

IBM Storage Scale
6.0.1

*Concepts, Planning, and Installation
Guide*



Note

Before using this information and the product it supports, read the information in [“Notices” on page 775](#).

This edition applies to Version 6 release 0 modification 1 of the following products, and to all subsequent releases and modifications until otherwise indicated in new editions:

- IBM Storage Scale Data Management Edition ordered through Passport Advantage® (product number 5737-F34)
- IBM Storage Scale Data Access Edition ordered through Passport Advantage (product number 5737-I39)
- IBM Storage Scale Erasure Code Edition ordered through Passport Advantage (product number 5737-J34)
- IBM Storage Scale Data Management Edition ordered through AAS (product numbers 5641-DM1, DM3, DM5)
- IBM Storage Scale Data Access Edition ordered through AAS (product numbers 5641-DA1, DA3, DA5)
- IBM Storage Scale Data Management Edition for IBM® ESS (product number 5765-DME)
- IBM Storage Scale Data Access Edition for IBM ESS (product number 5765-DAE)
- IBM Storage Scale Backup ordered through Passport Advantage® (product number 5900-AXJ)
- IBM Storage Scale Backup ordered through AAS (product numbers 5641-BU1, BU3, BU5)
- IBM Storage Scale Backup for IBM® Storage Scale System (product number 5765-BU1)

Significant changes or additions to the text and illustrations are indicated by a vertical line (|) to the left of the change.

IBM welcomes your comments; see the topic [“How to send your comments” on page xxxviii](#). When you send information to IBM, you grant IBM a nonexclusive right to use or distribute the information in any way it believes appropriate without incurring any obligation to you.

© **Copyright International Business Machines Corporation 2015, 2026.**

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

Contents

Figures.....	xi
Tables.....	xv
About this information.....	xix
Prerequisite and related information.....	xxxvii
Conventions used in this information.....	xxxvii
How to send your comments.....	xxxviii
Summary of changes.....	xli
Chapter 1. Introducing IBM Storage Scale.....	1
Strengths of IBM Storage Scale.....	1
Shared file system access among IBM Storage Scale clusters.....	1
Improved system performance.....	2
File consistency.....	2
Increased data availability.....	2
Enhanced system flexibility.....	3
Simplified storage management.....	3
Simplified administration.....	4
Basic structure of IBM Storage Scale.....	4
IBM Storage Scale administration commands.....	5
The GPFS kernel extension.....	5
The GPFS daemon.....	5
The GPFS open source portability layer.....	6
IBM Storage Scale cluster configurations.....	6
Signed kernel modules for UEFI secure boot on x86_64 and secure boot Linux on Z.....	8
GPFS architecture.....	9
Special management functions.....	10
Use of disk storage and file structure within a GPFS file system.....	13
GPFS and memory.....	15
GPFS and network communication.....	17
Application and user interaction with GPFS.....	19
NSD disk discovery.....	24
Failure recovery processing.....	26
Cluster configuration data files.....	26
GPFS backup data.....	27
Cluster configuration repository.....	27
GPUDirect Storage support for IBM Storage Scale.....	28
Protocols support overview: Integration of protocol access methods with GPFS.....	30
Cluster Export Services overview.....	31
NFS support overview.....	33
SMB support overview.....	37
S3 support overview.....	38
Swift Object storage support overview.....	52
Active File Management.....	54
AFM concepts.....	55
Active File Management (AFM) features.....	70
AFM limitations.....	128

AFM-based asynchronous disaster recovery (AFM DR)	130
Recovery time objective (RTO).....	132
Modes and concepts.....	133
AFM-based Asynchronous Disaster Recovery features.....	133
Online creation and linking of a dependent fileset in AFM DR.....	141
AFM DR limitations.....	144
AFM DR deployment considerations and best practices.....	146
AFM to cloud object storage.....	154
AFM to cloud object storage operation modes.....	156
Multi-site replication.....	162
AFM to cloud object storage parallel read data transfer.....	163
Connectivity to cloud object storage.....	166
Eviction in AFM to cloud object storage.....	168
Data eviction from an AFM to cloud object storage fileset after uploading.....	168
Audit messages support for the AFM to cloud object storage.....	170
Partial file or object caching for AFM to cloud object storage.....	170
AFM to cloud object storage directory object support.....	174
AFM to cloud object storage support for more than 2 K metadata.....	175
AFM to cloud object storage policy-based upload for manual updates mode	176
Support of Google Cloud Platform for AFM to cloud object storage.....	179
Symbolic links.....	179
Replication by using the manual updates mode of the AFM to cloud object storage.....	179
Microsoft Azure Blob support for the AFM to cloud object storage.....	181
Asynchronous auto prefetch.....	181
Asynchronous notifications and Amazon SQS integration.....	182
Glacier storage class support for the AFM to cloud object storage.....	183
AFM to cloud object storage limitations.....	184
Introduction to IBM Storage Scale on AWS.....	186
AWS Services.....	186
Regions and Availability Zones.....	187
IBM Storage Scale instance types and operating systems.....	187
IBM Storage Scale usage restrictions.....	188
Introduction to system health and troubleshooting	188
Introduction to performance monitoring.....	190
Data protection and disaster recovery in IBM Storage Scale.....	191
Data backup options in IBM Storage Scale.....	191
Data restore options in IBM Storage Scale.....	192
Data mirroring in IBM Storage Scale.....	192
Protecting file data using snapshots	192
Introduction to Scale Out Backup and Restore (SOBAR).....	193
Commands for data protection and recovery in IBM Storage Scale.....	193
IBM Storage Scale GUI.....	194
IBM Storage Scale management API.....	197
Functional overview.....	198
API requests.....	199
API responses.....	204
Asynchronous jobs.....	206
Accessing the IBM Storage Scale REST API endpoint details through Swagger and API explorer.....	206
List of IBM Storage Scale management API commands.....	210
IBM Storage Scale native REST API.....	216
Non-root administration daemon.....	218
API nodes for IBM Storage Scale native REST API	219
Authentication for IBM Storage Scale native REST API.....	220
Role-based access control.....	223
Long-running operations.....	225
Target nodes field in IBM Storage Scale native REST API.....	227
Pagination in IBM Storage Scale native REST API.....	228
Views and fields options in IBM Storage Scale native REST API.....	230

Base64-encoded fields in IBM Storage Scale native REST API.....	232
Update masks in IBM Storage Scale native REST API.....	233
List of IBM Storage Scale native REST API endpoints.....	234
Limitations of IBM Storage Scale native REST API	238
Introduction to transparent cloud tiering (discontinued).....	240
How transparent cloud tiering works.....	242
How cloud data sharing works.....	243
How Write Once Read Many (WORM) storage works.....	246
Supported cloud providers.....	246
Interoperability of transparent cloud tiering with other IBM Storage Scale features.....	247
Interoperability of cloud data sharing with other IBM Storage Scale features.....	249
File audit logging.....	250
Producers in file audit logging.....	250
The file audit logging fileset.....	251
File audit logging records.....	251
File audit logging events.....	252
JSON attributes in file audit logging.....	252
Remotely mounted file systems in file audit logging.....	255
Clustered watch folder.....	256
Producers in clustered watch folder.....	256
Interaction between the clustered watch folder and the external Kafka sink.....	256
Clustered watch folder events.....	257
JSON attributes in clustered watch folder.....	258
Overview of call home	261
Types of call home data upload.....	263
Inspecting call home data uploads.....	289
Benefits of enabling call home.....	292
Data privacy with call home.....	293
Call home monitors for PTF updates.....	294
IBM Storage Scale in an OpenStack cloud deployment.....	295
IBM Storage Scale product editions.....	298
IBM Storage Scale license designation.....	300
Capacity-based licensing.....	304
Dynamic pagepool.....	304
Dynamic pagepool limitations.....	304
Shared memory communications direct on Linux on Z.....	305
Plugins for NVIDIA Nsight Systems.....	305
Data tiering overview.....	306

Chapter 2. Planning for IBM Storage Scale.....309

Planning for GPFS.....	309
Hardware requirements	309
Software requirements.....	310
Recoverability considerations.....	313
GPFS cluster creation considerations.....	319
Disk considerations.....	328
File system creation considerations.....	345
Backup considerations for using IBM Storage Protect.....	362
Planning for quality of service for I/O operations.....	374
Planning for extended attributes.....	376
Planning for systemd.....	378
Planning for protocols.....	380
Scaling considerations.....	380
Authentication considerations.....	381
Planning for NFS.....	389
Planning for SMB.....	393
SMB best practices.....	398

Fileset considerations for creating protocol data exports.....	403
Planning for S3.....	406
Planning for CES HDFS.....	407
Planning for cloud services	407
Hardware requirements for cloud services	407
Software requirements for cloud services	407
Network considerations for cloud services	408
Cluster node considerations for cloud services	409
IBM Cloud Object Storage considerations.....	410
Firewall recommendations for cloud services.....	412
Performance considerations.....	413
Security considerations.....	416
Planning for maintenance activities.....	416
Backup considerations for Transparent cloud tiering.....	417
Quota support for tiering.....	418
Client-assisted recalls.....	419
Planning for IBM Storage Scale on public clouds.....	419
Planning for deployment patterns on the cloud.....	420
Planning for resources creation.....	421
Preparing the installer node.....	421
Planning for AWS cloud.....	423
Planning for GCP cloud.....	429
Planning for Microsoft Azure cloud.....	434
Planning for IBM Cloud.....	441
Installing IBM Storage Scale in a network restricted (air gap) setup with the cloudkit.....	446
Other considerations.....	448
Configuring the IBM Storage Scale environment in the AWS Cloud.....	450
Planning for AFM.....	452
Requirements for UID and GID on the cache and home clusters.....	452
Recommended workerThreads on a cache cluster.....	452
Inode limits to set at cache and home.....	452
Planning for AFM gateway nodes.....	453
General recommendations for AFM gateway node configuration.....	453
General guidelines and recommendations for AFM SW and IW mode fileset.....	454
Planning for AFM DR.....	455
Requirements for UID/GID on primary and secondary clusters.....	455
Recommended worker1Threads on primary cluster.....	455
NFS setup on the secondary cluster.....	455
General guidelines and recommendations for AFM-DR.....	455
Planning for AFM to cloud object storage.....	456
Firewall recommendations for AFM to cloud object storage.....	457
Planning for IBM Storage Scale native REST API.....	458
Planning for performance monitoring tool.....	461
Performance monitoring limitations.....	462
Planning for UEFI secure boot on x86_64 and secure boot on Linux on Z.....	462
Firewall recommendations.....	463
Considerations for GPFS applications.....	463
Security-Enhanced Linux support.....	463
Space requirements for call home data upload.....	463
Shared memory communications direct planning.....	464
Planning for data tiering.....	464
Chapter 3. Establishing and starting your IBM Storage Scale cluster.....	467
Chapter 4. Installing IBM Storage Scale on Linux nodes and deploying protocols. 469	
Deciding whether to install IBM Storage Scale and deploy protocols manually or with the installation toolkit.....	469

Installation prerequisites.....	470
Preparing the environment on Linux nodes.....	472
IBM Storage Scale packaging overview.....	473
Preparing to install the IBM Storage Scale software on Linux nodes.....	473
Accepting the electronic license agreement on Linux nodes.....	473
Extracting the IBM Storage Scale software on Linux nodes.....	473
Verifying signature of IBM Storage Scale packages.....	475
Extracting IBM Storage Scale patches (update SLES and Red Hat Enterprise Linux RPMs or Ubuntu Linux packages).....	476
Installing the IBM Storage Scale man pages on Linux nodes.....	476
For Linux on Z: Changing the kernel settings.....	477
Manually installing the IBM Storage Scale software packages on Linux nodes.....	477
Building the GPFS portability layer on Linux nodes.....	483
Manually installing IBM Storage Scale and deploying protocols on Linux nodes.....	486
Verifying the IBM Storage Scale installation on Ubuntu Linux nodes.....	493
Verifying the IBM Storage Scale installation on SLES and Red Hat Enterprise Linux nodes.....	494
Manually installing the performance monitoring tool.....	495
Manually installing IBM Storage Scale management GUI.....	497
Manually installing IBM Storage Scale bridge for Grafana with the RPM.....	505
Installing IBM Storage Scale on Linux nodes with the installation toolkit.....	509
Overview of the installation toolkit.....	509
Understanding the installation toolkit options.....	516
Limitations of the installation toolkit.....	518
Mixed operating system support with the installation toolkit.....	526
Preparing to use the installation toolkit.....	528
Using the installation toolkit to perform installation tasks: Explanations and examples.....	532
Populating cluster definition file with current cluster state using the installation toolkit.....	551
Upgrading protocol nodes in an IBM Storage Scale System cluster.....	555
Configuration of an IBM Storage Scale stretch cluster in an export services environment: a sample use case.....	556
Performing additional tasks using the installation toolkit.....	567
Protocol node IP further configuration.....	573
Chapter 5. Installing IBM Storage Scale on AIX nodes.....	575
Creating a file to ease the AIX installation process.....	575
Verifying the level of prerequisite software.....	575
Procedure for installing GPFS on AIX nodes.....	576
Accepting the electronic license agreement.....	576
Creating the GPFS directory.....	576
Creating the GPFS installation table of contents file.....	577
Installing the GPFS man pages.....	577
Installing GPFS over a network.....	577
Verifying the GPFS installation.....	578
Chapter 6. Installing IBM Storage Scale on Windows nodes.....	579
GPFS for Windows overview.....	579
GPFS support for Windows.....	580
GPFS limitations on Windows	581
File system name considerations.....	582
File name considerations.....	583
Case sensitivity.....	583
Antivirus software.....	584
Differences between GPFS and NTFS.....	585
Access control on GPFS file systems.....	586
Installing GPFS prerequisites.....	587
Configuring Windows.....	588
Installing Cygwin.....	588

Procedure for installing GPFS on Windows nodes.....	590
Running GPFS commands.....	591
Configuring a mixed Windows and UNIX (AIX or Linux) cluster.....	592
Configuring the Windows HPC server.....	595
Chapter 7. Installing IBM Storage Scale on public cloud by using cloudkit.....	597
Overview of the cloudkit.....	597
Overview of the cloudkit installation options.....	598
Deploying cloudkit.....	601
Limitations of cloudkit.....	602
Supported features of cloudkit.....	603
Accessing the cluster on the cloud.....	608
Deploying IBM Storage Scale on AWS.....	608
Option 1: Deploying IBM Storage Scale on a new Amazon VPC with a single Availability Zone	612
Option 2: Deploying IBM Storage Scale on a new Amazon VPC with multiple Availability Zones ..	615
Option 3: Deploying IBM Storage Scale on an existing Amazon VPC.....	619
Deploying AFM on AWS.....	623
Chapter 8. Installing transparent cloud tiering (discontinued).....	627
Creating a user-defined node class for transparent cloud tiering or cloud data sharing.....	627
Installation steps.....	628
Setting up a cloud services cluster.....	629
Adding a cloud services node to an existing cloud services cluster.....	630
Chapter 9. Installing and configuring IBM Storage Scale management API.....	633
Chapter 10. Installing IBM Storage Scale native REST API.....	635
Manually installing IBM Storage Scale native REST API with the RPM.....	635
Installing IBM Storage Scale native REST API with the installation toolkit.....	635
Chapter 11. Installing GPUDirect Storage for IBM Storage Scale.....	639
Chapter 12. Installation of Active File Management (AFM).....	641
Chapter 13. Installing AFM Disaster Recovery.....	643
Chapter 14. Installing call home.....	645
Chapter 15. Installing file audit logging.....	647
Requirements, limitations, and support for file audit logging.....	647
Requirements for using file audit logging with remotely mounted file systems.....	648
Chapter 16. Installing clustered watch folder.....	651
Requirements, limitations, and support for clustered watch folder.....	651
Requirements for using clustered watch folder with remotely mounted file systems.....	652
Manually installing clustered watch folder.....	652
Chapter 17. Installing the signed kernel modules for UEFI secure boot on x86_64.....	655
Chapter 18. Installing the signed kernel modules for secure boot on Linux on Z...	657
Chapter 19. Uninstalling IBM Storage Scale.....	659
Cleanup procedures required if reinstalling with the installation toolkit.....	660
Uninstalling the performance monitoring tool.....	666

Uninstalling the IBM Storage Scale management GUI.....	666
Removing nodes from management GUI-related node class.....	667
Permanently uninstall cloud services and clean up the environment.....	668

Chapter 20. Upgrading.....671

IBM Storage Scale supported upgrade paths.....	672
Online upgrade support for protocols and performance monitoring.....	673
Upgrade of IBM Storage Scale nodes.....	674
Upgrading IBM Storage Scale nodes that run on AIX.....	674
Upgrading IBM Storage Scale nodes that run on Windows.....	676
Upgrading IBM Storage Scale non-protocol nodes that run on Linux.....	677
Upgrading IBM Storage Scale protocol nodes.....	682
Upgrading IBM Storage Scale on cloud.....	688
Cluster upgrade.....	688
Upgrading IBM Storage Scale.....	690
Upgrading GPUDirect Storage.....	692
Upgrading AFM and AFM DR.....	692
Upgrading an AFM-DR system.....	694
Restarting replication to upgrade AFM and AFM DR.....	694
Upgrading object packages.....	695
Upgrading SMB packages.....	697
Upgrading the SMB package after upgrading OS.....	699
Upgrading NFS packages.....	699
Upgrading S3 packages.....	700
Upgrading call home.....	702
Call home configuration changes to be made while upgrading to IBM Storage Scale 5.0.x from IBM Storage Scale 4.2.1	702
Call home configuration changes to be made while upgrading to IBM Storage Scale 4.2.1 from IBM Storage Scale 4.2.0	703
Removing residual configuration files while upgrading to IBM Storage Scale 5.1.x.....	703
Upgrading the performance monitoring tool.....	704
Manually upgrading the performance monitoring tool.....	704
Upgrading signed kernel modules for UEFI secure boot on x86_64 and Linux on Z.....	705
Manually upgrading pmswift.....	706
Manually upgrading the IBM Storage Scale management GUI.....	708
Upgrading transparent cloud tiering (discontinued).....	710
Upgrading to Transparent cloud tiering 1.1.2 from Transparent cloud tiering 1.1.0 or 1.1.1	710
Upgrading to cloud services 1.1.2.1 from 1.1.2	711
Upgrading to cloud services 1.1.3 from 1.1.2	712
Upgrading to cloud services 1.1.3 from 1.1.0 and 1.1.1	712
Upgrading to cloud services 1.1.4 from 1.1.2.x	713
Upgrading to cloud services 1.1.4 from 1.1.3	714
Upgrading to cloud services 1.1.5 from 1.1.4	715
Upgrading to cloud services 1.1.6 from 1.1.5	717
Upgrading the cloud services sensors	718
Upgrading to IBM Cloud Object Storage software level 3.7.2 and above.....	722
Upgrade paths and commands for file audit logging and clustered watch folder.....	722
Upgrading IBM Storage Scale components with the installation toolkit.....	723
Upgrade process flow.....	725
Performing online upgrade by using the installation toolkit.....	731
Performing offline upgrade or excluding nodes from upgrade by using installation toolkit.....	733
Upgrade rerun after an upgrade failure.....	741
Upgrading Swift Object packages to IBM Storage Scale 5.1.8 or earlier by using the installation toolkit.....	741
Upgrading IBM Storage Scale on IBM Storage Archive Enterprise Edition (EE) nodes by using the installation toolkit.....	742
Upgrading IBM Storage Scale native REST API.....	744

Protocol authentication configuration changes during upgrade.....	744
Changing the IBM Storage Scale product edition.....	747
Changing the product edition by using the installation toolkit.....	748
Changing the product edition manually, one node at a time.....	749
Changing the product edition manually, all nodes at the same time.....	750
Changing Express Edition to Standard Edition.....	752
Completing the upgrade to a new level of IBM Storage Scale.....	753
Reverting to the previous level of IBM Storage Scale.....	758
Reverting to a previous level of GPFS when you have not issued mmchconfig release=LATEST....	759
Reverting to a previous level of GPFS when you <i>have</i> issued mmchconfig release=LATEST.....	760
Coexistence considerations.....	761
Compatibility considerations.....	761
Considerations for IBM Storage Protect for Space Management.....	761
Applying maintenance to your IBM Storage Scale system.....	761
Guidance for upgrading the operating system on IBM Storage Scale nodes.....	762
Guidance for Red Hat Enterprise Linux 8.x on IBM Storage Scale nodes.....	764
Instructions for removing object protocol packages when upgrading protocol nodes to Red Hat Enterprise Linux 8.x.....	767
Servicing IBM Storage Scale protocol nodes.....	768
Offline upgrade with complete cluster shutdown.....	769
Accessibility features for IBM Storage Scale.....	773
Accessibility features.....	773
Keyboard navigation.....	773
IBM and accessibility.....	773
Notices.....	775
Trademarks.....	776
Terms and conditions for product documentation.....	776
Glossary.....	779
Index.....	787

Figures

1. A cluster with disks that are SAN-attached to all nodes.....	7
2. A cluster with some nodes connected to disks.....	7
3. A multicluster configuration.....	8
4. GPFS files have a typical UNIX structure.....	14
5. IBM Storage Scale CES S3 architecture.....	41
6. IBM Storage Scale CES S3 control path IBM Storage Scale CES S3 control path	42
7. An example of IBM Storage Scale CES S3 data path.....	43
8. Sample of an AFM relationship	55
9. Global namespace implemented by using AFM.....	62
10. Read only mode.....	65
11. Single writer mode.....	65
12. Behaviors with local files.....	66
13. Independent writer mode.....	68
14. Sample setup of IBM Storage Protect for Space Management connected to home	109
15. IBM Storage Protect for Space Management connected to both home and cache.....	110
16. IW cache (Side A) to home (Side B).....	111
17. IW cache site (Side A) and home site (Side B).....	112
18. Asynchronous disaster recovery.....	131
19. AFM to cloud object storage.....	155
20. AFM to cloud object storage by using Azure Blob.....	156
21. IBM Storage Scale management API architecture.....	197
22. Option to download API description files in JSON format from the GUI node.....	207
23. REST API documentation opened in Swagger editor.....	207

24. Example for options that are available in the Swagger editor for each endpoint.....	209
25. High-level architecture.....	217
26. Default scalectl flow.....	218
27. scalectl flow with --url option.....	218
28. High-level architecture of a cluster with API nodes.....	220
29. transparent cloud tiering and cloud data sharing features.....	241
30. Exporting file system data to a cloud storage tier.....	244
31. Importing object storage data into the file system.....	245
32. WORM storage overview.....	246
33. Call home architecture.....	261
34. Guidance on which license to buy.....	302
35. GPFS configuration using node quorum	314
36. GPFS configuration using node quorum with tiebreaker disks.....	315
37. An example of a highly available SAN configuration for a GPFS file system.....	317
38. Configuration using GPFS replication for improved availability.....	317
39. High-level overview of protocol user authentication.....	383
40. High-level flow of authentication for File protocols.....	385
41. IBM Storage Scale for NFS architecture.....	390
42. Cloud deployment modes.....	420
43. Single availability zone architecture diagram.....	450
44. Multi availability zone architecture diagram.....	451
45. Synchronous mirroring by using GPFS replication	452
46. Synchronous mirroring with GPFS replication.....	557
47. Local Site CES Group.....	559
48. High-level IBM Storage Scale cluster architecture to connect from host.....	610

49. Example of an on-premises to AWS hybrid cloud networking architecture with AFM.....	623

Tables

1. IBM Storage Scale library information units.....	xx
2. Conventions.....	xxxviii
3. Supported devices on AIX.....	24
4. Supported devices on Linux.....	25
5. Supported devices on Windows.....	25
6. IAM user management operations.....	48
7. Comparison between NSD and NFS protocols.....	57
8. Conditions in which disabling AFM fileset fails, with corresponding messages	115
9. Supported function and supported OS/architecture.....	129
10. Supported function and supported OS/architecture.....	144
11. Parameters for the replication of a new file system by using the AFM to cloud object storage manual updates mode.....	180
12. Features associated with IBM Storage Scale GUI pages.....	194
13. Getting information about filesets.....	200
14. Creating a snapshot.....	200
15. Response codes.....	204
16. Operations supported for resources and resource elements in API endpoints	210
17. Operations supported for resources and resource elements in API endpoints	234
18. File audit logging events.....	252
19. JSON attributes in file audit logging.....	254
20. File access events that are supported by clustered watch folder.....	257
21. RAS events that trigger data collection.....	264
22. Data collected and uploaded by call home.....	270

23. Features that IBM Storage Scale supports for deploying the OpenStack cloud storage.....	297
24. Features in IBM Storage Scale editions.....	299
25. GPFS cluster creation options.....	319
26. Thin provisioning storage solutions in IBM Storage Scale.....	332
27. Supported operating systems and device connections.....	337
28. Supported devices.....	338
29. File system creation options.....	345
30. Correspondence between mount options and the -S option in mmcrfs and mmchfs.....	353
31. Comparison of mmbackup and IBM Storage Protect Backup-Archive client backup commands.....	365
32. General authentication support matrix.....	381
33. Tested NFS clients.....	391
34. Coherency option and description.....	401
35. OS support matrix.....	408
36. Recommended settings when you create a vault template on IBM Cloud Object Storage.....	411
37. Port requirements.....	412
38. workerThreads value per backend.....	452
39. Monitoring data consumption limits by pmcollector memory	461
40. GUI packages essential for each platform.....	498
41. Installation toolkit: List of features.....	511
42. spectrumscale command options for installing IBM Storage Scale and deploying protocols.....	517
43. Installation toolkit limitations.....	518
44. Operating systems supported with the installation toolkit in a mixed cluster.....	527
45. Validations by the installation toolkit in a mixed operating system cluster.....	527
46. Limitations of the config populate functionality.....	552
47. Generating short names for Windows.....	583

48. cloudkit command options.....	600
49. Supported public clouds.....	603
50. Operating systems that the cloudkit installer nodes support.....	604
51. IBM Storage Scale features.....	605
52. Feature matrix for cloudkit.....	606
53. File system configurations.....	612
54. NSD Configurations.....	612
55. Server Node Configurations.....	613
56. Compute Node Configurations.....	613
57. Network Configuration.....	614
58. Amazon EC2 Configuration.....	614
59. Personal Configuration.....	615
60. License Information.....	615
61. File System Configurations.....	615
62. NSD Configurations.....	615
63. Server Node Configurations.....	616
64. Compute Node Configurations.....	617
65. Network Configuration.....	617
66. Amazon EC2 Configuration.....	618
67. Personal Configuration.....	618
68. License Information.....	618
69. File System Configurations.....	619
70. NSD Configurations.....	619
71. Server Node Configurations.....	620
72. Compute Node Configurations.....	621

73. Network Configuration.....	621
74. Amazon EC2 Configuration.....	622
75. Personal Configuration.....	622
76. License Information.....	622
77. IBM Storage Scale supported online upgrade paths.....	672
78. Upgrade matrix.....	688
79. Supported cloud platforms.....	688
80. Other options and parameters with mmchconfig release=LATEST.....	754

About this information

This edition applies to IBM Storage Scale version 6.0.1 for AIX®, Linux®, and Windows.

IBM Storage Scale is a file management infrastructure, based on IBM General Parallel File System (GPFS) technology, which provides unmatched performance and reliability with scalable access to critical file data.

To find out which version of IBM Storage Scale is running on a particular AIX node, enter:

```
lslpp -l gpfs\*
```

To find out which version of IBM Storage Scale is running on a particular Linux node, enter:

```
rpm -qa | grep gpfs (for SLES and Red Hat Enterprise Linux)
```

```
dpkg -l | grep gpfs (for Ubuntu Linux)
```

To find out which version of IBM Storage Scale is running on a particular Windows node, open **Programs and Features** in the control panel. The IBM Storage Scale installed program name includes the version number.

Which IBM Storage Scale information unit provides the information you need?

The IBM Storage Scale library consists of the information units listed in [Table 1 on page xx](#).

To use these information units effectively, you must be familiar with IBM Storage Scale and the AIX, Linux, or Windows operating system, or all of them, depending on which operating systems are in use at your installation. Where necessary, these information units provide some background information relating to AIX, Linux, or Windows. However, more commonly they refer to the appropriate operating system documentation.

Note: Throughout this documentation, the term "Linux" refers to all supported distributions of Linux, unless otherwise specified.

