifications through an email server to support, and local users or services that monitor activity on the system. Call Home with email notifications requires configuration of at least one email server and

© Copyright IBM Corp. 2025

local user. You can add or delete custom email users and adjust the **Error** and **Inventory** toggles for the support center email.

To configure Call Home with email notifications, use the **Edit settings** option on the **Settings** > **Support** > **Call Home** panel in the management GUI. Make sure that you provide the system location and contact information.

Inventory information and heartbeats

Inventory data is a summary of the hardware components and configuration of a system and can be sent by using Call Home with cloud services or Call Home with email notifications. Service personnel can use this information to contact you when relevant software updates are available or when an issue that can affect your configuration is discovered. It is a good practice to enable inventory reporting.

If you are using Call Home with cloud services, the system sends a heartbeat notification to support as part of the Call Home function every 24 hours. Heartbeat notifications contain the same information about the system as inventory notifications, but they are not sent to any local user.

Because inventory information is sent by using the call home email function, you must enable call home email before you can send an inventory information email.

You can adjust the contact information, censor sensitive information, adjust the frequency of inventory emails, and enable or disable the sending of inventory data/heartbeats by using the management GUI. Use the **Settings** > **Support** > **Call Home** panel in the management GUI.

Storage Insights

You can also use Call Home with cloud services to send configuration, capacity, and performance metadata about your IBM Storage Virtualize systems to IBM Storage Insights for detailed analysis, monitoring, and reporting.

For more information on IBM Storage Insights, see .

Support assistance

Support assistance enables support personnel to access the system to complete troubleshooting and maintenance tasks.

You can configure either local support assistance, where support personnel visit your site to fix problems with the system, or local and remote support assistance. Remote support assistance allows support personnel to access the system remotely from the support center. Both local and remote support assistance use secure connections to protect data exchange between the support center and system. All actions that are completed with support assistance are recorded for auditing purposes. Local support assistance must be configured before remote support assistance is enabled.

Local support assistance

When you enable local support assistance, you can specify the IP address or domain name for the support connections. When support personnel log on to the systems with local support assistance, they are assigned either the Monitor role or the Restricted Administrator role.

The Monitor role can view, collect, and monitor logs and errors to determine the solution to problems on the system.

The Restricted Administrator role gives support personnel access to administrator tasks to help solve problems on the system. However, this role restricts these users from deleting volumes or pools, unmapping hosts, or creating, deleting, or changing users.

Roles limit access of the assigned user to specific tasks on the system. Users with the service role can set the time and date on the system, delete dump files, add and delete nodes, apply service, and shut down the system. They can also view objects and system configuration settings but cannot configure, modify, or manage the system or its resources. They also cannot read user data.

Remote support assistance

With remote support assistance, support personnel can access the system remotely through a secure connection from the support center. However, before you enable remote support assistance between the system and support, you first need to configure local support assistance. You must have either Call Home with cloud services or Call Home with email notifications configured. For more information about Call Home, see "Call Home" on page 519

During system initialization, you can optionally set up a service IP address and remote support assistance. If you did not configure a service IP address, go to **Settings** > **Network** > **Service IPs** to configure a service IP for each node on the system. If you use a firewall to protect your internal network you can configure a remote proxy server to allow access.

To prevent connection errors, ports 22 and 443 must be configured to support the service IP addresses remote support assistance.

When you enable remote support assistance, you can specify either IP address or domain name for support. 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**. In addition, you can define a shared-token that will be generated by the system and sent to the support center. If the system needs support services, support personnel can be authenticated onto the system with a challenge-response mechanism. After support personnel obtain the response code, it is entered to gain access to the system. Service personnel have three attempts to enter the correct response code. After three failed attempts, the system generates a new random challenge and support personnel must obtain a new response code.

Remote code load

Remote code load (RCL) is a service that allows remote support engineers to complete code updates on the storage system.

RCL is the process of having IBM support personnel securely connect to and update the microcode on the storage system. The RCL service is the preferred code delivery method, which proves to be both efficient and secure for IBM clients. RCL is fast and easy to coordinate because it does not require an onsite visit of an IBM services technician.

Prerequisites

If you are configuring remote support assistance, the following prerequisites are required for all configurations.

- Call home must be configured and functioning with a valid email server. To configure call home, select **Settings** > **Notifications** > **Email** in the management GUI or through system setup.
- Service IP addresses must be configured on each node on the system. To configure service IP addresses, select **Settings** > **Network** > **Service IPs** in the management GUI.
- A DNS server must be configured on your system. To configure a DNS server, select **Settings** > **System** > **DNS** in the management GUI.

  Note: DNS of your local system should allow for local and remote servers. It should not be configured to allow only a single external DNS server like Google 8.8.8.8.
- You can configure your firewall to allow traffic to pass directly from the system or you can route traffic through an HTTP proxy server within your environment.
- With the addition of the HTTP proxy support, Remote Support Proxy servers are no longer necessary, but they are still fully supported for existing configurations. Optionally, a Remote Support Proxy can be configured to consolidate firewall traffic from a number of storage systems. Remote upgrades cannot be completed through the Remote Support Proxy server.

Firewall configuration

The following network connections between IBM and the system are required to enable support assistance.

esupport.ibm.com

The esupport.ibm.com network connection is used to for the following actions:

- Uploading logs to the IBM Enhanced Customer Data Repository (ECUREP)
- Connecting to Call home with cloud services
- Downloading software from FixCentral

Note: The esupport.ibm.com network connection is fully certified to securely transmit data for Blue Diamond (HIPPA) users and General Data Protection Regulation (GDPR) protected users.

If you are using a firewall to route traffic instead of an HTTP proxy server, use the following information to configure a firewall rule.

Table 112. Firewall rule configuration

Source	Target	Port	Protocol	Direction
The service IP address of every node or node canister.	esupport.ibm.com	443	https	Outbound only

Remote Access

Use the following information to configure a firewall rule for remote access.

Table 113. Firewall rule configuration for remote access

Source	Target	Port	Protocol	Direction
The service IP address of every node or node canister	170.225.126.11, 170.225.126.12, 170.225.127.11, and 170.225.127.12	443 or 22	https or ssh	Outbound only

For remote access, port 443 is suggested as it provides more security. If the connection from port 443 remote access is unsuccessful, then port 22 is used as a backup. When you configure the firewall rules with port 443, the fallback of remote access from port 443 to port 22 is not applicable.

Configure support assistance

To configure support assistance, use the **Settings** > **Support** > **Support Assistance** > **Set Up Support Assistance** panel in the management GUI.

Support packages

If support assistance is configured on your system, you can either automatically or manually upload new support packages to the support center to help analyze and resolve errors on the system. To upload support packages, you must have at least one DNS server.

You can select individual logs to either download to review or send directly to the support center for analysis.

You can select from the following types of support packages:

Type 1: Standard logs

This option contains many different log files for both the underlying operating system and the software that includes these critical log files:

- Event log
- Audit logs

- Linux based logs that include `/var/log/messages`

These logs are sufficient to diagnose many problems, especially simple hardware replacement procedures.

Type 2: Standard logs plus one existing statesave
This option contains the standard logs for the system and an existing statesave from any of the nodes in the system. Statesaves can either be a *dump* or a *livedump*. A *dump* file is collected when the software restarts for any reason. This file can be used to understand why the software restarted. A *livedump* collects the current state of the software with minimal impact to I/O operations. The contents of a livedump are similar to the contents of a dump.

Type 3: Standard logs plus the most recent statesave from each node
Contains the standard logs for the system and the most recent statesave from each of the nodes on the system.

Type 4: Standard logs plus new statesave
This option contains the standard logs and a new statesave (livedump) for all the nodes in the system and packages them with the most recent logs.

When you are deciding what type of package to send, consider the speed of the upload and if more information is necessary later to resolve the issue. The standard logs are uploaded more rapidly since the support package is smaller than the other options. However, the standard logs might not include all the information that is needed to resolve the problem on the first attempt. Support personnel can require more data to resolve the problem. If a new statesave (livedump) is not generated shortly after the standard log, the livedump might not contain the data necessary to resolve the issue. When you generate the standard logs with the new statesave, you have the greatest chance of resolving the issue without more information. However, the support package is larger and can take longer to upload to the support center.

The following table shows the four types of support packages and the approximate size for those support packages for a basic configuration. The table describes common troubleshooting scenarios that require that type of support package. Actual sizes of support packages vary based on the number of volumes and MDisks, the software level, and the compressibility of the logs.

Table 114. Support packages

Description	Approximate Support Package Size: System with 1 I/O group and 30 volumes)	Approximate Support Package Size: System with 4 I/O groups and 250 volumes	Common Scenarios
Standard logs	10 MB	340 MB	• Simple hardware replacement • Critical performance issues[1] • All other issues
Standard logs plus one existing statesave	50 MB	520 MB	
Standard logs plus most recent statesave from each node	90 MB	790 MB	2030, 1196 or 1195 errors

Table 114. Support packages (continued)

Description	Approximate Support Package Size: System with 1 I/O group and 30 volumes)	Approximate Support Package Size: System with 4 I/O groups and 250 volumes	Common Scenarios
Standard logs plus new statesaves	90 MB	790 MB	• Compressed volumes issues • Host issues • External storage issues • General performance issues • Critical performance issues[1]

1. Critical performance issues require that both standard logs and a new statesave are generated and uploaded to the support center. Generate and upload the standard logs first and then create the new statesaves.

Uploading support packages

To upload support packages, use the **Settings** > **Support** > **Support Package** panel in the management GUI.

Chapter 21. System update

To obtain fixes or new features, you need to update the system software and firmware. For updates about known issues and new code levels, you can subscribe to the **My Notifications** service on the www.ibm.com website. See IBM support for details and to subscribe.

Most systems, including NVMe drives, can be fully updated using the main product software package. SAS drives need to be updated using . The full system update procedure restarts each node in turn to enable host I/O to continue through the partner node during the update.

Some parts of the system can be updated without restarting the nodes. This includes drive firmware updates, security patches and RTD models. Related packages are released independently of the main software package. Furthermore, some of these packages can be set to auto download and auto update. For more information, see .

Patches

Patches are one time updates that install without interrupting the I/O process. They are generally used to address minor bugs or security issues. are designed to occur periodically when a new versions of a feature become available, such as Ransomware Threat Detection.

Under some scenarios, IBM releases a patch that fixes or updates a part of the product. Patch installation does not cause any interruption to the IO processing.

Patches are made available on **FixCentral** with the help of direct download link (for more information, see) or as part of a general collection of available patches.

For patches that are made available by a direct download link

To start the installation process:

1. Go to **Settings** > **System** > **System Update**.
2. To install the patch, under **Manual updates** section, check **Use an existing package**.
3. Select the patch from the list of available packages on the system and click **Install Patch**.

For patches that are released as part of a general collection of available patches

To start the installation process:

1. Go to **Settings** > **System** > **System Update**.
2. To install the patches, select **Patches** tab, and click **Install** from the overflow menu. You can also click **Available patches** link on **System Update** panel.
3. On the Transfer and install page, enter your **IBMid** along with **Patch name**.
4. To view or edit the patches, select **Installed patches** link on the **System Update** panel.

The **Available patches** or **Installed patches** link is displayed only when minimum of one patch is installed on the system or at least one patch is available for patch installation.

Auto update

Patches subscribed for automatic updates persist across software upgrades, including CCUs. This means that once a patch is subscribed, it remains active and continues to receive updates, even after major system changes.

Use the **Auto updates** panel to manage auto update features, their current version, and update settings. Keep track of active updates, which ensures transparency in system behavior. If a patch fails to install, an event is generated to alert the user, who can then investigate further.

Note: Auto updates apply only to packages of type patch. They do not include full system updates.

There are two levels of auto update configuration:

© Copyright IBM Corp. 2025

System level

The system-level auto update setting can be configured to one of the following:

- **Auto Update**: automatically downloads and installs updates.
- **Auto Download**: downloads updates but requires manual installation.
- **None**: disables automatic updates entirely.

Per auto update level

The auto update level setting can be configured to one of the following:

- **Auto Update**: automatically downloads and installs updates.
- **Auto Download**: downloads updates but requires manual installation.
- **Use system setting**: inherits the system-level configuration.
- **None**: disables automatic updates entirely.

To ensure timely updates, it is recommended to configure the system level update setting to `Auto Update` setting and retain the default `Use System Setting` option for each update. This setup ensures that updates are automatically downloaded and installed in alignment with the system-wide policy, reducing the need for user intervention. This configuration ensures that patches for features are automatically downloaded and installed without requiring user action. System update settings can be configured by users with administrator permissions.

Preparing for a system update

Before you upgrade a system, complete the following actions:

1. Ensure Call Home connectivity and network settings are properly configured. For detailed instructions on setting up Call Home, see .
2. Decide the default behavior you want for automatic patch updates by setting the global policy using **`svctask chsystem -globalautoupdatepolicy`**.
3. After you enable automatic updates, allow time for the system to begin processing; updates might not start immediately.

Updating the system by using the GUI

To manage system updates (download, install, edit policies, or revert updates), go to **Settings** > **System** > **System Update** > **Auto updates** panel in the management GUI.

Manual update

Manual update is used to install the main product software package and drive bundles. As part of the manual update procedure, users also install and run the upgrade test utility package which performs system checks including whether the proposed upgrade path is supported. For more information about drive bundle upgrades, see .

For new product software packages, more information on the contents of the release can be found in the . During the product software 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 and drives. For more details on drives, see . Similar to the node updates, these updates are performed in a non-disruptive manner.

Note: Sometimes, firmware updates can take several hours to complete. For example, requirement of completing battery conditioning cycles before you apply the firmware. But this requirement is not a reason for any concern as these updates only take place once the main software upgrade process is complete and the cluster is running in a fully operational and redundant state.

While the system is updating, certain administrative actions are restricted until the upgrade process completes. Certain commands are permitted during the upgrade process, except during the final commit phase of the upgrade. Refer to the documentation for a command to see whether it is available during upgrade.

Note: The version of code on the config node determines whether the command is available during upgrade.

Preparing for a system update

Before you upgrade a system, complete the following actions:

1. Ensure that all errors in the event log are addressed through their associated fix procedures and marked as fixed.
2. Ensure that all host multipathing drivers are at supported levels and are correctly configured.
3. For older systems, where nodes are online for more than six months, consider performing a full reboot by using the service assistant to identify any potential latent hardware issues or issues in the system environment before performing the upgrade. Check that multipathing drivers have recovered paths after each reboot.

Updating the system by using the GUI

The update is started on the **System Update** panel in the GUI. The GUI guides the user through the procedure, including the preinstall checks that are performed by the upgrade test utility.

Troubleshooting issues during the upgrade

If issues occur during an upgrade, follow the fix procedures in the management GUI for guidance. Any user-driven node activities must only be done under the guidance of the fix procedure or IBM support to avoid any unnecessary impact to host applications.

System update duration

The following table shows the estimated time for an update on a four-node system:

Table 115. 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 10 minutes (for multi-pathing recovery)	10	29	59
Upgrade node A1	9-24	38	83
Upgrade node B3	9-24	47	107
Commit	10	57	117
Total	-	57	117

Service assistant manual update

Manual upgrade allows users to directly control the removal and upgrade of individual nodes. It is also used when automatic concurrent code update (CCU) has encountered problems midway through upgrade

and requires additional intervention from the user. Some of the checks from automatic CCU are removed, so manual upgrade should be used with caution.

During a manual system upgrade, each node is removed, individually updated, and added back to the cluster in turn. As each node is removed and then updated, which also includes a restart, the partner node in the I/O group processes I/O. Only one node should be updated at a time. After the final node is removed for upgrade, the system automatically commits the upgrade.

Start the update on the **Settings** > **System** > **System Update** in the management GUI. Select the appropriate test utility, upgrade package, and code level. When asked, select **Service Assistant Manual Update**.

Remove a node from the cluster, use the **Update Manually** panel to upload the necessary file, and start the update. After the node update is complete, it is added back to the system automatically. The update status indicates that a manual upgrade is in progress. Completion estimate is not provided because it depends on the pause between each node being removed and upgraded.

Update each node, one at a time, other than the configuration node by using this method. After all nodes, other than the configuration node, are updated and successfully returned to the cluster, remove the configuration node. Removal of the configuration node causes the remaining nodes, which are all on the uplevel version of the code, to commit the upgrade. One of the remaining nodes is elected as the new configuration node. After the commit is complete, a notification is sent to the GUI.

After the commit is complete, the removed node can be added back to the system. This node identifies that it is still on the downlevel build and it automatically updates the code level to conform with the rest of the cluster.

You can also start the manual update by using the **applysoftware** CLI command by using the **-prepare** option. Ensure that the **installsoftware** command is run with the **-pacedccu** option. Continue to run the upgrade test utility before proceeding with the update.

When the system is updating, certain administrative actions will be restricted until the upgrade process completes.

After the commit, the system might automatically update firmware levels on enclosure hardware components and firmware and drives. Similar to the node updates, these updates are performed in a nondisruptive manner.

Drive update

Drives can be updated automatically as part of the system update, or by using a drive package. 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.

Automatic drive updates

Every system update package includes drive firmware for all supported drives. After the system software is successfully updated, the drive firmware is automatically and systematically applied to all drives that have a use type of candidate, spare, or member.

Alternatively, drive-firmware packages can be downloaded and installed by an administrator. After drive firmware is installed on the system, the system applies the firmware at the following times outside of a system update:

- When a new drive is added to the system and brought into use
- When a failed drive is replaced with a replacement drive
- When the **applydrivesoftware** command is used.

Drive updates by using packages

To perform a drive firmware update using a drive firmware package, first obtain the latest drive firmware package and the latest upgrade-test utility from IBM FixCentral.

To apply the firmware using the management GUI, do the following:

1. Select **Pools > Internal Storage**.
2. Select the drives you want to update.
3. Select **Actions > Upgrade All** or **Upgrade** and follow the instructions to upload the package and test utility, and start the update.

Obtaining packages

The update packages are provided on **FixCentral**.

Package types

The following table shows different packages that are published to **FixCentral**:

*Table 116. Packages published to **FixCentral***

Fix Type	Example FixID	Description
Main product software	Storage-IBM_FlashSystem9100-8.6.0.2 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. The fix also contains FCM drive code. Other drive types must be upgraded using the Drive Bundle. **Note:** • Before installing this software, run the upgrade test utility.
Drive Bundle	Storage-IBM_FlashSystem9100-DriveMicrocode-231019 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. The fix contains various 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. **Note:** Before installing this firmware, run the upgrade test utility.

*Table 116. Packages published to **FixCentral** (continued)*		
Fix Type	**Example FixID**	**Description**
Software Upgrade Test Utility	Storage-IBM_FlashSystem9100-SwUpgradeTestUtility **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. It is essential to always use the latest upgrade test utility to protect against known issues. 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 I/O processing.
Generally available patches	These fixes are available from FixCentral without a direct download link.	Under some scenarios, IBM releases a patch that fixes or updates a part of the product, and can be updated without any interruption to I/O 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 .

To transfer a specific package to the system, select **Transfer** on the **System Update** 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.

If patch autoupdate is enabled, select **View available patches** on the **System Update** panel in the GUI to view and transfer generally available patches.

To obtain a package from **FixCentral** without using Call Home, access the following site 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 **System Update** 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.

Chapter 22. Management GUI

The system includes an easy-to-use management GUI to help you to monitor, manage, and configure your system.

You can access the management GUI by opening any supported web browser and entering the management IP address. In addition to simple setup, configuration, and management functions, the management GUI provides several functions that help filter and sort data on the system:

Automatic and manual resizing of main navigation
The main navigation menu automatically shrinks or expands based on the width of the browser viewport. If sufficient space is available, the navigation menu automatically expands to show the main navigation item names. When in the expanded mode, a toggle button is available to shrink or expand the navigation menu manually.

Filtering and sorting objects
On windows that include columns, you can sort each column by clicking the column heading. You can use the filtering feature to display items that include only text that you specify.

Selecting multiple objects
You can use the Ctrl key to select multiple items and choose actions on those objects by right-clicking them to display the actions menu. You can right-click any column heading to add or remove columns from the table.

Using preset options
The management GUI includes several preestablished configuration options to help you save time during the configuration process. For example, you can select from several preset options when you create a new volume. These preset options incorporate commonly used parameters.

Viewing help
Another useful interface feature is integrated help. You can access help for some fields and objects in the management GUI by hovering over the question mark icon near the object. Panel-specific help is available by clicking **Need Help** or by using the **Help** link in the upper-right corner of the management GUI.

Global search
The global search option enables you to find objects and menus in the system. This feature uses object attributes to search and navigates you to the selected object or menu that appears in the search result. The following objects are supported to appear in the search result:

- Partition
- Volume
- Volume group
- Snapshot
- Host
- Host port
- Host cluster
- Fibre Channel port
- IP address

Checking your web browser settings for the management GUI

You must have the latest supported web browser to access the management GUI. The management GUI supports the following HTML5-compliant browsers:

- Mozilla Firefox
- Mozilla Firefox Extended Support Release (ESR)

© Copyright IBM Corp. 2025

- Microsoft Edge
- Google Chrome

Note: The minimum viewport for browser window for the management GUI is 1024 x 768. Scaling browser view to less than those dimensions affects readability of data on management GUI pages.

IBM supports higher versions of the browsers if the vendors do not remove or disable function that the product relies upon. For browser levels higher than the versions that are certified with the product, customer support accepts usage-related and defect-related service requests. If the support center cannot re-create the issue, support might request the client to re-create the problem on a certified browser version. If a problem is identified with the browser, IBM might investigate potential solutions or work-arounds that the client can implement until a permanent solution becomes available.

The GUI requires the following configuration settings:

- JavaScript enabled
- Cookies enabled
- File Download enabled
- Disable or Replace Context menus (Mozilla Firefox only)

Refer to browser documentation on how to configure these settings.

When to use the management GUI

The management GUI is the primary tool that is used to service your system.

Regularly monitor the status of the system using the management GUI. If you suspect a problem, use the management GUI first to diagnose and resolve the problem.

Use the views that are available in the management GUI to verify the status of the system, the hardware devices, the physical storage, and the available volumes. The **Monitoring** > **Events** panel provides access to all problems that exist on the system. Use the **Recommended Actions** filter to display the most important events that need to be resolved.

If there is a service error code for the alert, you can run a fix procedure that assists you in resolving the problem. These fix procedures analyze the system and provide more information about the problem. They suggest actions to take and step you through the actions that automatically manage the system where necessary. Finally, they check that the problem is resolved.

If there is an error that is reported, always use the fix procedures within the management GUI to resolve the problem. Always use the fix procedures for both system configuration problems and hardware failures. The fix procedures analyze the system to ensure that the required changes do not cause volumes to be inaccessible to the hosts. The fix procedures automatically perform configuration changes that are required to return the system to its optimum state.

Chapter 23. FlashSystem.ai

FlashSystem.ai is an AI-driven assistant that enables natural-language administration of IBM FlashSystem storage. Administrators can request information or complete supported configuration tasks through plain-language instructions in the management GUI. The assistant analyses the request, asks for any information that is required to complete the task, and performs supported operations only after explicit user approval. This approach ensures that full administrative control remains with the organization.

About FlashSystem.ai

The assistant supports common tasks such as creating volumes, volume groups, hosts, host clusters, and snapshots. It can also assign snapshot or replication policies to volume groups to support disaster recovery workflows.

The assistant uses retrieval techniques to base responses from IBM FlashSystem product documentation, ensuring accurate and relevant answers to user queries.

Supported languages

FlashSystem.ai supports natural-language interaction in: English, Spanish, French, German, Portuguese (Brazil and Portugal), Italian, Dutch, Russian, Chinese (Simplified and Traditional), Japanese, and Korean.

English is used as the default language, and you can issue administrative commands or queries in any of the supported languages.

Supported versions

The FlashSystem.ai interface is available only on systems running 9.1.2.1 or later. Additional systems running 9.1.0 or later can be added to the container and managed through FlashSystem.ai, even though they do not display the interface themselves.

FlashSystem 5015 and 5045 systems are not supported.

FlashSystem.ai requires an active IBM Storage Insights connection, using a Storage Insights API key and tenant ID.

Capabilities

Natural-language administration

FlashSystem.ai supports natural-language interaction. You can use plain-language instructions in the management GUI to retrieve system information and perform selected administrative tasks, including the following operations and queries.

- Creating volumes, volume groups, hosts, and host clusters
- Creating snapshots
- Assigning snapshot or replication policies to volume groups
- Renaming volumes
- Increasing volume size
- Querying system configuration
- Asking follow-up questions to refine results

Supported administrative actions

You can use natural language instructions to initiate the following supported administrative actions:

- Create volumes
- Create hosts and host clusters
- Create volumes and attach to hosts

© Copyright IBM Corp. 2025

- Create volume groups
- Create snapshots
- Rename or expand volumes
- Move volumes between volume groups
- Modify snapshot and replication policies on volume groups
- Perform host-cluster operations (create clusters, map volumes)

Information retrieval capabilities

You can use natural language queries to retrieve the following information from managed systems:

- System overview
- Volumes and volume groups
- Hosts and host clusters
- Host–volume mappings
- Snapshots
- Snapshot and replication policies
- Storage pools
- Partitions
- Portsets

Filtering and aggregation are supported for all query results.

Context-aware responses

The assistant identifies the systems that you are managing and uses this context when it runs queries or other supported operations. By default, when you request information such as a list of volumes, the assistant retrieves data from the active system. You can also request information from any other system that is configured in the container by specifying the system name.

Workload placement recommendations

FlashSystem.ai integrates with IBM Storage Insights to provide workload placement recommendations. It enables natural language queries to identify optimal placement for new workloads by using fleet-wide analytics, historical performance trends, and capacity forecasts.

Systems must be registered in both FlashSystem.ai and IBM Storage Insights to use this feature. To receive workload placement recommendations in FlashSystem.ai, describe requirements in natural language. The requirements include capacity needs, performance characteristics (such as IOPS and latency), read and write patterns, block sizes, and snapshot requirements. FlashSystem.ai analyzes the storage fleet based on this input to recommend optimal workload placement options.

The time required for analysis depends on the number of systems in the fleet. Progress indicators and tool call panels display the status during execution. During this time, user interaction is unavailable. Starting a new chat session cancels any ongoing analysis.