Table 1. IBM Storage Scale library information units

Information unit	Type of information	Intended users
<i>IBM Storage Scale: Concepts, Planning, and Installation Guide</i>	This guide provides the following information: Product overview • Strengths of IBM Storage Scale • Basic structure of IBM Storage Scale • IBM Storage Scale cluster configurations • Signed kernel modules for UEFI secure boot nodes • GPFS architecture • Protocols support overview: Integration of protocol access methods with GPFS • Active File Management • AFM-based Asynchronous Disaster Recovery (AFM DR) • Introduction to AFM to cloud object storage • Introduction to IBM Storage Scale on AWS • Introduction to system health and troubleshooting • Introduction to performance monitoring • Data protection and disaster recovery in IBM Storage Scale • Introduction to IBM Storage Scale GUI • IBM Storage Scale native REST API • IBM Storage Scale management API • Introduction to transparent cloud tiering (discontinued) • Introduction to file audit logging • Introduction to clustered watch folder • Overview of call home • IBM Storage Scale in an OpenStack cloud deployment • IBM Storage Scale product editions	System administrators, analysts, installers, planners, and programmers of IBM Storage Scale clusters who are very experienced with the operating systems on which each IBM Storage Scale cluster is based

Table 1. IBM Storage Scale library information units (continued)

Information unit	Type of information	Intended users
<i>IBM Storage Scale: Concepts, Planning, and Installation Guide</i>	• IBM Storage Scale license designation • Capacity-based licensing • Dynamic pagepool • Shared memory communications direct on Linux on Z • Plugins for NVIDIA Nsight Systems • Data tiering overview Planning • Planning for GPFS • Planning for protocols • Planning for cloud services • Planning for IBM Storage Scale on public clouds • Planning for AFM • Planning for AFM DR • Planning for AFM to cloud object storage • Planning for IBM Storage Scale native REST API • Planning for performance monitoring tool • Planning for UEFI secure boot nodes • Firewall recommendations • Considerations for GPFS applications • Security-Enhanced Linux support • Space requirements for call home data upload • Shared memory communications direct planning • Planning for data tiering	

Table 1. IBM Storage Scale library information units (continued)

Information unit	Type of information	Intended users
<i>IBM Storage Scale: Concepts, Planning, and Installation Guide</i>	• IBM Storage Scale license designation • Capacity-based licensing • Dynamic pagepool • Shared memory communications direct on Linux on Z • Plugins for NVIDIA Nsight Systems • Data tiering overview Planning • Planning for GPFS • Planning for protocols • Planning for cloud services • Planning for IBM Storage Scale on public clouds • Planning for AFM • Planning for AFM DR • Planning for AFM to cloud object storage • Planning for IBM Storage Scale native REST API • Planning for performance monitoring tool • Planning for UEFI secure boot nodes • Firewall recommendations • Considerations for GPFS applications • Security-Enhanced Linux support • Space requirements for call home data upload • Shared memory communications direct planning • Planning for data tiering	

Table 1. IBM Storage Scale library information units (continued)

Information unit	Type of information	Intended users
<i>IBM Storage Scale: Concepts, Planning, and Installation Guide</i>	Installing • Establishing and starting your IBM Storage Scale cluster • Installing IBM Storage Scale on Linux nodes and deploying protocols • Installing IBM Storage Scale on public cloud by using cloudkit • Installing IBM Storage Scale on AIX nodes • Installing IBM Storage Scale on Windows nodes • Installing transparent cloud tiering (discontinued) • Installing and configuring IBM Storage Scale management API • Installing IBM Storage Scale native REST API • Installing GPUDirect storage for IBM Storage Scale • Installation of Active File Management (AFM) • Installing AFM Disaster Recovery • Installing call home • Installing file audit logging • Installing clustered watch folder • Installing the signed kernel modules for UEFI secure boot nodes • Installing the signed kernel modules for secure boot nodes • Uninstalling IBM Storage Scale	System administrators, analysts, installers, planners, and programmers of IBM Storage Scale clusters who are very experienced with the operating systems on which each IBM Storage Scale cluster is based

Table 1. IBM Storage Scale library information units (continued)

Information unit	Type of information	Intended users
<i>IBM Storage Scale: Concepts, Planning, and Installation Guide</i>	Upgrading • IBM Storage Scale supported upgrade paths • Online upgrade support for protocols and performance monitoring • Upgrading IBM Storage Scale nodes • Upgrading IBM Storage Scale non-protocol Linux nodes • Upgrading IBM Storage Scale protocol nodes • Upgrading IBM Storage Scale on cloud • Upgrading GPUDirect Storage • Upgrading AFM and AFM DR • Upgrading object packages • Upgrading SMB packages • Upgrading NFS packages • Upgrading S3 packages • Upgrading call home • Upgrading the performance monitoring tool • Upgrading signed kernel modules for UEFI secure boot nodes on Linux on Z • Manually upgrading pmswift • Manually upgrading the IBM Storage Scale management GUI • Upgrading transparent cloud tiering (discontinued) • Upgrading to IBM Cloud Object Storage software level 3.7.2 and above • Upgrade paths and commands for file audit logging and clustered watch folder • Upgrading IBM Storage Scale components with the installation toolkit • Protocol authentication configuration changes during upgrade • Changing the IBM Storage Scale product edition	System administrators, analysts, installers, planners, and programmers of IBM Storage Scale clusters who are very experienced with the operating systems on which each IBM Storage Scale cluster is based

Table 1. IBM Storage Scale library information units (continued)

Information unit	Type of information	Intended users
<i>IBM Storage Scale: Concepts, Planning, and Installation Guide</i>	• Completing the upgrade to a new level of IBM Storage Scale • Reverting to the previous level of IBM Storage Scale • Coexistence considerations • Compatibility considerations • Considerations for IBM Storage Protect for Space Management • Applying maintenance to your IBM Storage Scale system • Guidance for upgrading the operating system on IBM Storage Scale nodes • Considerations for upgrading from an operating system not supported in IBM Storage Scale 5.1.x.x • Servicing IBM Storage Scale protocol nodes • Offline upgrade with complete cluster shutdown	

Table 1. IBM Storage Scale library information units (continued)

Information unit	Type of information	Intended users
<i>IBM Storage Scale: Administration Guide</i>	This guide provides the following information: Configuring • Configuring the GPFS cluster • Configuring GPUDirect Storage for IBM Storage Scale • Configuring the CES and protocols • Configuring and tuning your system for GPFS • Parameters for performance tuning and optimization • Ensuring high availability of the GUI service • Configuring and tuning your system for cloud services • Configuring IBM Power Systems for IBM Storage Scale • Configuring file audit logging • Configuring clustered watch folder • Configuring the cloudkit • Configuring Active File Management • Configuring AFM-based DR • Configuring AFM to cloud object storage • Configuring IBM Storage Scale native REST API • Tuning for Kernel NFS backend on AFM and AFM DR • Configuring call home • Integrating IBM Storage Scale Cinder driver with Red Hat OpenStack Platform 16.1 • Configuring Multi-Rail over TCP (MROT) • Dynamic pagepool configuration • Configuring shared memory communications direct • Configuring fileset multitenancy by using the GUI • Configuring the plugins for NVIDIA Nsight Systems • Configuring data tiering	System administrators or programmers of IBM Storage Scale systems

Table 1. IBM Storage Scale library information units (continued)

Information unit	Type of information	Intended users
IBM Storage Scale: Administration Guide	Administering • Performing GPFS administration tasks • Performing parallel copy with mmxcp command • Managing shared memory communications direct • Protecting file data: IBM Storage Scale safeguarded copy • Verifying network operation with the mmnetverify command • Managing file systems • File system format changes between versions of IBM Storage Scale • Managing disks • Managing protocol services • Managing protocol user authentication • Managing protocol data exports • Managing S3 protocol • Managing Swift Object storage	System administrators or programmers of IBM Storage Scale systems

Table 1. IBM Storage Scale library information units (continued)

Information unit	Type of information	Intended users
<i>IBM Storage Scale: Administration Guide</i>	• Managing GPFS quotas • Managing GUI users • Managing GPFS access control lists • Native NFS and GPFS • Accessing a remote GPFS file system • Information lifecycle management for IBM Storage Scale • Creating and maintaining snapshots of file systems • Creating and managing file clones • Scale Out Backup and Restore (SOBAR) • Data Mirroring and Replication • Implementing a clustered NFS environment on Linux • Implementing Cluster Export Services • Identity management on Windows / RFC 2307 attributes • Protocols cluster disaster recovery • File Placement Optimizer • Encryption • Managing certificates to secure communications between GUI web server and web browsers • Securing protocol data • Cloud services: Transparent cloud tiering and Cloud data sharing • Managing file audit logging • RDMA tuning • Configuring Mellanox Memory Translation Table (MTT) for GPFS RDMA VERBS Operation • Administering cloudkit • Administering AFM • Administering AFM through GUI • Administering AFM DR	System administrators or programmers of IBM Storage Scale systems

Table 1. IBM Storage Scale library information units (continued)

Information unit	Type of information	Intended users
<i>IBM Storage Scale: Administration Guide</i>	• Administering AFM to cloud object storage • Migrating existing IBM Storage Scale clusters management from the legacy mm CLI to the IBM Storage Scale native REST API • Remote cluster support with IBM Storage Scale native REST API • Highly available write cache (HAWC) • Local read-only cache • Miscellaneous advanced administration topics • GUI limitations	System administrators or programmers of IBM Storage Scale systems
<i>IBM Storage Scale: Problem Determination Guide</i>	This guide provides the following information: Monitoring • Monitoring system health by using IBM Storage Scale GUI • Monitoring system health by using the mmhealth command • Dynamic page pool monitoring • Performance monitoring • Monitoring GPUDirect storage • Monitoring events through callbacks • Monitoring capacity through GUI • Monitoring AFM and AFM DR • Monitoring profiling sessions with the plugins for NVIDIA Nsight Systems • Monitoring AFM to cloud object storage • Monitoring IBM Storage Scale native REST API • GPFS SNMP support • Monitoring the IBM Storage Scale system by using call home • Monitoring remote cluster through GUI • Monitoring file audit logging • Monitoring clustered watch folder • Monitoring local read-only cache Troubleshooting • Best practices for troubleshooting • Understanding the system limitations • Collecting details of the issues • Managing deadlocks • Installation and configuration issues	System administrators of GPFS systems who are experienced with the subsystems used to manage disks and who are familiar with the concepts presented in the <i>IBM Storage Scale: Concepts, Planning, and Installation Guide</i>

Table 1. IBM Storage Scale library information units (continued)

Information unit	Type of information	Intended users
<i>IBM Storage Scale: Problem Determination Guide</i>	• Upgrade issues • CCR issues • Network issues • File system issues • Disk issues • GPUDirect Storage troubleshooting • Security issues • Protocol issues • Disaster recovery issues • Performance issues • GUI and monitoring issues • AFM issues • AFM DR issues • AFM to cloud object storage issues • Troubleshooting IBM Storage Scale native REST API • Transparent cloud tiering issues • File audit logging issues • Cloudkit issues • Troubleshooting mmwatch • Troubleshooting plugin issues for NVIDIA Nsight Systems • Troubleshooting shared memory communications direct issues • Maintenance procedures • Recovery procedures • Support for troubleshooting • References	

Table 1. IBM Storage Scale library information units (continued)

Information unit	Type of information	Intended users
<i>IBM Storage Scale: Command and Programming Reference Guide</i>	This guide provides the following information: Command reference • Command exceptions for AIX and Windows • cloudkit command • gpfs.snap command • mmaddcallback command • mmaddddisk command • mmaddnode command • mmadquery command • mmafmconfig command • mmafmcosaccess command • mmafmcosconfig command • mmafmcosctl command • mmafmcoskeys command • mmafmctl command • mmafmlocal command • mmapplypolicy command • mmaudit command • mmauth command • mmbackup command • mmbackupconfig command • mmbuildgpl command • mmcachectl command • mmcallhome command • mmces command • mmchattr command • mmchcluster command • mmchconfig command • mmchdisk command • mmcheckquota command • mmchfileset command • mmchfs command • mmchlicense command • mmchmgr command • mmchnode command • mmchnodeclass command • mmchnsd command • mmchpolicy command • mmchpool command	• System administrators of IBM Storage Scale systems • Application programmers who are experienced with IBM Storage Scale systems and familiar with the terminology and concepts in the XD SM standard

Table 1. IBM Storage Scale library information units (continued)

Information unit	Type of information	Intended users
<i>IBM Storage Scale: Command and Programming Reference Guide</i>	• mmchqos command • mmclidecode command • mmclone command • mmcloudgateway command • mmcloudworkflows utility • mmcrcluster command • mmcrfileset command • mmcrfs command • mmcrnodeclass command • mmcrnsd command • mmcrsnapshot command • mmdefedquota command • mmdefquotaoff command • mmdefquotaon command • mmdefragfs command • mmdelacl command • mmdelcallback command • mmdeldisk command • mmdelfileset command • mmdelfs command • mmdelnodde command • mmdelnoddeclass command • mmdelnodde command • mmdelsnapshot command • mmdf command • mmdiag command • mmdsh command • mmeditaccl command • mmedquota command • mmexportfs command • mmfsck command • mmfsckx command • mmfsctl command • mmgetacl command • mmgetstate command • mmhadoopctl command • mmhdfs command • mmhealth command • mmimgbackup command	• System administrators of IBM Storage Scale systems • Application programmers who are experienced with IBM Storage Scale systems and familiar with the terminology and concepts in the XDSM standard

Table 1. IBM Storage Scale library information units (continued)

Information unit	Type of information	Intended users
<i>IBM Storage Scale: Command and Programming Reference Guide</i>	• mmimgrestore command • mmimportfs command • mmkeyserv command • mmlinkfileset command • mmlsattr command • mmlscallback command • mmlscluster command • mmlsconfig command • mmlsdisk command • mmlsfileset command • mmlsfs command • mmlslicense command • mmlsmgr command • mmlsmount command • mmlsnodeclass command • mmlsnd command • mmlspolicy command • mmlspool command • mmlsqos command • mmlsquota command • mmlssnapshot command • mmmigratefs command • mmmount command • mmnetverify command • mmnfs command • mmnsddiscover command • mmobj command • mmperfmon command • mmpmon command • mmprotocoltrace command • mmpsnap command • mmpstat command • mmptop command • mmputacl command • mmqos command • mmquotaoff command • mmquotaon command	• System administrators of IBM Storage Scale systems • Application programmers who are experienced with IBM Storage Scale systems and familiar with the terminology and concepts in the XD SM standard

Table 1. IBM Storage Scale library information units (continued)

Information unit	Type of information	Intended users
<i>IBM Storage Scale: Command and Programming Reference Guide</i>	• mmreclaimspace command • mmremoteccluster command • mmremotefs command • mmrepquota command • mmrestoreconfig command • mmrestorefs command • mmrestrictedctl command • mmrestripefile command • mmrestripefs command • mmrpldisk command • mmsdrrestore command • mmsetquota command • mmshutdown command • mmsmb command • mmsnapdir command • mmstartup command • mmstartpolicy command • mms3 command • mmtracectl command • mmumount command • mmunlinkfileset command • mmuserauth command • mmwatch command • mmwinservctl command • mmxcp command • scalectl command • spectrumscale command Programming reference • IBM Storage Scale Data Management API for GPFS information • GPFS programming interfaces • GPFS user exits • IBM Storage Scale management API endpoints • IBM Storage Scale native REST API endpoints • Considerations for GPFS applications	• System administrators of IBM Storage Scale systems • Application programmers who are experienced with IBM Storage Scale systems and familiar with the terminology and concepts in the XD SM standard

Table 1. IBM Storage Scale library information units (continued)

Information unit	Type of information	Intended users
IBM Storage Scale: Big Data and Analytics Guide	This guide provides the following information: Big data and analytics support Summary of changes IBM Storage Scale support for Hadoop • Overview • Planning • Installing • Upgrading • Configuring • Administering • Troubleshooting IBM Storage Scale Hadoop performance tuning guide • Overview • Performance overview • Hadoop performance planning over IBM Storage Scale • Performance guide Cloudera Data Platform (CDP) Private Cloud Base • Overview • Planning • Installing • Configuring • Administering • Monitoring • Upgrading • Limitations • Problem determination	• System administrators of IBM Storage Scale systems • Application programmers who are experienced with IBM Storage Scale systems and familiar with the terminology and concepts in the XD SM standard
IBM Storage Scale: Big Data and Analytics Guide	Cloudera HDP 3.X Note: Starting with IBM Storage Scale 5.1.1.2, Cloudera Hortonworks Data Platform (HDP) is no longer supported. <i>This content is kept in the documentation set merely for reference.</i> • Planning • Installation • Upgrading and uninstallation • Configuration • Administration • Limitations • Problem determination	• System administrators of IBM Storage Scale systems • Application programmers who are experienced with IBM Storage Scale systems and familiar with the terminology and concepts in the XD SM standard

Table 1. IBM Storage Scale library information units (continued)

Information unit	Type of information	Intended users
<i>IBM Storage Scale: Big Data and Analytics Guide</i>	Apache Hadoop • Apache Hadoop 3.0.x support • Enabling Kerberos with Apache Hadoop and CES HDFS • HDFS clients configuration • MapReduce/YARN clients configuration • Add HDFS to CES HDFS nodes	• System administrators of IBM Storage Scale systems • Application programmers who are experienced with IBM Storage Scale systems and familiar with the terminology and concepts in the XD SM standard
<i>IBM Storage Scale Erasure Code Edition Guide</i>	IBM Storage Scale Erasure Code Edition • Summary of changes • Introduction to IBM Storage Scale Erasure Code Edition • Planning for IBM Storage Scale Erasure Code Edition • Installing IBM Storage Scale Erasure Code Edition • Uninstalling IBM Storage Scale Erasure Code Edition • Creating an IBM Storage Scale Erasure Code Edition storage environment • Using IBM Storage Scale Erasure Code Edition for data mirroring and replication • Deploying IBM Storage Scale Erasure Code Edition on VMware infrastructure • Upgrading IBM Storage Scale Erasure Code Edition • Incorporating IBM Storage Scale Erasure Code Edition in an Elastic Storage Server (ESS) cluster • Incorporating IBM Elastic Storage Server (ESS) building block in an IBM Storage Scale Erasure Code Edition cluster • Administering IBM Storage Scale Erasure Code Edition • Troubleshooting • IBM Spectrum® Scale RAID Administration	• System administrators of IBM Storage Scale systems • Application programmers who are experienced with IBM Storage Scale systems and familiar with the terminology and concepts in the XD SM standard

Table 1. IBM Storage Scale library information units (continued)

Information unit	Type of information	Intended users
IBM Storage Scale Container Native Storage Access	This guide provides the following information: • Overview • Planning • Installation prerequisites • Installing the IBM Storage Scale container native operator and cluster • Upgrading • Configuring IBM Storage Scale Container Storage Interface (CSI) driver • Using IBM Storage Scale GUI • Maintenance of a deployed cluster • Cleaning up the container native cluster • Monitoring • Troubleshooting • References	• System administrators of IBM Storage Scale systems • Application programmers who are experienced with IBM Storage Scale systems and familiar with the terminology and concepts in the XD SM standard
IBM Storage Scale Container Storage Interface Driver Guide	This guide provides the following information: • Summary of changes • Introduction • Planning • Installation • Upgrading • Configurations • Using IBM Storage Scale Container Storage Interface Driver • Managing IBM Storage Scale when used with IBM Storage Scale Container Storage Interface driver • Cleanup • Limitations • Troubleshooting	• System administrators of IBM Storage Scale systems • Application programmers who are experienced with IBM Storage Scale systems and familiar with the terminology and concepts in the XD SM standard

Prerequisite and related information

For updates to this information, see [IBM Storage Scale in IBM Documentation](#).

For the latest support information, see the [IBM Storage Scale FAQ in IBM Documentation](#).

Conventions used in this information

Table 2 on page xxxviii describes the typographic conventions used in this information. UNIX file name conventions are used throughout this information.

Note: Users of IBM Storage Scale for Windows must be aware that on Windows, UNIX-style file names need to be converted appropriately. For example, the GPFS cluster configuration data is stored in the `/var/mmfs/gen/mmsdrfs` file. On Windows, the UNIX namespace starts under

the %SystemDrive%\cygwin64 directory, so the GPFS cluster configuration data is stored in the C:\cygwin64\var\mmfs\gen\mmsdrfs file.

Table 2. Conventions

Convention	Usage
bold	Bold words or characters represent system elements that you must use literally, such as commands, flags, values, and selected menu options. Depending on the context, bold typeface sometimes represents path names, directories, or file names.
<u>bold underlined</u>	<u>bold underlined</u> keywords are defaults. These take effect if you do not specify a different keyword.
constant width	Examples and information that the system displays appear in constant-width typeface. Depending on the context, constant-width typeface sometimes represents path names, directories, or file names.
<i>italic</i>	<i>Italic</i> words or characters represent variable values that you must supply. <i>Italics</i> are also used for information unit titles, for the first use of a glossary term, and for general emphasis in text.
<key>	Angle brackets (less-than and greater-than) enclose the name of a key on the keyboard. For example, <Enter> refers to the key on your terminal or workstation that is labeled with the word <i>Enter</i> .
\	In command examples, a backslash indicates that the command or coding example continues on the next line. For example: <pre>mkcondition -r IBM.FileSystem -e "PercentTotUsed > 90" \ -E "PercentTotUsed < 85" -m p "FileSystem space used"</pre>
{item}	Braces enclose a list from which you must choose an item in format and syntax descriptions.
[item]	Brackets enclose optional items in format and syntax descriptions.
<Ctrl-x>	The notation <Ctrl-x> indicates a control character sequence. For example, <Ctrl-c> means that you hold down the control key while pressing <c>.
item...	Ellipses indicate that you can repeat the preceding item one or more times.
	In <i>synopsis</i> statements, vertical lines separate a list of choices. In other words, a vertical line means <i>Or</i> . In the left margin of the document, vertical lines indicate technical changes to the information.

Note: CLI options that accept a list of option values delimit with a comma and no space between values. As an example, to display the state on three nodes use `mmgetstate -N NodeA,NodeB,NodeC`. Exceptions to this syntax are listed specifically within the command.

How to send your comments

Your feedback is important in helping us to produce accurate, high-quality information. If you have any comments about this information or any other IBM Storage Scale documentation, send your comments to the following e-mail address:

mhvrcfs@us.ibm.com

Include the publication title and order number, and, if applicable, the specific location of the information about which you have comments (for example, a page number or a table number).

To contact the IBM Storage Scale development organization, send your comments to the following e-mail address:

`scale@us.ibm.com`

Summary of changes

Consult the changes made to the IBM Storage Scale licensed program and the IBM Storage Scale library. Within each information unit in the library, a vertical line (|) to the left of text and illustrations indicates technical changes or additions that are made to the previous edition of the information.

Summary of changes for IBM Storage Scale 6.0.1 as updated, June 2026

This release of the IBM Storage Scale licensed program and the IBM Storage Scale library includes the following improvements. All improvements are available after an upgrade, unless otherwise specified.

- [Commands, data types, and programming APIs](#)
- [Stabilized, deprecated, and discontinued features](#)

IBM Storage Scale Big Data and Analytics changes

For information on changes in IBM Storage Scale Big Data and Analytics support and HDFS protocol, see *Summary of changes* in the [IBM Storage Scale Big Data and Analytics](#) documentation.

IBM Storage Scale Erasure Code Edition changes

For information on changes in IBM Storage Scale Erasure Code Edition, see *Summary of changes* in the [IBM Storage Scale Erasure Code Edition](#) documentation.

IBM Storage Scale Container Native Storage

For information on changes in IBM Storage Scale Container Native Storage, see *What's new?* in the [IBM Storage Scale Container Native](#) documentation.

IBM Storage Scale Container Storage Interface driver

For information on changes in IBM Storage Scale Container Storage Interface driver, see *Summary of changes* in the [IBM Storage Scale Container Storage Interface Driver](#) documentation.

AFM changes

Snapshot and data management

- Added support for writable snapshots on AFM filesets for flexibility in data operations. For more information, see the *Writable snapshots in the AFM local update mode* section in the *IBM Storage Scale: Concepts, Planning, and Installation Guide*.
- Introduced compound remove operations to improve efficiency in file deletions within AFM filesets. For more information, see the *AFM fast remove for large directories* section in the *IBM Storage Scale: Concepts, Planning, and Installation Guide*.
- Enabled out-of-band cache updates (**outbandupdate**) for NFS backends to keep cache data synchronized. For more information, see the *AFM cache synchronization with outband updates* section in the *IBM Storage Scale: Concepts, Planning, and Installation Guide* and the *Synchronizing AFM cache by using outband updates* section in the *IBM Storage Scale: Administration Guide*.

Prefetch and performance improvements

- Enhanced multi-prefetch capabilities by using the **--multi-prefetch** parameter for faster and more efficient data downloads. For more information, see the *mmafmctl* command in the *IBM Storage Scale: Command and Programming Reference Guide*.
- Added support for prefetch with the **--wait** parameter to improve execution control. For more information, see the *mmafmctl* command in the *IBM Storage Scale: Command and Programming Reference Guide*.

- Supported the prefetch operations for files and directories by using a single list to simplify bulk operations.
- Enabled asynchronous prefetch (async prefetch) for AFM NFS and GPFS backend target filesets to support CAST. For more information, see the *Asynchronous auto prefetch for AFM NFS-based filesets* in the *IBM Storage Scale: Concepts, Planning, and Installation Guide*.

Resiliency and availability

- Added failover mapping to handle home NFS outages for system resiliency and continuity improvement. For more information, see *The failover mapping for AFM NFS targets* section in the *IBM Storage Scale: Concepts, Planning, and Installation Guide*.

Configuration enhancements

- Enabled configuration of AFM with nonstandard ports for NFS targets for deployment flexibility. For more information, see the *Nonstandard NFS ports for AFM* section in the *IBM Storage Scale: Concepts, Planning, and Installation Guide*.

Usability and management

- Improved migration precheck and overall usability enhancements by enabling the `afmMigrate` parameter. For more information, see the *Data migration by using AFM migration enhancements* section in the *IBM Storage Scale: Command and Programming Reference Guide*.
- Expanded coverage for new parameters for the system tuning. For more information about, `afmAllowReadBypass`, `afmReadBypassMapSize`, `afmMaxAsyncReadCount`, `afmMaxPMounts`, `afmPrefetchFileSize`, and `afmPrefetchSize`, see the *mmchconfig* command in the *IBM Storage Scale: Command and Programming Reference Guide*.

Snapshot and data management

- Added support for writable snapshots on AFM filesets for flexibility in data operations. For more information, see the *Writable snapshots in the AFM local update mode* section in the *IBM Storage Scale: Concepts, Planning, and Installation Guide*.
- Introduced compound remove operations to improve efficiency in file deletions within AFM filesets. For more information, see the *AFM fast remove for large directories* section in the *IBM Storage Scale: Concepts, Planning, and Installation Guide*.
- Enabled out-of-band cache updates (**outbandupdate**) for NFS backends to keep cache data synchronized. For more information, see the *AFM cache synchronization with outband updates* section in the *IBM Storage Scale: Concepts, Planning, and Installation Guide* and the *Synchronizing AFM cache by using outband updates* section in the *IBM Storage Scale: Administration Guide*.

Prefetch and performance improvements

- Enhanced multi-prefetch capabilities by using the **--multi-prefetch** parameter for faster and more efficient data downloads. For more information, see the *mmafmctl* command in the *IBM Storage Scale: Command and Programming Reference Guide*.
- Added support for prefetch with the **--wait** parameter to improve execution control. For more information, see the *mmafmctl* command in the *IBM Storage Scale: Command and Programming Reference Guide*.
- Introduced ability to prefetch files and directories by using a single list to simplify bulk operations.
- Enabled asynchronous prefetch (async prefetch) for AFM NFS and GPFS backend target filesets to support CAST. For more information, see the *Asynchronous auto prefetch for AFM NFS-based filesets* in the *IBM Storage Scale: Concepts, Planning, and Installation Guide*.

Resiliency and availability

- Added failover mapping to handle home NFS outages for system resiliency and continuity improvement. For more information, see *The failover mapping for AFM NFS targets* section in the *IBM Storage Scale: Concepts, Planning, and Installation Guide*.

Configuration enhancements

- Enabled configuration of AFM with nonstandard ports for NFS targets for deployment flexibility. For more information, see the *Nonstandard NFS ports for AFM* section in the *IBM Storage Scale: Concepts, Planning, and Installation Guide*.

IBM Storage Scale native REST API support

Starting with IBM Storage Scale 6.0.1.0, you can perform the AFM and AFM Object Storage operations through the `scalectl` commands and the corresponding native REST API endpoints. For more information, see *scalectl afm command*, *scalectl afm command*, AFM endpoints, and AFM to cloud object storage endpoints in the *IBM Storage Scale: Command and Programming Reference Guide*.

Call home changes

- Review mode for call home uploads
- Only for IBM Storage Scale Systems, new option to automatically collect `gpfs.snap` for automatically created service tickets.

Changes to the `mmchconfig` command

Starting with IBM Storage Scale 6.0.1.0, the `seqDiscardThreshold` parameter of the `mmchconfig` command is no longer supported. The `PagePool steal` algorithm now works based on reaccess history and sequence detection, which turns obsolete the threshold-based approach of `seqDiscardThreshold`. The new algorithm automatically handles buffer management without requiring manual threshold configuration.

Cloudkit changes

The following features are supported:

- The `cloudkit` command can deploy an IBM Storage Scale cluster on IBM Cloud.
- Protocol support for Google Cloud Platform clusters.

Data tiering

Added support for data tiering, which enables placing blocks from a single file across different storage tiers based on file offset ranges. For more information, see the *Data tiering overview* section in the *IBM Storage Scale: Administration Guide*.

File system core improvements

- Starting with IBM Storage Scale 6.0.1.0, the `no` value is set as default for the `mmapRangeLock` attribute of the `mmchconfig` command. This default setting helps to avoid possible deadlocks when the file is concurrently accessed by using mixed `mmap` operations (`mmap` and non-`mmap` operations). For more information, see [mmchconfig command](#) in the *IBM Storage Scale: Command and Programming Reference Guide*.
- Starting with IBM Storage Scale 6.0.1.0, the `mmchconfig` command supports the `restripePauseEnabled` and `restripePauseSetTimeWindow` parameters, which can pause a restripe process when direct input/output (DIO) applications are detected. This new capability prevents potential starvation for DIO threads when a command like `mmchdisk`, `mmdeldisk`, `mmadddisk`, or `mmrestripefs` is used. For more information, see [mmchconfig command](#) in the *IBM Storage Scale: Command and Programming Reference Guide*.
- Starting with IBM Storage Scale 6.0.1.0, to prevent applications from inadvertently using the `fcntl(F_SETFL, O_DIRECT)` call without receiving an indication that it is not supported, the code now enforces this limitation and rejects the call. `fcntl(F_SETFL, O_DIRECT)` has never been supported on a GPFS file system. For more information, see [Considerations for the use of direct I/O \(O_DIRECT\)](#) in the *IBM Storage Scale: Command and Programming Reference Guide*.
- For new clusters with a minimum release level of 6.0.1.0 or later, the default value of the `maxMBps` parameter is 25000. For existing clusters, the default value of this parameter is 2048 and the system administrator can change it to 25000 if the minimum release level of the cluster is at least 6.0.1.0. For more information, see [mmchconfig command](#) in the *IBM Storage Scale: Command and Programming Reference Guide*.