Note: IBM Storage Insights does not support workload placement recommendations for SAN Volume Controller (SVC) systems or systems outside the specified tenant because the required historical data is unavailable. If a workload placement request includes such systems, they are excluded from consideration and the workload placement operation continues with the remaining eligible systems.

Proactive AI-initiated notifications

FlashSystem.ai supports proactive AI-initiated notifications by initiating conversations when critical issues require attention. This capability enables early identification and resolution of potential issues before they impact the environment.

The system monitors storage and IBM Storage Insights for security, compliance, and capacity issues, presenting actionable notifications in the chat interface at sign-in.

To enable email notifications, the tenant administrator email ID must be configured in IBM Storage Insights.

At sign-in, relevant notifications are displayed in the chat interface.

- Select an action to process a notification.
- Select **Mute this notification** for a notification type to suppress similar notifications.
- Enter **`/unmute`** to view and unmute notification categories.
- Up to three high-priority notifications are displayed at a time.
- Selecting an action removes all other notifications from the view.

If no action is taken, the same notifications continue to appear in subsequent chat sessions until they are addressed or replaced by higher-priority notifications. Starting a new chat session displays the next set of highest-priority notifications that are pending.

Assistant behavior and approval model

For all supported create and modify actions, FlashSystem.ai requires explicit user approval and follows these rules:

- The assistant presents a human-readable confirmation dialog before each operation it runs.
- The dialogs explain the action in plain language.
- You must explicitly approve each configuration action before it runs.
- All changes are performed through on-premises API calls; nothing is executed in the cloud.
- The assistant is available only when the AI container is deployed and registered.

You can view results in a table format that supports filtering and aggregation for easier analysis.

Guardrails for safe operation

FlashSystem.ai enforces guardrails that block destructive, security-related, and other high-risk actions. The assistant does not support the following operations:

- Deleting system objects
- Reducing the size of objects
- Reversing replication relationships
- Restoring data from snapshots
- Removing snapshot or replication policies
- Performing any security-related tasks
- Creating objects that require special user handling (for example, partitions, replication policies)

The assistant does not have access to any security-related information on the system.

Human-in-the-loop administrative actions

All configuration actions run on-premises. The assistant can perform supported actions on the user's behalf and uses human-in-the-loop prompt boxes to confirm approval before each action is run.

Access, support, and licensing

User access

FlashSystem.ai uses role-based access control to define which actions each user is permitted to perform:

- A dedicated Administrator account must be created on each system and provided to the AI container.
- This account must be provisioned manually by the customer.

Licensing

FlashSystem.ai is included with the system at no additional cost. No WatsonX.ai subscription is required.

Note: Responses generated in FlashSystem.ai are provided for convenience and can be inaccurate or incomplete. Review and verify critical information before acting on or relying on the generated content.

Upgrading FlashSystem.ai

If you have already deployed an existing instance of FlashSystem.ai, use these instructions to update your deployment to a new version concurrently without losing any previous configuration.

Procedure

1. Log in to the IBM Cloud Registry and download the latest FlashSystem.ai Helm chart.

```
helm registry login icr.io --username iamapikey --password <apikey>
helm pull oci://icr.io/flashsystem-ai-release/flashsystem-ai:1.2.0 --untar
```

2. Change to the extracted package directory.

   Note: For OpenShift deployments, ensure that the OpenShift (oc) is installed. If both `kubectl` and `oc` are installed on the host where you run the upgrade script, `remove kubectl` before running the script.

3. Run the upgrade script.

```
chmod +rx scripts/upgrade.sh
scripts/upgrade.sh
```

   During the upgrade, the following conditions are expected:

   - The upgrade process requires preserving configuration values from the existing deployment.
   - If your existing installation uses a namespace other than `flashsystem-app`, specify the namespace when prompted.
   - During a rolling update, a new FlashSystem.ai pod is started with the updated image before the existing pod is removed.
     - You can monitor the rollout status by running the **`kubectl rollout status`** command, which is displayed when the `upgrade.sh` script runs.
     - Depending on the environment, this process typically takes 2–5 minutes.
     - As a result, it is normal to temporarily see two FlashSystem.ai pods running simultaneously during the upgrade.
     - After the new pod is verified, the pod using the previous image is removed.

4. Verify that the upgrade completed successfully and that FlashSystem.ai is accessible.

Planning for the FlashSystem.ai container

The FlashSystem.ai container enables the FlashSystem.ai assistant to interpret natural-language queries and perform supported administrative tasks on IBM FlashSystem storage. The container runs in a customer-managed, on-premises Kubernetes or OpenShift environment, and integrates securely with IBM FlashSystem systems and IBM Storage Insights.

Before users can access FlashSystem.ai, you must deploy and configure the FlashSystem.ai container. Review the requirements in this topic before you begin the deployment:

Environment requirements

Secure communication

To ensure secure interaction between FlashSystem systems and the container:

- Use TLS-protected communication end to end.
- Install certificates trusted by both the FlashSystem and the FlashSystem.ai container.
- Enable only the network paths necessary for FlashSystem.ai.
- The Kubernetes ingress controller must expose a valid TLS endpoint and support WebSocket connections.

Deployment characteristics

- Deploy the container in a customer-managed, on-premises Kubernetes or OpenShift environment.
- GPUs not required.
- A single container instance can support multiple FlashSystem systems.
- One container per grid is recommended.

Connectivity

FlashSystem.ai uses REST API and WebSocket protocols. Your environment must allow the following connections:

Table 117. Connectivity requirements

Direction	Protocol	Requirement
FlashSystem ↔ Container	• HTTPS 443 • WSS 443	• Outbound to ingress endpoint • FlashSystem REST API • Bi-directional communication
Container → IBM Storage Insights	HTTPS	• Required for AI processing • Unidirectional communication

FlashSystem.ai does not connect to IBM Cloud services directly.

Required credentials

Before you begin the deployment, obtain the required credentials.

IBM Storage Insights credentials

- Storage Insights API key
- Storage Insights tenant ID

IBM FlashSystem credentials

- One administrator service account per FlashSystem

IBM Cloud credentials

- IBM Cloud login credentials

Obtaining IBM Storage Insights credentials

Log in to the IBM Storage Insights user interface to obtain the tenant ID and REST API key.

Obtain a REST API key

1. Click **Settings** icon in the top-right corner of the IBM Storage Insights toolbar.
2. In the left navigation, select **User access management**.

3. On the **User access management** page, select **Manage REST API key access**.

Creating and copying a new REST API key

1. From the **Manage REST API key access** page, select the user name in the **Name** column for whom you want to create a REST API key
2. Click **Create REST API Key**
3. Enter the required information
4. Click **Create**
5. **Copy** the generated API key and save it in a secure location

Copying an existing REST API key

1. Select the user whose API key you want to use
2. Select the REST API key
3. Verify that the key status is **Active**
4. Click **Copy** to copy the API key

Obtain the tenant ID

Use one of the following methods to obtain the tenant ID for your IBM Storage Insights instance:

1. Locate the tenant ID in the welcome email that you received when your IBM Storage Insights instance was provisioned.
2. Log in to IBM Storage Insights and copy the tenant ID from the instance URL. The tenant ID is the unique identifier that appears in the URL, as shown in the following example: If your URL is `https://insights.ibm.com/ui/01ecc493-7859-1560-a1de-f604c4ea1592/...`, then `01ecc493-7859-1560-a1de-f604c4ea1592` is your tenant ID.

Note: Save the **REST API key** and **Tenant ID** before you begin the FlashSystem.ai deployment. You must provide these values during the container configuration process.

Certificate configuration for FlashSystem.ai

Determine the certificate configuration that matches your OpenShift or Kubernetes environment for secure communication with FlashSystem.ai.

FlashSystem.ai runs in a container in your OpenShift or Kubernetes environment. To support secure communication, certificates must be configured for connections between users' browsers, the FlashSystem.ai container, and the FlashSystem system. Certificate management varies across OpenShift and Kubernetes environments. Select the procedure that matches your environment.

FlashSystem.ai uses two TLS layers:

Frontend TLS
The certificate that the Kubernetes ingress controller or OpenShift route presents to external clients.

Backend TLS
The certificate that secures communication between the ingress controller or OpenShift router and the FlashSystem.ai application.

FlashSystem.ai supports the following certificate configuration modes:

Wildcard
Uses an existing ingress or router certificate that covers the FlashSystem.ai hostname.

Custom
Uses dedicated certificates that you provide for FlashSystem.ai.

cert-manager
Uses cert-manager to issue and manage certificates automatically.

Quick selection guide

Table 118.

Certification mode	Environment
Wildcard	An existing wildcard certificate or a certificate that covers the FlashSystem.ai hostname is available
Custom	Dedicated certificates are provided for FlashSystem.ai
cert-manager	cert-manager is available and manages certificates in the cluster

Wildcard

Select `Wildcard` mode if your ingress controller or OpenShift router already provides a certificate that covers the FlashSystem.ai hostname. Examples include:

- A wildcard certificate such as `*.example.com`
- A certificate whose `CN` or `SAN` includes the exact hostname that you will use

Requirements

- Use the existing ingress or router certificate for frontend TLS.
- Create the backend TLS secret before deployment.

Custom

Select `Custom` mode if you want to provide dedicated certificates for FlashSystem.ai.

Requirements

- Create the frontend TLS secret.
- Create the backend TLS secret.

cert-manager

Select `cert-manager` mode if your Kubernetes or OpenShift environment uses cert-manager to manage certificates automatically.

Requirements

- Ensure that cert-manager is installed and configured.
- No TLS secrets are created manually before deployment.

Note: After deployment, wait until the certificate resources are ready before using FlashSystem.ai.

Set up FlashSystem.ai container

Set up the FlashSystem.ai container and configure your environment for FlashSystem.ai.

Choose an installation path

After you select your platform and TLS certificate mode, complete the setup for the platform, follow one installation section only.

Set up Kubernetes with wildcard certificate

Use this installation path when you deploy on Kubernetes and your ingress-nginx controller provides a default TLS certificate that covers your application hostname.

In Wildcard mode:

- The ingress controller uses its default certificate for frontend TLS.
- You create a backend TLS secret for TLS re-encryption between ingress-nginx and the application pod

Before you begin

- Complete the common setup steps from “Set up FlashSystem.ai container” on page 539.
- Confirm that the ingress controller default certificate matches your hostname (for example, `*.example.com` or a certificate with the hostname in the CN or SAN).
- Decide the namespace, application name, and backend secret name.

Default values used in the examples:

- Namespace: `flashsystem-app`
- Application name: `flashsystem-app`
- Backend TLS secret: `flashsystem-app-tls`

Installing TLS certificates

1. **Create the Backend TLS Secret**

   Create a TLS secret that the application uses for backend TLS. You do not need a frontend TLS secret in this mode.

   Use one of the following methods to create the TLS secret:

   - “Method 1: Use the helper script” on page 540
   - “Method 2: Create the secret manually with OpenSSL” on page 540

   Method 1: Use the helper script

   Run the script from the chart directory:

   ```
   ./scripts/create-backend-cert.sh [namespace] [secret-name] [appname]
   ```

   If you do not provide values, the script uses these defaults:

   - Namespace: `flashsystem-app`
   - Secret name: `flashsystem-app-tls`
   - App name: `flashsystem-app`

   Method 2: Create the secret manually with OpenSSL

   Create a self-signed certificate.

   ```
   NAMESPACE="flashsystem-app"
   APPNAME="flashsystem-app"

   openssl req -x509 -nodes -days 365 -newkey rsa:2048 \
     -keyout tls.key -out tls.crt \
     -subj "/CN=${APPNAME}-service.${NAMESPACE}.svc.cluster.local"
   ```

   Create the Kubernetes TLS secret

   ```
   kubectl create secret tls flashsystem-app-tls \
     --key=tls.key \
     --cert=tls.crt \
     -n flashsystem-app
   ```

Use the service DNS name in the certificate.

Verify that the secret exists

```
kubectl get secret flashsystem-app-tls -n flashsystem-app
```

2. **Configure the Helm chart values file**
   Configure the Helm chart values file before you deploy FlashSystem.ai.

   Configure `values.yaml`

   Run the script from the extracted Helm chart directory:

   ```
   scripts/configure-values.sh
   ```

   Note: If you have not already created a namespace for FlashSystem.ai, enter Y when the script prompts you to create a namespace during the Helm installation. If you already created a namespace, enter N and specify the existing namespace.

   The script:

   - Updates the `values.yaml` file with deployment-specific settings.
   - Uses the default resource names that are referenced throughout this deployment guide.
   - Prompts for alternative values if you changed any resource names from the defaults.
   - Creates the namespace, if required.

3. **Install the Helm chart**
   Install the chart in the target namespace.

   ```
   helm install flashsystem-app . -f values.yaml -n flashsystem-app
   ```

4. **Verify the deployment**

   Confirm that the release is installed

   ```
   helm list -n flashsystem-app
   ```

   Confirm that the Ingress resource exists

   ```
   kubectl get ingress -n flashsystem-app
   ```

   Confirm that the pods are ready

   ```
   kubectl get pods -n flashsystem-app
   ```

Next step

“Configure FlashSystem systems for FlashSystem.ai” on page 551.

Set up Kubernetes with custom certificate

Use this installation path when you deploy on Kubernetes and you want to provide a dedicated frontend TLS certificate for this application.

In Custom mode:

- You create a backend TLS secret for TLS re-encryption between ingress-nginx and the application pod.
- You create a frontend TLS secret that ingress-nginx presents to external

Before you begin

1. Complete the common setup steps from “Set up FlashSystem.ai container” on page 539.

2. Confirm that ingress-nginx is installed and working.
3. Ensure that your chosen hostname resolves to the Ingress endpoint.
4. Decide the namespace, application name, and secret names.

Default values used in the examples:

- Namespace: `flashsystem-app`
- Application name: `flashsystem-app`
- Backend TLS secret: `flashsystem-app-tls`
- Frontend TLS secret: `my-custom-frontend-cert`

Installing TLS certificates

1. **Create the backend TLS secret**

   Use one of the following methods to create the TLS secret:

   - “Method 1: Use the helper script” on page 542
   - “Method 2: Create the secret manually with OpenSSL” on page 542

   Method 1: Use the helper script

   Run the script from the chart directory:

   ```
   ./scripts/create-backend-cert.sh [namespace] [secret-name] [appname]
   ```

   If you do not provide values, the script uses these defaults:

   - Namespace: `flashsystem-app`
   - Secret name: `flashsystem-app-tls`
   - App name: `flashsystem-app`

   Method 2: Create the secret manually with OpenSSL

   The backend certificate secures traffic between ingress-nginx and the application pod. This traffic stays inside the cluster. A self-signed certificate is sufficient. Set the certificate common name (CN) to the in-cluster service DNS name:

   ```
   <appName>-service.<namespace>.svc.cluster.local
   ```

   Create a self-signed certificate.

   ```
   NAMESPACE="flashsystem-app"
   APPNAME="flashsystem-app"

   openssl req -x509 -nodes -days 365 -newkey rsa:2048 \
     -keyout tls.key -out tls.crt \
     -subj "/CN=${APPNAME}-service.${NAMESPACE}.svc.cluster.local"
   ```

   Create the Kubernetes TLS secret

   ```
   kubectl create secret tls flashsystem-app-tls \
     --key=tls.key \
     --cert=tls.crt \
     -n flashsystem-app
   ```

Verify that the secret exists

```
kubectl get secret flashsystem-app-tls -n flashsystem-app
```

2. **Create the frontend TLS secret**

   The frontend certificate is presented to users who connect to the application.

   Use one of the following methods to create the frontend TLS secret:

   - “Method 1: Use a certificate from your certificate authority (CA)” on page 543
   - “Method 2: Create a self-signed certificate (testing only)” on page 543

   Method 1: Use a certificate from your certificate authority (CA)

   Create the secret from your existing key and certificate files.

   ```
   kubectl create secret tls my-custom-frontend-cert \
     --key=frontend.key \
     --cert=frontend.crt \
     -n flashsystem-app
   ```

   Note: For the **AI container configuration** page, the certificate file specified by `--cert=frontend.crt` must contain the complete certificate chain, including any intermediate CA certificates. Do not provide only the leaf certificate. Otherwise, browsers might display certificate trust warnings for the FlashSystem.ai website.

   Method 2: Create a self-signed certificate (testing only)

   Create this certificate only for non-production testing. Use a CA-issued certificate in production.

   Create the certificate. Include your chosen DNS name in the CN and SAN.

   ```
   APPLICATION_DNS_NAME="<APPLICATION_DNS_NAME>"

   openssl req -x509 -nodes -days 365 -newkey rsa:2048 \
     -keyout frontend.key -out frontend.crt \
     -subj "/CN=${APPLICATION_DNS_NAME}" \
     -addext "subjectAltName = DNS:${APPLICATION_DNS_NAME}"
   ```

   Create the Kubernetes TLS secret

   ```
   kubectl create secret tls my-custom-frontend-cert \
     --key=frontend.key \
     --cert=frontend.crt \
     -n flashsystem-app
   ```

Verify that the secret exists.

```
kubectl get secret my-custom-frontend-cert -n flashsystem-app
```

3. **Configure the Helm chart values file**

   Configure the Helm chart values file before you deploy FlashSystem.ai.

   Configure `values.yaml`

   Run the script from the extracted Helm chart directory:

   ```
   scripts/configure-values.sh
   ```

   Note: If you have not already created a namespace for FlashSystem.ai, enter Y when the script prompts you to create a namespace during the Helm installation. If you already created a namespace, enter N and specify the existing namespace.

   The script:

   - Updates the `values.yaml` file with deployment-specific settings.
   - Uses the default resource names that are referenced throughout this deployment guide.
   - Prompts for alternative values if you changed any resource names from the defaults.
   - Creates the namespace, if required.

4. **Install the Helm chart**
   Install the chart in the target namespace.

   ```
   helm install flashsystem-app . -f values.yaml -n flashsystem-app
   ```

5. **Verify the deployment**

   Confirm that the release is installed

   ```
   helm list -n flashsystem-app
   ```

   Confirm that the Ingress resource exists

   ```
   kubectl get ingress -n flashsystem-app
   ```

   Confirm that the pods are ready

   ```
   kubectl get pods -n flashsystem-app
   ```

Next step

"Configure FlashSystem systems for FlashSystem.ai" on page 551.

Set up Kubernetes with cert-manager

Use this installation path when you deploy on Kubernetes and you use cert-manager to issue and renew certificates.

In cert-manager mode:

- cert-manager creates the frontend and backend certificates.
- The chart creates the certificate resources.
- You do not create TLS secrets before you install the chart

Before you begin

- Complete the common setup steps from "Set up FlashSystem.ai container" on page 539.
- Confirm that cert-manager is installed in the cluster.
- Identify the Issuer or ClusterIssuer that cert-manager uses in your cluster.
- Ensure that your chosen hostname resolves to the Ingress endpoint.

Configure Helm values

1. **Configure the Helm chart values file**
   Configure the Helm chart values file before you deploy FlashSystem.ai.

   Configure `values.yaml`

   Run the script from the extracted Helm chart directory:

   ```
   scripts/configure-values.sh
   ```

   Note: If you have not already created a namespace for FlashSystem.ai, enter Y when the script prompts you to create a namespace during the Helm installation. If you already created a namespace, enter N and specify the existing namespace.

   The script:

   - Updates the `values.yaml` file with deployment-specific settings.
   - Uses the default resource names that are referenced throughout this deployment guide.

- Prompts for alternative values if you changed any resource names from the defaults.
- Creates the namespace, if required.

2. **Install the Helm chart**
   Install the chart in the target namespace.

```
helm install flashsystem-app . -f values.yaml -n flashsystem-app
```

3. **Wait for certificates to become ready**
   Cert-manager creates the TLS secrets asynchronously. Wait until both certificates are ready.

```
kubectl get certificates -n flashsystem-app

kubectl wait --for=condition=Ready \
  certificate/flashsystem-app-frontend-tls \
  -n flashsystem-app --timeout=300s

kubectl wait --for=condition=Ready \
  certificate/flashsystem-app-backend-tls \
  -n flashsystem-app --timeout=300s
```

   If the certificates do not become ready, review the cert-manager controller logs.

```
kubectl logs -n cert-manager -l app=cert-manager
```

4. **Verify the deployment**

   Confirm that the release is installed

```
helm list -n flashsystem-app
```

   Confirm that the Ingress resource exists

```
kubectl get ingress -n flashsystem-app
```

   Confirm that the pods are ready

```
kubectl get pods -n flashsystem-app
```

Next step

“Configure FlashSystem systems for FlashSystem.ai” on page 551.

Set up OpenShift with wildcard certificate

Use this installation path when you deploy on OpenShift and the default router TLS certificate already covers your application hostname.

In wildcard mode:

- The OpenShift router uses its default certificate for frontend TLS.
- You create a backend TLS secret for TLS re-encryption between the router and the application pod.

Before you begin

- Complete the common setup steps from “Set up FlashSystem.ai container” on page 539..
- Confirm that the router default certificate matches your hostname (for example, `*.apps.example.com` or a certificate with the hostname in the CN or SAN).
- Decide the project name, application name, and backend secret name.

Default values used in the examples:

- Project: `flashsystem-app`

- Application name: `flashsystem-app`
- Backend TLS secret: `flashsystem-app-tls`

Installing TLS certificates

1. **Create the backend TLS secret**
   Create a TLS secret that the application uses for backend TLS. You do not need a frontend TLS secret in this mode.

   Use one of the following methods to create the TLS secret:

   - “Method 1: Use the helper script” on page 546
   - “Method 2: Create the secret manually with OpenSSL” on page 546

   Method 1: Use the helper script
   Run the script from the chart directory.

   ```
   ./scripts/create-backend-cert.sh [namespace] [secret-name] [appname]
   ```

   If you do not provide values, the script uses these defaults:

   - Namespace (project): `flashsystem-app`
   - Secret name: `flashsystem-app-tls`
   - App name: `flashsystem-app`

   Method 2: Create the secret manually with OpenSSL

   Create a self-signed certificate

   ```
   NAMESPACE="flashsystem-app"
   APPNAME="flashsystem-app"

   openssl req -x509 -nodes -days 365 -newkey rsa:2048 \
     -keyout tls.key -out tls.crt \
     -subj "/CN=${APPNAME}-service.${NAMESPACE}.svc.cluster.local"
   ```

   Create the OpenShift TLS secret.

   ```
   oc create secret tls flashsystem-app-tls \
     --key=tls.key \
     --cert=tls.crt \
     -n flashsystem-app
   ```

   Verify that the secret exists

   ```
   oc get secret flashsystem-app-tls -n flashsystem-app
   ```

2. **Configure the Helm chart values file**
   Configure the Helm chart values file before you deploy FlashSystem.ai.

   Configure `values.yaml`

   Run the script from the extracted Helm chart directory:

   ```
   scripts/configure-values.sh
   ```

   Note: If you have not already created a namespace for FlashSystem.ai, enter Y when the script prompts you to create a namespace during the Helm installation. If you already created a namespace, enter N and specify the existing namespace.

   The script:

   - Updates the `values.yaml` file with deployment-specific settings.
   - Uses the default resource names that are referenced throughout this deployment guide.

- Prompts for alternative values if you changed any resource names from the defaults.
- Creates the namespace, if required.

3. **Install the Helm chart**
   Install the chart in the target project.

   ```
   helm install flashsystem-app . -f values.yaml -n flashsystem-app
   ```

4. **Verify the deployment**

   Confirm that the release is installed

   ```
   helm list -n flashsystem-app
   ```

   Confirm that the Route resource exists

   ```
   oc get route -n flashsystem-app
   ```

   Confirm that the pods are ready

   ```
   oc get pods -n flashsystem-app
   ```

Next step

"Configure FlashSystem systems for FlashSystem.ai" on page 551.

Set up Open Shift with custom certificate

Use this installation path when you deploy on OpenShift and you want to provide a dedicated frontend TLS certificate for this application.

In Custom mode:

- You create a backend TLS secret for TLS re-encryption between the OpenShift router and the application pod.
- You create a frontend TLS secret that the router presents to external clients.

Before you begin

- Complete the common setup steps from "Set up FlashSystem.ai container" on page 539.
- Confirm the project name, application name, and the secret names that you will use.
- Ensure that your chosen hostname resolves to the Route endpoint.

Default values used in the examples:

- Project: `flashsystem-app`
- Application name: `flashsystem-app`
- Backend TLS secret: `flashsystem-app-tls`
- Frontend TLS secret: `my-custom-frontend-cert`

Installing TLS certificates

1. **Create the backend TLS secret**

   Use one of the following methods to create the TLS secret:

   - "Method 1: Use the helper script" on page 548
   - "Method 2: Create the secret manually with OpenSS" on page 548

Method 1: Use the helper script

Run the script from the chart directory.

```
./scripts/create-backend-cert.sh [namespace] [secret-name] [appname]
```

If you do not provide values, the script uses these defaults:

- Namespace (project): `flashsystem-app`
- Secret name: `flashsystem-app-tls`
- App name: `flashsystem-app`

Method 2: Create the secret manually with OpenSS

The backend certificate secures traffic between the OpenShift router and the application pod. This traffic stays inside the cluster. A self-signed certificate is sufficient. Set the certificate common name (CN) to the in-cluster service DNS name:

```
<appName>-service.<project>.svc.cluster.local
```

Create a self-signed certificate

```
NAMESPACE="flashsystem-app"
APPNAME="flashsystem-app"

openssl req -x509 -nodes -days 365 -newkey rsa:2048 \
  -keyout tls.key -out tls.crt \
  -subj "/CN=${APPNAME}-service.${NAMESPACE}.svc.cluster.local"
```

Create the OpenShift TLS secret

```
oc create secret tls flashsystem-app-tls \
  --key=tls.key \
  --cert=tls.crt \
  -n flashsystem-app
```

Verify that the secret exists

```
oc get secret flashsystem-app-tls -n flashsystem-app
```

2. **Create the frontend TLS secret**

The frontend certificate is presented to users who connect to the application. Choose one method.