- A new parameter is introduced with IBM Storage Scale 6.0.1.0 for the **mmbackup** command: **--ignore-migration-state**. When this parameter is specified, if the only change of a file is the migration state, that file is not selected as an update candidate. For more information, see [mmbackup command](#) in the *IBM Storage Scale: Command and Programming Reference Guide*.
- Through improvements to the **mmfsckx** command, IBM Storage Scale 6.0.1 introduces enhanced checks in **mmfsckx**. The **mmfsckx** command can now detect and report inconsistencies with directory entries. This function only detects and reports the directory entries inconsistencies but does not repair them. For more information, see [mmfsckx command](#) in the *IBM Storage Scale: Command and Programming Reference Guide*.
- Cluster communication from remote clusters is strictly inspected on the home cluster. Invalid remote procedure calls (RPCs) are rejected and warnings are logged in the home cluster. For more information, see [Accessing a remote GPFS file system](#) in the *IBM Storage Scale: Administration Guide*.
- The **mmauth** command now supports the **--force** option. For more information, see [mmauth command](#) in the *IBM Storage Scale: Command and Programming Reference Guide*.
- Thanks to scaling changes made to the IBM Storage Scale feature known as remote fileset access control (RFAC), the limit of 31 clients per cluster is eliminated. For more information, see [Fileset access control for remote clusters](#) in the *IBM Storage Scale: Administration Guide*.

First-time data capture

Only for IBM Storage Scale Systems and tailored for specific hardware, more targeted diagnostics data is collected.

IBM Storage Scale native REST API updates

- Supports AFM and AFM Object Storage operations through the **scalectl** commands and the corresponding native REST API endpoints. For more information, see [scalectl afm command](#), [scalectl afmc command](#), AFM endpoints, and AFMCOS endpoints in the *IBM Storage Scale: Command and Programming Reference Guide*.
- Supports applying file system policies through the **scalectl policy apply** command and the corresponding native REST API with support for RBAC domain filtering and GRPC communication. For more information, see [scalectl policy command](#) and [gpfs_igetfilesetdomainid\(\)](#) subroutine endpoint in the *IBM Storage Scale: Command and Programming Reference Guide*.
- Full coexistence support for the **mmaddnode** command within native REST API environments by removing all prior restrictions at release level 6.0.1.0 or later. This support includes auto-migration of newly added nodes for concurrent management through classic IBM Storage Scale mm-commands and the native REST API.
- Enables file system rebalance and restripe operations through the **scalectl filesystem rebalance** and **scalectl filesystem restripe** commands and the corresponding native REST API endpoints. For more information, see [scalectl filesystem command](#) and file system endpoints in the *IBM Storage Scale: Command and Programming Reference Guide*.
- Supports parallel copy and file synchronization operations through the **scalectl utils xcp** command and the corresponding native REST API endpoints. For more information, see [scalectl utils command](#) and XCP endpoints in the *IBM Storage Scale: Command and Programming Reference Guide*.
- Added support for managing default quotas and per-fileset quotas through the IBM Storage Scale native REST API and corresponding **scalectl filesystem quota** commands. It includes support for listing and updating quotas, enabling or disabling default quotas and quota enforcement, checking quota usage, reporting quota discrepancies, and filtering quota entries by quota name. For more information, see [scalectl filesystem command](#), [scalectl fileset command](#) and quota endpoints in the *IBM Storage Scale: Command and Programming Reference Guide*.
- Supports creating, listing, retrieving, and deleting directories within a file system through the **scalectl filesystem directory** commands and the corresponding native REST API endpoints. For more information, see [scalectl filesystem command](#) and file system endpoints in the *IBM Storage Scale: Command and Programming Reference Guide*.

- Includes the `--fast-tier` option in the `scalectl filesystem disk add` command for configuring fast tier storage pools. For more information, see *scalectl filesystem command* in the *IBM Storage Scale: Command and Programming Reference Guide*.
- Accepts the unique identifier (UID) value in addition to the file system name for the `scalectl filesystem get` command and the corresponding native REST API file system GET endpoint. For more information, see *scalectl filesystem command* and `/scalemgmt/v3/filesystems/{filesystem}: GET` in the *IBM Storage Scale: Command and Programming Reference Guide*.
- Retrieves time zone, operating system, architecture, and IBM Storage Scale version information in the JSON output of the `scalectl node status` command and the corresponding native REST API endpoints. For more information, see `/scalemgmt/v3/nodes/status: GET` in the *IBM Storage Scale: Command and Programming Reference Guide*.
- Added support for installing the `scalectl` binary as a separate `gpfs.scalectl` package, enabling use of `scalectl` on remote systems outside the cluster without requiring the full `gpfs.scaleapi` package dependencies.

Note: The `gpfs.scaleapi` package continues to provide the core REST API components and depends on `gpfs.scalectl`, so `scalectl` remains available on IBM Storage Scale nodes.

Installation toolkit changes

The following features and characteristics are supported:

- Support for ansible-core versions earlier than 2.24.
- Support for Python 3.10 and later.
- Added support for installing the IBM Storage Scale Grafana bridge on scale collector nodes.
- Added support for installing and configuring NFS core dump on NFS protocol nodes.
- Addition of the `nfs_core_dump` enable option to support NFS core dump collection on NFS protocol nodes for troubleshooting NFS-related issues.

Management API changes

New endpoints are added to support content-aware storage and AFM.

For more information, see *IBM Storage Scale management API endpoints* in *IBM Storage Scale: Command and Programming Reference Guide*.

Management GUI Changes

- Modernization of the GUI toward IBM's Carbon standard, with enhancements to the Nodes and Events pages.
- Fileset-level multitenancy.
- AFM overview implementation.
- S3 implementation.
- Enhancements for creating and deleting snapshots.
- Improvements to support the creation of many filesets through the REST API.
- Support for Azure Linux.

For more information, see *Introduction to IBM Storage Scale GUI* in *IBM Storage Scale: Concepts, Planning, and Installation Guide*.

Operating systems support

- Starting with 6.0.1.0, you can run IBM Storage Scale on Red Hat Enterprise Linux 9.8 (RHEL 9.8).
- Starting with 6.0.1.0, you can run IBM Storage Scale on RHEL 10.2.
- Starting with 6.0.1.0, you can run IBM Storage Scale on SUSE Linux Enterprise Server 16 (SLES 16).

For more information, see [2.1 What is supported on IBM Storage Scale for AIX, Linux, and Windows?](#)

Performance monitoring changes

- New sensor `GPFSSnapshot` to better analyze every phase of the snapshot creation or deletion.

- New sensor CesS3 to monitor S3 protocol performance.
- New CLI options `--legend` and `--top` for the `mmperfmon` command.
- New default measurements: `nsdds_lat_disk_write_sum` and `nsdds_lat_disk_read_sum`.
- More descriptive error messages for `mmperfmon`, `mmpstat`, `mmptop` commands.
- New default threshold added: `Neterr_Rule`.
- Only for IBM Storage Scale Systems, added anomaly detection to monitor and report probable anomalies in cluster behavior.
- Support for the NVIDIA Nsight plug-in.
- For easier deployment of the Grafana bridge, new RPM and DEB packages: `gpfs.grafana-bridge`.
- Multi-tenancy observability improvements to allow tenants to query their own storage metrics.

Protocol changes

First-time data capture

- Enhanced this feature for periodic statistics collection, hang diagnostics, core dump generation, and tracing tools for server and client-side analysis. Also, supported data capture for default Samba config, CTDB config.
- Added new scripts such as `ganeshatracectl` and `nfs_client_debug_script.py` to capture metrics such as NFS operations, memory usage, stack traces, network state, and packet. For more information, see *The NFS client and server runtime diagnostic* section in the *IBM Storage Scale: Problem Determination Guide*.

For more information, see the *First time data capture for CES NFS* section in the *IBM Storage Scale: Concepts, Planning, and Installation Guide*.

Samba support

- Updated Samba on IBM Storage Scale to the upstream Samba 4.22.4.
- Added The SMB multichannel feature in Samba.
- Improved domain controller failover time for multiple DCs in Samba.

S3 protocol

Support for S3 over RDMA

Started S3 over RDMA support. This feature enables high-performance object access to data in GPFS file systems through RDMA transport for S3 data traffic. For more information, see the *S3 over RDMA overview* section in *IBM Storage Scale: Concepts, Planning, and Installation Guide*.

Support for S3 IAM user management

Started S3 Identity Access Management (IAM) support for S3 user management. For more information, see the *IAM users for S3 access control* section in *IBM Storage Scale: Concepts, Planning, and Installation Guide*.

CES S3 sensor for performance

Added support to collect performance metrics for S3 operations. For more information, see the *CES S3 metrics* section in *IBM Storage Scale: Problem Determination Guide*.

Support for LDAP

AUTH by using LDAP is supported.

System health monitoring changes

- Extended AFM replication monitoring, which adds dedicated events with tailored user actions.
- Added ARP cache monitoring.
- `mmhealth cluster show <component> -v` now shows events like `mmhealth node show`.
- The `mmhealth` command introduces the alias `--not-ok` option for the `--unhealthy` option.

- For the distinction between administrator and daemon names in the `mmhealth` command, the human-readable output is now switched to administrator names only. The `-Y` output stays consistent but adds the new field `adminname`.
- If more than three events are active, the output for `mmhealth node show` displays ellipsis after the third event.
- `Autoexpel` now allows to specify a custom exclusion list.

File-system-level quality of service (QoS) for I/O operations

QoS is a storage management feature that helps to ensure that critical applications consistently receive the performance they require, while maximizing the efficiency of the storage infrastructure.

Commands, data types, and programming APIs

The following section lists the modifications to the documented commands, structures, and subroutines:

New commands

- `scalectl afm`
- `scalectl afmcos`
- `scalectl utils`

Changed commands

- `gpfs.snap`
- `mmadddisk`
- `mmaddnode`
- `mmafmctl`
- `mmafmconfig`
- `mmauth`
- `mmbackup`
- `mmcachectl`
- `mmcallhome`
- `mmchattr`
- `mmchconfig`
- `mmchfileset`
- `mmchfirmware`
- `mmchfs`
- `mmchpolicy`
- `mmchvdisk server`
- `mmcrfileset`
- `mmcrfs`
- `mmcrsnapshot`
- `mmdelfs`
- `mmdumpperfdata`
- `mmfsckx`
- `mmhealth`
- `mmfsfs`
- `mmfslicensem`
- `mmperfmon`
- `mmrestorefs`

- **mms3**
- **mmsmb**
- **mmvdisk pdisk**
- **mmvdisk vdisket**
- **scalectl**
- **scalectl filesystem**
- **scalectl fileset**
- **scalectl policy**

Chapter 1. Introducing IBM Storage Scale

IBM Storage Scale is a clustered file system that provides concurrent access to a single file system or set of file systems from multiple nodes. The nodes can be SAN-attached, network-attached, a mixture of SAN-attached and network-attached, or in a shared-nothing cluster configuration. This enables high-performance access to this common set of data to support a scale-out solution or to provide a high availability platform.

IBM Storage Scale has many features, beyond common data access, including data replication, policy-based storage management, and multisite operations. You can create a cluster of AIX nodes, Linux nodes, Windows server nodes, or a mix of all three. IBM Storage Scale can run on virtualized instances that provides common data access in environments, and uses logical partitioning or other hypervisors. Multiple IBM Storage Scale clusters can share data within a location or across wide area network (WAN) connections.

Strengths of IBM Storage Scale

IBM Storage Scale provides a global namespace, shared file system access among IBM Storage Scale clusters, simultaneous file access from multiple nodes, high recoverability, and data availability through replication, the ability to make changes while a file system is mounted, and simplified administration even in large environments.

Shared file system access among IBM Storage Scale clusters

IBM Storage Scale allows you to share data between separate clusters within a location or across a WAN.

IBM Storage Scale clusters are independently managed, but IBM Storage Scale also shares data access through remote cluster mounts. This is known as a multicluster environment. When multiple clusters are configured to access the same IBM Storage Scale file system, IBM Global Security Kit (GSKit) is used to authenticate and check authorization for all network connections.

With GSKit, all messages within and across clusters are authenticated. You can also configure a cipherList to cause messages to be encrypted for transmissions.

The multicluster environment has the following features:

- The cluster that is hosting the file system can specify different security levels for each cluster authorized to mount a particular file system.
- The local cluster can remain active while changing security keys. Periodic changing of keys is necessary for a variety of reasons:
 - The number of keys should remain small to facilitate good performance.
 - Key changes prevent use or continued use of compromised keys.
 - As a matter of policy, some institutions require security keys to be changed periodically.

IBM Storage Scale uses public key authentication in a manner similar to the host-based authentication mechanism of OpenSSH. Each cluster has a pair of these keys that identify the cluster. In addition, each cluster also has an `authorized_keys` list. Each line in the `authorized_keys` list contains the public key of one remote cluster and a list of the file systems that the cluster is authorized to mount. For details about multicluster (remote mount) file system access, see *Accessing a remote GPFS file system* in *IBM Storage Scale: Administration Guide*.

For more information, see *Active file management* in the *IBM Storage Scale: Administration Guide*.

Improved system performance

GPFS file systems can improve system performance in a number of ways.

- Allowing multiple processes or applications on all nodes in the cluster simultaneous access to the same files. That is, it allows concurrent reads and writes from multiple nodes.
- Increasing aggregate bandwidth of your file system by spreading reads and writes across multiple disks.
- Balancing the load evenly across all disks to maximize their combined throughput, eliminating storage hotspots.
- Supporting very large file and file system sizes.
- Allowing concurrent reads and writes from multiple nodes.
- Provides sophisticated token management that handles fast and fine-grained access to cluster, file system, and file resources.
- Allowing for the specification of multiple networks for GPFS daemon communication and for GPFS administration command usage within your cluster.

Achieving high throughput to a single, large file that requires striping the data across multiple disks and disk controllers.

GPFS implements data striping in the file system instead of relying on data striping in a separate volume manager layer. As GPFS manages its own data striping, it enables the GPFS to achieve fault tolerance and balance load across adapters, storage controllers, and disks. Large files in GPFS are divided into equal sized blocks, and the consecutive blocks are placed on different disks in a round-robin manner.

GPFS automatically detects common data access patterns and automatically begins prefetching data accordingly. The prefetching and caching provide high throughput and fast response times. Some of the recognized I/O patterns include sequential, reverse sequential, and various forms of strided access patterns.

File consistency

IBM Storage Scale provides concurrent access to clients across the cluster by utilizing sophisticated token management. This provides concurrent and detailed access to IBM Storage Scale features, file systems, and file resources.

For more information, see [“GPFS architecture” on page 9](#).

Increased data availability

GPFS provides multiple features that improve the reliability of your file system. This includes automatic features like file system logging and configurable features like intelligently mounting file systems on startup to providing tools for flexible synchronous replication.

GPFS allows you to organize your storage hardware into *failure groups*. A failure group is defined as a set of disks that share a common point of failure that could cause them all to become simultaneously unavailable. Failure groups are defined by the system administrator, so care needs to be taken when defining disks to ensure proper failure group isolation. When used in conjunction with the *replication* feature of GPFS, the creation of multiple failure groups provides for increased file availability should a group of disks fail. Replication in GPFS ensures that there is a copy of each block of replicated data and metadata on disks in different failure groups. In this case, should a set of disks become unavailable, GPFS fails over to the replicated copies in another failure group.

During configuration, you assign a replication factor to indicate the total number of copies of data and metadata you wish to store. Currently, the maximum replication factor is 3. Replication allows you to set different levels of protection for each file or one level for an entire file system. Since replication uses additional disk space and requires extra write time, you should consider the impact of replication on your application, especially if the replication is occurring over a WAN. To reduce the overhead involved with the replication of data, you can also choose to replicate only metadata as a means of providing additional file system protection. For more information on GPFS replication, see [“File system replication parameters” on page 355](#).

GPFS is a logging file system. It creates separate logs for each file system. GPFS automatically replicates recovery logs if multiple failure groups are available. When used in conjunction with geographic based replication, disaster recovery abilities are provided. For more information on failure groups, see [“Network Shared Disk \(NSD\) creation considerations”](#) on page 329. For more information on disaster recovery with GPFS, see *Data Mirroring and Replication* in *IBM Storage Scale: Administration Guide*.

Once your file system is created, it can be configured to mount whenever the GPFS daemon is started. This feature assures that whenever the system and disks are up, the file system is available. When utilizing shared file system access among GPFS clusters to reduce overall GPFS control traffic, you might decide to mount the file system when it is first accessed. This is done either by running the `mmremotefs` command or the `mmchfs` command by using the `-A automount` option. GPFS mount traffic might be reduced by using automatic mounts instead of mounting at GPFS startup. Automatic mounts only produce additional control traffic at the point that the file system is first used by an application or user. Mounting at GPFS startup on the other hand produces additional control traffic at every GPFS startup. Thus, startup of hundreds of nodes at once is better served by using automatic mounts. However, when exporting the file system through Network File System (NFS) mounts, it might be useful to mount the file system when GPFS is started.

Enhanced system flexibility

With GPFS, your system resources are not frozen. You can add or delete disks while the file system is mounted.

When the time is favorable and system demand is low, you can rebalance the file system across all currently configured disks.

With the QoS capability, you can prevent I/O-intensive, long running administration commands from dominating file system performance and significantly delaying other tasks.

You can also add or delete nodes without having to stop and restart the GPFS daemon on all nodes.

Note: GPFS allows many quorum nodes to facilitate maintaining quorum and continued cluster operation. GPFS also allows a tiebreaker disk configuration to further enhance cluster availability. For additional information, refer to [“Quorum”](#) on page 314.

If the physical connection to the disk is broken, GPFS dynamically switches disk access to the server nodes and continues to provide data through NSD server nodes. GPFS falls back to local disk access when it discovers that the path is repaired.

After GPFS is configured for your system, depending on your applications, hardware, and workload, you can re-configure GPFS to increase throughput. You can set up your GPFS environment for your current applications and users, secure in the knowledge that you can expand in the future without jeopardizing your data. GPFS capacity can grow as your hardware expands.

Simplified storage management

IBM Storage Scale can help you achieve information lifecycle management (ILM) through powerful policy-driven, automated tiered storage management.

IBM Storage Scale provides storage management based on the definition and use of:

- Storage pools
- Policies
- Filesets

Storage pools

A storage pool is a collection of disks or RAID configurations with similar properties that are managed together as a group. Storage pools provide a method to partition storage within a file system. While you plan how to configure your storage, consider factors such as:

- Improved price performance by matching the cost of storage to the value of the data.

- Improved performance by reducing the contention for premium storage and impact of slower devices.
- Improved reliability by providing replication based on need and better failure containment.

Policies

Files are assigned to a storage pool based on defined policies. Policies provide for:

Placement policies

Placing files in a specific storage pool when the files are created.

File management policies

- Migrating files from one storage pool to another.
- Deleting files based on file characteristics.
- Changing the replication status of files.
- Snapshot metadata scans and file list creation.
- Compressing static files.

Filesets

Filesets provide a method for partitioning a file system and allow administrative operations at a finer granularity than the entire file system. For example, filesets allow you to:

- Define data block and inode quotas at the fileset level.
- Apply policy rules to specific fileset.
- Create snapshots at the fileset level.

For more information on storage pools, filesets, and policies, see *Information lifecycle management for IBM Storage Scale* in *IBM Storage Scale: Administration Guide*.

Simplified administration

GPFS offers many of the standard file system interfaces allowing most applications to execute without modification.

Operating system utilities complement the GPFS utilities, which means you can continue to also use the commands you have always used for ordinary file operations. For more information, see [“Considerations for GPFS applications”](#) on page 463.

GPFS administration commands are similar in name and function to UNIX and Linux file system commands with one important difference: the GPFS commands operate on multiple nodes. A single GPFS command can perform an administration function across the entire cluster. For more information, see the individual commands in the *IBM Storage Scale: Command and Programming Reference Guide*.

GPFS commands save configuration and file system information in one or more files. These are collectively known as GPFS cluster configuration data files. GPFS maintains the consistency of its configuration files across the cluster, and this provides accurate and consistent confirmation information. For more information, see [“Cluster configuration data files”](#) on page 26.

Basic structure of IBM Storage Scale

IBM Storage Scale is a clustered file system that is defined over one or more nodes. On each node in the cluster, IBM Storage Scale consists of three basic components: administration commands, a kernel extension, and a multithreaded daemon.

For more information, see [“GPFS architecture”](#) on page 9.

IBM Storage Scale administration commands

IBM Storage Scale administration commands are programs and scripts that control the configuration and operation of IBM Storage Scale. They typically begin with the letters "mm", like in the **mmcrcluster** command.

By default, you can issue IBM Storage Scale administration commands from any node in a cluster. If the execution of an administration command requires tasks to be performed on other nodes in the cluster, then the command automatically sends orders to the nodes to perform the tasks.

For this inter-node communication to succeed, one of the following setups is required:

- All the nodes in the cluster must be configured to allow passwordless remote shell communications.
- IBM Storage Scale sudo wrapper scripts must be configured on all the nodes of the cluster.

For more information about these two setups, See *Requirements for administering a GPFS file system* in the *IBM Storage Scale: Administration Guide*.

For descriptions of the commands, see *Command reference* in the *IBM Storage Scale: Command and Programming Reference Guide*.

The GPFS kernel extension

The GPFS kernel extension provides the interfaces to the operating system vnode and virtual file system (VFS) layer to register GPFS as a native file system.

Structurally, applications make file system calls to the operating system, which presents them to the GPFS file system kernel extension. GPFS uses the standard mechanism for the operating system. In this way, GPFS appears to applications as just another file system. The GPFS kernel extension will either fulfill these requests by using resources that are already available in the system, or sends a message to the GPFS daemon to complete the request.

The GPFS daemon

The GPFS daemon performs all I/O operations and buffer management for GPFS. This includes read-ahead for sequential reads and write-behind for all writes that are not specified as synchronous. I/O operations are protected by GPFS token management, which ensures consistency of data across all nodes in the cluster.

The daemon is a multithreaded process with some threads dedicated to specific functions. Dedicated threads for services requiring priority attention are not used for or blocked by routine work. In addition to managing local I/O, the daemon also communicates with instances of the daemon on other nodes to coordinate configuration changes, recovery, and parallel updates of the same data structures. Specific functions that execute in the daemon are as follows:

1. Allocation of disk space to new files and newly extended files, which is done in coordination with the [file system manager](#).
2. Management of directories including creation of new directories, insertion, and removal of entries into existing directories, and searching of directories that require I/O.
3. Allocation of appropriate locks to protect the integrity of data and metadata. Locks affect data that might be accessed from multiple nodes require interaction with the token management function.
4. Initiation of actual disk I/O on threads of the daemon.
5. Management of user security and quotas in conjunction with the file system manager.

The GPFS Network Shared Disk (NSD) component provides a method for cluster-wide disk naming and high-speed access to data for applications running on nodes that do not have direct access to the disks.

The NSDs in your cluster can be physically attached to all nodes or serve their data through an NSD server that provides a virtual connection. You are allowed to specify up to eight NSD servers for each NSD. If one server fails, the next server on the list takes control from the failed node.

For a given NSD, each of its NSD servers must have physical access to the same NSD. However, different servers can serve I/O to different non-intersecting sets of clients. The existing subnet functions in GPFS determine which NSD server should serve a particular GPFS client.

Note: GPFS assumes that nodes within a subnet are connected using high-speed networks. For more information on subnet configuration, see [“Using public and private IP addresses for GPFS nodes”](#) on page 18.

GPFS determines whether a node has physical or virtual connectivity to an underlying NSD through a sequence of commands that are invoked from the GPFS daemon. The NSD discovery, occurs at initial GPFS startup and whenever a file system is mounted.

Note: To manually cause this discovery action, use the `mmnsddiscover` command. For more information, see *mmnsddiscover* command in *IBM Storage Scale: Command and Programming Reference Guide*.

This is the default order of access used during NSD discovery:

1. Local block device interfaces for SAN, SCSI, IDE, or DASD disks
2. NSD servers

This order can be changed with the `useNSDserver` mount option.

You must define NSD servers for the disks. In a SAN configuration where NSD servers are defined, if the physical connection is broken, GPFS dynamically switches to the server nodes and continues to provide data. GPFS falls back to local disk access when the discovered path is repaired. This is the default behavior, and it can be changed with the `use NSD server` file system mount option.

For more information, see [“Disk considerations”](#) on page 328 and [“NSD disk discovery”](#) on page 24.

The GPFS open source portability layer

On Linux platforms, GPFS uses a loadable kernel module that enables the GPFS daemon to interact with the Linux kernel. The source code is provided for the portability layer so that the GPFS portability can be built and installed on a wide variety of Linux kernel versions and configuration.

When installing GPFS on Linux, you build a portability module based on your particular hardware platform and Linux distribution to enable communication between the Linux kernel and GPFS. For more information, see [“Building the GPFS portability layer on Linux nodes”](#) on page 483.

IBM Storage Scale cluster configurations

An IBM Storage Scale cluster can be configured in various ways. The cluster can be a heterogeneous mix of hardware platforms and operating systems.

IBM Storage Scale clusters can contain a mix of all supported node types including Linux, AIX, and Windows Server and these operating systems can run on various hardware platforms, such as IBM POWER®, x86-based servers, and IBM Z®. These nodes can all be attached to a common set of SAN storage or through a mix of SAN and network attached nodes. Nodes can all be in a single cluster, or data can be shared across multiple clusters. A cluster can be contained in a single data center or spread across geographical locations. To determine which cluster configuration is best for your application, start by determining the following:

- Application I/O performance and reliability requirements.
- Properties of the underlying storage hardware.
- Administration, security, and ownership considerations.

Understanding these requirements helps you determine which nodes require direct access to the disks and which nodes must access the disks over a network connection through an NSD server.

There are four basic IBM Storage Scale configurations:

- When all nodes are attached to a common set of Logical Unit Numbers (LUNs).
- When some nodes are NSD clients.

- When a cluster is spread across multiple sites.
- When data is shared between clusters.

All nodes attached to a common set of LUNs

In this type of configuration, all of the nodes in the cluster are connected to a common set of LUNs (for example, over a SAN). The following factors must be considered while defining this configuration:

- The maximum number of nodes accessing a LUN that you want to support.
- You cannot mix different operating systems with IBM Storage Scale to directly access the same set of LUNs on SAN.

For example, see [Figure 1 on page 7](#).




Figure 1. A cluster with disks that are SAN-attached to all nodes

Some nodes are NSD clients

In this type of configuration, only some nodes are connected to disks. Other nodes access the disks using the NSD path.

For an example, see [Figure 2 on page 7](#).




Figure 2. A cluster with some nodes connected to disks

IBM Storage Scale servers and clients

You can configure an IBM Storage Scale cluster in which some nodes have a direct attachment to the disks and others access the disks through other IBM Storage Scale nodes. This configuration is often used in large clusters or to provide a cost effective and high-performance solution.

When an IBM Storage Scale node is providing access to a disk for another IBM Storage Scale node, the node that provides access is called an NSD server. The node that accesses the data through an NSD server is called an IBM Storage Scale client.

Sharing data across multiple IBM Storage Scale clusters

IBM Storage Scale allows you to share data across multiple IBM Storage Scale clusters. After a file system is mounted in another IBM Storage Scale cluster, all access to the data is the same as if you were in the host cluster. You can connect multiple clusters within the same data center or across long distances over a WAN. In a multicluster configuration, each cluster can be placed in a separate administrative group simplifying administration or provide a common view of data across multiple organizations.




Figure 3. A multicluster configuration

Note: For more information, see *Accessing a remote GPFS file system* in *IBM Storage Scale: Administration Guide*.

Signed kernel modules for UEFI secure boot on x86_64 and secure boot Linux on Z

Starting with IBM Storage Scale 5.1.9.0, the secure boot that is defined by the Unified Extensible Firmware Interface (UEFI) is supported on x86_64; 5.2.1.0 onward, secure boot is supported on Linux on Z. The secure boot is a verification mechanism that ensures the code that is launched by a computer firmware can be trusted.

For IBM Storage Scale, using the secure boot means that the kernel modules are cryptographically signed by IBM so that their integrity can be verified when the system starts.

IBM provides the public key for verification at [Fix Central](#). Use a support ticket to request the signed kernel modules.

The correct RPM package with signed kernel modules for your system can be determined from the name of the RPM because the name incorporates the necessary information.

On x86_64, for example, if the name of the RPM package is `gpfs.gplbin-5.14.0-284.25.1.el9_2.x86_64-6.0.1.0.x86_64.rpm`:

- "5.14.0-284.25.1" is the kernel for which the modules in the RPM are built.
- "el9_2.x86_64" stands for RHEL 9.2 on x86_64.

- "6.0.1.0" refers to the IBM Storage Scale release.

On Linux on Z, for example, if the name of the RPM package is `gpfs.gplbin-5.14.0-427.18.1.el9_4.s390x-6.0.1-0.s390x.rpm`:

- "5.14.0-427.18.1" is the kernel for which the modules in the RPM are built.
- "el9_4.s390x" stands for RHEL 9.4 on s390x (Linux on Z).
- "6.0.1-0" refers to the IBM Storage Scale release.

The following information of the RPM package must be an exact match with the information of the system where the RPM is to be installed:

- Kernel level
- RHEL operating system level
- Architecture
- All four digits of the IBM Storage Scale release

If the match is not exact, the signed modules in the RPM package are not accepted and IBM Storage Scale does not start. This remains true if you upgrade the kernel level, which means that you must upgrade to a new version of the IBM Storage Scale RPM with the signed kernel modules.

Notes:

- It is sufficient to upgrade just the `gpl.bin` RPM rather than all IBM Storage Scale RPMs. You need to upgrade all IBM Storage Scale RPMs only if a `gpl.bin` RPM is unavailable for the needed level.
- If you installed the `gpfs.gplbin` and the `gpfs.gpl`, which contain the source of the kernel modules, do not run the `mmbuildgpl` command. Because the `mmbuildgpl` command overwrites the signed kernel modules with unsigned kernel modules; and IBM Storage Scale does not work anymore until you install the `gpfs.gplbin` RPM again.

IBM Storage Scale 5.1.9.0 supports kernel modules for RHEL 9.2 on the x86_64 platform; 5.2.1.0 onward, kernel modules for RHEL 9.4 on the s390x platform are supported. New signed modules will be available for all succeeding kernel updates and IBM Storage Scale PTFs in these cases:

- If a new IBM Storage Scale package is installed, make sure that you have the matching RPM for the signed kernel modules.
- If the kernel on a node that runs IBM Storage Scale is updated, make sure that you also upgrade the signed kernel modules with a matching version.