- "Method 1: Use a certificate from your certificate authority (CA)" on page 548
- "Method 2: Create a self-signed certificate (testing only)" on page 548

Method 1: Use a certificate from your certificate authority (CA)

Create the secret from your existing key and certificate files.

```
oc create secret tls my-custom-frontend-cert \
  --key=frontend.key \
  --cert=frontend.crt \
  -n flashsystem-app
```

Method 2: Create a self-signed certificate (testing only)

Create this certificate only for non-production testing. Use a CA-issued certificate in production.

Create the certificate. Include your chosen DNS name in the CN and SAN.

```
APPLICATION_DNS_NAME="<APPLICATION_DNS_NAME>"

openssl req -x509 -nodes -days 365 -newkey rsa:2048 \
  -keyout frontend.key -out frontend.crt \
  -subj "/CN=${APPLICATION_DNS_NAME}" \
  -addext "subjectAltName = DNS:${APPLICATION_DNS_NAME}"
```

Create the OpenShift TLS secret

```
oc create secret tls my-custom-frontend-cert \
  --key=frontend.key \
  --cert=frontend.crt \
  -n flashsystem-app
```

Verify that the secret exists

```
oc get secret my-custom-frontend-cert -n flashsystem-app
```

3. **Configure the Helm chart values file**
   Configure the Helm chart values file before you deploy FlashSystem.ai.

   Configure `values.yaml`

   Run the script from the extracted Helm chart directory:

   ```
   scripts/configure-values.sh
   ```

   Note: If you have not already created a namespace for FlashSystem.ai, enter Y when the script prompts you to create a namespace during the Helm installation. If you already created a namespace, enter N and specify the existing namespace.

   The script:

   - Updates the `values.yaml` file with deployment-specific settings.
   - Uses the default resource names that are referenced throughout this deployment guide.
   - Prompts for alternative values if you changed any resource names from the defaults.
   - Creates the namespace, if required.

4. **Install the Helm chart**
   Install the chart in the target project

   ```
   helm install flashsystem-app . -f values.yaml -n flashsystem-app
   ```

5. **Verify the deployment**

   Confirm that the release is installed

   ```
   helm list -n flashsystem-app
   ```

   Confirm that the Route resource exists

   ```
   oc get route -n flashsystem-app
   ```

   Confirm that the pods are ready

   ```
   oc get pods -n flashsystem-app
   ```

Next step

"Configure FlashSystem systems for FlashSystem.ai" on page 551.

Set up OpenShift with cert-manager certificate

Use this installation path when you deploy on OpenShift and you want cert-manager to issue and renew the front-end and backend certificates automatically.

In this mode:

- Cert-manager creates the certificates and the related TLS secrets.
- You do not create TLS secrets before you install the chart.

- On OpenShift, you must update the Route after the certificates are ready.

Before you begin

1. Complete the common setup steps from “Set up FlashSystem.ai container” on page 539.
2. Verify that cert-manager is installed in the cluster.
3. Identify the Issuer or ClusterIssuer that your cluster uses.
4. Ensure that your Route hostname resolves to the OpenShift router.

Configure Helm values

1. **Configure the Helm chart values file**
   Configure the Helm chart values file before you deploy FlashSystem.ai.

   Configure `values.yaml`

   Run the script from the extracted Helm chart directory:

   ```
   scripts/configure-values.sh
   ```

   Note: If you have not already created a namespace for FlashSystem.ai, enter Y when the script prompts you to create a namespace during the Helm installation. If you already created a namespace, enter N and specify the existing namespace.

   The script:

   - Updates the `values.yaml` file with deployment-specific settings.
   - Uses the default resource names that are referenced throughout this deployment guide.
   - Prompts for alternative values if you changed any resource names from the defaults.
   - Creates the namespace, if required.

2. **Install the Helm chart**
   Install the chart in the target project:

   ```
   helm install flashsystem-app . -f values.yaml -n flashsystem-app
   ```

3. **Wait for certificates to become ready**
   Cert-manager creates the TLS secrets asynchronously. Wait until both certificates are ready:

   ```
   oc get certificates -n flashsystem-app

   oc wait --for=condition=Ready \
     certificate/flashsystem-app-frontend-tls \
     -n flashsystem-app --timeout=300s

   oc wait --for=condition=Ready \
     certificate/flashsystem-app-backend-tls \
     -n flashsystem-app --timeout=300s
   ```

   If the certificates do not become ready, review the cert-manager logs:

   ```
   oc logs -n cert-manager -l app=cert-manager
   ```

4. **Update the Route after certificate issuance**
   OpenShift Routes store certificate content in the Route specification. The initial **`helm install`** can create the Route before cert-manager creates the secrets. After the certificates are ready, run `helm upgrade` to update the Route with the issued certificate:

   ```
   helm upgrade flashsystem-app . -f values.yaml -n flashsystem-ap
   ```

   Verify that the Route includes TLS configuration:

```
oc describe route flashsystem-app-route -n flashsystem-app | grep -A 5 "TLS"
```

Next step

"Configure FlashSystem systems for FlashSystem.ai" on page 551.

Configure FlashSystem systems for FlashSystem.ai

Configure FlashSystem.ai after container setup. Register FlashSystem systems, establish AI container connectivity, and access FlashSystem.ai from the management GUI.

Using FlashSystem.ai in the management GUI

After registration is complete and successful connectivity to the container:

- The FlashSystem.ai icon appears in the management GUI.
- Selecting the icon opens the AI assistant panel.
- Each interaction creates a session that is associated with the logged-in user.
- Sessions end when the page is refreshed or when the user logs out.

Chapter 24. Ansible collection

The Ansible collection provides a series of Ansible modules and plugins for interacting with the IBM Storage Virtualize family products.

For more information, refer to https://galaxy.ansible.com/ui/repo/published/ibm/storage_virtualize/docs/.

© Copyright IBM Corp. 2025

Chapter 25. Storage Virtualize RESTful API

The Storage Virtualize Representational State Transfer (REST) model application programming interface (API) consists of command targets that are used to create, read, update, and delete system resources. These command targets allow command parameters to pass through unedited to the Storage Virtualize command-line interface, which handles parsing parameter specifications for validity and error reporting. Use Hypertext Transfer Protocol Secure (HTTPS) to successfully communicate with the RESTful API server.

The RESTful API server does not consider transport security (such as SSL), but instead assumes that requests are initiated from a local, secured server. The HTTPS protocol provides privacy through data encryption. The RESTful API provides security by requiring command authentication by using tokens. The token timeout is a user configurable option and the default value is 60 minutes. However, the token timeout value can be viewed and changed by using **lssecurity** or **chsecurity** commands.

Uniform Resource Locators (URLs) target different node objects on the system. The HTTPS POST method acts on command targets that are specified in the URL. To make changes or view information about different objects on the system, you must create and send a request to the system. You are required to provide certain elements for the RESTful API server to receive and transform the request into a command, as described in the next section.

Making an HTTPS request

To interact with the system by using the RESTful API, make an HTTPS command request with a valid configuration node URL destination. Open TCP port 7443 and include the keyword rest. For IBM Storage Virtualize for Public Cloud on Amazon Web Services (AWS), you must manually open TCP port 7443 in the AWS security group before you use the RESTful API. The following URL formats are used for all requests:

```
https://system_ip:7443/rest/v1/target/ID
```

Where:

- *system_ip* is the IP address to which you are sending requests.
- The port number is always 7443 for the Storage Virtualize RESTful API.
- *rest* is a keyword.
- *v1* is version 1 of API.
- *target* is the target object of commands, which includes any object IDs, names, and parameters.
- *ID* is the object ID.

```
https://system_ip:7443/rest/v1/command
```

Where:

- *system_ip* is the system IP address, which is the address that is taken by the configuration node of the system.
- The port number is always 7443 for the Storage Virtualize RESTful API.
- rest is a keyword.
- *v1* is the version 1 of the API.
- *command* is the target command object (such as **auth** or **lseventlog** with any parameters).

  The *command* specification follows this format:

```
command_name,method="POST",headers={'parameter_name': 'parameter_value',
'parameter_name': 'parameter_value',...}
```

© Copyright IBM Corp. 2025

Why we use POST exclusively: All of the Storage Virtualize RESTful API command targets are named after Storage Virtualize commands, with names that already reflect Create, Read, Update, and Delete actions. MK commands make (create) resources, LS commands list (read) resources, CH commands change (update) resources, and RM commands remove (delete) resources. Using HTTP methods on such commands is redundant.

All commands, including LS commands, accept at least one named parameter because the Storage Virtualize RESTful API is POST method only. It implements a convention of appending positional parameters to the URI, and packing named parameters into the request body as JavaScript Object Notation (JSON) strings.

Why we do not use GET: Most HTTP servers reject GET requests with body data. To support GET, means appending named parameters to the URI as a query string. A query string does not support arbitrary data. A valid URI cannot contain any spaces and other reserved characters that might be present in the data. Therefore, you must URL-encode named parameter data before you append it to the URI. In addition to that inconvenience, the key-value structure of named parameter data must be conveyed in some way. The query string would likely be a URL-encoded JSON object. Supporting the GET method means that parameter data represented in your programming language of choice must pass through two separate encoders before the data is sent as part of an HTTP request.

The more elegant method for the command targets was to adopt POST exclusively.

As mentioned, in addition to the URL and the name of the command target, other information is required in the request line and in the body of the HTTP request that regards the action to take on the specified object. In the request line, include the POST HTTP method. Include any required parameters (such as RAID level or IP address) in the body of the request. For more information, refer to `https://<system-ip>:7443/rest/explorer/` for OpenApi explorer guide.

Provide any required parameters as valid JSON in the HTTP body, such as shown in the following example:

```
{'X-Auth-Username': 'superuser'}
```

The request is routed to port 7443 on the specified destination (which must be the configuration node of the system), where the request is received by the RESTful API server. The server runs the command, collects any resulting output, then creates an HTTP response like the following example:

Response header

```
content-length: 273
content-type: application/json;  charset=UTF-8
date: Mon,21 Jun 2021 10:47:18 GMT
server: nginx
strict-transport-security:  max-age=31536000; includeSubDomains
```

Response body

```
{
  "key": "value"
}
```

API limits

To handle the incoming requests from client, the web server is configured with some rate limits so that API does not fail. The API returns 429 HTTP error code to the client if that limit is reached. The following table lists the rate limits that are configured on the web server.

Table 119.

Rate Limit	Type	Requests
Max active connections per cluster	RESTful API	40

Table 119. (continued)

Rate Limit	Type	Requests
Max requests/sec to auth endpoint	RESTful API	3
Max requests/sec to command endpoints	RESTful API	10
Number of simultaneous CLIs in progress	System	1

Authentication overview

Aside from data encryption, the HTTPS server requires authentication of a valid username and password for each API session. Use two authentication header fields to specify your credentials: X-Auth-Username and X-Auth-Password.

The **/auth** endpoint uses the POST method for authentication request. The **/auth** endpoint returns JWT(JSON Web Token) authentication token that authenticates every command that a user runs. The JWT token has the token expiry encoded into it.

Initial authentication requires that you POST the authentication target (/auth) with the username and password. The RESTful API server returns a JWT token. The system supports configuring a session between 10 minutes to 2 hours. However, the default single session lasts a maximum of 1 hour. If the maximum allotted time is reached, an error code 403 occurs that indicate the loss of authorization. Also, the token expiry is configurable and can be set between 10 - 120minutes. Use the **/auth** command target to authorize again with the username and password.

For example, the following command passes the authentication command to target node IP 192.168.10.109 at port 7443:

```
https://192.168.10.109:7443/rest/v1/auth, method="POST",
   headers={'X-Auth-Username': 'superuser', 'X-Auth-Password': 'passw0rd'}
```

The HTTP request that was sent to the API server looks like:

```
:authority: 192.168.10.109:7443
:method: POST
:path: /rest/v1/auth
:scheme: https
accept: application/json
content-length: 0
x-auth-password: passw0rd
x-auth-username: superuser
```

Where:

- The host is directing the HTTP request to the port (7443) and IP address (192.168.10.109) on the system .
- POST is the only HTTPS method that the Storage Virtualize RESTful API supports
- *rest* is a keyword.
- *v1* is the version 1 of the API.
- **/auth** is the authentication request command.
- The *accept* specifies the content type (application/json).
- The content length specifies the message size.
- The *x-auth-password* and *x-auth-username* field is used to authenticate and retrieve token.

The successful **auth** command returns a token similar to the following example:

```
{"token":
eyJ0eXAiOiJKV1QiLCJhbGciOiJIUzUxMiJ9.eyJpYXQiOjE2MTg4NDc4ODAsImV4cCI6MTYxODg1
MTQ4MCwianRpIjoiYTMzY2VlYmRkZGM3YmFhOTUwMmZmNGQ5ZmEwMzBhZTIiLCJzdiI6eyJ1c2VyI
joic3VwZXJ1c2VyIn19.1SXD8DoUYY-BO_cjGL9XZSnjTyBSRfp25YSWbr0Y9k5UaQrmz8xtp2XU
TR5l0PZKhqii4eU2AYf7gaTd3kfxWw}
```

The returned token string can be decoded by using any JWT decode library freely available on the web. The following is the example of the decoded JWT token string:

```
HEADER:ALGORITHM & TOKEN TYPE
{
  "typ": "JWT",
  "alg": "HS512"
}
PAYLOAD:DATA
{
  "iat": 1618847880,
  "exp": 1618851480,
  "jti": "a33ceebdddc7baa9502ff4d9fa030ae2",
  "sv": {
    "user": "superuser"
  }

}
```

API authentication management

To authenticate, use the user ID and the authentication (auth) target.