If a `gpfs.gplbin` package matching your kernel version is not included in the installation package or available on IBM [Fix Central](#), contact IBM Support to request a compatible version.

GPFS architecture

Use this information to understand the architecture of GPFS.

Interaction between nodes at the file system level is limited to the locks and control flows that are required to maintain data and metadata integrity in the parallel environment.

A discussion of GPFS architecture includes:

- [“Special management functions” on page 10](#)
- [“Use of disk storage and file structure within a GPFS file system” on page 13](#)
- [“GPFS and memory” on page 15](#)
- [“GPFS and network communication” on page 17](#)
- [“Application and user interaction with GPFS” on page 19](#)
- [“NSD disk discovery” on page 24](#)
- [“Failure recovery processing” on page 26](#)
- [“Cluster configuration data files” on page 26](#)

- [“GPFS backup data” on page 27](#)

Special management functions

In general, GPFS performs the same functions on all nodes. It handles application requests on the node where the application exists. This provides maximum affinity of the data to the application.

In the following cases, one node provides a more global function affecting the operation of multiple nodes. These are nodes acting as:

1. [“GPFS cluster manager” on page 10](#)
2. [“File system manager” on page 10](#)
3. [“Metanode” on page 12](#)
4. [“CES node \(protocol node\)” on page 12](#)
5. [“AFM gateway node” on page 12](#)

Note:

- For quorum nodes, see [“Quorum” on page 314](#).
- For Cloud services nodes, see *Designating the Cloud services nodes* in the *IBM Storage Scale: Administration Guide*.

GPFS cluster manager

There is one GPFS cluster per manager, which is selected from a set of quorum nodes that are designated for the cluster.

For more information, see [“Quorum” on page 314](#).

The cluster manager performs the following tasks:

- Monitors disk leases.
- Detects failures and manages recovery from node failure within the cluster.

The cluster manager determines whether a quorum of nodes exists to allow the GPFS daemon to start and for file system usage to continue.

- Distributes certain configuration changes that must be known to nodes in remote clusters.
- Selects the file system manager node.

The cluster manager prevents multiple nodes from assuming the role of file system manager to avoid data corruption. The [token management](#) service resides on the file system manager node and any other nodes that you specify. For more information, see *Using multiple token servers* in *IBM Storage Scale: Administration Guide*.

- Handles UID mapping requests from remote cluster nodes.
- Aggregates health information from all nodes in the cluster.

To identify the cluster manager, issue the **mmlsmgr -c** command. For more information, see *mmlsmgr command* in *IBM Storage Scale: Command and Programming Reference Guide*.

To change the cluster manager, issue the **mmchmgr -c** command. For more information, see *mmchmgr command* in *IBM Storage Scale: Command and Programming Reference Guide*.

File system manager

There is one file system manager per file system, which handles all of the nodes by using the file system.

The services that are provided by the file system manager include:

1. File system configuration

Processes the following file system changes:

- Adding disks
- Changing disk availability
- Repairing the file system

Mount and unmount processing is performed on both the file system manager and the node that are requesting the service.

2. Management of disk space allocation

Controls which regions of disks are allocated to each node, allowing effective parallel allocation of space.

3. Token management

The file system manager node might also perform the duties of the token manager server. If you have explicitly designated some of the nodes in your cluster as file system manager nodes, then the token server load is distributed among all of the designated manager nodes. For more information, see *Using multiple token servers in IBM Storage Scale: Administration Guide*.

The token management server coordinates access to files on shared disks by granting tokens that convey the right to read or write the data or metadata of a file. This service ensures the consistency of the file system data and metadata when different nodes access the same file. The status of each token is held in two places:

- On the token management server
- On the token management client holding the token

The first time a node accesses a file that it must send a request to the token management server to obtain a corresponding read or write token. After the token is granted, a node might continue to read or write to the file without requiring additional interaction with the token management server. This continues until an application on another node attempts to read or write to the same region in the file.

The normal flow for a token is:

- A message to the token management server.

The token management server then either returns a granted token or a list of the nodes that are holding conflicting tokens.

- The token management function at the requesting node then has the responsibility to communicate with all nodes holding a conflicting token and get them to relinquish the token.

This relieves the token server of having to deal with all nodes holding conflicting tokens. For a node to relinquish a token, the daemon must give it up. First, the daemon must release any locks that are held using this token, which might involve waiting for I/O to complete.

4. Quota management

In a quota-enabled file system, the file system manager node automatically assumes quota management responsibilities whenever the GPFS file system is mounted. Quota management involves:

- Allocating disk blocks to nodes that are writing to the file system
- Comparing the allocated space to the quota limits at regular intervals

Notes:

- To reduce the number of space requests from nodes writing to the file system, the quota manager allocates more disk blocks than requested (see [“Enabling quotas” on page 357](#)). This allows nodes to write to the file system without having to go to the quota manager and check quota limits each time they write to the file system.
- Quota-enabled file systems with more than 100,000 users or groups must avoid designating nodes as manager. This situation is avoided because nodes might have low memory or they might get heavily loaded due to high-memory demands for quota manager operations.

The file system manager is selected by the cluster manager. If a file system manager fails for any reason, then a new file system manager is selected by the cluster manager. During this transition, all functions continue without disruption, except for the time required to accomplish the takeover.

Depending on the application workload, the memory and CPU requirements for the services that are provided by the file system manager can make it undesirable to run a resource intensive application on the same node as the file system manager. GPFS allows you to control the pool of nodes from which the file system manager is selected by using the following commands:

- The `mmcrcluster` command when creating your cluster.
- The `mmaddnode` command when adding nodes to your cluster.
- The `mmchnode` command to change a node's designation at any time.

These preferences are honored except in certain failure situations where multiple failures occur. For more information, see *Multiple file system manager failures* in *IBM Storage Scale: Problem Determination Guide*. You might list which node is currently assigned as the file system manager by issuing the `mm1smgr` command or change which node is assigned to this task by using the `mmchmgr` command.

Metanode

There is one metanode per open file. A file refers to a file system object that also includes a directory. The metanode is responsible for maintaining file metadata integrity.

In almost all cases, the node that has had the file open for the longest continuous period is the metanode. All nodes accessing a file can read and write data directly, but updates to metadata are written only by the metanode. The metanode for each file can move to any node to meet application requirements, except when multiple remote clusters access the same file. In such a scenario, the metanode is placed on a home cluster node. The *home cluster* is also known as the *owning cluster*.

CES node (protocol node)

Only nodes that are designated as CES nodes can serve integrated protocol function.

Nodes in the cluster can be designated to be CES nodes by using the `mmchnode --ces-enable Node` command. Each CES node serves each of the protocols (NFS, SMB, Block, HDFS namenode) that are enabled. CES IP addresses that are assigned for protocol serving can failover to any of CES nodes that are up based on the configured failback policy. CES functionality can be designated only on nodes that are running on supported operating systems, and all the CES nodes must have the same platform (either all Intel or all POWER Big Endian).

For information about supported operating systems and their required minimum kernel levels, see [IBM Storage Scale FAQ in IBM Documentation](#).

For more information about Cluster Export Services, see *Implementing Cluster Export Services* and *Configuring the CES and protocol configuration* in *IBM Storage Scale: Administration Guide*.

Related concepts

[“Protocols support overview: Integration of protocol access methods with GPFS” on page 30](#)

IBM Storage Scale provides extra protocol access methods. Providing these additional file and object access methods and integrating them with GPFS offers several benefits. It enables users to consolidate various sources of data efficiently in one global namespace. It provides a unified data management solution and enables not only efficient space utilization but also avoids making unnecessary data moves just because access methods might be different.

AFM gateway node

Each cache fileset in a cluster is served by one of the nodes that is designated as the gateway in the cluster.

The gateway node that is mapped to a fileset is called the primary gateway of the fileset. The primary gateway acts as the owner of the fileset and communicates with the home cluster. While you are using

parallel data transfers, the primary gateway might also communicate with other participating gateway nodes for data transfer to home clusters.

All other nodes in the cluster, including other gateways, become the application nodes of the fileset. Therefore, any node in the cache cluster can function as a gateway node and an application node for different filesets based on configuration of the node.

Use of disk storage and file structure within a GPFS file system

A file system (or stripe group) consists of a set of disks that store file data, file metadata, and supporting entities, such as quota files and recovery logs.

When a disk is assigned to a file system, a file system descriptor is written on each disk. The *file system descriptor* is written at a fixed position on each of the disks in the file system and is used by GPFS to identify this disk and its place in a file system. The file system descriptor contains file system specifications and information about the state of the file system.

Within each file system, files are written to disk as in other UNIX file systems, using inodes, indirect blocks, and data blocks. Inodes and indirect blocks are considered *metadata*, as distinguished from data, or actual file content. You can control which disks GPFS uses for storing metadata when you create the file system with the `mmcrfs` command or when you modify the file system with the `mmchdisk` command.

The metadata for each file is stored in the inode and contains information such as file size and time of last modification. The inode also sets aside space to track the location of the data of the file. On file systems that are created in IBM Storage Scale, if the file is small enough that its data can fit within this space, the data can be stored in the inode itself. This method is called data-in-inode and improves the performance and space utilization of workloads that use many small files.

Otherwise, the data of the file must be placed in data blocks, and the inode is used to find the location of these blocks. The location-tracking space of the inode is then used to store the addresses of these data blocks. If the file is large enough, the addresses of all of its data blocks cannot be stored in the inode itself, and the inode points instead to one or more levels of indirect blocks. These trees of additional metadata space for a file can hold all of the data block addresses for large files. The number of levels that are required to store the addresses of the data block is referred to as the *indirection level* of the file.

To summarize, on file systems that are created in IBM Storage Scale, a file typically starts out with data-in-inode. When it outgrows this stage, the inode stores direct pointers to data blocks; this arrangement is considered a zero level of indirection. When more data blocks are needed, the indirection level is increased by adding an indirect block and moving the direct pointers there; the inode then points to this indirect block. Subsequent levels of indirect blocks are added as the file grows. The dynamic nature of the indirect block structure allows file sizes to grow up to the file system size.

For security reasons, encrypted files skip the data-in-inode stage. They always begin at indirection level zero.




Figure 4. GPFS files have a typical UNIX structure

File system limitations:

1. The maximum number of mounted file systems within a GPFS cluster is 256.
2. The supported file system size depends on the version of GPFS that is installed.
3. The maximum number of files within a file system cannot exceed the architectural limit.

For the latest information on these file system limitations, see the [IBM Storage Scale FAQ in IBM Documentation](#).

GPFS uses the file system descriptor to find all of the disks that make up the file system's stripe group, including their size and order. Once the file system descriptor is processed, it is possible to address any block in the file system. In particular, it is possible to find the first inode, which describes the *inode file*, and a small number of inodes that contain the rest of the file system information. The inode file is a collection of fixed-length records that represent a single file, directory, or link. The unit of locking is the single inode. Specifically, there are fixed inodes within the inode file for the following components:

- Root directory of the file system.
- *Block allocation map*, which is a collection of bits that represent the availability of disk space within the disks of the file system. One unit in the allocation map represents a subblock. A subblock is the smallest unit of contiguous disk space that can be allocated to a file. Block size, subblock size, and the number of subblocks per block are set when the file system is created and cannot be changed afterward. For more information, see the *mmcrfs* command in the *IBM Storage Scale: Command and Programming Reference Guide*. The allocation map is broken into regions that reside on disk sector boundaries. The number of regions is set at file system creation time by the parameter that specifies how many nodes access this file system. The regions are separately locked and as a result, different nodes can be allocating or deallocating space that is represented by different regions independently and concurrently.
- *Inode allocation map*, which represents the availability of inodes within the inode file. The *Inode allocation map* is located in the *inode allocation file*, and represents all the files, directories, and links that can be created. The *mmchfs* command can be used to change the maximum number of files that can be created in the file system up to the architectural limit.

The data contents of each of these files are taken from the data space on the disks. These files are considered metadata and are allocated only on disks where metadata is allowed. For more information, see the *mmcrfs* command in the *IBM Storage Scale: Command and Programming Reference Guide*.

Quota files

For each file system, IBM Storage Scale maintains quota information in internal quota files: a quota file for users, a quota file for groups, and a quota file for filesets.

Quota files are created when quotas are enabled by the `-Q yes` option, which can be either set in the `mmcrfs` command or in the `mmchfs` command. Quota files are maintained until quotas are disabled by the `mmchfs -Q no` command. To determine whether quotas are enabled on a file system, issue the following command:

```
mmfsfs Device -Q
```

where *Device* is the file system.

Quota files are internal product files and are not available for access by users. Quota file information is backed up by the `mmbackupconfig` command and is restored by the `mmrestoreconfig` command.

The quota files contain the following information:

- The quota file for users contains the block limits, file limits, actual usage, and grace period for each user.
- The quota file for groups contains the block limits, file limits, actual usage, and grace period for each group.
- The quota file for filesets contains the block limits, file limits, actual usage, and grace period for each fileset.

GPFS recovery logs

GPFS recovery logs are created at file system creation. Additional recovery logs are automatically created as needed. The file system manager assigns a recovery log to each node that accesses the file system.

GPFS maintains the atomicity of the on-disk structures of a file through a combination of rigid sequencing of operations and logging. The data structures maintained are the inode, indirect blocks, the allocation map, and the data blocks. Data blocks are written to disk before any control structure that references the data is written to disk. This ensures that the previous contents of a data block can never be seen in a new file. Allocation blocks, inodes, and indirect blocks are written and logged in such a way that there can never be a pointer to a block marked unallocated that is not recoverable from a log.

Recovery logs are replicated only if default metadata replication is turned on (`-m 2`) or if explicitly enabled for logs (`--log-replicas 2`). You can check to see if log replication is enabled for a file system by using the `mmfsfs` command and looking at the value of the `-m` and `--log-replicas` parameters. If both are set, the `--log-replicas` value takes precedence over the `-m` value for log replication.

There are certain failure cases where blocks are marked allocated but not yet assigned to a file, and these can be recovered by running the `mmfsck` command in online or offline mode. Log recovery is run as part of:

1. The recovery of a node failure affecting the objects that the failed node might have had locked.
2. A mount after the file system has been unmounted everywhere.

GPFS and memory

GPFS uses three areas of memory: memory allocated from the kernel heap, memory allocated within the daemon segment, and shared segments accessed from both the daemon and the kernel.

Memory allocated from the kernel heap

GPFS uses kernel memory for control structures such as vnodes and related structures that establish the necessary relationship with the operating system.

Memory allocated within the daemon segment

GPFS uses daemon segment memory for file system manager functions. Because of that, the file system manager node requires more daemon memory since token states for the entire file system are initially stored there. File system manager functions requiring daemon memory include:

- Structures that persist for the execution of a command
- Structures that persist for I/O operations
- States related to other nodes

The file system manager is a token manager, and other nodes might assume token management responsibilities. Therefore, any manager node can consume additional memory for token management. For more information, see *Using multiple token servers* in *IBM Storage Scale: Administration Guide*.

Shared segments accessed from both the daemon and the kernel

Shared segments consist of both pinned and unpinned memory that is allocated at daemon startup. The initial values are the system defaults. However, you can change these values later by using the `mmchconfig` command. See [“Cluster configuration file” on page 327](#).

The pinned memory is called the *pagepool* and is configured by setting the `pagepool` cluster configuration attribute. This pinned area of memory is used for storing file data and for optimizing the performance of various data access patterns. In a non-pinned area of the shared segment, GPFS keeps information about open and recently opened files. This information is held in two forms:

1. A full inode cache
2. A stat cache

Pinned memory

GPFS uses pinned memory (also called page pool memory) for storing file data and metadata in support of I/O operations.

With some access patterns, increasing the amount of page pool memory can increase I/O performance. Increased page pool memory can be useful in the following cases:

- There are frequent writes that can be overlapped with application execution.
- There is frequent reuse of file data that can fit in the page pool.
- The I/O pattern contains various sequential reads large enough that the prefetching of data improves performance.

Pinned memory regions cannot be swapped out to disk, which means that GPFS always consumes at least the value of the `pagepool` attribute in system memory. So, consider the memory requirements of GPFS and other applications running on the node when determining a value for the `pagepool` attribute.

Non-pinned memory

There are two levels of cache used to store file metadata.

Inode cache

The inode cache contains copies of inodes for open files and for some recently used files that are no longer open. The **`maxFilesToCache`** parameter controls the number of inodes cached by GPFS. Every open file on a node consumes a space in the inode cache. Additional space in the inode cache is used to store the inodes for recently used files in case another application needs that data.

The number of open files can exceed the value defined by the **`maxFilesToCache`** parameter to enable applications to operate. However, when the **`maxFilesToCache`** number is exceeded, there is no more caching of recently open files, and only open file inode data is kept in the cache.

Stat cache

The stat cache contains enough information to respond to inquiries about the file and open it, but not enough information to read from it or write to it. There is sufficient data from the inode in the stat cache to respond to a `stat()` call (for example, when issuing the `ls -l` command on a UNIX node). A stat cache entry consumes significantly less memory than a full inode. Stat cache entries are kept for the following:

- Recently accessed files
- Directories recently accessed by a number of `stat()` calls

To set the value of the stat cache, use the **maxStatCache** attribute of the **mmchconfig** command.

Notes:

1. GPFS prefetches data for stat cache entries if a pattern of use indicates this is productive (for example, if a number of `ls -l` commands issued for a large directory).
2. In versions of IBM Storage Scale earlier than 5.0.2, the stat cache is not effective on the Linux platform unless the Local Read-Only Cache (LROC) is configured. For more information, see the description of the **maxStatCache** parameter in the topic *mmchconfig command* in the *IBM Storage Scale: Command and Programming Reference Guide*. The size of the GPFS shared segment can limit the maximum setting of **maxStatCache**.
3. Each entry in the inode cache and the stat cache requires appropriate tokens:
 - a. To ensure the cached information remains correct
 - b. For the storage of tokens on the file system manager node
4. Depending on the usage pattern, system performance might degrade when an information update requires revoking a token. This happens when two or more nodes share the same information and the most recent information is moved to a different location. When the current node needs to access the updated information, the token manager must revoke the token from the current node before that node can access the information in the new location.

GPFS and network communication

Within the GPFS cluster, you can specify different networks for GPFS daemon communication and for GPFS command usage.

You can select the different networks by issuing **mmaddnode**, **mmchnode**, and **mmcrcluster** commands. In these commands, the node descriptor allows you to specify separate node interfaces for those functions on each node. The correct operation of GPFS is directly dependent upon these selections.

GPFS might not work properly if there is a firewall enabled on the nodes within the cluster. To ensure proper operation, you must either configure the firewall to allow the appropriate ports or disable the firewall. For more information, see *GPFS port usage* in *IBM Storage Scale: Administration Guide* and *mmnetverify command* in *IBM Storage Scale: Command and Programming Reference Guide*.

GPFS daemon communication

In a cluster environment, the GPFS daemon depends on the correct operation of TCP/IP.

These dependencies exist because of the following reasons:

- The communication path between nodes must be built at the first attempt to communicate.
- Each node in the cluster is required to communicate with the cluster manager and the file system manager during startup and mount processing.
- Once a connection is established, it must remain active until the GPFS daemon is shut down on the nodes.

Note: Establishing other communication paths depends upon application usage among nodes.

The daemon also uses sockets to communicate with other instances of the file system on other nodes. Specifically, the daemon on each node communicates with the file system manager for allocation of logs, allocation segments, and quotas, as well as for various recovery and configuration flows. GPFS requires an active internode communications path between all nodes in a cluster for locking, metadata coordination, administration commands, and other internal functions. The existence of this path is necessary for the correct operation of GPFS. The instance of the GPFS daemon on a node goes down if it senses that this communication is not available to it. If communication is not available to another node, then one of the two nodes exit GPFS.

Using public and private IP addresses for GPFS nodes

GPFS permits the system administrator to set up a cluster such that both public and private IP addresses are in use. For example, if a cluster has an internal network connecting some of its nodes, it is advantageous to use private IP addresses to communicate on this network, and public IP addresses to communicate to resources outside of this network.

Public IP addresses are those that can be used to access the node from any other location for which a connection exists. Private IP addresses might be used only for communications between nodes directly connected to each other with a communications adapter. Private IP addresses are assigned to each node at hardware setup time, and must be in a specific address range (IP addresses on a 10.0.0.0, 172.16.0.0, or 192.168.0.0 subnet). For more information on private IP addresses, refer to [RFC 1597 - Address Allocation for Private Internets \(www.ip-doc.com/rfc/rfc1597\)](http://www.ip-doc.com/rfc/rfc1597).

The `subnets` operand on the `mmchconfig` command specifies an ordered list of subnets available to GPFS for private TCP/IP communications. Each listed subnet might have a list of cluster names (allowing shell-style wild cards) that specifies other GPFS clusters that have direct access to the same subnet.

When the GPFS daemon starts on a node, it obtains a list of its own IP addresses and associated subnet masks from its local IP configuration. For each IP address, GPFS checks whether that address is on one of the subnets specified on the `subnets` configuration parameter. It records the list of its matching IP addresses and subnet masks, and then listens for connections on any of these addresses. If any IP address for the node (specified when the cluster was created or when the node was added to the cluster), is not specified with the `subnets` configuration parameter, GPFS automatically adds it to the end of the node's IP address list.

Therefore, when you use public IP addresses for a node, there is no need to explicitly list the public IP subnet with the `subnets` configuration parameter. For example, the normal way to configure a system would be to use host names that resolve to the external Ethernet IP address in the `mmcrcluster` command, and then, if the system administrator wants GPFS to use the High Performance Switch within the cluster, add one `subnets` configuration parameter for the HPS subnet. It is acceptable to add two `subnets` configuration parameters, one for the HPS and one for the external Ethernet, making sure that they are in that order. In this case, it does not matter which of each node's two addresses was specified when the cluster was created or when the node was added to the cluster. For example, to add remote access to an existing cluster that was using only switch addresses, it is sufficient to add two `subnets` configuration parameters.

When a node joins a cluster (its own cluster on startup, or another cluster when mounting a file system owned by another cluster), the node sends its list of IP addresses (ordered according to the order of `subnets` configuration parameters) to the cluster manager node, which forwards the list to all other nodes as part of the join protocol. No other additional information needs to be propagated.

When a node attempts to establish a connection to another node, GPFS determines the destination IP address to use according to the following procedure:

1. For each of its own IP addresses, it searches the other node's list of IP addresses for an address that is on the same subnet.
 - For normal public IP addresses this is done by comparing IP address values ANDed with the node's subnet mask for its IP address.
 - For private IP addresses GPFS assumes that two IP addresses are on the same subnet only if the two nodes are within the same cluster, or if the other node is in one of the clusters that are explicitly listed in the `subnets` configuration parameter.
2. If the two nodes have more than one IP address pair on a common subnet, GPFS uses the first one found according to the order of `subnets` specified in the initiating node's configuration parameters.
3. If a given pair of nodes do not share a common subnet defined by the `subnets` configuration parameter, IBM Storage Scale uses the network based on the daemon address of the node, which is automatically added as the last entry in each node's network address list.

Note: When you use subnets, both the interface corresponding to the daemon address and the interface that matches the subnet settings must be operational.

For more information and an example, see *Using remote access with public and private IP addresses* in *IBM Storage Scale: Administration Guide*.

Network communication and GPFS administration commands

Socket communications are used to process GPFS administration commands. Depending on the nature of the command, GPFS might process commands either on the node by issuing the command or on the file system manager. The actual command processor merely assembles the input parameters and sends them along to the daemon on the local node by using a socket.

Some GPFS commands permit you to specify a separate administrative network name. You make this specification by using the `AdminNodeName` field of the node descriptor. For additional information, see *IBM Storage Scale: Command and Programming Reference Guide* for descriptions of these commands:

- `mmaddnode`
- `mmchnode`
- `mmcrcluster`

If the command changes the state of a file system or its configuration, the command is processed at the file system manager. The results of the change are sent to all nodes and the status of the command processing is returned to the node, and eventually, to the process issuing the command. For example, a command to add a disk to a file system originates on a user process and the following process takes place:

1. Is sent to the daemon and validated.
2. If acceptable, it is forwarded to the file system manager, which updates the file system descriptors.
3. All nodes that have this file system are notified of the need to refresh their cached copies of the file system descriptor.
4. The return code is forwarded to the originating daemon and then to the originating user process.

Note: This chain of communication might allow faults that are related to the processing of a command to occur on nodes other than the node on which the command was issued.

Application and user interaction with GPFS

There are four ways to interact with a GPFS file system.

You can interact with a GPFS file system by using any of the following interfaces:

- The IBM Storage Scale management GUI as described in [“IBM Storage Scale GUI” on page 194](#).
- Operating system commands, which run at GPFS daemon initialization time or at file system mount time (as described in [“Operating system commands” on page 19](#)).
- Operating system calls, such as **open()**, from an application requiring access to a file that is controlled by GPFS (as described in [“Operating system calls” on page 20](#)).
- GPFS commands as described in *IBM Storage Scale: Command and Programming Reference Guide*.
- GPFS programming interfaces as described in *IBM Storage Scale: Command and Programming Reference Guide*.
- IBM Storage Scale management API as described in [“IBM Storage Scale management API” on page 197](#).

Operating system commands

Operating system commands operate on GPFS data during the following scenarios.

- The initialization of the GPFS daemon
- The mounting of a file system

Initialization of the GPFS daemon

GPFS daemon initialization can be done automatically as part of the node startup sequence, or manually by using the `mmstartup` command.

The daemon startup process loads the necessary kernel extensions, if they are not previously loaded by another process. The daemon then waits for the cluster manager to declare that a quorum exists. When quorum is achieved, the cluster manager changes the state of the group from *initializing* to *active*. You can see the transition to active state when the `mmfsd ready` message appears in the GPFS log file (`/var/adm/ras/mmfs.log.latest`) or by running the `mmgetstate` command. When this state changes from *initializing* to *active*, the daemon is ready to accept mount requests.

The mounting of a file system

GPFS file systems are mounted by using the GPFS `mmmount` command.

On AIX or Linux, you can also use the operating system's `mount` command. GPFS mount processing builds the structures required to provide a path to the data and is performed on both the node requesting the mount and the file system manager node. If there is no file system manager, a call is made to the cluster manager, which appoints one. The file system manager ensures that the file system is ready to be mounted. The file system manager ensures that each of the following is true:

- There are no conflicting utilities being run by the `mmfsck` command, which checks and repairs a file system.
- There are no conflicting utilities being run by the `mmcheckquota` command, which checks file system user, group and fileset quotas.
- All of the disks are available.
- Any necessary file system log processing is completed to ensure that metadata on the file system is consistent.

On the local node, the control structures required for a mounted file system are initialized and the token management function domains are created. In addition, paths to each of the disks that make up the file system are opened. Part of mount processing involves unfencing the disks, which might be necessary if this node was previously failed. This is done automatically without user intervention. If insufficient disks are up, the mount fails. That is, in a replicated system if two disks are down in different failure groups, the mount fails. In a non-replicated system, one disk down causes the mount to fail.

Operating system calls

The most common interface to files residing in a GPFS file system is through normal file system calls to the operating system.

When a file is accessed, the operating system submits the request to the GPFS kernel extension, which attempts to satisfy the application request using data already in memory. If this can be accomplished, control is returned to the application through the operating system interface. If the data is not available in memory, the request is transferred for execution by a daemon thread. The daemon threads wait for work in a system call in the kernel, and are scheduled as necessary. Services available at the daemon level include the acquisition of tokens and disk I/O.

Operating system calls operate on GPFS data during:

- Opening of a file
- Reading of data
- Writing of data

Opening a GPFS file

The open of a file residing in a GPFS file system involves the application making a call to the operating system specifying the name of the file.

Processing of a file open involves two stages:

1. Identifying the file specified by the application.

2. Building the required data structures based on the inode.

The kernel extension code processes the directory search. If the required information is not in memory, the daemon is called to acquire the necessary tokens for the directory or part of the directory needed to resolve the lookup, then reads the directory from disk into memory.

The lookup process occurs one directory at a time in response to calls from the operating system. In the final stage of open, the inode for the file is read from disk and connected to the operating system vnode structure. This requires acquiring locks on the inode and a lock that indicates the presence to the metanode. The metanode is discovered or created any time a file is opened.

- If no other node has this file open, this node becomes the metanode.
- If another node has a previous open, then that node is the metanode and this node interfaces directly with the metanode for metadata operations.

If the open involves the creation of a new file, the appropriate locks are obtained on the parent directory and the inode allocation file block. The directory entry is created, an inode is selected and initialized and then open processing is completed.

Reading file data

The GPFS read function is invoked in response to a read system call.

File read processing falls into three levels of complexity based on system activity and status:

1. Buffers are available in memory.
2. Tokens are available in memory but data must be read.
3. Data and tokens must be acquired.

At the completion of a read, a determination of the need for prefetching is made. GPFS computes a desired read-ahead for each open file based on the performance of the disks, the data access pattern and the rate at which the application is reading data. If additional prefetching is needed, a message is sent to the daemon that processes it asynchronously with the completion of the current read.

Buffer and locks available in memory

The simplest read operation occurs when the data is already available in memory because either it is prefetched or because it has been read recently by another read call.

In either case, the buffer is locally locked and the data is copied to the application data area. The lock is released when the copy is complete. Note that no token communication is required because possession of the buffer implies that the node at least has a read token that includes the buffer. After the copying, prefetching is initiated if appropriate.

Tokens available locally but data must be read

The second, more complex, type of read operation is necessary when the data is not in memory.

This occurs under three conditions:

- The token is acquired on a previous read that found no contention.
- The buffer is stolen for other uses.
- On some random read operations.

In the first of a series of reads, the token is not available locally, but in the second read it might be available.

In such situations, the buffer is not found and must be read from disk. No token activity has occurred because the node has a sufficiently strong token to lock the required region of the file locally. A message is sent to the daemon, which is handled on one of the waiting daemon threads. The daemon allocates a buffer, locks the file range that is required so the token cannot be stolen for the duration of the I/O, and initiates the I/O to the device holding the data. The originating thread waits for this to complete and is posted by the daemon upon completion.

Data and tokens must be acquired

The third and most complex read operation requires that tokens and data be acquired on the application node.

The kernel code determines that the data is not available locally and sends a message to the daemon waiting after posting the message. The daemon thread determines that it does not have the required tokens to perform the operation. In that case, a token acquire request is sent to the token management server. The requested token specifies a required length of that range of the file, which is needed for this buffer. If the file is being accessed sequentially, a desired range of data, starting at this point of this read and extending to the end of the file, is specified. In the event that no conflicts exist, the desired range is granted, eliminating the need for additional token calls on subsequent reads. After the minimum token needed is acquired, the flow proceeds with the token management process described in [“File system manager”](#) on page 10.

Writing file data

The write processing is initiated by a system call to the operating system, which calls GPFS when the write involves data in a GPFS file system.

GPFS moves data from a user buffer into a file system buffer synchronously with the application `write` call, but defers the actual write to disk. This asynchronous I/O technique allows better scheduling of the disk I/O operations and improved performance. The file system buffers come from the memory allocated based on the `pagepool` parameter in the `mmchconfig` command. Increasing the size of the `pagepool` might allow more writes to be deferred, which can improve performance in certain workloads.

A block of data is scheduled to be written to a disk when:

- The application has specified a synchronous write.
- The system needs the memory used to buffer the written data.
- The file token is revoked.
- The last byte of a block of a file is being written sequentially.
- A system sync command is run.

Until one of these occurs, the data remains in GPFS memory.

The write processing falls into three levels of complexity based on system activity and status:

- Buffer available in memory.
- Tokens available locally but data must be read.
- Data and tokens must be acquired.

Metadata changes are flushed under a subset of the same conditions. They can be written either directly, if this node is the metanode, or through the metanode, which merges changes from multiple nodes. This last case occurs most frequently if processes on multiple nodes are creating new data blocks in the same region of the file.

Buffer available in memory

The simplest path involves a case where a buffer already exists for this block of the file, but might not have a strong token.

This occurs if a previous `write` call accessed the block and it is still resident in memory. The write token already exists from the prior call. In this case, the data is copied from the application buffer to the GPFS buffer. If this is a sequential write and the last byte has been written, an asynchronous message is sent to the daemon to schedule the buffer for writing to disk. This operation occurs on the daemon thread overlapped with the execution of the application.

Token available locally but data must be read

There are two situations in which the token might exist but the buffer does not exist.

- The buffer is recently stolen to satisfy other needs for buffer space.
- A previous write obtained a desired range token for more than it needed.

In either case, the kernel extension determines that the buffer is not available, suspends the application thread, and sends a message to a daemon service thread requesting the buffer. If the write call is for a full file system block, an empty buffer is allocated since the entire block is replaced. If the write call is for less than a full block and the rest of the block exists, the existing version of the block must be read and overlaid. If the write call creates a new block in the file, the daemon searches the allocation map for a block that is free and assigns it to the file. With both a buffer assigned and a block on the disk associated with the buffer, the write proceeds as it would in [“Buffer available in memory” on page 22](#).

Data and tokens must be acquired

The third, and most complex path through write occurs when neither the buffer nor the token exists at the local node.

Prior to the allocation of a buffer, a token is acquired for the area of the file that is needed. As was true for read, if sequential operations are occurring, a token covering a larger range than is needed is obtained if no conflicts exist. If necessary, the token management function revokes the needed token from another node holding the token. Having acquired and locked the necessary token, the write continues as in [“Token available locally but data must be read” on page 22](#).

The stat() system call

The stat() system call returns data on the size and parameters associated with a file. The call is issued by the `ls -l` command and other similar functions.

The data required to satisfy the stat() system call is contained in the inode. GPFS processing of the stat() system call differs from other file systems in that it supports handling of stat() calls on all nodes without funneling the calls through a server.

This requires that GPFS obtain tokens that protect the accuracy of the metadata. To maximize parallelism, GPFS locks inodes individually and fetches individual inodes. In cases where a pattern can be detected, such as an attempt to stat() all of the files in a larger directory, inodes are fetched in parallel in anticipation of their use.

Inodes are cached within GPFS in two forms:

- Full inode
- Limited stat cache form

The full inode is required to perform data I/O against the file.

The stat cache form is smaller than the full inode, but is sufficient to open the file and satisfy a stat() call. It is intended to aid functions such as `ls -l`, `du`, and certain backup programs that scan entire directories looking for modification times and file sizes.

These caches and the requirement for individual tokens on inodes are the reason why a second invocation of directory scanning applications might run faster than the first.

GPFS command processing

GPFS commands fall into two categories: those that are processed locally and those that are processed at the file system manager for the file system involved in the command.

The [file system manager](#) is used to process any command that alters the state of the file system. When commands are issued and the file system is not mounted, a file system manager is appointed for the task. The `mmchdisk` command and the `mmfsck` command represent two typical types of commands that are processed at the file system manager.

The mmchdisk command

The `mmchdisk` command is issued when a failure that caused the unavailability of one or more disks has been corrected. The need for the command can be determined by the output of the `mmfsdisk` command.

The `mmchdisk` command performs four major functions:

- It changes the availability of the disk to **recovering**, and to **up** when all processing is complete. All GPFS utilities honor an availability of **down** and do not use the disk. **recovering** means that recovery has not been completed but the user has authorized use of the disk.
- It restores any replicas of data and metadata to their correct value. This involves scanning all metadata in the system and copying the latest to the **recovering** disk. Note that this involves scanning large amounts of data and potentially rewriting all data on the disk. This can take a long time for a large file system with a great deal of metadata to be scanned.
- It stops or suspends usage of a disk. This merely involves updating a disk state and should run quickly.
- Change metadata of disk attributes.

Subsequent invocations of the **mmchdisk** command attempts to restore the replicated data on any disk left in with an availability of **recovering**.

For more information, see *mmchdisk command* in *IBM Storage Scale: Command and Programming Reference Guide*.

The mmfsck command

The **mmfsck** command repairs file system structures.

The **mmfsck** command operates in two modes:

- Online
- Offline

For performance reasons, GPFS logging allows the condition where disk blocks are marked **used** but not actually part of a file after a node failure. The online version of **mmfsck** cleans up that condition. Running **mmfsck -o -n** scans the file system to determine whether correction might be useful. The online version of **mmfsck** runs on the file system manager and scans all inodes and indirect blocks looking for disk blocks that are allocated but not used. If authorized to repair the file system, it releases the blocks. If not authorized to repair the file system, it reports the condition to standard output on the invoking node.

The offline version of **mmfsck** is the last line of defense for a file system that cannot be used. It is most often be needed in the case where GPFS recovery logs are not available because of disk media failures. The **mmfsck** command runs on the file system manager and reports status to the invoking node. It is mutually incompatible with any other use of the file system and checks for any running commands or any nodes with the file system mounted. It exits if any are found. It also exits if any disks are **down** and require the use of **mmchdisk** to change them to **up** or **recovering**. The **mmfsck** command performs a full file system scan looking for metadata inconsistencies. This process can be lengthy on large file systems. It seeks permission from the user to repair any problems that are found, which might result in the removal of corrupted files or directories. The processing of this command is similar to those for other file systems.

For more information, see *mmfsck command* in *IBM Storage Scale: Command and Programming Reference Guide*.

NSD disk discovery

When the GPFS daemon starts on a node, it discovers the disks that are defined as NSDs by reading a disk descriptor that is written on each disk owned by GPFS. This enables the NSDs to be found regardless of the current operating system device name that is assigned to the disk.

On UNIX, NSD discovery is done by the GPFS shell script `/usr/lpp/mmfs/bin/mmdevdiscover`, which generates a list of available disk devices that appear in the node's local `/dev` file system. On Linux, any devices that are designated to be an IBM Storage Scale NSD must also exist in the `/proc/partitions` directory. Each NSD is assigned a device type as shown in the following tables:

Table 3. Supported devices on AIX		
Device name	GPFS device type	Description
vpath	vpath	IBM virtual path disk

<i>Table 3. Supported devices on AIX (continued)</i>		
Device name	GPFS device type	Description
hdisk	hdisk	AIX hard disk
hdiskpower	powerdisk	EMC power path disk
dlnmfdrv	dlnmfdrv	Hitachi dlnm

<i>Table 4. Supported devices on Linux</i>		
Device name	GPFS device type	Description
dm-	dmm	Device-Mapper Multipath (DMM)
vpath	vpath	IBM virtual path disk
sd/hd	Generic	Device having no unique failover or multipathing characteristic (predominantly Linux devices).
emcpower	powerdisk	EMC power path disk
dasd*	dasd	Direct access storage devices (DASD) (for Linux on z Systems)
vd*	Generic	Device having no unique failover or multipathing characteristic (predominantly Linux devices).
sd*	Generic	SCSI disk.
hd*	Generic	Disk on IDE controller.
scini*	Generic	Device having no unique failover or multipathing characteristic (predominantly Linux devices).
pmem*	Generic	Device having no unique failover or multipathing characteristic (predominantly Linux devices).
nvm*	Generic	NVMe disk.

<i>Table 5. Supported devices on Windows</i>		
Device Name	GPFS Device Type	Description
A Number 0-n	gpt	GPFS partition on Windows disk

To override or enhance NSD discovery, you can create a script and name it `/var/mmfs/etc/nsddevices`. The user-created `nsddevices` script, if it exists, is executed before the default discovery process.

On Windows, NSDs have a GUID Partition Table (GPT) with a single GPFS partition. NSD discovery is done by scanning the system for a list of disks that contain a GPFS partition.

On all platforms, the list of disk devices is then used by the GPFS daemon to determine whether a device interface on the local node maps to an NSD name recorded in the configuration database. The process of mapping a device interface to an NSD involves GPFS opening each device in turn and reading any NSD volume identifier potentially recorded at sector two of the disk.

If GPFS discovers that an NSD volume identifier read from a disk device matches the volume identifier recorded with the NSD name in the GPFS configuration database, I/O for the local node proceeds over the local device interface.

If no device mapping appears on the local node for a particular NSD, I/O proceeds over the IP network to the first NSD server specified in the server list for that NSD. If the first NSD server in the server list is not available, I/O proceeds sequentially through the server list until it finds an available NSD server.

Consult the **/usr/lpp/mmfs/samples/nsddevices.sample** file for configuration guidelines on how to provide additional disk discovery capabilities unique to your configuration.

Failure recovery processing

GPFS failure recovery processing occurs automatically. Therefore, though not necessary but some familiarity with its internal functions is useful when failures are observed.

Only one state change, such as the loss or initialization of a node, can be processed at a time and subsequent changes are queued. Therefore, the entire failure processing must be completed before the failed node can rejoin the group. All failures are processed first, which means that GPFS handles all failures before it completes any recovery.

GPFS recovers from a node failure by using join or leave processing messages that are sent explicitly by the cluster manager node. The cluster manager node observes that a node fails when it no longer receives heartbeat messages from the node. The join or leave processing messages are broadcast to the entire group of nodes that run GPFS, and each node updates its status for the failing or joining node. Failure of the cluster manager node results in a new cluster manager that is elected by the cluster. Then, the newly elected cluster configuration manager node processes the failure message for the failed cluster manager.

GPFS is notified of a node failure or that the GPFS daemon has failed on a node. It then starts the recovery for each of the file systems that were mounted on the failed node. If necessary, new file system managers are selected for any file systems that no longer have one.

The file system manager for each file system ensures that the failed node can no longer access the disks that comprise the file system. If the file system manager is newly appointed as a result of this failure, it rebuilds a token state by querying the other nodes in the group. After this rebuilding process is completed, the actual recovery of the failed node log begins. This recovery also rebuilds the metadata that was being modified at the time of the failure to a consistent state. Sometimes, blocks are allocated that are not part of any file and are effectively lost until `mmfsck` is run, online, or offline. After log recovery is complete, the locks that are held by the failed nodes are released for this file system. When this activity is completed for all file systems, failure processing is done. The last step of this process allows a failed node to rejoin the cluster.

Cluster configuration data files

GPFS commands store configuration and file system information in one or more files collectively known as configuration data files of the GPFS cluster. These files are not intended to be modified manually.

The GPFS administration commands are designed to keep these files synchronized between each other and with the GPFS system files on each node in the cluster. The GPFS commands constantly update the GPFS cluster configuration data files and any user modification made to this information might be lost without warning. On AIX nodes this includes the GPFS file system stanzas in **/etc/filesystems** and on Linux nodes the lists in **/etc/fstab**.

The GPFS cluster configuration data is stored in the `/var/mmfs/gen/mmsdrfs` file. The first record in the `mmsdrfs` file contains a generation number. Whenever a GPFS command causes something to change in the cluster or any of the file systems, this change is reflected in the `mmsdrfs` file and the generation number is increased by one. The latest generation number is always recorded in the `mmsdrfs` file on CCR.

The master copy of configuration data files is stored redundantly on all quorum nodes. This method of storing configuration data is called the cluster configuration repository (CCR) and is the default for new clusters created on IBM Storage Scale. Existing clusters can be converted to the new repository type using the `--ccr-enable` option of the `mmchcluster` command.

Using CCR has the advantage that full read/write access to the configuration data remains available as long as a majority of quorum nodes are accessible. For example, in a cluster with five quorum nodes, commands that update the `mmsdrfs` file continues to work normally, even if any two of the five quorum nodes have failed. In a two-node cluster with tiebreaker disks, it is still possible to run commands that change the `mmsdrfs` file if one of the two nodes has failed, as long as the surviving node has access to the tiebreaker disks. In general, full configuration command functionality remains available as long as enough nodes and, if specified, tiebreaker disks are accessible for GPFS to reach quorum. The CCR also has the advantage that it allows changing the tiebreaker disk configuration, including switching between node-based quorum and node quorum with tiebreaker, without first shutting down GPFS on all of the nodes. For more information about CCR, see [“Cluster configuration repository”](#) on page 27.

Based on the information in the GPFS cluster configuration data, the GPFS commands generate and maintain a number of system files on each of the nodes in the GPFS cluster.

Linux

`/etc/fstab`

On Linux nodes, contains lists for all GPFS file systems that exist in the cluster.

AIX

`/etc/filesystems`

On AIX nodes, contains lists for all GPFS file systems that exist in the cluster.

All GPFS nodes

`/var/mmfs/gen/mmfsNodeData`

Contains GPFS cluster configuration data pertaining to the node.

`/var/mmfs/gen/mmsdrfs`

Contains a local copy of the `mmsdrfs` file from CCR.

`/var/mmfs/gen/mmfs.cfg`

Contains GPFS daemon startup parameters.

GPFS backup data

The GPFS `mmbbackup` command creates several files during command execution. Some of the files are temporary and deleted at the end of the backup operation. There are other files that remain in the root directory of the fileset or the file system, and they must not be deleted.

The `mmbbackup` command creates other files that begin with `.mmbbackupShadow.*`. These files are associated with the `mmbbackup` command and are required for proper backups to be complete, so do not manually delete or change them.

For more information, see *mmbbackup command* in *IBM Storage Scale: Command and Programming Reference Guide*.

Cluster configuration repository

The cluster configuration repository (CCR) of IBM Storage Scale is a fault tolerant configuration store that is used by nearly all IBM Storage Scale components, including GPFS, GUI, system health, and Cluster Export Services (CES) to name a few. It is not meant to be used directly by the customer. It offers an interface to store configuration files and flat key-value pairs.

The CCR state consists of files and directories under `/var/mmfs/ccr`. The consistency of CCR state replicas is maintained by using a majority consensus algorithm. That is, the CCR state from a majority of quorum nodes or tiebreaker disks must be available for the CCR to function properly. For example, a cluster with just one or two quorum nodes and no tiebreaker disks has zero fault tolerance. If CCR state is unavailable on one of the quorum nodes, the cluster reports a quorum loss. If the CCR state is unavailable on both quorum nodes, the cluster becomes inoperable. You can create a fault tolerance by assigning more quorum nodes or tiebreaker disks.

Tiebreaker disks can be configured while GPFS is up and running. Increasing the fault tolerance by leveraging up to eight quorum nodes and up to three tiebreaker disks helps to keep the CCR state. That means the CCR remains fully functional as long as a majority of quorum nodes or tiebreaker disks are available.

Common CCR functions

The CCR has the following functions:

- Provides a PUT and GET interface for storing configuration files and flat key-value pairs redundantly across all quorum nodes. CCR uses a Paxos based algorithm to keep the stored data consistent among the quorum nodes. That is, ensure that all quorum nodes agree on the most recent version of each configuration file or key-value pair.
- Updates the CCR configuration when the number of quorum nodes or tiebreaker disks are changing.
- Creates a CCR backup file and initializes the entire quorum nodes from a CCR backup by using the **mmsdrrestore -F <CCR_BACKUP_FILE> -a** command.
- Elects the cluster manager by running the Paxos protocol among the available quorum nodes.

If the cluster is configured with tiebreaker disks and the network between the current cluster manager and the remaining quorum nodes is not working, then one of the quorum nodes writes the so-called challenges to the tiebreaker disks. The current cluster manager must answer within a specific time limit to keep this role. Otherwise, the challenger node becomes the new cluster manager.

- Provides monitor and debug commands that can be used for analysis when CCR is not functional. For more information about how to use these commands, see *CCR issues* in *IBM Storage Scale: Problem Determination Guide*.

GPUDirect Storage support for IBM Storage Scale

IBM Storage Scale's support for NVIDIA's GPUDirect Storage (GDS) enables a direct path between GPU memory and storage. This solution addresses the need for higher throughput and lower latencies. File system storage is directly connected to the GPU buffers to reduce latency and load on the CPU. For IBM Storage Scale, this means that data can be read or written directly from / to an NSD server's pagepool and it is sent to the GPU buffer of the IBM Storage Scale clients by using RDMA. IBM Storage Scale with GDS requires an InfiniBand or RoCE fabric. In IBM Storage Scale, the **mmdiag** command is enhanced to print diagnostic information for GPUDirect Storage.

IBM Storage Scale supports Nvidia GDS over RDMA over Converged Ethernet (RoCE). It allows to achieve low latencies and high throughput when reading the data from an NSD server pagepool into a GPU buffer. It requires high speed Ethernet fabric with GDS capable hardware and the CUDA environment must be installed on GDS clients.

GDS is useful where significant I/O is involved and the CPU stands as a bottleneck to overall system performance. This happens when the CPU cycles are heavily used for managing data transfers into and out of CPU memory and CPU DRAM bandwidth is used. The addition of GDS enhances the ability of IBM Storage Scale and all-flash solutions like ESS 3200 to avoid the most of those bottlenecks.

You need to install CUDA, which is provided by NVIDIA, on the IBM Storage Scale client. GDS enables the CUDA developer to copy data directly between IBM Storage Scale storage and the GPU memory by using RDMA. GDS eliminates the buffer copies in system (CPU) memory, bypasses the CPU memory, and can place data directly into the GPU application memory. GDS delivers benefits such as increased data transfer rate, lower latency, and reduced CPU utilization.

I/O requests triggered by calling the APIs `cuFileRead(.)` or `cuFileWrite(.)` from a CUDA application are run as RDMA requests started on the NSD servers copying data into or from the GPU buffers. In case the preconditions for an RDMA are not satisfied, the read and write operations are transparently handled in compatibility mode.

Note: GDS writes follow Direct I/O semantics. Therefore, IBM Storage Scale does not serialize concurrent GDS writes (or concurrent DirectIO writes) to overlapping regions of a file. Overlapping writes must be serialized by the application.

Compatibility mode

For certain types of I/O, GDS cannot use the direct RDMA from the pagepool into the GPU buffer. In those cases, the buffered I/O path is taken, which gets the data correctly into the GPU but it does not produce any performance improvements. This is called *compatibility mode* for GDS. The types of I/O that switches GDS into compatibility mode contents are as follows:

- Files with size less than 4096 bytes.
- Sparse files or files with preallocated storage. For example, `fallocate()` and `gpfs_prealloc()`.
- Encrypted files.
- Memory-mapped files.
- Compressed files or files that are marked for deferred compression. For more information on compression, see *File compression* in *IBM Storage Scale: Administration Guide*.
- Files in snapshots or clones.
- Direct I/O is disabled by using the `mmchconfig disableDIO = true` option. The default value of the **disableDIO** parameter is `false`.
- For a full list of cases for which the compatibility mode is used and how to diagnose these, see *Restriction counters* in The *GPUDirect Storage troubleshooting* topic in *IBM Storage Scale: Problem Determination Guide*.

IBM Storage Scale supports recovery for failing GDS RDMA operations by returning the failed read request to CUDA and CUDA retries the failed read request in the compatibility mode.

Other limitations

The following limitations are also applicable for the GDS support:

- IBM Storage Scale does not support GDS in the following scenarios:
 - NVIDIA GDS in asynchronous "poll" mode. The NVIDIA GDS lib implicitly converts a poll mode request on a file in an IBM Storage Scale mount to a synchronous GDS I/O request.
 - Reading a file with GDS read concurrently with a buffered read does not deliver full GDS performance for the GDS thread. This limitation holds whether the concurrent threads are part of the same or different user application. In this context, *buffered read* is considered as a nonGDS and indirect I/O.
 - Files that use data tiering, including Transparent Cloud Tiering.

Related concepts

[“Planning for GPUDirect Storage” on page 343](#)

IBM Storage Scale support for GPUDirect Storage (GDS) enables a direct path between GPU memory and storage. You need to ensure that certain conditions are met before you start installing the feature.

Related tasks

[“Installing GPUDirect Storage for IBM Storage Scale” on page 639](#)

IBM Storage Scale support for GPUDirect Storage (GDS) enables a direct path between GPU memory and storage.

[“Upgrading GPUDirect Storage” on page 692](#)

The recommended version for use of GDS with IBM Storage Scale is 5.1.6 as it supports accelerated writes.

Related information

Protocols support overview: Integration of protocol access methods with GPFS

IBM Storage Scale provides extra protocol access methods. Providing these additional file and object access methods and integrating them with GPFS offers several benefits. It enables users to consolidate various sources of data efficiently in one global namespace. It provides a unified data management solution and enables not only efficient space utilization but also avoids making unnecessary data moves just because access methods might be different.

Protocol access methods that are integrated with GPFS are NFS, SMB, HDFS, and S3. While each of these server functions (NFS, SMB, HDFS, and S3) uses open source technologies, this integration adds value by providing scaling and by providing high availability by using the clustering technology in GPFS.

For information about HDFS protocol support, see *CES HDFS in IBM Storage Scale Big Data and Analytics Support* documentation.

- The integration of file and object serving with GPFS allows the ability to create NFS exports, SMB shares, and S3 buckets containers that have data in GPFS file systems for access by client systems that do not run GPFS.
- Some nodes (at least two recommended) in the GPFS cluster must be designated as protocol nodes (also called CES nodes) from which (non-GPFS) clients can access data that resides in and is managed by GPFS by using the appropriate protocol artifacts (exports, shares, buckets, or containers).
- The protocol nodes need to have GPFS server license designations.
- The protocol nodes must be configured with "external" network addresses that are used to access the protocol artifacts from clients. The (external) network addresses are different from the GPFS cluster address that is used to add the protocol nodes to the GPFS cluster.
- The CES nodes allow SMB, NFS, HDFS, and S3 clients to access the file system data using the protocol exports on the configured public IP addresses. The CES framework allows network addresses that are associated with protocol nodes to fail over to other protocol nodes when a protocol node fails.

All protocol nodes in a cluster must be running on the same operating system and they must be of the same CPU architecture. The other nodes in the GPFS cluster might be on other platforms and operating systems.

For information about supported operating systems for protocol nodes and their required minimum kernel levels, see [IBM Storage Scale FAQ in IBM Documentation](#).

Like GPFS, the protocol serving functionality is also delivered only as software.

- The CES framework provides access to data managed by GPFS through more access methods.
- While the protocol function provides several aspects of NAS file serving, the delivery is not a NAS appliance.
- Role-based access control of the command line interface is not offered.
- Further, the types of workloads that are suited for this delivery continue to be workloads that require the scaling or consolidation aspects that are associated with traditional GPFS.

Note: Some NAS workloads might not be suited for delivery in the current release. For instance, extensive use of snapshots, or support for many SMB users. For SMB limitations, see [“SMB limitations” on page 394](#).

For more information, see [IBM Storage Scale FAQ in IBM Documentation](#).

Note: Even though some of the components that are provided are open source, the specific packages that are provided must be used. If there are existing versions of these open source packages on your system, they must be removed before you install this software.

The protocols feature provides following additional commands for administering the protocols, along with the enhancements to the existing commands:

- The commands for managing these functions includes `mmces`, `mmuserauth`, `mmnfs`, `mmsmb`, `mmobj`, and `mmperfmon`.
- In addition, `mmdumpperfdata` and `mmprotocoltrace` are provided to help with data collection and tracing.
- Existing GPFS commands are expanded with some options for protocols include `mm1scluster`, `mmchnode`, and `mmchconfig`.
- The `gpfs.snap` command is enhanced to include data by gathering about the protocols to help with problem determination.

For information on the use of CES including administering and managing the protocols, see the *Implementing Cluster Export Services* chapter of *IBM Storage Scale: Administration Guide*.

In addition to the installation toolkit, IBM Storage Scale also includes a performance monitoring toolkit. Sensors to collect performance information are installed on all protocol nodes, having one of these nodes as a designated collector node. The `mmperfmon query` command can be used to view the performance counters that have been collected.

The `mmhealth` command can be used to monitor the health of the node and the services hosted on that node.

Cluster Export Services overview

Cluster Export Services (CES) provide highly available file and object services to a GPFS cluster by using Network File System (NFS), Simple Storage Service (S3), Hadoop Distributed File System (HDFS), or Server Message Block (SMB) protocols.

IBM Storage Scale includes Cluster Export Services (CES) infrastructure to support the integration of the NFS, S3, HDFS, and SMB servers.

- The NFS server supports NFS v3 and the mandatory features in NFS v4.0 and NFS v4.1.
- The SMB server supports SMB 2, SMB 2.1, and the mandatory features of SMB 3.0 and 3.11.
- The S3 server supports S3 APIs to access the data stored in IBM Storage Scale file systems.
- CES supports HDFS protocols. HDFS Transparency is supported from IBM Storage Scale version 3.1.1. The IBM Storage Scale Big Data Analytics Integration Toolkit for HDFS Transparency or Toolkit for HDFS must be installed. The Toolkit for HDFS is supported from version 1.0.0.0. See [Support Matrix](#) under *CES HDFS* in *Big data and analytics support documentation*.

The CES infrastructure is responsible for:

- Managing the setup for high-availability clustering that is used by the protocols.
- Monitoring the health of these protocols on the protocol nodes and raising events or alerts during failures.
- Managing the addresses that are used for accessing these protocols by including failover and failback of these addresses because of protocol node failures.

High availability

With GPFS, you can configure a subset of nodes in the cluster to provide a highly available solution for exporting GPFS file systems by using the Network File System (NFS), Server Message Block (SMB), Simple Storage Service (S3), and Hadoop Distributed File System (HDFS) protocols. The participating nodes are designated as Cluster Export Services (CES) nodes or protocol nodes. The set of CES nodes is frequently referred to as the *CES cluster*.

A set of IP addresses, the *CES address pool*, is defined and distributed among the CES nodes. As nodes enter and leave the GPFS cluster, the addresses in the pool can be redistributed among the CES nodes to provide high availability. GPFS nodes that are not CES nodes that are leaving or entering the GPFS cluster do not redistribute the IP addresses. However, clients that use the SMB, NFS, and S3 protocols must not share the IP address for these protocols with the IP addresses that are used by the HDFS service to avoid

impacting the clients of other protocols during an HDFS failover. However, applications that use the CES IP addresses are not impacted. SMB failover is not transparent.

Monitoring

CES monitors the state of the protocol services. Monitoring ensures that the CES addresses are assigned to the appropriate node and that the processes that implement the services in the CES cluster are running correctly. Upon failure detection, CES marks the node as temporarily unable to participate in the CES cluster and reassigns the IP addresses to another node.

Protocol support

CES supports the following export protocols: NFS, SMB, S3, and HDFS. Each protocol can be enabled or disabled in the cluster. If a protocol is enabled in the CES cluster, all CES nodes serve that protocol.

The following are examples of enabling and disabling protocol services by using the **mmces** command:

mmces service enable nfs

Enables the NFS protocol in the CES cluster.

mmces service disable s3

Disables the S3 protocol in the CES cluster.

Commands

To set or configure CES options, use the following commands:

mmces

Manages the CES address pool and other CES cluster configuration options.

mmhdfs

Manages the HDFS configuration operations.

mmnfs

Manages NFS exports and sets the NFS configuration attributes.

mms3

Manages Simple Storage Service (S3) configuration, accounts, and buckets.

mm smb

Manages SMB exports and sets the SMB configuration attributes.

mmuserauth

Configures the authentication methods that are used by the protocols.

For more information, see *mmces command*, *mmhdfs*, *mmnfs command*, *mms3 command*, *mmobj command*, *mm smb command*, and *mmuserauth command* in [IBM Storage Scale: Command and Programming Reference Guide](#).

Restrictions

- For an up-to-date list of supported operating systems, specific distributions, and other dependencies for CES, refer to the [IBM Storage Scale FAQ in IBM Documentation](#).
- A CES cluster can contain a maximum of 32 CES nodes. All the nodes in a CES cluster must be running on the same operating system and they must be of the same CPU architecture. If the SMB protocol is enabled, then the CES cluster is limited to a total of 16 CES nodes.
- Each CES node must have network adapters capable of supporting all IP addresses in the CES address pool. The primary address of these adapters must not be used as a CES address.
- The S3 protocol is tested with 10 CES nodes in the cluster.
- Do not configure CES IP addresses in a subnet that is included in the '**subnets=**' parameter. Because the CES IP addresses move across nodes, if this restriction is not heeded, the **mmfsd** server could fail to connect to another node.

NFS support overview

The NFS support for IBM Storage Scale enables clients to access the GPFS file system by using NFS clients with their inherent NFS semantics.