Use the authentication command target (auth) to authenticate by using the POST method. In this process, you provide a username and password through the X-Auth-Username and X-AuthPassword header fields. The beginning of the session is the only instance where you must enter the username and password. Upon successfully entering the username and password, an authentication JWT token is displayed, which has an expiry encoded into it. The token is a base64 encoded string that you include in future commands by using the X-Auth-Token header. The X-Auth-Token header in combination with the JWT authentication token replaces the username and password for each action. The JWT token authorizing hours are user configurable. The token remains authorized only for the configured time even if the token is unused. The default authorized time for the JWT token is one hour. After the configured time is reached, an error code of 403 occurs that indicates the loss of authorization. Use the **/auth** command target to reauthenticate with the username and password.

The system supports configuring a session between 10 minutes to 2 hours. However, the default single session lasts a maximum of 1 hour. If the maximum allotted time is reached, an error code 403 occurs that indicate the loss of authorization.

REST API response

The following example shows the REST response for **lsuser** command:

```
[
    {
        "id": "0",
        "name": "superuser",
        "password": "yes",
        "ssh_key": "no",
        "remote": "no",
        "usergrp_id": "0",
        "usergrp_name": "SecurityAdmin",
        "owner_id": "",
        "owner_name": "",
        "locked": "no",
        "password_change_required": "no"
    }
]
```

The following example shows **lsuser** command that uses the SSH protocol :

```
lsuser
```

The resulting output:

```
id name password ssh_key remote usergrp_id usergrp_name owner_id owner_name locked password_change_required
0 superuser yes   no      no     0          SecurityAdmin                    no        no
```

Command target parameters

Refer to `https://<system_ip>:7443/rest/explorer/` for OpenApi explorer for command targets parameters.

Storage Virtualize Performance Statistics

The IBM Storage Virtualize REST model API replaces the Common Information Model (CIM) methods to extract general performance statistics from a Storage Virtualize system for new users.

Overview

Application programmers, who are migrating away from CIM to REST API, and new REST API users can complete the following tasks to extract system performance statistics:

- Understand the now decommissioned CIM-based performance statistics for the IBM Storage Virtualize systems.
- Discover and connect to the REST API to list and download performance data files.
- Retrieve and extract metrics and apply calculations that were previously applied by the CIM agent.

Introduction

The CIM agent is a set of standards that are developed by the Distributed Management Task Force (DMTF). CIM includes an open approach to the design and implementation of storage systems, applications, databases, networks, and devices.

Newer systems increasingly rely on the REST architecture to manage and monitor systems, as it optimally utilises HTTP, thereby providing an advantage over CIM. The HTTP-based infrastructure is widely available and comes with built-in features, including caches, proxies, and scalability. REST systems are stateless, making REST operations simple and performance efficient.

IBM Storage Virtualize support a REST API implementation that is tuned for performance and scalability. Therefore, existing systems that rely on CIM agent functionality must upgrade to use the REST interface as the underlying delivery mechanism to obtain performance metrics.

The following sections provide information related to the migration process, calculations that are performed internally by the CIM agent, and coding examples. This information aims at simplifying the migration process such that the level of complexity involved in implementing and managing the existing scripts is greatly reduced.

Performance data files

The performance statistics of the IBM Storage Virtualize are stored in the performance data files.

The IBM Storage Virtualize system collects performance statistics and stores them on the management node for a certain time, usually statistics are updated at 5 minutes interval which is configurable via the chsystem command and can be viewed using the lssystem command. The performance data files are accessible over the REST API endpoints. You can list and download the files by using a scripting language, for example Python or Linux utility like cURL. The REST API accepts response type JSON (JavaScript Object Notation) which is a language-independent data format and has built-in support in

various scripting languages. You can filter the JSON response from the REST API server on the client and enable the client end programs to remove any noise from the data.

Note: The following examples use Python programming language, and the code and variables that are declared in previous examples.

Authenticating

Acquire a token to complete operations on the system by using the REST API. To request an API token, make an HTTP POST call to `/auth` endpoint and provide a username and password in the request header. See the following example.

Request

```
import requests

url = "https://example.com:7443/rest/v1"
username = "myuser"
password = "mypassword"

headers = {
    "X-Auth-Username": username,
    "X-Auth-Password": password
}

response = requests.post(url + "/auth", headers=headers, verify=False)
token = response.json()['token']

response.json()
```

Response

```
{"token":
"eyJ0eXAiOiJKV1QiLCJhbGciOiJIUzUxMiJ9.eyJpYXQiOjE3MDg1MjY4NzYsImV4cCI6MTcwODUzMDQ3NiwianRpIjoiZT
EzMzQ4YmUzZGRkMWRjZTQ4YzhkYzE5Y2RhMTFjYzMiLCJzdiI6eyJ1c2VyIjoic3VwZXJ1c2VyIn19.I--
KYnXUsFQ4wbtw_zmfkHrZ6V4bsWJgZA1j3WeHchw_kJVFRrPAhrjzJlYybIIDBYA05OSnmIZ9YWjpFY9uUQ"}
```

Listing files

The XML statistics files that contain the performance data are stored in `/dumps/iostats/` directory on the config node. The `/lsdumps` API endpoint identifies the files that are accessible for download through the API. To list the contents of `/dumps/iostats/`, use the `prefix` parameter in the request body. See the following example.

Request

```
headers = {
    "X-Auth-Token": token,
}
data = {
    "prefix": "/dumps/iostats/"
}
response = requests.post(url + "/lsdumps", json=data, headers=headers, verify=False)
response.json()
```

Response

```
[{'id': '0', 'filename': 'Nm_stats_F304745-2_240221_151627'},
 {'id': '1', 'filename': 'Ng_stats_F304745-2_240221_151627'},
 {'id': '2', 'filename': 'Nn_stats_F304745-2_240221_151627'},
 {'id': '3', 'filename': 'Nv_stats_F304745-2_240221_151627'},
…
]
```

Filtering on file names

The system collects statistics over an interval and creates files that can be downloaded. For each collection interval, the system creates five statistics files:

- Nm_stats for managed disks (MDisks)
- Nv_stats for volumes and volume copies
- Ng_stats for volume groups
- Nn_stats for nodes
- Nd_stats for drives

The file name consists of several parts, as shown in the following description:

<category>_<panel_name>_<yymmdd>_<hhmmss>

For example, Nv_stats_F304745-2_240221_151627

The naming convention enables a filtering capability that a client-side script might use to download:

- Specific type of files
- Files that were created at a specific time interval.

See the following example.

Request

```
import re

headers = {
    "X-Auth-Token": token,
}
data = {
    "prefix": "/dumps/iostats/"
}
response = requests.post(url + "/lsdumps", json=data, headers=headers, verify=False)
files = [row['filename'] for row in response.json()]

# Filter on category
nm_files = [name for name in files if re.match("Nm_stats", name)]
nv_files = [name for name in files if re.match("Nv_stats", name)]
ng_files = [name for name in files if re.match("Ng_stats", name)]
nn_files = [name for name in files if re.match("Nn_stats", name)]
nd_files = [name for name in files if re.match("Nd_stats", name)]
# Filter on time interval
tm_files = [name for name in files if re.match(".*_142127$", name)]
# Filter on panel name
pn_files = [name for name in files if re.match(".*F304745-2.*", name)]
# Make a selection
target_file = nm_files[0]
# Print
pn_files[:2]
```

Response

```
['Nm_stats_F304745-2_240222_143127', 'Nn_stats_F304745-2_240222_143127']
```

Downloading files

You can download the files that can be listed by using the /lsdumps API endpoint, through the /download endpoint. The /download interface accepts two parameters:

- prefix: Name of the target directory
- filename: Name of the target file

The contents of the downloaded file can be processed either in memory or stored on the host for post processing.

Note: The performance data files are constantly updated on the system. If an error occurs during the file download that means the file is rotated. Usually, you can fetch the latest list of files, followed by another download request. If not, reevaluate statistics intervals and collection intervals to ensure that the list and download is completed within the network and system limits.

See the following example.

Request

```
headers = {
    "X-Auth-Token": token,
}
data = {
    "prefix": "/dumps/iostats",
    "filename": target_file
}
response = requests.post(url + "/download", json=data, headers=headers, verify=False)

if response.status_code != 200:
    raise Exception("File not found")

response.content[:100]
```

Response

```
b'<?xml version="1.0" encoding="utf-8" ?>\n<diskStatsColl\nxmlns="http://ibm.com/storage/
management/perf'
```

Getting performance statistics

You can extract various data metrics from the system by using code.

The following code snippet provides a reference for authenticating a user session, getting a list of the statistics files, and downloading a statistics file for individual calculations. The example code is also used in later sections to perform the actual calculations for various types of metrics that can be pulled from the system.

Note: The following examples use Python programming language, and the code and variables that are declared in previous examples.

```
import re
import xml.etree.ElementTree as ET
import requests
import json
import pprint

def get_token(url, username, password):
    response = requests.post(
        url + "/auth",
        headers={"X-Auth-Username": username, "X-Auth-Password": password},
        verify=False)
    return response.json()['token']

def get_namespace(element):
    m = re.match(r'\{.*\}', element.tag)
    return m.group(0) if m else ''

def get_directory_listing(url, token):
    response = requests.post(
        url + "/lsdumps",
        json={"prefix": "/dumps/iostats"},
        headers={"X-Auth-Token": token},
        verify=False)
    if response.status_code != 200:
        raise Exception("Invalid directory listing")
    return response.json()

def get_file_contents(url, token, filename):
    response = requests.post(
        url + "/download",
        json={"prefix": "/dumps/iostats", "filename": filename},
        headers={"X-Auth-Token": token},
        verify=False)
    if response.status_code != 200:
        raise Exception("File not found")
    return response.content

def get_latest_stats(url, username, password):
    token = get_token(url, username, password)
    files = [row['filename'] for row in get_directory_listing(url, token)]
    nm_files = [name for name in files if re.match("Nm_stats", name)]
    nv_files = [name for name in files if re.match("Nv_stats", name)]
```

```
    ng_files = [name for name in files if re.match("Ng_stats", name)]
    nn_files = [name for name in files if re.match("Nn_stats", name)]
    nd_files = [name for name in files if re.match("Nd_stats", name)]
    nh_files = [name for name in files if re.match("Nh_stats", name)]

    nm_xml = ET.fromstring(get_file_contents(url, token, nm_files[0]))
    nv_xml = ET.fromstring(get_file_contents(url, token, nv_files[0]))
    ng_xml = ET.fromstring(get_file_contents(url, token, ng_files[0]))
    nn_xml = ET.fromstring(get_file_contents(url, token, nn_files[0]))
    nd_xml = ET.fromstring(get_file_contents(url, token, nd_files[0]))
    nh_xml = ET.fromstring(get_file_contents(url, token, nh_files[0]))

    xml = {
        "Nm_stats": nm_xml,
        "Nv_stats": nv_xml,
        "Ng_stats": ng_xml,
        "Nn_stats": nn_xml,
        "Nd_stats": nd_xml,
        "Nh_stats": nh_xml
    }
    return xml

def populate_fields(xml, fields):
    for item in fields.items():
        item[1]['value'] = int(
            xml[item[1]['filename']].
            find(get_namespace(xml[item[1]['filename']]) + item[1]['element']).
            attrib[item[1]['field']]
        )
```

For details on different types of metrics, see the following sections:

Node statistics

The following table lists:

- Performance statistics calculations previously provided by the CIM agent
- Respective fields in the statistics file that is reported for individual nodes.

Table 120. Calculations

CIMOM field	Calculation
TotalIOs	READ_OPERATION_TAG + WRITE_OPERATION_TAG
ReadIOs	READ_OPERATION_TAG
ReadHitIOs	READ_HITCACHE_TAG
WriteIOs	WRITE_OPERATION_TAG
WriteHitIOs	WRITE_HITCACHE_TAG

Var	File name	XML field	Description
DEVICE_ID_TAG	Nn_stats	"idx"	Indicates the enclosure or node ID
READ_OPERATION_TAG	Nn_stats	"ro"	Indicates the number of Read I/Os
WRITE_OPERATION_TAG	Nn_stats	"wo"	Indicates the number of Write I/Os
READ_HITCACHE_TAG	Nv_stats	"ctrh"	Indicates the number of reads received from other components and processed in read-cache hit

Var	File name	XML field	Description
WRITE_HITCACHE_TAG	Nv_stats	"ctwfw"	Indicates the number of writes received from other components and processed in write-cache hit

The following code example shows an illustration of the calculations in a scripting form.