The following features are provided:

Clustered NFS support

NFS clients can connect to any of the protocol nodes and get access to the exports defined. A clustered registry makes sure that all nodes see the same configuration data. It does not matter to the client to which of the CES nodes the connections are established. Moreover, the state of opened files is also shared among the CES nodes so that data integrity is maintained. On failures, clients can reconnect to another cluster node as the IP addresses of failing nodes are automatically transferred to another healthy cluster node. The supported protocol levels are NFS version 3 (NFSv3) and NFS version 4 (NFSv4.0, NFSv4.1).

Export management commands

With the **mmnfs export** command, IBM Storage Scale provides a comprehensive entry point to manage all NFS-related export tasks such as creating, changing, and deleting NFS exports. The **mmnfs export** command follows the notions of supported options, that is, a limited set of NFS-related options that proved useful. For more information, see the **mmnfs** command in the *IBM Storage Scale: Command and Programming Reference Guide*.

Export configuration commands

With the **mmnfs config** command, IBM Storage Scale provides a tool for administering the global NFS configuration. You can use this command to set and list default settings such as the port number for the NFS service, the default access mode for exported file systems, and the NFS server log level. This command follows the notions of supported options, that is, a limited set of NFS-related options that proved useful. For more information, see the **mmnfs** command in the *IBM Storage Scale: Command and Programming Reference Guide*.

NFS monitoring

The monitoring framework detects issues with the NFS services and triggers failover if an unrecoverable error occurs. Moreover, the **mmces** command provides a quick access to current and past system states and these states are useful to diagnose issues with the NFS services on the CES nodes. You can check the node performance by using the **mmhealth node show** command. Issues that are detected and causing failover are, for example, GPFS daemon failures, node failures, or NFS service failures. For more information, see the *mmces* command and the *mmhealth* command in *IBM Storage Scale: Command and Programming Reference Guide*.

NFS performance metrics

The NFS services provide performance metrics that are collected by the performance monitor framework. The **mmperfmon query** tool provides access to the most important NFS metrics through predefined queries. For more information, see the **mmperfmon** command in the *IBM Storage Scale: Command and Programming Reference Guide*.

Cross-protocol integration with SMB

IBM Storage Scale enables concurrent access to the same file data by using NFS, SMB, and native POSIX access (limitations apply).

Authentication and ID mapping

You can configure NFS services to authenticate against the most popular authentication services such as Microsoft Active Directory and LDAP. Mapping Microsoft security identifiers (SIDs) to the POSIX

user and group IDs on the file server can either be done automatically or by using the external ID-mapping service like RFC 2307. If none of the offered authentication and mapping schemes match the environmental requirements, the option to establish a user-defined configuration is available. The **mmuserauth service create** command can be used to set up all authentication-related settings. For more information, see the **mmuserauth** command in the *IBM Storage Scale: Command and Programming Reference Guide*.

First-time data capture for CES NFS

First-time data capture (FTDC) enhancements collect diagnostics and performance data from Cluster Export Services (CES) NFS environments to reduce investigation time and improve problem analysis.

Overview

In large-scale NFS environments, diagnosing issues often requires collecting diagnostic data from multiple components, including CES NFS services, protocol statistics, server state information, and client activity logs. Manually gathering this data can be time-consuming and can delay troubleshooting.

To simplify problem determination, IBM Storage Scale CES NFS provides enhancements to the FTDC framework for NFS-Ganesha environments. These new configuration options help automate the collection of key diagnostic and performance data during issue investigation.

FTDC use

- Automatically collects diagnostic information during failures or performance issues.
- Improves visibility into CES NFS server behavior and workloads.
- Standardizes data collection for troubleshooting and support.
- Reduces the time required to analyze production issues.

Scope

For the following components, FTDC enhancements apply only to diagnostic data collection and monitoring behavior:

- IBM Storage Scale 6.0.1.
- CES NFS environments that use NFS-Ganesha.
- Systems where the **mmsysmon** service is active.

This feature includes two main functions.

NFS statistics collection

You can collect NFS-Ganesha performance and memory statistics through **mmsysmon** by using the **ganesha_stats** utility.

To configure FTDC for CES NFS, see the *Configuring NFS statistics collection* section in the *IBM Storage Scale: Administration Guide*.

Enhanced debug data collection for hang scenarios

You can collect a Ganesha core dump, GPFS trace data, or both when NFS-Ganesha becomes unresponsive.

To configure FTDC for CES NFS, see the *Configuring debug data collection for NFS hung scenarios* section in the *IBM Storage Scale: Administration Guide*.

NFS statistics collection

You can use NFS statistics collection in IBM Storage Scale CES NFS to gather periodic NFS-Ganesha performance and memory statistics for performance analysis and workload investigation.

Overview

Periodic NFS-Ganesha performance and memory statistics can be collected through `mmsysmon` by using the `ganesha_stats` utility. These statistics help you analyze NFS performance and operational behavior more effectively.

The collected NFS statistics provide details about protocol activity, authentication operations, client workload behavior, memory utilization, inode cache performance, and GPFS FSAL metrics. Use these details to investigate performance issues, identify workload patterns, and understand overall server-side behavior.

Configuration parameters

The following configuration parameters control NFS statistics collection in `mmsysmonitor.conf`.

collectnfsperf

Enables periodic collection of general NFS server-side counters and is enabled by default.

Use to investigate general NFS performance issues or workload behavior. Because this option is enabled by default, it is suitable for baseline visibility and routine troubleshooting.

collectnfsclientperf

Collects IOPS and operation data for each connected NFS client IP address and is disabled by default.

Use to analyze activity from specific clients or identify heavy workloads. This setting collects metrics by connected client IP address.

This option runs one `ganesha_stats` command per connected client, which creates significant CPU and I/O overhead. Keep this option disabled unless you are actively investigating a specific client-related issue.

These metrics help identify clients that generate high request volumes or unusual workload patterns.

collectnfsmemperf

Enable investigation of memory usage growth or suspected memory leaks in NFS-Ganesha. The default value is **true**.

Ganesha memory usage tracking. Collects detailed memory statistics. These metrics provide visibility into memory consumption across NFS-Ganesha components and can help identify memory growth or potential memory leaks.

perf_collection_interval

Controls how frequently NFS performance statistics are collected, in seconds. The default value is 900 seconds.

The `perf_collection_interval` value determines the resolution of your collected data and the associated system overhead.

- 900 seconds for the default monitoring cadence for steady-state observation and is appropriate for baseline data collection or nonurgent investigation with minimal overhead.
- 300 seconds for increased resolution for active investigations and can help with more granular trending without a significant overhead increase.
- 60 seconds for high-resolution capture for acute issues and must be used only during active debugging of intermittent or fast-moving problems. Revert the setting after capture because this interval has higher overhead.

Performance statistics are written to the `/var/adm/ras/nfs_ganesha_perf_stats_<hostname>.log` file.

Operational considerations



CAUTION: Keep `collectnfsclientperf` disabled unless you are actively investigating a specific client-related issue because this option can create significant CPU and I/O overhead.



CAUTION: Use a 60-second interval only for short-term debugging and revert the setting after you capture enough data.

Mapping between configuration options and `mmchconfig` parameters

Use the `mmchconfig` command to update configuration options without restarting the `mmsysmon` service. Map the following options to the options in the `mmchconfig` command:

- `collectnfsperf` to `mmhealth-nfs-collectnfsperf`
- `collectnfsclientperf` to `mmhealth-nfs-collectnfsclientperf`
- `collectnfsmemperf` to `mmhealth-nfs-collectnfsmemperf`
- `perf_collection_interval` to `mmhealth-nfs-perf_collection_interval`

For more information, see the `mmchconfig` command in the *IBM Storage Scale: Command and Programming Reference Guide*.

To configure FTDC for CES NFS, see the *Configuring NFS statistics collection* section in the *IBM Storage Scale: Administration Guide*.

Enhanced debug data collection for hang scenarios

You can use enhanced debug data collection in IBM Storage Scale CES NFS to gather additional diagnostic information when NFS-Ganesha becomes unresponsive or hung.

Overview

Additional diagnostic data can be collected when NFS-Ganesha becomes unresponsive whereas the service process remains active. These diagnostics investigate hang conditions and identify root causes more efficiently.

Depending on the configuration, the system can collect lightweight process diagnostic information, NFS-Ganesha core dumps, and GPFS traces.

Configuration parameters

The following parameters in the `mmsysmonitor.conf` file control additional debug data collection during NFS hang scenarios:

`collectnfscoreonhang`

Controls whether the `sysmon` service captures a Ganesha core dump through `gcore` when NFS is detected as unresponsive. This flag changes what is collected on a hang event, not what is collected periodically. The default value is `false`.

When the option is enabled, the collection can include network statistics, Ganesha statistics, cluster-wide waiter information, GPFS statistics, process limits, stack traces, and the core dump through `gcore`.

`collectgpfstraceonhang`

Controls whether GPFS trace data is collected when NFS is detected as hung or unresponsive. The default value is `true`.

This flag is enabled by default to capture GPFS deadlock traces and internal state during hang scenarios.

This trace data helps to identify GPFS-level issues that can contribute to NFS unresponsiveness, such as lock contention, thread deadlocks, or resource exhaustion.

Unresponsive event handling and data collection

The `mmsysmon` service continuously monitors NFS-Ganesha health by checking process responsiveness through NFS NULL requests and I/O activity. If the NFS-Ganesha process becomes unresponsive or does not respond, the `sysmon` service triggers an `nfs_unresponsive` event.

1. Health checks run periodically at regular intervals.
2. If NFS-Ganesha does not respond within the `UnrespTimeout` period of 70 seconds, the service is marked as unresponsive.
3. The `maxwaittime` setting of 240 seconds defines the maximum time to wait before declaring NFS-Ganesha as hung.
4. After an unresponsive condition is detected, the `nfsdebug.snap` script is triggered automatically.
5. The script collects diagnostic data based on the enabled `collectnfscoreonhang` and `collectgpfstraceonhang` configuration flags.
6. Data collection starts immediately after detection to capture the NFS-Ganesha state during the `nfs_hang` condition.

Operational considerations

The flags `collectnfscoreonhang` and `collectgpfstraceonhang` can be enabled or disabled independently based on troubleshooting needs.



Warning: Ensure that you have sufficient free space in `/var/adm/ras/` and `/tmp/` before you enable these options, and monitor disk usage during active troubleshooting sessions.

To configure FTDC for CES NFS, see the *Configuring debug data collection for NFS hung scenarios* section in the *IBM Storage Scale: Administration Guide*.

SMB support overview

The SMB support for IBM Storage Scale allows clients to access the GPFS file system by using SMB clients with their inherent SMB semantics.

The following features are provided:

- **Clustered SMB support**

SMB clients can connect to any of the protocol nodes and get access to the shares defined. A clustered registry makes sure that all nodes see the same configuration data, that is, for the client it does not matter to which of the CES nodes the connection is established. Moreover, the state of opened files (share modes, open modes, access masks, and locks) is also shared among the CES nodes so that data integrity is maintained. On failures, clients can reconnect to another cluster node as the IP addresses of failing nodes are transferred to another healthy cluster node. The supported protocol levels are SMB2 and the base functionality of SMB3 (dialect negotiation, secure negotiation, encryption of data on the wire).

- **Export management command**

With the `mmsmb` command, IBM Storage Scale provides a comprehensive entry point to manage all SMB-related configuration tasks like creating, changing, and deleting SMB shares and administering the global configuration. The `mmsmb` command follows the notions of supported options, that is, a limited set of SMB-related options that are proven useful. Moreover, the Microsoft Management Console can be used to administer SMB shares.

- **SMB monitoring**

The monitoring framework detects issues with the SMB services and triggers failover if there is an unrecoverable error. Moreover, the `mmhealth` and `mmces` commands provide quick access to current and past system states and aid to diagnose issues with the SMB services on the CES nodes. Issues that are detected and causing failover are, for example, GPFS daemon failures, node failures, or SMB service failures.

- **Integrated installation**

The SMB services can be installed by the integrated installer (toolkit) together with the CES framework, and the NFS, S3, and Object protocols.

- **SMB performance metrics**

The SMB services provide two sets of performance metrics that are collected by the performance monitor framework. Thus, not just the current metrics can be retrieved but also historic data is available (at some lower granularity). The two sets of metrics are global SMB metrics (like number of connects and disconnects) and the metrics for each SMB request (number, time, throughput). The **mmperfmon query** tool provides access to all SMB metrics by using predefined queries. Moreover, metrics for the clustered file metadata base CTDB are collected and exposed by using the **mmperfmon query** command.

- **Cross-protocol integration with NFS, S3, and SMB**

IBM Storage Scale allows concurrent access to the same file data by using unified file (such as, SMB, S3, and NFS), and object interfaces.

- **Authentication and ID mapping**

The SMB services can be configured to authenticate against the most popular authentication services MS Active Directory and LDAP. Mapping MS security identifiers (SIDs) to the POSIX user and group IDs on the file server can either be done automatically by using the automatic or by using an external ID mapping service like RFC 2307. If none of the offered authentication and mapping schemes matches the environmental requirements the option to establish a user-defined configuration is available. The **mmuserauth service create** command can be used to set up all authentication-related settings.

S3 support overview

IBM Storage Scale supports the S3 access protocol. It provides high-performance and scalable S3 object access to data, which is stored in IBM Storage Scale file systems. It is integrated with Cluster Export Services (CES) to provide highly available S3 access service in an IBM Storage Scale cluster. The IBM Storage Scale S3 is also sometimes called CES S3.

The IBM Storage Scale S3 features are as follows:

Clustered S3 support

S3 clients can connect to any of the protocol nodes to get access to the S3 buckets. The S3 configuration data is stored in the CES shared root, and this makes sure that all nodes see the same configuration data. It does not matter to the client to which of the CES nodes the connections are established. On failures, clients can reconnect to another protocol node as the IP addresses of failing nodes are automatically transferred to another healthy protocol node.

S3 service configuration

With the **mms3 config** command, IBM Storage Scale provides a command for administering the global S3 configuration. You can use this command to set and list default settings such as the port number for the S3 service, the S3 service log level, number of S3 endpoint forks. For a list of configurable options, see the *mms3 command* in the *IBM Storage Scale: Command and Programming Reference Guide*.

S3 accounts management

IBM Storage Scale S3 allows administrators to have S3 accounts to manage S3 access keys for S3 clients and their respective UIDs and GIDs.

With the **mms3 account** command, IBM Storage Scale provides a comprehensive entry point to create, delete, and list S3 accounts at run time of the S3 service. For more information, see the *mms3 command* in the *IBM Storage Scale: Command and Programming Reference Guide*.

S3 buckets management

IBM Storage Scale S3 maps each S3 bucket to a directory in the IBM Storage Scale file system. With the **mms3 bucket** command, IBM Storage Scale provides a comprehensive entry point to create, delete, and list S3 buckets at run time of the S3 service. Buckets can be declared for any directory in the GPFS file system, including a fileset junction. No directories that are located only locally on a server node can be exported because they cannot be used in a failover situation. For more information, see the *mms3 command* in the *IBM Storage Scale: Command and Programming Reference Guide*.

Supported S3 APIs

IBM Storage Scale S3 allows S3 clients to manage S3 objects and buckets by using S3 API requests. To learn about the supported S3 APIs, see the *S3 objects API* and *S3 buckets API* sections in the *IBM Storage Scale: Administration Guide*.

S3 monitoring

The monitoring framework monitors the S3 service to detect issues with the S3 services and triggers failover if an unrecoverable error occurs. If a problem is found, the CES addresses of the node are reassigned, and the node state is set to the failed state. When the problem is corrected, the node resumes normal operation. Moreover, the **mmhealth** command provides quick access to current and past system states and these states are useful to diagnose issues with the S3 services on the CES nodes. You can check the node state by using the **mmhealth node show** command. Issues that are detected and causing failover are, for example, GPFS daemon failures, node failures, or S3 service failures. For more information, see the **mmhealth** command in the IBM Storage Scale 5.2.1 documentation.

S3 failover

When a CES node leaves the cluster, the CES addresses assigned to that node are redistributed among the remaining nodes. Remote clients that access the GPFS file system might see a pause in service while the internal state information is passed to the new servers.

S3 configuration backup and restore

S3 supports automatic or periodic backup of the S3 configuration from CES Share root directory to CCR. Backups are scheduled every 10 minutes on CES nodes where S3 is enabled. S3 also supports manual backup and restore of S3 configuration data. For more information, see the Backup and Restore of S3 configuration data section.

ILM support including tiering to tape support

IBM Storage Scale S3 supports two solutions to tier S3 buckets and objects stored using IBM Storage Scale S3 to tape. The supported solutions are as follows:

- IBM Storage Archive Enterprise Edition (EE)
- IBM Storage Protect for Space Management

Limitation of tiering to tape support limitation

Customers who are planning to use tiering to tape support with IBM Storage Scale must first review their deployments with IBM Storage Scale development. Contact IBM Storage Scale support at scale@us.ibm.com to get your use case reviewed.

Note: This limitation does not apply to base S3 protocol support. This limitation applies only when you are planning to use tiering to tape support with that S3 protocol that is available from the IBM Storage Scale 5.2.1 and later.

S3 versioning support

Customers who are planning to use S3 versioning with IBM Storage Scale must first review their deployments with IBM Storage Scale development. Contact IBM Storage Scale support at scale@us.ibm.com to get your use case reviewed.

Note: This limitation does not apply to base S3 protocol support. This limitation applies only when you are planning to use S3 versioning support with the S3 protocol that is available on IBM Storage Scale 5.2.2 and later.

S3 over RDMA transport support

Starting with IBM Storage Scale 6.0.1.0, S3 over RDMA is supported. This feature enables high-performance object access to data stored on IBM Storage Scale file systems by transporting S3 data traffic over RDMA transports. For more information, refer to [“S3 over RDMA overview and deployment” on page 49](#).

Customers who are planning to use S3 over RDMA transport with IBM Storage Scale must review your deployments with the IBM Storage Scale development team. Contact IBM Storage Scale support at scale@us.ibm.com to get your use case reviewed.

S3 bucket notification support

Use the S3 bucket notifications feature to receive notifications for specific events in your S3 bucket. To receive notifications, add a notification connection configuration by using the **mms3** command that identifies the destination where the S3 server sends notifications. For more details about the events supported, refer to [“S3 bucket notification” on page 46](#).

S3 architecture

The S3 support for IBM Storage Scale enables S3 access to data, which is stored in IBM Storage Scale file systems. The S3 (NooBaa) server runs on the CES nodes in an IBM Storage Scale CES cluster. You can use the **mms3** commands to manage S3 configuration, S3 accounts, and S3 buckets.

The S3 applications use the S3 APIs to access data in IBM Storage Scale. The S3 applications run outside the IBM Storage Scale cluster on client systems by using any underlying infrastructure.





Figure 5. IBM Storage Scale CES S3 architecture

Control path architecture

For the control path, the **mmces** command is enhanced to support management of S3 service in a CES cluster. CES S3 also introduces the new **mms3** command to manage S3 configuration, accounts, and buckets. The **mmhealth** command is enhanced to display the results of the background monitoring for the health of S3 service.

Figure 6. IBM Storage Scale CES S3 control path



Data path architecture

The following figure illustrates how object data is represented at each layer in the infrastructure.




Figure 7. An example of IBM Storage Scale CES S3 data path

The NooBaa endpoint processes provide S3 access to data in IBM Storage Scale.

IBM Storage Scale S3 makes directories in IBM Storage Scale file systems accessible as S3 buckets. In this example, the S3 bucket `proj1-bucket` is mapped to the directory `proj1-data` on the `/ibm/fs1` file system. The `mms3` command can report all mappings of S3 buckets to file system directories.

```
mms3 bucket list proj1-bucket
```

A sample output is as follows:

```
Name Filesystem Path Bucket Owner
-----
proj1-bucket /ibm/fs1/proj1-data s3user
```

Multi-protocol data sharing with S3, NFS, SMB, POSIX, and IBM Storage Scale Container Storage Interface driver

Multi-protocol data sharing for file and object access allows use cases where you can access instance of same data by using object and file interfaces.

Some of the key unified file and object access use cases are as follows:

- Accessing object by using file interfaces and accessing file by using object interfaces (S3 APIs) help legacy applications that are designed for file to start integrating into the object world.
- It allows files exported using NFS, SMB or IBM Spectrum Scale CSI, or files available on POSIX, to be accessible as objects using HTTP to the end clients.
- Multi-protocol access for file and object that is available in different environments allows supporting and sharing data with multiple access options. For more information about the NFS protocol, see the *Configuring the CES and protocol configuration* section in the *IBM Storage Scale: Administration Guide*.

The unified file and object access allows users to access the same data as an object and as a file. Data can be stored and retrieved through IBM Storage Scale as objects by using S3 APIs or as files from POSIX, NFS, SMB interfaces, or through IBM Storage Scale Container Storage Interface driver. The unified file and object access provides the following capabilities:

- Ingest data by using the S3 APIs, and access this data from the file interface.
- Ingest data by using the file interface, and access this data using the S3 APIs.
- Ingest data by using IBM Storage Scale Container Storage Interface driver, and access this data from the file interfaces or S3 APIs.

Note: This feature is tested with basic authentication only. It is not tested with any external authentication mechanism on IBM Storage Scale.

For more details about multi-protocol data access in an IBM Storage Scale cluster, refer to <https://community.ibm.com/community/user/storage/blogs/madhu-punjabi/2024/08/19/multi-protocol-data-access-with-ibm-storage-scale>.

Load balancing

Use multiple protocol nodes for higher throughput and performance of S3 by allowing more requests in parallel and distributing the workload among all protocol nodes.

Certain steps might be necessary to ensure that client requests are distributed among all protocol nodes.

If a hostname is resolved to a single protocol IP address (such as 192.168.1.1), then all client requests are sent to a protocol node that is associated with the specific protocol address. The requests must be distributed among all the protocol nodes instead of a single protocol node. To accomplish this you can use a load balancer.

By using a load balancer, such as HAProxy configure the destination for an endpoint to resolve to the load balancer. The load balancer forwards incoming requests to one of the protocol nodes (based on a balancing strategy). By configuring the load balancer with all of the protocol IP addresses, the client requests can be sent to the entire set of protocol nodes.

Another common solution is to use a DNS Round Robin, and configure the DNS server to associate a single host name with a set of addresses. When a client does a DNS lookup of the hostname, one address from the set is returned. By configuring the DNS Round Robin service to associate the set of protocol IP addresses with host name, client requests are distributed among all protocol nodes. As each client initiates a request to the endpoint hostname, the DNS Round Robin service returns an IP address by cycling through the set of protocol IP addresses.

Reference performance numbers for an IBM Storage Scale CES S3 cluster

The performance of an IBM Storage Scale CES S3 cluster is highly dependent on an underlying infrastructure and workload.

IBM published the following benchmark results for IBM Storage Scale CES S3:

- COSBench using objects with a size of 1 GB running against a three-node IBM Storage Scale System CES S3 cluster and using IBM Storage Scale System 3200 as the back-end storage:
 - 63+ GB/s aggregated throughput for read workloads
 - 25+ GB/s aggregated throughput for write workloads
- COSBench using small objects of different sizes ranging from 4 KB to 8 MB, running against a three-node IBM Storage Scale System CES S3 cluster and using IBM Storage Scale 3200 as the back-end storage:
 - 56K ops/s read for 4K objects
 - 5.6K ops/s write for 4K objects

For more information, see [IBM Storage Scale CES S3 evaluation of large and small objects using COSBench](#).

Migration from Swift Object or IBM Storage Scale Data Access Services S3 to IBM Storage Scale S3

Migration from Swift Object to IBM Storage Scale S3

- CES Swift Object protocol feature is not supported from IBM Storage Scale 5.1.9 onwards. CES Swift Object is replaced by IBM Storage Scale S3.
- IBM Storage Scale 5.1.8 is the last release that has CES Swift Object protocol.
- IBM Storage Scale 5.2.1.x tolerates the update of a CES node from IBM Storage Scale 5.1.8 or 5.1.9.
 - *Tolerate* means:
 - The CES node will be updated to IBM Storage Scale 5.2.1.x.
 - Swift Object support will not be updated as part of the IBM Storage Scale 5.2.1.x update.
 - You may continue to use the version of the Swift Object protocol that was provided in IBM Storage Scale 5.1.8 on the CES node in IBM Storage Scale 5.2.1.x or higher cluster.
 - IBM will provide usage and known defect support for the version of Swift Object that was provided in IBM Storage Scale 5.1.8 until you migrate to a supported object solution that IBM Storage Scale provides.
- Contact IBM for further details and migration planning by sending an email to scale@us.ibm.com.

Migration from IBM Storage Scale Data Access Services S3 to IBM Storage Scale S3

IBM Storage Scale Data Access Services S3 is not supported from IBM Storage Scale 5.2.0. IBM Storage Scale Data Access Services S3 is replaced by IBM Storage Scale S3. Contact IBM for further details and migration planning by sending an email to scale@us.ibm.com.

S3 buckets API

IBM Storage Scale S3 maps each S3 bucket to a directory in the IBM Storage Scale file system. IBM Storage Scale S3 allows S3 clients to manage S3 buckets by using the following API requests:

- S3 GetBucketTagging
- S3 PutBucketTagging
- S3 DeleteBucketTagging
- S3 CreateBucket
- S3 ListObjects
- S3 ListObjectsV2
- S3 DeleteBucket
- S3 HeadBucket
- S3 ListBuckets
- S3 ListMultipartUploads
- S3 PutBucketVersioning (RPQ)
- S3 GetBucketVersioning
- S3 ListObjectVersions
- S3 GetBucketPolicy
- S3 PutBucketPolicy
- S3 GetBucketPolicyStatus
- S3 SetBucketPolicy
- S3 PutBucketNotificationConfiguration

- S3 GetBucketNotificationConfiguration

S3 bucket notification

Note:

- The supported event types are s3:ObjectCreated and s3:ObjectRemoved.
- Support for configuring webhooks to get notifications.
- S3 ObjectCreated:*
- S3 ObjectCreated:Put
- S3 ObjectCreated:Post
- S3 ObjectCreated:Copy
- S3 ObjectCreated:CompleteMultipartUpload
- S3 ObjectRemoved:*
- S3 ObjectRemoved>Delete
- S3 ObjectRemoved>DeleteMarkerCreated

Other

- Bucket policies
- Support MD5-based ETags
- Upload and download objects with pre-signed URLs.

S3 versioning support

Note:

- After you modify the S3 versioning state (Enabled or Suspended), the new state might change after 60 or more seconds.
- If you perform any operations on a bucket during modification of the versioning state (Enabled or Suspended), unexpected issues might occur.
- When you use S3 buckets with the Enabled versioning state, a maximum of 50 parallel operations (tested limit) can be performed on a bucket from a CES node at any specified time.

S3 objects API

IBM Storage Scale allows S3 users to manage objects by using the following API requests:

- S3 PutObject
- S3 GetObject
- S3 HeadObject
- S3 CopyObject
- S3 DeleteObject
- S3 DeleteObjects
- S3 CreateMultipartUpload
- S3 CompleteMultipartUpload
- S3 AbortMultipartUpload
- S3 ListMultipartUploads
- S3 UploadPart
- S3 UploadPartCopy
- S3 ListParts

- S3 DeleteObjectTagging
- S3 GetObjectTagging
- S3 ListObjectVersions

IBM Storage Scale S3 allows S3 applications to store user-defined object metadata in addition to the object data itself.

IAM users for S3 access control

Identity and access management (IAM) users grant specific S3 permissions without exposure of primary account credentials, which reduces security risk and enables granular access control through policies.

IAM users are secondary identities that operate under an S3 account. Each IAM user receives unique credentials and permissions that bucket policies defined. This separation prevents credential share and limits access to only required operations.

Identity types in S3

S3 access control uses two identity types:

S3 accounts

Primary identities that administrators created by using the **mms3** command. S3 accounts own all resources (buckets and objects) and authenticate with access keys and secret keys. Each S3 account can create multiple IAM users.

IAM users

Secondary identities created by S3 accounts through the AWS CLI. IAM users inherit the parent account context but operate with restricted permissions. Access is controlled entirely through bucket policies, not through UID or GID mappings.

Access control mechanism

When an S3 account creates an IAM user, the system generates unique access credentials. The IAM user authenticates S3 requests with these credentials. The following components determine permissions:

- Bucket policies that explicitly allow or deny operations.
- Resource ownership (parent S3 account owns all resources).
- Access control lists on individual objects.

IAM users can create buckets and objects. These resources belong to the parent S3 account. Access to resources, which other users created, requires explicit permission through bucket policies.

Common use cases

IAM users address these access control requirements:

Application integration

Applications receive read-only or write-only access to specific buckets without full account credentials.

Team access

Team members receive separate credentials with permissions that match to their responsibilities.

Service integration

External services or backup systems access specific resources without broader account permissions.

Access tracking

Operations are attributed to individual IAM users for audit logs and compliance reports.

Security model

IAM users reduce security risk through credential isolation and permission boundaries:

- Each IAM user has unique credentials that can be revoked independently.
- Bucket policies enforce the minimum required permissions.
- Compromised IAM credentials affect only the resources that are explicitly granted in policies.
- Access keys can be rotated without effect on other users.
- Deleted IAM users lose all access immediately.

Scope and requirements

The requirements for IAM user support are as follows:

- The S3 service must be configured and operational.
- You must have administrative access to create S3 accounts with the **mms3** command.
- AWS CLI must be installed on client systems.
- Network connectivity must be between clients and S3 endpoints.

IAM users apply when access control extends beyond simple bucket ownership or when multiple identities require different permission levels for the same resources.

IAM user management operations

The S3 service supports a subset of AWS IAM API operations for IAM users and access credentials management. Use operations to create, modify, and delete IAM users, manage access keys, and track key usage. All operations are run by using the AWS CLI with the appropriate endpoint configuration.

These operations manage the lifecycle of IAM users in an S3 account.

Table 6. IAM user management operations		
Operation	Purpose	Command
CreateUser	Creates a new IAM user under the authenticated S3 account. The IAM user inherits the parent account context and can be assigned specific permissions through bucket policies.	<code>aws iam create-user --user-name <i>username</i></code>
GetUser	Retrieves detailed information about a specific IAM user, including the user ID (principal ID), ARN, creation date, and path. Use this operation to obtain the principal ID required for bucket policies.	<code>aws iam get-user --user-name <i>username</i></code>
ListUsers	Lists all IAM users associated with the authenticated S3 account. Returns user names, user IDs, ARNs, and creation dates for all users.	<code>aws iam list-users</code>
UpdateUser	Modifies IAM user properties such as the user name or path. Use this operation to rename users or reorganize user hierarchies.	<code>aws iam update-user --user-name <i>old-name</i> --new-user-name <i>new-name</i></code>
DeleteUser	Deletes an IAM user from the S3 account. You must delete all access keys associated with the user before deleting the user.	<code>aws iam delete-user --user-name <i>username</i></code>

Table 6. IAM user management operations (continued)		
Operation	Purpose	Command
CreateAccessKey	Creates an access key pair (access key ID and secret access key pair) for an IAM user. The secret key is returned only during creation and cannot be retrieved later. Use this operation to generate initial credentials for new IAM users or to create more keys for key rotation.	<pre>aws iam create-access-key -- user-name <i>username</i></pre>
ListAccessKeys	Lists all access keys that are associated with an IAM user. Returns the access key ID, status (Active or Inactive), and creation date for each key. Use this operation to audit existing credentials or identify keys for rotation.	<pre>aws iam list-access-keys -- user-name <i>username</i></pre>
GetAccessKeyLastUsed	Retrieves the last usage information for an access key, including the date and time of last use, the service that was accessed, and the AWS region. Use this operation to identify inactive keys for removal or to verify that credentials are in active use.	<pre>aws iam get-access-key- last-used --access-key-id AKIAIOSFODNN7EXAMPLE</pre>
UpdateAccessKey	Changes the status of an access key between Active and Inactive. Inactive keys cannot authenticate requests but remain in the system. Use this operation to temporarily disable credentials without deletion or to reactivate previously disabled keys.	<pre>aws iam update-access-key --user-name <i>username</i> --access- key-id AKIAIOSFODNN7EXAMPLE -- status <i>Inactive</i></pre>
DeleteAccessKey	Permanently deletes an access key from an IAM user. The key can no longer authenticate requests after deletion. Use this operation to remove compromised credentials or to complete key rotation after new keys are deployed.	<pre>aws iam delete-access-key --user-name <i>username</i> --access- key-id AKIAIOSFODNN7EXAMPLE</pre>

See the *Creating and managing IAM user for S3 access* in *IBM Storage Scale: Administration Guide* to create and manage IAM users.

S3 over RDMA overview and deployment

IBM Storage Scale 6.0.1.0 and later support S3 over RDMA and enables high-performance object access to IBM Storage Scale file systems through RDMA fabrics for S3 data transport.

Overview

IBM Storage Scale 6.0.1.0 and later supports S3 over Remote Direct Memory Access (RDMA). This capability enables high-performance object access to the file systems by allowing S3 data traffic to use RDMA transports.


By using RDMA, you can reduce CPU overhead, lower latency, and improve throughput for demanding object storage workloads, such as AI and machine learning, analytics, and high-performance data pipelines.

The solution is designed to support both traditional CPU-based workloads and GPU-accelerated environments.

Note: If you want to use S3 over RDMA transport with IBM Storage Scale, review your deployments with the IBM Storage Scale development team. Contact IBM Storage Scale support at scale@us.ibm.com to get your use case reviewed.

Architecture

The following figure shows how an S3 client that is running on a CPU- or GPU-based system communicates with IBM Storage Scale CES nodes by using RDMA, enabling efficient S3 header exchange over CES IP and high-speed data transfer directly between client and storage.



The S3 over RDMA architecture separates control paths and data paths to isolate management and data flow operations.

Control path (HTTP/HTTPS)

Standard S3 GET and PUT request headers are modified to include specific metadata (for example, x-amz-rdma-token), before transmitting to the S3 server.. This path handles authentication, coordination, and transfer setup.

Data path (RDMA)

The actual data payload is transferred directly between the IBM Storage Scale CES node and GPU memory or system memory by using RDMA operations.

Benefits

S3 over RDMA provides the following advantages:

- Lower network latency compared to TCP/IP

- Reduced CPU usage during data transfer
- Higher throughput for large object operations
- Improved performance and consistency under heavy workloads
- Better scalability for AI, analytics, and HPC environments

Configuration considerations

The following configuration considerations apply to a CES cluster where the S3 service with RDMA is enabled. Each CES node must meet the following requirements:

- Operating system: Red Hat Enterprise Linux 9.x or 10.x.
- Architecture: x86_64.
- By default, up to 64 GB RAM for NooBaa service that is configured with up to eight endpoint processes.
- Use the **mms3 config** command to adjust the number of S3 NooBaa endpoint processes based on workload requirements. Make sure that sufficient processes to manage the workload.

Supported RDMA network technologies

IBM Storage Scale supports the following RDMA network technologies for S3 communication:

- RoCE (RDMA over Converged Ethernet)
- InfiniBand

You can use these RDMA interfaces for high-speed object data transfer between S3 clients and the IBM Storage Scale CES nodes.

GPU workload support

S3 over RDMA supports both CPU-based workloads and GPU-based workloads.

In GPU-enabled environments, this feature can use GPUDirect RDMA capabilities when the hardware and software stack support this feature. The GPUDirect RDMA capabilities allow direct data movement between GPU memory and the network interface, reduces unnecessary CPU copies, and improves performance for AI and accelerated computing workloads.

Deployment considerations

To use S3 over RDMA successfully, you must configure and validate RDMA interfaces across all participating CES nodes on an IBM Storage Scale cluster. For more information, see [Configuring InfiniBand and RDMA networks](#). To enable RDMA transport with the S3 service, set the `S3_RDMA_ENABLED` parameter with the **mms3** command.

The RDMA IP addresses for the S3 service are configured by using the `host_customization` option of the **mms3** command and are used to establish RDMA communication channels for object data transfer.

Note: RDMA IP addresses manage S3 data transfers, while CES IP addresses (TCP/IP based) manage S3 request or response headers for transfers.

Make sure that the following requirements are met:

- RDMA interfaces are active and operational.
- RoCE or InfiniBand fabrics are correctly configured.
- Required RDMA connectivity exists between all participating CES nodes on an IBM Storage Scale cluster.
- RDMA kernel modules and drivers are installed and validated.

Installation of NVIDIA CUDA toolkit and cuObject libraries

To use S3 over RDMA, you need to install NVIDIA cuObject libraries. For more information, see [NVIDIA cuObject: GPUDirect Storage for Objects](#).

Install the cuObject library on the S3 client systems and IBM Storage Scale CES cluster nodes. Install the CUDA toolkit to install the cuObject client library. GDS enables the CUDA developer to copy data directly between IBM Storage Scale storage and the GPU memory by using RDMA. GDS eliminates the buffer copies in system (CPU) memory, bypasses the CPU memory, and can place data directly into the GPU application memory. GDS delivers benefits such as increased data transfer rate, lower latency, and reduced CPU usage.

- To download the CUDA Toolkit, see [CUDA downloads](#).
- For CUDA installation details, see [NVIDIA GPUDirect Storage Installation and Troubleshooting Guide](#).
- For information about the CUDA configuration, see the *Configuring Redirect Storage* section in the *IBM Storage Scale: Administration Guide*.
- For information about the NVIDIA cuObject server library installation on the IBM Storage Scale cluster nodes, see [cuObject server Downloads](#).

Important: At IBM, the following versions were used for testing:

- CUDA toolkit: v13.2
- cuObject server library: v1.0.0

Limitations

- S3 over RDMA does not support High Availability (HA) configurations.

If an RDMA communication path is not available, automatic failover for RDMA traffic is not provided as part of this RPQ feature. So, plan deployments and make sure that operational monitoring of RDMA network components.

- All RDMA IP addresses specified in the `host_customization` configuration must remain reachable and available for successful data transfer operations.

If one or more configured RDMA IPs are not available, S3 data transfer operations might fail or experience degraded functions.

Recommendation:

- Use stable RDMA network configurations.
- Avoid dynamic IP reassignment.
- Continuously monitor RDMA link health.
- Validate end-to-end RDMA connectivity before workload deployment.

Swift Object storage support overview

Important:

- CES Swift Object protocol feature is not supported from IBM Storage Scale 5.2.0 onwards.
- IBM Storage Scale 5.1.8 is the last release that has CES Swift Object protocol.
- IBM Storage Scale 5.2.0 will tolerate the update of a CES node from IBM Storage Scale 5.1.8.
 - *Tolerate* means:
 - The CES node will be updated to 5.2.0.
 - Swift Object support will not be updated as part of the 5.2.0 update.
 - You may continue to use the version of Swift Object protocol that was provided in IBM Storage Scale 5.1.8 on the CES 5.2.0 node.
 - IBM will provide usage and known defect support for the version of Swift Object that was provided in IBM Storage Scale 5.1.8 until you migrate to a supported object solution that IBM Storage Scale provides.
 - CES Swift Object is replaced with IBM Storage Scale S3. For more details, refer the *S3 support overview* section the *IBM Storage Scale: Concepts, Planning, and Installation Guide*.

- For more information about Swift Object in IBM Storage Scale, refer to the [IBM Storage Scale 5.2.0 documentation](#).
- Contact IBM for further details and migration planning.

Secure communication between the proxy server and other backend servers

Use this feature to establish secure communication between the proxy server and the backend Swift Object Storage servers.

By default, object-server, object-server-sof, container-server, and account-server do not have authentication for the requests that they are serving. Processes, including the proxy-server that are connecting to these servers over their listening ports, can send requests that can result into updating the database and altering the object data on disk. Extra security between these servers can be enabled. Requesting process signs a request with a secret key kept in `swift.conf`. This key is verified by the serving object, container, or account server. To enable this feature, set:

```
mmobj config change --ccrfile swift.conf --section node_communication --property secure --value true
```

The signing middleware is added to proxy-server and the validating middleware is added to object-server, object-server-sof, container-server, and account-server. If the secret key is not present in `swift.conf`, it is randomly chosen and set to `key secure_communication_secret` under `node_communication` section. In a multi-region environment, this key must be reset and kept common in all the clusters.

To revert to the original configuration, set:

```
mmobj config change --ccrfile swift.conf --section node_communication --property secure --value false
```

Note: Disable SSH access on the protocol nodes on the IBM Storage Scale cluster for the users that have the same UID and GID as the local swift user.


Swift Object heatmap data tiering

Swift Object heatmap data tiering policies can be applied to data that is frequently accessed.

For more information, see the topic *File heat: Tracking file access temperature* in the *IBM Storage Scale: Administration Guide*.

With the Swift Object heatmap data tiering policies, frequently accessed Swift Objects are stored in higher-performing storage pools such as the SSD-based pool. The policy also stores Swift Objects that are not accessed frequently in slower disk-based storage pools.

Note: Data that is not required for regular access is moved from slower disk-based storage pools to tapes.



The Swift Object heatmap policy

The Swift Object heatmap policy can be applied to Swift Object data files in the file system.

The policy manages the placement of Swift Objects in the gold, silver, and system pools. The policy also specifies the maximum threshold so that the capacity in the storage tiers is not over-utilized.

Enabling heatmap tracking

Enable heatmap tracking on the IBM Storage Scale file system by running the following command:

```
mmchconfig fileheatperiodminutes=1440,fileheatlosspercent=10
mmchconfig: Command successfully completed
mmchconfig: Propagating the cluster configuration data to all
affected nodes. This is an asynchronous process.
```

Active File Management

Active File Management (AFM) is a feature of IBM Storage Scale. You can use AFM to share data across clusters.

Introduction to Active File Management (AFM)

Active File Management (AFM) enables sharing of data across clusters, even if the networks are unreliable or have high latency.

The following figure is a sample of an AFM relationship.




Figure 8. Sample of an AFM relationship

You can use AFM to create associations between IBM Storage Scale clusters or between IBM Storage Scale clusters and NFS data source. With AFM, you can implement a single namespace view across sites around the world by making your global namespace truly global.

By using AFM, you can build a common namespace across locations, and automate the flow of file data. You can duplicate data for disaster recovery purposes without suffering from wide area network (WAN) latencies.

AFM can be enabled on GPFS-independent filesets only. A dependent fileset can be linked under an AFM fileset, but only up to one level below the AFM-independent fileset. The dependent fileset does not allow nested dependent filesets under the AFM-independent fileset.

Individual files in the AFM filesets can be compressed. Compressing files saves disk space. For more information, see *File compression* in the *IBM Storage Scale: Administration Guide*.

Namespace replication with AFM occurs asynchronously so that applications can operate continuously on an AFM fileset without network bandwidth constraints.

Note: AFM does not offer any feature to check consistency of files across source and destination. However, after files are replicated, you can use any third-party utility to check consistency.

AFM concepts

The following topics provide an overview of AFM and describe the concepts.

Cache and home

An AFM fileset can be enabled on a GPFS-independent fileset. Each fileset has a distinct set of AFM attributes. An IBM Storage Scale cluster that contains AFM filesets is called a cache cluster. A cache cluster has a relationship with another remote site called the home, where either the cache or the home can be the data source or destination.

AFM constantly maintains an active relationship between the cache and the home. Changes are managed per fileset results in modular, scalable architecture capable of supporting billions of files and petabytes of data. Each AFM-enabled fileset is associated with a single home path.

AFM uses an NFSv3, NFSv4, or NSD (GPFS multi-cluster) protocol for the communication between the home and cache sites. A home export path is either an NFSv3 or NFSv4 exported path or a multi-cluster or remote file system, which is mounted at the IBM Storage Scale cache cluster. This path is used by an AFM or AFM-DR fileset as a target path for data synchronization between sites. For AFM RO-mode filesets,

the target path at the home NFS server can be exported as 'Read-Only' or 'Read/Write'. However, for AFM LU/SW/IW and AFM-DR mode filesets, the target NFS export path must be 'Read/Write'.

Example:

- Export as Read-Only for an AFM RO-mode fileset.

```
/ibm/gpfs0/homeDataSource GatewayIP/  
*(ro,nohide,insecure,no_subtree_check,sync,no_root_squash,fsid=1001)
```

- Export as Read/Write for an AFM-LU/SW/IW and AFM-DR mode fileset.

```
/ibm/gpfs0/homeDataSource GatewayIP/  
*(rw,nohide,insecure,no_subtree_check,sync,no_root_squash,fsid=1001)
```

Architecturally, AFM works with any NFS home export. However, if the home is a GPFS file system then AFM supports replicating ACLs, EA, and sparse files. You must issue the **mmafmconfig enable <ExportPath>** command on the home GPFS cluster to enable files replication.

Note: If you do not run the **mmafmconfig enable** command and configure an AFM relationship, ACLs and extended attributes are not supported. Sparse file information is not maintained. The following error message appears in the `/var/adm/ras/mmfs.log.latest` file: `AFM: AFM is not enabled for fileset cachefs in filesystem gpfs. Some features will not be supported, see documentation for enabling AFM and unsupported features.`

Any cluster can be a home cluster, a cache cluster, or both. In typical setup, a home is in an IBM Storage Scale cluster and a cache is defined in another IBM Storage Scale cluster. Multiple AFM-enabled filesets can be defined in one cache cluster, and each cache has a relationship with targets with a home, or different cluster.

In IW, RO, and LU modes, multiple caches might point to the same home. But in SW mode, only one-to-one relationship between cache and home is supported. AFM can also be configured as a subscription service, where home is the data feed and all caches can subscribe to this data feed.

Within a single cache cluster, application nodes experience POSIX semantics. File locking across cache and home is not supported.

When you perform operations on AFM filesets, ensure that the operations are supported on home over the chosen protocol. Because the operations that are done from cache are replayed on the remote as normal file system operations. When you use the NSD protocol with UID remapping, operations such as `chown` (change ownership) are not supported.

Caching modes

The caching modes in AFM are Read-Only (RO), Single-Writer (SW), Local-Update (LU), and Independent-Writer (IW).

- In a Read-Only mode, files in the cache fileset cannot be modified.
- In a Single-Writer mode, cache pushes all changes to home but never checks home for any updates.
- Independent-Writer mode allows both reads and writes, it pushes changes to home and checks home for file updates.
- Local-Update is read only but files can be changed in the cache fileset, though the changes are never pushed to home and any changed file no longer checks home for updated versions.

Communication between cache and home

AFM uses the NFSv3, NFSv4, and NSD protocols for communication between clusters. AFM supports both kernel NFS (NFSv3, NFSv4.1, and NFSv4.2) and CES NFS (Ganesha) (NFSv3 and NFSv4.1).

Communication for caching between clusters is performed by one or more nodes on the cache cluster that are designated as gateway nodes by using the **mmchnode** command. The target path on the home server must be NFS-exported on one or more nodes in the home cluster, or the home file system must be mounted on the cache cluster by using the NSD protocol. For filesets in the AFM RO mode, the target path

at home NFS export can be a read-only or read/write export. The exported path or the home-mounted path of the home cluster is used in the cache cluster as the target while creating a cache fileset.

You can export paths with Kerberos-enabled security levels from home to cache. At cache, configure Kerberos clients at gateway nodes and run the **mmchconfig afmEnableNFSec=yes -i** command. Gateway nodes can mount Kerberos-enabled home exports with security levels like `sys`, `krb5`, `krb5i`, `krb5p` after the **afmEnableNFSec** parameter is set to `yes` at the cluster level.

To configure Kerberos for AFM-DR filesets, the user must enable the read/write access to the AFM-DR secondary. The read/write access can be enabled by setting the **afmSecondaryRW** parameter value to `'yes'`.

Run the following command at the AFM-DR secondary:

```
# mmchconfig afmSecondaryRW=yes -i
```

AFM maintains a relationship between cache and home by monitoring home availability and reconnects cache and home.

Note: If a home fileset is removed or deleted, you must remove the export path first and restart NFS at home.

The backend protocol - NFS versus NSD

The NSD protocol is a stateful protocol. The NFSv3 protocol is a stateless protocol, which is resilient to the low bandwidth and lossy networks.

The current recommended transport protocol for AFM data transfers is NFS due to the tolerance of NFS to unstable network connections. It is recommended to use NFS first and shift to the NSD protocol only if NFS does not meet the performance requirements even with multiple primary gateways and use of parallel data transfers. The implications of using the NSD protocol between the cache and home cluster are:

1. Network availability fluctuations and instability can affect the NSD protocol connection to the home on the cache cluster primary gateways. This network issue can lead to frequent data access interruptions from the home cluster, and can even cause the connection to the home cluster to stop responding. In these cases, it might be necessary to restart the GPFS daemon on the primary gateway, and possibly even restart the primary gateway server.
2. IBM Storage Scale instability issues on the home cluster can affect the cache cluster. The AFM fileset in the cache cluster does not respond because of the instability, and you might also need to restart the IBM Storage Scale service on the home and cache clusters.

For more information about setting up primary gateway nodes that communicate with multiple NFS servers at home, see [“Parallel data transfers” on page 77](#).

The following table summarizes the differences between NSD and NFS protocols on various parameters:

Table 7. Comparison between NSD and NFS protocols		
	NSD	NFS
Usability	Customers are familiar with NSD use in multi-cluster environments. Configuration does not require NFS knowledge or tuning, but requires NSD tuning.	Configuration requires NFS knowledge and performance tuning for both NFS and TCP over WAN.

Table 7. Comparison between NSD and NFS protocols (continued)

	NSD	NFS
Performance	By default, uses all primary gateway nodes for parallel data transfers. Large file data transfer performance is better than NFS from a single primary gateway node as it can use the inherent parallelism of striping to multiple NSDs.	Parallel data transfers can be achieved by creating mapping between primary gateway nodes and NFS servers on the home. In summary, while both NFS and NSD can do similar forms of parallelism, generally NSD achieves higher performance.
Security	Encryption is built in, which can be turned on optionally.	Supports Kerberos-enabled exported paths from home to cache. The afmEnableNFSSec parameter must be set to yes on cache.
Firewall configuration	Special ports might not be required.	Must be configured for the traffic to pass through.
Stability	Performs well if the network is stable and has low latency.	More resilient to failures within the network such as packet drops that readily happen over WAN and it is more resilient, protecting the cache cluster from being affected by home cluster issues.

Considerations when you use the NSD protocol for AFM data transfers

The NSD protocol is more sensitive to packet drops and network latency than NFS. If the network does not respond, or if the packets are dropped, the NSD mount on cache cluster stops responding, causing the cache cluster also to stop responding. More causes of issues when the NSD protocol is used:

- Deadlock in the home cluster - This deadlock might cause the NSD mounts on the cache cluster to stop responding for some time. Due to a non-responsive NSD mount, AFM fileset at cache that uses these NSD mounts as target might be in the 'unmounted' state. After the home cluster is responsive, the home cluster tries queued operations again.
- Cluster reconfiguration or higher resource consumption on the home cluster - This configuration and resource consumption might cause a temporary loss of communication between home and cache cluster. If the home cluster does not respond within AFM wait timeout intervals, AFM filesets at cache that use these NSD mounts as target might be in the 'unmounted' state. After the home cluster is responsive, the home cluster tries queued operations again.
- When a new primary gateway node joins the cluster, the old primary gateway node transfers the fileset to a new primary gateway node. If the remote file system is not mounted on the new primary gateway node, the fileset remains in the 'unmounted' state. After the remote file system is mounted at gateway node, the fileset automatically moves to the Active state.
- Remote file system cannot be unmounted unless replication is stopped, or primary gateway node is restarted. AFM puts a hold on remote mount, not allowing the file system to be unmounted.
- Creating an AFM association, by using GPFS protocol, to the same local file system is not supported.

Note:

- If the NSD mount on the gateway node is unresponsive, AFM does not synchronize data with home or secondary. The file system might be unmounted on the gateway node. A message `AFM: Remote filesystem remotefs is panicked due to unresponsive messages on fileset <fileset_name>`, re-mount the filesystem after it becomes responsive. `mmcommon preunmount` invoked. File system: `fs1` Reason: `SGPanic` is written to `mmfs.log`. After the home or secondary is responsive, you must restore the NSD mount on the gateway node.

- When NSD backend is used, **afmDIO** parameter is set to 0 by default. For NFS backend, **afmDIO** parameter value must be set to 2.

Nonstandard NFS ports for AFM

Active File Management (AFM) uses port 2049 by default for NFS communication between cache and home (or primary and secondary) clusters. However, in environments where network policies restrict this port, AFM supports configuration by using nondefault NFS ports.

AFM fileset mode with an NFS backend can operate on any port in the range 1024-65535, if:

- The nondefault port is configured on the home/export cluster in the `/etc/nfs.conf` file under the `[nfsd]` section.
- The same port is specified in the AFM target definition or mapping configuration.

This configuration ensures that AFM communication operations, including mounting, data transfer, and metadata synchronization, continue to function in secured or segmented network environments.

Nonstandard ports are supported for:

- NFSv3 and NFSv4
- IPv4 and IPv6

For more information about configuration, see the *Configuring AFM with nonstandard NFS ports* section in the *IBM Storage Scale: Administration Guide*.

Primary gateway and afmHashVersion

Each cache fileset in a cluster is served by one of the nodes that is designated as a gateway in the cluster. The gateway node that is mapped to a fileset is called the primary gateway of the fileset. The primary gateway acts as the owner of the fileset and communicates with the home cluster.

All other nodes in the cluster, including other gateways, become the application nodes of the fileset. Therefore, any node in the cache cluster can function as a gateway node and an application node for different filesets based on configuration of the node.

Application nodes communicate with the primary gateway for a fileset through internal network requests. The gateway function is highly available and can be scaled out. When a gateway node fails, all cache filesets that are owned by this gateway node are moved to another gateway node. This gateway node runs AFM recovery and takes over as the primary gateway of the filesets. When the old primary gateway returns after node failure, the original primary gateway resumes managing its fileset, and AFM transfers the queues for all the filesets from the current primary gateway to the old primary gateway.

The primary gateway can be configured to take help from other gateway nodes for parallel data movement of large files to and from home, if the home is an IBM Storage Scale cluster for increasing performance during data transfer. It is not recommended to assign the manager role to the gateway node because the assigned role can hamper the gateway node performance. The manager role can be assigned to other non-gateway nodes.

The gateway node designation is supported only on the Linux operating system. The Windows or AIX nodes cannot be designated as gateway.

The recommended setup configuration of a primary gateway node is as follows:

- The memory and CPU requirement for a gateway node depends on the number of assigned AFM or AFM-DR filesets and files inside the filesets. It is recommended to have a minimum 128 GB of memory for the gateway node.
- Up to 20 AFM or AFM-DR filesets can be assigned to a gateway node.
- A gateway node is a dedicated node in the cluster without any other designations such as NSD server, CES protocol nodes, quorum, and manager.
- The following table contains the cluster level parameters:

Parameter	Value
Pagepool	8G
afmHardMemThreshold	40G
afmNumFlushThreads	8
afmDIO	2
maxFilestoCache	10000
afmMaxParallelRecoveries	3

To set these parameters, use the **mmchconfig** command. For example,

```
# mmchconfig pagepool=8G
```

For more information about configuration parameters, see *Parameters for performance tuning and optimization* in *IBM Storage Scale: Administration Guide*.

This configuration recommendation is based on observations in a controlled test environment by running moderate or reasonable workload. The parameters values might vary based on the setup or workload.

Note: AFM gateway node is licensed as a server node.

See “Parallel data transfers” on page 77 to set up gateway nodes that communicate to multiple NFS servers at home.

afmHashVersion

In a cluster with multiple gateway nodes and AFM cache filesets, AFM uses a hashing algorithm to elect the primary gateway for each of the filesets. This algorithm can be selected by changing the cluster level config tunable **afmHashVersion**. Issue the following command to change the hashing algorithm.

```
# mmchconfig afmHashVersion=version
```

Note: You need to shut down the cluster and then issue this command because you cannot set **afmHashVersion** on an active cluster.

AFM Hash Version 2

The hashing algorithm is set by default on an IBM Storage Scale 4.1 cluster. This version is based on the static fileset ID mapping. On an upgraded cluster, the old hashing algorithm is effective by default. You can change the hashing version by changing the value of the **afmHashVersion** tunable.

AFM Hash Version 4

This version is an improved hash version '2'. If you have many filesets and if these filesets are not evenly distributed across gateway nodes, you can set the **afmHashVersion** parameter to 4 for the improvement. This hashing version is dynamic, which ensures balanced mapping of AFM filesets and gateway nodes.

Note: Do not set the **afmHashVersion** parameter to 4, if AFM has heavy data transfer operations from the gateway node to filesets. You are not recommended to set this version because you might face some deadlock issues in some cases. You can choose version '5' over version '4'.

AFM Hash Version 5

On the new cluster, the AFM Hash Version 5 is the default version. For better load balancing of AFM or AFM-DR filesets across gateway nodes, you can set the **afmHashVersion** parameter to '5'. This version is recommended for the load balancing of the filesets. AFM hash version '5' does not cause any known performance degradation in comparison with earlier versions. This option must set only at the cache or primary cluster and at the client cluster if it has remote mounts, that is, the AFM cache cluster. However,

do not set this option at the home or secondary cluster because the home or primary cluster does not have AFM or AFM-DR filesets.

AFM hash version 5 supports manual assignment of the selected gateway node to an AFM or AFM-DR fileset by using the **mmafmctl stop** command and the **mmafmctl start** command. You can change the default assigned gateway node of an AFM or AFM-DR fileset by using the **mmchfileset** or **mmcrfileset** command. When the specified gateway node is down or not available, then AFM internally assigns a gateway node from the available gateway node list to the fileset. For filesets that do not have the **afmGateway** parameter set, the gateway node is assigned by using the hash version 2.

To assign the **afmGateway** parameter to a fileset, see the *mmchfileset* command or *mmcrfileset* command in the *IBM Storage Scale: Command and Programming Reference Guide*.

Stop replication on an AFM or AFM-DR fileset before you set the **afmGateway** parameter and start replication after the parameter is set.

Complete the following steps to assign a gateway node by using the **mmchfileset** command:

1. Stop replication on an AFM or AFM-DR fileset by issuing the following command:

```
# mmafmctl fs stop -j AFM/AFM-DR Fileset
```

2. Set the **afmGateway** parameter by issuing the following command:

```
# mmchfileset fs AFM/AFM-DR Fileset -p afmGateway=NewGatewayNode
```

3. Start replication on the AFM or AFM-DR fileset by issuing the following command:

```
# mmafmctl fs start -j AFM/AFM-DR Fileset
```

To assign a gateway node by using the **mmcrfileset** command, issue the following command:

```
# mmcrfileset fs fileset -p afmMode=mode,afmTarget=Target,afmGateway=Gateway --inode-space new
```

For more information about start and stop replication on a fileset, see [“Stop and start replication on a fileset”](#) on page 118.

Note:

- To set the **afmHashVersion** parameter, IBM Storage Scale cluster must be shut down and start after this cluster level tunable is set for a new hash algorithm to take effect.
- To change the **afmHashVersion** parameter to ‘5’, all the nodes in the AFM cache and the client cluster must be upgraded to the minimum 5.0.2 level.

Global namespace

You can combine the home and cache entities to create a global namespace.

Any client node can use the same path to connect to the data within any of the IBM Storage Scale clusters that are part of the namespace.

In such a global namespace, the following AFM features can improve application performance on a remote site:

1. When a file is being read into a cache, the data can be read after it arrives in the cache.
2. Multiple cache filesets can share a single file system. Data can be transferred between sites in parallel.

AFM also performs on networks that are unreliable, or have high latency. The following example is of a global namespace, which implemented by using AFM, with three different sites. An IBM Storage Scale client node from any site sees all of the data from all of the sites. Each site has a single file system. Each site is the home for two of the sub directories and cache filesets that point to the data that originates at the other sites. Every node in all three clusters has direct access to the global namespace.

The following figure shows global namespace that is implemented by using AFM.




Figure 9. Global namespace implemented by using AFM

Revalidation

As users traverse the directory tree of an AFM cache fileset, files and directory metadata information from the home cluster is checked and updated as needed on the cache cluster. This process is called AFM revalidation.