Request

```
def get_latest_cimom_sample_node_stats(url, username, password):
    xml = get_latest_stats(url, username, password)

    fields = {
        "READ_OPERATION_TAG": {
            "filename": "Nn_stats",
            "element": "node",
            "field": "ro",
            "value": 0},
        "WRITE_OPERATION_TAG": {
            "filename": "Nn_stats",
            "element": "node",
            "field": "wo",
            "value": 0},
        "READ_HITCACHE_TAG": {
            "filename": "Nv_stats",
            "element": "vdsk",
            "field": "ctrh",
            "value": 0},
        "WRITE_HITCACHE_TAG": {
            "filename": "Nv_stats",
            "element": "vdsk",
            "field": "ctwfw",
            "value": 0}
    }
    populate_fields(xml, fields)
    cimom_fields = {
        "TotalIOs": fields['READ_OPERATION_TAG']['value'] + fields['WRITE_OPERATION_TAG']
['value'],
        "ReadIOs": fields['READ_OPERATION_TAG']['value'],
        "ReadHitIOs": fields['READ_HITCACHE_TAG']['value'],
        "WriteIOs": fields['WRITE_OPERATION_TAG']['value'],
        "WriteHitIOs": fields['WRITE_HITCACHE_TAG']['value']
    }
    return cimom_fields

pprint.pprint(get_latest_cimom_node_stats(url, username, password))
```

Response

```
{'ReadHitIOs': 0,
 'ReadIOs': 19291610,
 'TotalIOs': 39009988,
 'WriteHitIOs': 0,
 'WriteIOs': 19718378}
```

Backend (MDisk) volumes

The following table lists:

- Performance statistics calculations accomplished by the CIM agent
- Respective fields in the statistics file that is reported for individual mdisks.

Table 121. Calculations

CIMOM field	Calculation
deviceId	DEVICE_ID_TAG

Table 121. Calculations (continued)	
CIMOM field	**Calculation**
TotalIOs	READ_OPERATION_TAG + WRITE_OPERATION_TAG
KBytesTransferred(sizeUnits)	(READ_BLOCK_TAG + WRITE_BLOCK_TAG)*sizeUnits/1024
IOTimeCounter	READ_TIME_TAG + WRITE_TIME_TAG
ReadIOs	READ_OPERATION_TAG
ReadIOTimeCounter	READ_TIME_TAG
KBytesRead(sizeUnits)	READ_BLOCK_TAG*sizeUnits/1024
WriteIOs	WRITE_OPERATION_TAG
WriteIOTimeCounter	WRITE_TIME_TAG
KBytesWritten(sizeUnits)	WRITE_BLOCK_TAG*sizeUnits/1024

Var	File name	XML field	Description
DEVICE_ID_TAG	Nm_stats	"idx"	Indicates the name of the MDisk for which the statistics apply
READ_OPERATION_TAG	Nm_stats	"ro"	Indicates the cumulative number of MDisk read operations that were processed (since the node started)
WRITE_OPERATION_TAG	Nm_stats	"wo"	Indicates the cumulative number of MDisk write operations that were processed (since the node started)
READ_BLOCK_TAG	Nm_stats	"rb"	Indicates the cumulative number of blocks of data that was read (since the node started)
WRITE_BLOCK_TAG	Nm_stats	"wb"	Indicates the cumulative number of blocks of data written (since the node started)
READ_TIME_TAG	Nm_stats	"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 host read command is issued and stopped when the command

Var	File name	XML field	Description
			completes successfully. The elapsed time is added to the cumulative counter
WRITE_TIME_TAG	Nm_stats	"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 host write command is issued and stopped when the command completes successfully. The elapsed time is added to the cumulative counter

The following code example shows an illustration of the calculations in a scripting form.

Request

```
def get_latest_cimom_sample_mdisk_stats(url, username, password):
    xml = get_latest_stats(url, username, password)

    fields = {
        "READ_OPERATION_TAG": {
            "filename": "Nm_stats",
            "element": "mdsk",
            "field": "ro",
            "value": 0},
        "WRITE_OPERATION_TAG": {
            "filename": "Nm_stats",
            "element": "mdsk",
            "field": "wo",
            "value": 0},
        "READ_BLOCK_TAG": {
            "filename": "Nm_stats",
            "element": "mdsk",
            "field": "rb",
            "value": 0},
        "WRITE_BLOCK_TAG": {
            "filename": "Nm_stats",
            "element": "mdsk",
            "field": "wb",
            "value": 0},
        "READ_TIME_TAG": {
            "filename": "Nm_stats",
            "element": "mdsk",
            "field": "re",
            "value": 0},
        "WRITE_TIME_TAG": {
            "filename": "Nm_stats",
            "element": "mdsk",
            "field": "re",
            "value": 0}
    }
    populate_fields(xml, fields)
    cimom_fields = {
        "TotalIOs": fields['READ_OPERATION_TAG']['value'] + fields['WRITE_OPERATION_TAG']
['value'],
        "KBytesTransferred": (fields['READ_BLOCK_TAG']['value'] + fields['WRITE_BLOCK_TAG']
['value']) / 1024,
        "IOTimeCounter": fields['READ_TIME_TAG']['value'] + fields['WRITE_TIME_TAG']['value'],
        "ReadIOs": fields['READ_OPERATION_TAG']['value'],
        "ReadIOTimeCounter": fields['READ_TIME_TAG']['value'],
        "KBytesRead": fields['READ_BLOCK_TAG']['value'] / 1024,
```

```
        "WriteIOs": fields['WRITE_OPERATION_TAG']['value'],
        "WriteIOTimeCounter": fields['WRITE_TIME_TAG']['value'],
        "KBytesWritten": fields['WRITE_BLOCK_TAG']['value'] / 1024
    }
    return cimom_fields
pprint.pprint(get_latest_cimom_sample_mdisk_stats(url, username, password))
```

Response

```
{'IOTimeCounter': 0,
 'KBytesRead': 0.0,
 'KBytesTransferred': 1632.0,
 'KBytesWritten': 1632.0,
 'ReadIOTimeCounter': 0,
 'ReadIOs': 0,
 'TotalIOs': 18,
 'WriteIOTimeCounter': 0,
 'WriteIOs': 18}
```

Volume (VDisk) statistics

The following table lists:

- Performance statistics calculations accomplished by the CIM agent
- Respective fields in the statistics file that is reported for individual vdisks.

Table 122. Calculations

CIMOM field	Calculation
deviceId	DEVICE_ID_TAG
TotalIOs	READ_OPERATION_TAG + WRITE_OPERATION_TAG
KBytesTransferred(sizeUnits)	(READ_BLOCK_TAG + WRITE_BLOCK_TAG)*sizeUnits/1024
IOTimeCounter	READ_TIME_TAG + WRITE_TIME_TAG
ReadIOs	READ_OPERATION_TAG
ReadHitIOs	READ_HITCACHE_TAG
ReadIOTimeCounter	READ_TIME_TAG
KBytesRead(sizeUnits)	READ_BLOCK_TAG*sizeUnits/1024
WriteIOs	WRITE_OPERATION_TAG
WriteHitIOs	WRITE_HITCACHE_TAG
WriteIOTimeCounter	WRITE_TIME_TAG
KBytesWritten(sizeUnits)	WRITE_BLOCK_TAG*sizeUnits/1024

Var	File name	XML field	Description
DEVICE_ID_TAG	Nv_stats	"idx"	Indicates the volume for which the statistics apply
READ_OPERATION_TAG	Nv_stats	"ro"	Indicates the cumulative number of volume read operations that were processed (since the node started)

Var	File name	XML field	Description
WRITE_OPERATION_TAG	Nv_stats	"wo"	Indicates the cumulative number of volumes write operations that were processed (since the node started).
READ_BLOCK_TAG	Nv_stats	"rb"	Indicates the cumulative number of blocks of data read (since the node started)
WRITE_BLOCK_TAG	Nv_stats	"wb"	Indicates the cumulative number of blocks of data written (since the node started)
READ_TIME_TAG	Nv_stats	"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.
WRITE_TIME_TAG	Nv_stats	"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.
READ_HITCACHE_TAG	Nv_stats	"ctrh"	Indicates the number of reads received from other components and processed in read-cache
WRITE_HITCACHE_TAG	Nv_stats	"ctwfw"	Indicates the number of writes received from other components and processed in write-cache

The following code example shows an illustration of the calculations in a scripting form.

Request

```
def get_latest_cimom_sample_vdisk_stats(url, username, password):
    xml = get_latest_stats(url, username, password)

    fields = {
        "READ_OPERATION_TAG": {
            "filename": "Nv_stats",
            "element": "vdsk",
            "field": "ro",
            "value": 0},
        "WRITE_OPERATION_TAG": {
            "filename": "Nv_stats",
            "element": "vdsk",
            "field": "wo",
            "value": 0},
        "READ_BLOCK_TAG": {
            "filename": "Nv_stats",
            "element": "vdsk",
            "field": "rb",
            "value": 0},
        "WRITE_BLOCK_TAG": {
            "filename": "Nv_stats",
            "element": "vdsk",
            "field": "wb",
            "value": 0},
        "READ_TIME_TAG": {
            "filename": "Nv_stats",
            "element": "vdsk",
            "field": "rl",
            "value": 0},
        "WRITE_TIME_TAG": {
            "filename": "Nv_stats",
            "element": "vdsk",
            "field": "wl",
            "value": 0},
        "READ_HITCACHE_TAG": {
            "filename": "Nv_stats",
            "element": "vdsk",
            "field": "ctrh",
            "value": 0},
        "WRITE_HITCACHE_TAG": {
            "filename": "Nv_stats",
            "element": "vdsk",
            "field": "ctwfw",
            "value": 0}
    }
    populate_fields(xml, fields)
    cimom_fields = {
        'TotalIOs': fields['READ_OPERATION_TAG']['value'] + fields['WRITE_OPERATION_TAG']['value'],
        'KBytesTransferred': (fields['READ_BLOCK_TAG']['value'] + fields['WRITE_BLOCK_TAG']['value'])
* 1024 / 1024,
        'IOTimeCounter': fields['READ_TIME_TAG']['value'] + fields['WRITE_TIME_TAG']['value'],
        'ReadIOs': fields['READ_OPERATION_TAG']['value'],
        'ReadHitIOs': fields['READ_HITCACHE_TAG']['value'],
        'ReadIOTimeCounter': fields['READ_TIME_TAG']['value'],
        'KBytesRead': fields['READ_BLOCK_TAG']['value'] * 1024 / 1024,
        'WriteIOs': fields['WRITE_OPERATION_TAG']['value'],
        'WriteHitIOs': fields['WRITE_HITCACHE_TAG']['value'],
        'WriteIOTimeCounter': fields['WRITE_TIME_TAG']['value'],
        'KBytesWritten': fields['WRITE_BLOCK_TAG']['value'] * 1024 / 1024
    }
    return cimom_fields

pprint.pprint(get_latest_cimom_sample_vdisk_stats(url, username, password))
```

Response

```
{'IOTimeCounter': 0,
'KBytesRead': 0.0,
'KBytesTransferred': 1632.0,
'KBytesWritten': 1853.0,
'ReadHitIOs': 0,
'ReadIOTimeCounter': 3,
'ReadIOs': 178,
'TotalIOs': 178,
'WriteHitIOs': 54,
```

```
 'WriteIOTimeCounter': 2,
 'WriteIOs': 54}
```

Disk (drive) statistics

The following table lists:

- Performance statistics calculations accomplished by the CIM agent
- Respective fields in the statistics file that is reported for individual drives.

Table 123. Calculations

CIMOM field	Calculation
deviceId	DEVICE_ID_TAG
TotalIOs	READ_OPERATION_TAG + WRITE_OPERATION_TAG
KBytesTransferred(sizeUnits)	(READ_BLOCK_TAG + WRITE_BLOCK_TAG)*sizeUnits/1024
IOTimeCounter	READ_TIME_TAG + WRITE_TIME_TAG
ReadIOs	READ_OPERATION_TAG
ReadIOTimeCounter	READ_TIME_TAG
KBytesRead(sizeUnits)	READ_BLOCK_TAG*sizeUnits/1024
WriteIOs	WRITE_OPERATION_TAG
WriteIOTimeCounter	WRITE_TIME_TAG
KBytesWritten(sizeUnits)	WRITE_BLOCK_TAG*sizeUnits/1024

Var	File name	XML field	Description
DEVICE_ID_TAG	Nd_stats	"idx"	Indicates the name of the MDisk for which the statistics apply
READ_OPERATION_TAG	Nd_stats	"ro"	Indicates the cumulative number of MDisk read operations that were processed (since the node started)
WRITE_OPERATION_TAG	Nd_stats	"wo"	Indicates the cumulative number of MDisk write operations that were processed (since the node started)
READ_BLOCK_TAG	Nd_stats	"rb"	Indicates the cumulative number of blocks of data that was read (since the node started).
WRITE_BLOCK_TAG	Nd_stats	"wb"	Indicates the cumulative number of blocks of data written (since the node started)

Var	File name	XML field	Description
READ_TIME_TAG	Nd_stats	"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
WRITE_TIME_TAG	Nd_stats	"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

The following code example shows an illustration of the calculations in a scripting form.

Request

```
def get_latest_cimom_sample_drive_stats(url, username, password):
    xml = get_latest_stats(url, username, password)

    fields = {
        "READ_OPERATION_TAG": {
            "filename": "Nd_stats",
            "element": "mdsk",
            "field": "ro",
            "value": 0},
        "WRITE_OPERATION_TAG": {
            "filename": "Nd_stats",
            "element": "mdsk",
            "field": "wo",
            "value": 0},
        "READ_BLOCK_TAG": {
            "filename": "Nd_stats",
            "element": "mdsk",
            "field": "rb",
            "value": 0},
        "WRITE_BLOCK_TAG": {
            "filename": "Nd_stats",
            "element": "mdsk",
            "field": "wb",
            "value": 0},
        "READ_TIME_TAG": {
            "filename": "Nd_stats",
            "element": "mdsk",
            "field": "re",
            "value": 0},
```

```
        "WRITE_TIME_TAG": {
            "filename": "Nd_stats",
            "element": "mdsk",
            "field": "we",
            "value": 0}
    }
    populate_fields(xml, fields)
    cimom_fields = {
        'TotalIOs': fields['READ_OPERATION_TAG']['value'] + fields['WRITE_OPERATION_TAG']['value'],
        'KBytesTransferred': (fields['READ_BLOCK_TAG']['value'] + fields['WRITE_BLOCK_TAG']['value'])
* 1024 / 1024,
        'IOTimeCounter': fields['READ_TIME_TAG']['value'] + fields['WRITE_TIME_TAG']['value'],
        'ReadIOs': fields['READ_OPERATION_TAG']['value'],
        'ReadIOTimeCounter': fields['READ_TIME_TAG']['value'],
        'KBytesRead': fields['READ_BLOCK_TAG']['value'] * 1024 / 1024,
        'WriteIOs': fields['WRITE_OPERATION_TAG']['value'],
        'WriteIOTimeCounter': fields['WRITE_TIME_TAG']['value'],
        'KBytesWritten': fields['WRITE_BLOCK_TAG']['value'] * 1024 / 1024

    }
    return cimom_fields

pprint.pprint(get_latest_cimom_sample_drive_stats(url, username, password))
```

Response

```
{'IOTimeCounter': 409969,
'KBytesRead': 356062720.0,
'KBytesTransferred': 356064872.0,
'KBytesWritten': 2152.0,
'ReadIOTimeCounter': 409940,
'ReadIOs': 695437,
'TotalIOs': 695451,
'WriteIOTimeCounter': 29,
'WriteIOs': 14}
```

Fibre Channel (FC) statistics

The following table lists:

- Performance statistics calculations accomplished by the CIM agent
- Respective fields in the statistics file that is reported for individual FC ports.

Table 124. Calculations

CIMOM field	**Calculation**
portType	PORT_TYPE_TAG
deviceId	PORT_FC_WWPN_TAG
portId	PORT_ID_TAG
nodeId	NODE_ID
LinkFailures	LINK_FAILURES
LossOfSignalCounter	LOSS_SIGNAL
LossOfSyncCounter	LOSS_SYNC
PrimitiveSeqProtocolErrCount	PRIMITIVE_SEQ_PROTOCOL_ERRCOUNT
CRCErrors	INVALID_CRC
InvalidTransmissionWords	INVALID_WORD
BBCreditZeroTime	ZERO_BUFFER_CREDIT_COUNTER
writeBytesTransmitted	BYTES_TO_CONTROLLER + BYTES_TO_CLUSTERNODE + BYTES_TO_HOST + BYTES_TO_OTHER_CLUSTERNODE

Table 124. Calculations (continued)

CIMOM field	Calculation
BytesReceived	BYTES_FROM_CLUSTERNODE + BYTES_FROM_CONTROLLER + BYTES_FROM_HOST + BYTES_FROM_OTHER_CLUSTERNODE
DataTransmissionDelayCount	DATA_TRANS_DELAY_COUNT
DataTransmissionDelayTime	DATA_TRANS_DELAY_TIME
DataTransmissionDelayMeasured	DATA_TRANS_DELAY_MEASURED

Var	File name	XML field	Description
PORT_ID_TAG	Nn_stats	"id"	Indicates the port identifier for the node
PORT_TYPE_TAG	Nn_stats	"type"	
PORT_FC_WWPN_TAG	Nn_stats	"wwpn"	Indicates the worldwide port name for the node
LINK_FAILURES	Nn_stats	"lf"	Indicates a link failure count
LOSS_SYNC	Nn_stats	"lsy"	Indicates the loss-of-synchronization count
LOSS_SIGNAL	Nn_stats	"lsi"	Indicates the lost-of-signal count
PRIMITIVE_SEQ_PROTOCOL_ERRCOUNT	Nn_stats	"pspe"	Indicates the primitive sequence-protocol error count
INVALID_WORD	Nn_stats	"itw"	Indicates the number of transmission word counts that are not valid
INVALID_CRC	Nn_stats	"icrc"	Indicates the number of CRC that is not valid
ZERO_BUFFER_CREDIT_COUNTER	Nn_stats	"bbcz"	Indicates the total time in microseconds for which the buffer credit counter was at zero. This statistic is only reported by 8 Gbps Fibre Channel ports. For other port types, this statistic is 0
BYTES_FROM_CLUSTERNODE	Nn_stats	"lnbr"	Indicates the bytes received to other nodes in the same cluster
BYTES_FROM_CONTROLLER	Nn_stats	"cbr"	Indicates the bytes received from controllers
BYTES_FROM_HOST	Nn_stats	"hbr"	Indicates the bytes received from hosts

Var	File name	XML field	Description
BYTES_FROM_OTHER_CLUSTERNODE	Nn_stats	"rmbr"	Indicates the bytes received to other nodes in the other clusters
BYTES_TO_CLUSTERNODE	Nn_stats	"lnbt"	Indicates the bytes transmitted to other nodes in the same cluster
BYTES_TO_HOST	Nn_stats	"hbt"	Indicates the bytes transmitted to hosts
BYTES_TO_CONTROLLER	Nn_stats	"cbt"	Indicates the bytes transmitted to disk controllers
BYTES_TO_OTHER_CLUSTERNODE	Nn_stats	"rmbt"	Indicates the bytes transmitted to other nodes in the other clusters
DATA_TRANS_DELAY_MEASURED	Nn_stats	"dtdm"	Indicates the number of transfers that had their data transmission delay measured
DATA_TRANS_DELAY_COUNT	Nn_stats	"dtdc"	Indicates the number of transfers that experienced excessive data transmission delay
DATA_TRANS_DELAY_TIME	Nn_stats	"dtdt"	Indicates the total time in microseconds for which data transmission was excessively delayed

The following code example shows an illustration of the calculations in a scripting form.

Request

```
def get_latest_cimom_sample_fc_port_stats(url, username, password):
    xml = get_latest_stats(url, username, password)

    fields = {
        "LINK_FAILURES": {
            "filename": "Nn_stats",
            "element": "port",
            "field": "lf",
            "value": 0},
        "LOSS_SYNC": {
            "filename": "Nn_stats",
            "element": "port",
            "field": "lsy",
            "value": 0},
        "LOSS_SIGNAL": {
            "filename": "Nn_stats",
            "element": "port",
            "field": "lsi",
            "value": 0},
        "PRIMITIVE_SEQ_PROTOCOL_ERRCOUNT": {
            "filename": "Nn_stats",
            "element": "port",
            "field": "pspe",
            "value": 0},
        "INVALID_WORD": {
            "filename": "Nn_stats",
            "element": "port",
```

```
            "field": "itw",
            "value": 0},
        "INVALID_CRC": {
            "filename": "Nn_stats",
            "element": "port",
            "field": "icrc",
            "value": 0},
        "ZERO_BUFFER_CREDIT_COUNTER": {
            "filename": "Nn_stats",
            "element": "port",
            "field": "bbcz",
            "value": 0},
        "BYTES_FROM_CLUSTERNODE": {
            "filename": "Nn_stats",
            "element": "port",
            "field": "lnbr",
            "value": 0},
        "BYTES_FROM_CONTROLLER": {
            "filename": "Nn_stats",
            "element": "port",
            "field": "cbr",
            "value": 0},
        "BYTES_FROM_HOST": {
            "filename": "Nn_stats",
            "element": "port",
            "field": "hbr",
            "value": 0},
        "BYTES_FROM_OTHER_CLUSTERNODE": {
            "filename": "Nn_stats",
            "element": "port",
            "field": "rmbr",
            "value": 0},
        "BYTES_TO_CLUSTERNODE": {
            "filename": "Nn_stats",
            "element": "port",
            "field": "lnbt",
            "value": 0},
        "BYTES_TO_HOST": {
            "filename": "Nn_stats",
            "element": "port",
            "field": "hbt",
            "value": 0},
        "BYTES_TO_CONTROLLER": {
            "filename": "Nn_stats",
            "element": "port",
            "field": "cbt",
            "value": 0},
        "BYTES_TO_OTHER_CLUSTERNODE": {
            "filename": "Nn_stats",
            "element": "port",
            "field": "rmbt",
            "value": 0},
        "DATA_TRANS_DELAY_MEASURED": {
            "filename": "Nn_stats",
            "element": "port",
            "field": "dtdm",
            "value": 0},
        "DATA_TRANS_DELAY_COUNT": {
            "filename": "Nn_stats",
            "element": "port",
            "field": "dtdc",
            "value": 0},
        "DATA_TRANS_DELAY_TIME": {
            "filename": "Nn_stats",
            "element": "port",
            "field": "dtdt",
            "value": 0}
    }
    populate_fields(xml, fields)
    cimom_fields = {
        'LinkFailures': fields['LINK_FAILURES']['value'],
        'LossOfSignalCounter': fields['LOSS_SIGNAL']['value'],
        'LossOfSyncCounter': fields['LOSS_SYNC']['value'],
        'PrimitiveSeqProtocolErrCount': fields['PRIMITIVE_SEQ_PROTOCOL_ERRCOUNT']['value'],
        'CRCErrors': fields['INVALID_CRC']['value'],
        'InvalidTransmissionWords': fields['INVALID_WORD']['value'],
        'BBCreditZeroTime': fields['ZERO_BUFFER_CREDIT_COUNTER']['value'],
        'writeBytesTransmitted': fields['BYTES_TO_CONTROLLER']['value'] +
fields['BYTES_TO_CLUSTERNODE']['value'] + fields['BYTES_TO_HOST']['value'] +
fields['BYTES_TO_OTHER_CLUSTERNODE']['value'],
        'BytesReceived': fields['BYTES_FROM_CLUSTERNODE']['value'] + fields['BYTES_FROM_CONTROLLER']
['value'] + fields['BYTES_FROM_HOST']['value'] + fields['BYTES_FROM_OTHER_CLUSTERNODE']
```

```
['value'],
        'DataTransmissionDelayCount': fields['DATA_TRANS_DELAY_COUNT']['value'],
        'DataTransmissionDelayTime': fields['DATA_TRANS_DELAY_TIME']['value'],
        'DataTransmissionDelayMeasured': fields['DATA_TRANS_DELAY_MEASURED']['value']
    }
    return cimom_fields

pprint.pprint(get_latest_cimom_sample_fc_port_stats(url, username, password))
```

Response

```
{'BBCreditZeroTime': 0,
'BytesReceived': 1710209,
'CRCErrors': 0,
'DataTransmissionDelayCount': 32,
'DataTransmissionDelayMeasured': 20098258,
'DataTransmissionDelayTime': 998,
'InvalidTransmissionWords': 7,
'LinkFailures': 6,
'LossOfSignalCounter': 0,
'LossOfSyncCounter': 6,
'PrimitiveSeqProtocolErrCount': 0,
'writeBytesTransmitted': 628974}
```

Chapter 26. Conventions used in this IBM Documentation

Ensure that you are familiar with the notational conventions that are used in this IBM Documentation.

Notifications

This topic describes the meaning of notices and how to search, synchronize, and navigate by keyboard within the IBM Documentation.

Notices

The following notices are used in the product documentation and convey the following meanings:

Note: These notices provide important tips, guidance, or advice.



Attention: These notices indicate possible damage to programs, devices, or data. An attention notice is displayed before the instruction or situation in which damage could occur.



CAUTION: These notices indicate situations that can be potentially hazardous to you. A caution notice is displayed before the description of a potentially hazardous procedural step or situation.



DANGER: These notices indicate situations that can be potentially lethal or extremely hazardous to you. A danger notice is displayed before a description of a potentially lethal or extremely hazardous procedural step or situation.

Emphasis

Different typefaces are used in this guide to show emphasis.

The following typefaces are used to show emphasis.

Emphasis	Meaning
Boldface	Text in **boldface** represents menu items.
`Bold monospace`	Text in **`bold monospace`** represents command names.
Italics	Text in *italics* is used to emphasize a word. In command syntax, it is used for variables for which you supply actual values, such as a default directory or the name of a system.
`Monospace`	Text in `monospace` identifies the data or commands that you type, samples of command output, examples of program code or messages from the system, or names of command flags, parameters, arguments, and name-value pairs.

Numbering conventions

A specific numbering convention is used for this product.

The following numbering conventions are used:

- 1 kilobyte (KB) is equal to 1024 bytes.
- 1 megabyte (MB) is equal to 1 048 576 bytes.

© Copyright IBM Corp. 2025

- 1 gigabyte (GB) is equal to 1 073 741 824 bytes.
- 1 terabyte (TB) is equal to 1 099 511 627 776 bytes.
- 1 petabyte (PB) is equal to 1 125 899 906 842 624 bytes.

Note: The management GUI uses the same numerical values, but displays those values in the KiB, MiB, GiB, TiB, or PiB style of presentation.