Revalidation performance is dependent upon network latency and bandwidth available between the cache and the home. Revalidations are done per node, and not per fileset. If a file or directory is revalidated from one node on the cache cluster, the same fileset goes through another revalidation operation when accessed from another node on the same cache cluster. You can modify the refresh intervals by using the following command:

```
mmchfileset fs1 sw1 -p afmFileOpenRefreshInterval=160
```

In this example, the **afmFileOpenRefreshInterval** parameter is set to 160 for the **sw1** fileset in the **fs1** file system.

Revalidation intervals can be adjusted to support the workload and network latency. Setting a parameter by using the **mmchconfig** command sets the default for all filesets. Parameters set by using the **mmchfileset** command affect only a particular fileset and override the global defaults. You can enable, modify, or disable any of the intervals based on the application needs, though it is recommended to use default values for most cases.

If file or directory refresh intervals are disabled for a fileset, the refresh intervals can be enabled by using the **mmchfileset** command. Enabling requires the fileset to be unlinked, and then linked again.

For more information, see the *mmchfileset* command in the *IBM Storage Scale: Command and Programming Reference Guide*.

It is recommended not to set the revalidation intervals to 0, as a revalidation request is continuously sent to home, thus resulting in performance degradation. You must set the revalidation interval to as large as possible, depending on how frequently home gets updated, and at what interval the cache needs the updated data.

For more information about revalidation intervals, see the *mmcrfileset* command in the *IBM Storage Scale: Command and Programming Reference Guide*.

The revalidation intervals are defined by the following configuration parameters. These parameters are tunable at the cluster and fileset level and can be changed by using the **mmchconfig** command and the **mmchfileset** command at the cluster and the file level:

1. **afmFileLookupRefreshInterval**: The frequency of revalidation that is triggered by a lookup operation on a file such as ls or stat, from the cache.
2. **afmDirLookupRefreshInterval**: The frequency of revalidation that is triggered by a lookup operation on a directory such as ls or stat, in the cache.
3. **afmFileOpenRefreshInterval**: The frequency of revalidations that are triggered by read or write operations on a file in the cache. Open requests on that file are served from the cache fileset until the **afmFileOpenRefreshInterval** expires after which the open requests are sent to home.
4. **afmDirOpenRefreshInterval**: The frequency of revalidations that are triggered by read or update operations on a directory from the cache. Open requests on files or subdirectories on that directory are served from the cache fileset until the **afmDirOpenRefreshInterval** expires after which the open requests are sent to home.
5. **afmRevalOpWaitTimeout**: The time for which AFM waits for completion of revalidation to get response from the home cluster. The **afmRevalOpWaitTimeout** parameter can be set only for a cluster and not for an individual fileset.

RO, LU, and IW modes of filesets revalidate regularly with home. The SW mode populates metadata only one time during first access and does not revalidate with home thereafter. To revalidate, AFM sends a message to the home cluster to find out whether the metadata of that file or directory is modified since the last time it was revalidated. If so, the latest file metadata or data information, depending on the type of request, at home is reflected on the cache.

Revalidation in asynchronous mode

Synchronization with home before revalidation delays response to the applications querying data. To overcome this delay, you can perform cache data refresh operation in asynchronous mode. You must set the **afmRefreshAsync** parameter to 'yes'.

For more information about the **afmRefreshAsync** parameter, see *mmchconfig* command in the *IBM Storage Scale: Command and Programming Reference Guide* and *Configuration parameters for AFM* in the *IBM Storage Scale: Administration Guide*.

Cached and uncached files

A `readdir` operation on a directory populates the metadata of the directory in the cache, but this operation does not populate contents of each file within the directory. A read operation on file generates a request to the home to make contents available in the cache. The file contents do not need to be in the cache to start reading it.

AFM allows data to be pre-populated before actual read operation by using the **mmafmctl prefetch** command. For more information about pre-populating data, see [“Prefetch” on page 82](#).

A file whose contents are available in the cache is called a cached file. A file whose contents are not yet present in the cache is called an uncached file. An uncached file cannot be evicted, resynched with home, or failed over to a new home. See the sections on Asynchronous operations and delay, for AFM eviction and syncing to home.

Files that have all blocks read, or the entire file contents that is fetched, are marked as cached. AFM does whole file caching by default. By default, reading more than three blocks of a file drives AFM to cache the full file in the background for performance. Sometimes, the whole file might not need to be cached. For example, some applications read only a few bytes of a file to detect the file mime type. In such cases, you can configure partial file read behavior on cache fileset. For more information about partial file caching, see [“Partial file caching” on page 81](#).

When a write operation is performed directly on an uncached file in IW, SW, or LU modes, AFM fetches first the entire file to the cache (cached) by default and later allows the local write to proceed. This behavior is for in place writes within the size of the file. If the append that is being performed to the file is more than the last offset, the behavior is slightly different:

- The append on files in IW or SW filesets appends only data to the file at the cache. This append can leave the rest of the file data from the home as uncached.

- The append on files in the LU mode must mark the file as dirty, and the append disconnects the appended file with an appended file of the cache. So that AFM fetches the entire file even for an append more than the file size.

Likewise, the files that originate in the cache on IW, SW, or LU filesets are marked as cached by default because these files are created and written to at the cache.

Synchronous or asynchronous operations

AFM operations are serviced either synchronously or asynchronously. Reads and revalidations are synchronous operations, and update operations from the cache are asynchronous.

Synchronous operations require an application request to be blocked until the operation completes at the home cluster. File data is available while AFM queues are processed asynchronously in the background.

If synchronous operations depend on the results of one or more updates or asynchronous operations, AFM prioritizes the dependent asynchronous operations before the execution of the synchronous operations. For example, synchronous operations like file lookup cause dependent asynchronous operations to be flushed, overriding asynchronous delay (**afmAsyncDelay**).

AFM deploys a filtering algorithm for most optimal flushing performance by analyzing the queued requests. For example, if a write and a delete are in a queue on the same file, the write request is dropped from the queue and the delete is run at home. Similarly, when **mkdir** and **rmdir** of the same directory name are in a queue, both requests are dropped.

AFM sends data from cache to home as root. If home has fileset quotas at home, it must be set such that it can accommodate all writes from cache.

The **afmRefreshAsync** parameter changes the synchronous behavior. This parameter causes a **readdir**, which must be a synchronous command, to behave asynchronously. That is, like a follow-up **readdir** for an already cached directory. The **Async readdir** goes in the queue like any other asynchronous operation, where the cache serves existing readdir contents locally.

This **afmRefreshAsync** is useful for IW, RO, and non-dirty or unchanged LU modes that need to perform revalidations frequently with the home. The applications must not be affected with such a **readdir** on the application path.

Asynchronous delay

All update operations from the writable cache filesets are on the primary gateway. Queues in the primary gateway are pushed to home asynchronously based on the **afmAsyncDelay** interval.

This interval can be modified by using the **mmchfileset** command. You can force a flush of all the pending updates without waiting for asynchronous delay by using the **mmafmctl** command with the **flushPending** option. The **flushPending** option can be set at the fileset or file system level, or with a file that contains the list of files to be flushed.

For more information, see **mmafmctl** command in *IBM Storage Scale: Command and Programming Reference Guide*.

Operations with AFM modes

The operations with AFM modes are read only (RO), single writer (SW), local updates (LU), and independent writer (IW).

1. Read only (RO)

The following figure illustrates the RO mode.




Figure 10. Read only mode

In this mode, data in the cache is read-only. Creating or modifying files in the cache fileset is not allowed. If a file is renamed at the home for an RO fileset, the file is re-created in cache and is assigned a new inode number in the cache. If the file is in use by an application while it is re-created (deleted and re-created with the same name) at home, it gets re-created in cache.

2. Single writer (SW)

The following figure illustrates the SW mode.



Figure 11. Single writer mode

In this mode, only the cache fileset does all the writing and the cache does not check home for file or directory updates. The administrator needs to guarantee no writes are done in the home cluster. AFM does not enforce this check.

An SW home can have some pre-existing data. An SW cache can cache and update this data. Update of an uncached file from SW home caches the file. However, the truncate or append operations on an uncached file does not fetch the contents of the file into the cache, but queues the truncate or append operation to the home.

3. Local updates (LU)

This mode is similar to the RO mode, although you can create and modify files in the cache fileset. Updates in the cache are considered local to the cache and get decoupled from the corresponding file at home. Local updates are never pushed back to home. After modifications in a file, the file is no longer compared to the version at home during revalidation to verify that it is up-to-date. Changes of this file on home do not have an impact on the cached copy of the file and vice versa.

Behaviors with local files: In AFM, LU mode files have one of the following states:

- **Uncached:** Files on the home for which metadata but no data is copied into the cache are shown in the cache as uncached. The file is not resident on cache, but only on the home. Changes on the home are reflected in the cache.
- **Cached:** If an uncached file is read in the cache or pre-fetched, the state of the file changes to cached. In the cached state, all changes to the file on the home are reflected in the cache. The file is resident on the cache.
- **Local:** File data or metadata that is modified at cache becomes local to the cache. The cached files relationship to the file in the home is broken. Changes on the home are not reflected in the cache anymore and file changes are not copied to the home.

Operations in the cache that trigger the transitions between these states are shown below:




Figure 12. Behaviors with local files

The following tables summarize AFM LU mode behavior with local files:

File operation at home	Uncached file	Cached file	Local file
New Create	Reflects	Reflects	Reflects
Change data or attributes/Rename/Delete	Reflects	Reflects	Does not reflect

File operation at home	Uncached file	Cached file	Local file
File lookup revalidation	Pulls the metadata	Pulls the latest metadata	Does not pull the latest metadata
Read	Pulls the data	Pulls the latest data	Shows data in the local file – does not pull the latest data from the home
Prefetch	Prefetches the data	Prefetches the latest data	Does not prefetch the latest data from home
Change data	Pulls the data into cache and updates the changes	Pulls the latest to cache and updates the changes	Updates the local copy of the file and does not pull the latest data from the home before the update
Change attributes/metadata except delete EA	Updates the file and marks file as local	Updates the file and marks file as local	Updates the file – does not pull the latest metadata from the home before the update

Directories in a local-update cache:

Directories can become local in the LU mode with the following directory operations from the cache:

- Deleting files or sub-directories in the cache directory
- Creating new files or sub-directories in the cache directory
- Renaming files or sub-directories in the cache directory

The following file operations on cache do not cause directories on the cache to become local:

- Update the uncached or cached file
- Add attributes to the file
- Migrate the cached or local files to tape
- Recall of the cached or local files from tape

A local directory in an LU cache can contain local, cached, and uncached files and directories. A local directory is not revalidated with home. Therefore, operations on the directory at home do not show up on the cache.

The following table summarizes AFM LU mode behavior with local directories:

Operation at home	Cache behavior - Normal directory	Cache behavior - Local directory
Create a file or subdirectory	Reflects in cache after next directory lookup	Does not show in cache after next directory lookup. Contents can be read or prefetched if the file name is known in the cache. After read or prefetch, the new file or directory shows up in the cache for future file lookup
Update data or attributes of an existing file or subdirectory	Reflects in cache after next file lookup	Reflects in cache after next file lookup
Rename an existing file or subdirectory	Reflects in cache after next file lookup	Does not show in cache after next directory lookup. On reading the file, contents from renamed file that is fetched into the cache into the old file name, old file name always prevails in the cache. New file name can arrive at the cache with prefetch, in which case both names will co-exist and point to the same inode at home.
Delete an existing file or subdirectory	Reflects in the cache after next directory lookup	Deleted upon the next file lookup

Directory operation from cache	Cache behavior - Normal directory	Cache behavior - Local directory
Directory revalidation	Pulls the latest metadata	Does not pull the latest metadata
Read a non-local file in the directory	Pulls the latest data	Pulls the latest data
Prefetch a non-local file	Prefetches the latest data	Prefetches the latest data
Create a file or directory	Updates and marks the directory local	Updates the local directory
Change data of a non-local file	Pulls the data into the cache, updates the file and marks file as local	Pulls the data into the cache, updates the file and marks the file as local
Change metadata of a non-local file except delete of EA	Updates the file and marks the file as local	Updates the file and marks the file as local

Appending to, truncating, or writing to an uncached file in LU mode fetches the entire file to the cache before making the change locally.

If you expect revalidation, change the LU fileset root directory with caution as this might cause the fileset to be marked as local, and context with the home is lost. For example, running **chmod** or **chown**

of the LU fileset root directory causes all sub directories from fileset root to be out of synchronization with the home.

4. Independent writer (IW)

The following figure illustrates the IW mode.




Figure 13. Independent writer mode

This mode allows multiple cache filesets to point to the same home. There is no synchronous locking between clusters while updating home. Each cache reads from home and makes updates to the home independently of each other, based on the revalidation intervals and asynchronous delay.

This mode is used to access different files from each IW cache site, such as, the unique users at each site updating files in their home directory. While this mode allows multiple cache clusters to modify the same set of files, this should be only done by advanced users. This is because there is no locking or ordering between updates. Updates are propagated to the home in an asynchronous manner and can be delayed due to network disconnections. Therefore, conflicting updates from multiple cache sites can cause the data at the home site to be undetermined.

Writing in-place on a pre-existing uncached file at the home pulls the complete contents of the file into cache. However, truncate and append operations on an uncached file do not fetch the contents of the file into cache, but queues the truncate and append operations to home.

When many IW cache filesets point to the same NFS home, the number of NFS threads at home can be tuned for better efficiency. Increasing the revalidation intervals for the IW cache filesets might reduce the frequency of periodic refreshes from home and improve cache performance.

Note: Do not create hard links at home or in the cache when IW is used, as IW fallback does not preserve hard links.

The following file attributes are maintained locally in cache: Control attributes, direct I/O, replication factors, fileset quotas, storage pool flags, special file types such as FIFO, socket, block, character device. Hard links can be created at cache. Hard links at home are displayed as individual files in cache.

Filesets to the same home target

The following are the possibilities of filesets that point to home as the target:

- One or more RO filesets can point to the same target.
- One or more LU filesets can point to the same target.
- One or more RO and LU filesets can point to the same target.
- SW, RO/LU filesets can point to the same target.
- Only one SW and one or more RO or LU filesets can point to the same target.
- One or more IW, RO, and LU filesets can point to the same target.
- More than one SW fileset must not point to the same target. More than one SW fileset can technically point to the same target but can be only a writer. The other must behave like a reader fileset.
- One or more IW filesets can point to the same target.
- SW and IW filesets must not point to the same target.

Conversion of mode

An Active File Management (AFM) cache can be converted to another mode in some conditions to meet all types of requirements of data management by unlinking the fileset and by using **mmchfileset -p afmmode** command.

The mode conversion limitations are as follows:

- A single-writer (SW) or independent-writer (IW) fileset with pending requests in the queue cannot be converted.
- Local-update (LU) filesets cannot be converted to any other mode.
- Before you convert a fileset from the IW mode to the SW mode, ensure that all the remaining IW filesets are converted to either RO or LU. This conversion is necessary to avoid home conflicts that are caused by writers from other IW filesets.
- You cannot change the mode or disable AFM when the fileset is linked.

Internal AFM directories

AFM uses `.afm`, `.ptrash`, `.pconflicts`, and `.afm/.afmtrash` as internal directories.

AFM uses the following internal directories.

Important: Do not alter or remove any of these directories.

- The `.afm` directory: Present in the cache and the home.

The **mmafmconfig enable** command creates this directory in the home cluster. The **mmafmconfig disable** command removes the directory from home. The `.afm` directory is created in the cache when you create a fileset, and is valid throughout the lifetime of the cache fileset.

- The `.ptrash` directory: Present in the cache.

The existing `.ptrash` directory is applicable to the RO and LU modes.

For these modes, when a directory along with files or directories is moved to a cache fileset, and the same directory is removed from the home fileset, then the cache fileset directory is moved to the `.ptrash` directory.

- The `.pconflicts` directory: Present in the cache.
- The `.afm/.afmtrash` directory: Present in the home.

This directory is created at the home when the **mmafmconfig enable** command is run, or AFM creates it automatically from the cache after the relationship is initialized. Directories that are conflicting at the home are moved to this cache directory to resume replication successfully.

If an SW or IW mode fileset has deletion conflict for a file, the conflict is resolved by removing the file from the cache and moving it to `.ptrash`. A conflict occurs when a file in the home is removed but still exists in the cache because the cached version has outstanding changes. A file with outstanding changes in the cache not yet copied to home is called dirty. Files that are moved from the `.ptrash` or `.pconflicts` directory by using the **mv** command are treated as local files to the cache fileset and are not queued to the home. However, if these files are copied, for example, by using the **cp** command to the cache, they are queued to the home as new files.

Active File Management (AFM) features

The following sections list features of Active File Management (AFM).

AFM Network File System version 4 support

AFM supports NFSv3 and NFSv4 protocols for communication between the home and the AFM cache. An NFS client must be enabled on the AFM gateway node to replicate data between the home and cache on the gateway node.

NFSv4 support

The AFM gateway node must mount the remote exports by using NFSv3 or NFSv4. AFM does not differentiate between an actual NFS version on mounted remote exports except where ACLs from a third-party file system are migrated by using NFSv4. NFSv4 is more secure and improves the replication performance even on a high latency network.

You can specify an NFS version for the gateway node by changing the **afmNFSVersion** parameter value at the cluster level. When this value is set, AFM communicates between the home and the cache clusters by using specified versions for all NFS targeted filesets on the file systems. The default value of this parameter is 3. The allowed values of the **afmNFSVersion** parameter are listed in the following table:

NFS service	NFS version
Kernel NFS	3, 4.1, 4.2
CES NFS	3, 4.1

You can change an existing value of the **afmNFSVersion** parameter by issuing the following command:

```
mmchconfig afmNFSVersion=4.1 -i
```

The new NFS version takes effect immediately.

Enabling NFSv4 at the AFM fileset level

NFSv4 can be enabled at the cluster level or at the fileset level.

To enable a specific NFS version at the cluster level, set the value of **afmNFSVersion** parameter by issuing the following command:

```
mmchconfig afmNFSVersion=4.1 -i
```

In this example, the configured NFS version is 4.1.

After the version is set, AFM uses the specified NFS version for communication between home or secondary and cache or primary sites for all the filesets across the file systems.

Configuring NFSv4 at the AFM fileset level

You can configure an individual AFM fileset to use NFSv4, so that AFM overrides the value set at cluster level in the **afmNFSVersion** parameter.

Configure **afmNFSV4** to yes or no as needed.

- To configure an individual AFM fileset to use NFSv4, set the **afmNFSV4** parameter to yes:

```
afmNFSV4=yes
```

Note: With this configuration at fileset level, AFM ignores the **afmNFSVersion** value set at the cluster level and uses NFS v4.2 in the AFM fileset. If NFS v4.2 is not available, AFM uses NFS v4.1. If neither NFS v4.2 nor NFS v4.1 are available, the fileset is changed to the unmounted state.

- To change this configuration, so that the AFM fileset uses the NFS version that is specified at cluster level, set the **afmNFSV4** parameter to no:

```
afmNFSV4=no
```

Example: Enabling NFSv4 during fileset creation

To enable NFSv4 on a specific fileset during the time of fileset creation, specify **afmNFSV4=yes** with the **mmcrfileset** command.

```
mmcrfileset fs1 testfset -p afmnfsv4=yes
...
```

Example: Enabling NFSv4 on an existing fileset

To enable NFSv4 on an existing fileset, follow these steps.

1. Stop the fileset.

```
mmafmctl fs1 stop -j <fileset name>
```

2. Specify **afmNFSV4=yes** with the **mmchfileset** command.

```
mmchfileset fs1 <fileset name> -p afmnfsv4=yes
```

3. Restart the fileset.

```
mmafmctl fs1 start -j <fileset name>
```

Example: Disabling NFSv4 on an existing fileset

1. Stop the fileset.

```
mmafmctl fs1 stop -j <fileset name>
```

2. Specify **afmNFSV4=yes** with the **mmchfileset** command.

```
mmchfileset fs1 <fileset name> -p afmnfsv4=no
```

3. Restart the fileset.

```
mmafmctl fs1 start -j <fileset name>
```

For more information, see [“NFSv4 support” on page 70](#).

Migrating existing AFM filesets with an old NFS version as a target

After the **afmNFSVersion** value is changed from 3 to 4, migrate existing AFM filesets with NFS targets by setting NFSv3 to NFSv4. Complete the following steps on all AFM filesets to change an NFS version on all file systems:

1. Check the existing NFS version on all file systems.

```
mmlsconfig afmNFSVersion
```

A sample output is as follows:

```
afmNFSVersion 3
```

2. Change the NFS version to 4.1 on all file systems.

```
mmchconfig afmNFSVersion=4.1 -i
```

3. Stop all AFM filesets, which are using NFS as a target.

```
mmafmctl fs1 stop -j afmNFSFileset1
```

4. Start all AFM filesets, which are using NFS as a target.

```
mmafmctl fs1 start -j afmNFSFileset1
```

An AFM fileset is enabled to work with NFSv4 exports from the home server.

Restriction: AFM filesets can use NFSv3 or NFSv4. You cannot set both NFS versions simultaneously for a cluster.

For more information about the **afmNFSVersion** parameter, see the *mmchconfig command* in *IBM Storage Scale: Command and Programming Reference Guide*. For more information, see the *Enabling AFM Network File System version 4* section in *IBM Storage Scale: Administration Guide*.

Migrating NFSv4 ACL from a third-party file server

After the NFSv4 is configured on a cluster, AFM can also download the NFSv4 ACLs of the files or directories from a third-party file server (non-GPFS) to an IBM Storage Scale AFM fileset. AFM can sync both data and NFSv4 ACLs from the third-party file server to the AFM cache. After the migration is complete, you can convert this cache to a GPFS file system by disabling AFM or you can replicate AFM cache data to another target.

For the migration of NFS v4 ACLs from a non-GPFS or a third-party file server, configure the **afmSyncNFSv4ACL** parameter on the cache. This parameter must be set when home is a non-GPFS home or a third-party file server. This parameter is a cluster level parameter.

```
mmchconfig afmSyncNFSv4ACL=yes -i  
mmlsconfig afmSyncNFSv4ACL
```

AFM pulls the NFS v4 ACLs of files or directories from the non-GPFS or third-party file server.

NFSv4 ACL conversion examples

1. Display ACL that is set on an external file system on the home.

```
# getfacl /ext4/dir1/1.txt
```

A sample output is as follows:

```
getfacl: Removing leading '/' from absolute path names  
# file: ext4/dir1/1.txt  
# owner: root  
# group: root  
user::rw-  
user:user12:rw-  
group::r--  
group:user12:rw-  
mask::rw-  
other::r--
```

2. A single writer AFM mode fileset is created and data is cached. Check the directory contents.

```
# cd /gpfs/gpfs1/sw1
```

```
# ls -l
```

A sample output is as follows:

```
total 0  
-rw-rwxr--+ 1 root root 3 Apr  
-rw-rwxr--+ 1 root root 3 Apr  
8 15:08 1.txt  
8 15:08 2.txt
```

3. Display NFSV4 ACL on the cache by issuing the **getfacl** command.

```
# getfacl 1.txt
```

A sample output is as follows:

```
# file: 1.txt
# owner: root
# group: root
user::rw-
user:user12:rw-
group::r--
mask::rwx
group:user12:rw-
other::r--
```

4. Display NFSV4 ACL on the cache by issuing the **mmgetacl** command.

```
# mmgetacl -k nfs4 1.txt
```

A sample output is as follows:

```
#NFSv4 ACL
#owner:root
#group:root
special:owner@:--x:deny
(-)READ/LIST (-)WRITE/CREATE (-)APPEND/MKDIR (-)SYNCHRONIZE (-)READ_ACL (-)READ_ATTR
(-)READ_NAMED
(-)DELETE
(X)DELETE_CHILD (-)CHOWN
(X)EXEC/SEARCH (-)WRITE_ACL (-)WRITE_ATTR
(-)WRITE_NAMED
special:owner@:rw-c:allow
(X)READ/LIST (X)WRITE/CREATE (X)APPEND/MKDIR (X)SYNCHRONIZE (X)READ_ACL (X)READ_ATTR
(X)READ_NAMED
(-)DELETE
(-)DELETE_CHILD (X)CHOWN
(-)EXEC/SEARCH (X)WRITE_ACL (X)WRITE_ATTR
(X)WRITE_NAMED
user:user12:rw-:allow
(X)READ/LIST (X)WRITE/CREATE (X)APPEND/MKDIR (X)SYNCHRONIZE (X)READ_ACL (X)READ_ATTR
(X)READ_NAMED
(-)DELETE
(-)DELETE_CHILD (-)CHOWN
(X)EXEC/SEARCH (-)WRITE_ACL (-)WRITE_ATTR
(-)WRITE_NAMED
group:user12:rw-:allow
(X)READ/LIST (X)WRITE/CREATE (X)APPEND/MKDIR (X)SYNCHRONIZE (X)READ_ACL (X)READ_ATTR
(X)READ_NAMED
(-)DELETE
(-)DELETE_CHILD (-)CHOWN
(X)EXEC/SEARCH (-)WRITE_ACL (-)WRITE_ATTR
(-)WRITE_NAMED
special:group@:r--:allow
(X)READ/LIST (-)WRITE/CREATE (-)APPEND/MKDIR (X)SYNCHRONIZE (X)READ_ACL (X)READ_ATTR
(X)READ_NAMED
(-)DELETE
(-)DELETE_CHILD (-)CHOWN
(-)EXEC/SEARCH (-)WRITE_ACL (-)WRITE_ATTR
(-)WRITE_NAMED
special:everyone@:r--:allow
(X)READ/LIST (-)WRITE/CREATE (-)APPEND/MKDIR (X)SYNCHRONIZE (X)READ_ACL (X)READ_ATTR
(X)READ_NAMED
(-)DELETE
(-)DELETE_CHILD (-)CHOWN
(-)EXEC/SEARCH (-)WRITE_ACL (-)WRITE_ATTR
(-)WRITE_NAMED
```

Active file management-hierarchical storage management integration for the data migration

Active file management (AFM) replicates data from a hierarchical storage management (HSM)-configured file system on a home cluster to a file system a cache cluster without recalling data, which is migrated to a tape.

Many files with large data might be migrated from the home file system to the tape. These files in the home file system can be in the following states:

- **Resident:** Data is on a disk.

- **Pre-Migrated:** Data is on a disk and a tape.
- **Migrated:** Data is only on a tape.

Instead of recalling data on the tape, IBM Storage Scale integrates only the file or inode metadata by adding **afmIncludeDmapiCache** and **afmIncludeDmapiHome** data management API (DMAPI) flags to locate data on the tape. For this data migration, consider the following points:

- DMAPI is enabled at the cluster level on both the home and cache clusters.
- Both clusters are connected to the same HSM server, so that tape metadata remains consistent and valid on the cache side.

For more information about the flags, see the *mmchconfig* command in the *IBM Storage Scale: Command and Programming Reference Guide*.

Fast create

When the fast create feature is enabled on AFM or AFM-DR filesets, this feature reduces the number of operations that are generated at the time of file creation at the cache. AFM needs to replicate each file creation operation to the home as a separate operation.

Every time when a file is created, operations such as create, truncate, write, chmod, chown, or set ACL/EA are generated at the file system level. Each operation is replicated to the target site as a separate operation. The applications that generate such workload are tar, git, or make, which increases stress on the gateway node to queue all file creation operations as a separate operation for each file creation. Each operation that is performed from an application node for a file creation triggers a separate RPC send or acknowledge to or from the gateway node.

On the clusters where the fast create feature is enabled, AFM does not send all operations. However, AFM combines all operations of a file creation and sends these combined operations to the target. The target performs all operations that are required so that both sites are in sync. This feature helps you to avoid queuing all operations and send only one operation for every file creation. Thus, the fast create feature improves the performance by saving many RPC exchanges between the application node and the gateway node and also saves memory and network bandwidth.

For the performance improvement, the **afmAsyncDelay** parameter must be tuned when the **afmFastCreate** parameter is set on an AFM or AFM-DR fileset. The tuning of the **afmAsyncDelay** parameter depends on the application and network workload. If you set the **afmAsyncDelay** parameter value too high, more memory is consumed on the gateway node. To change the value of the **afmAsyncDelay** parameter, issue the following command:

```
# mmchfileset fs -j fileset -p afmAsyncDelay=Value
```

This feature can be enabled or disabled on any individual AFM or AFM-DR fileset. To enable this feature, all the nodes in the cluster must be upgraded to the latest release level and the latest file system format level. Before you enable this feature on a fileset, you must stop the fileset and then set this parameter. After you enable this feature, you need to start the fileset to continue the replication. For more information about stop and start the fileset, see [“Stop and start replication on a fileset” on page 118](#).

Complete the following steps to enable this feature on an AFM or AFM-DR fileset:

1. Update the cluster release level to the latest level and the file system version to the latest version. For more information about the cluster release level, see [“Completing the upgrade to a new level of IBM Storage Scale” on page 753](#), and for more information about file system version, see *File system format changes between versions of IBM Storage Scale* in the *IBM Storage Scale: Administration Guide*.
2. Ensure that the fileset state is active and operations not are in the queue. To check the fileset state, issue following command:

```
# mmafmctl fs getstate -j fileset
```

3. To stop the fileset, issue the following command:

```
# mmafmctl fs stop -j fileset
```

Note: If you set the **afmFastCreate** parameter without stopping or unlinking a fileset,

```
# mmchfileset fs fileset -p afmFastCreate=yes
```

the following error is displayed:

```
mmchfileset: [E] Fileset cannot be changed, either fileset is linked or not stopped.  
mmchfileset: Command failed. Examine previous error messages to determine cause.
```

4. To enable or disable the fast create at the fileset level, issue the following command:

```
# mmchfileset fs -j fileset -p afmFastCreate=yes|no
```

5. Start the fileset by issuing the following command:

```
# mmafmctl fs start -j fileset
```

6. Ensure that the fast create is enabled successfully by issuing the following command:

```
# mmlsfileset fs fileset -L --afm
```

For example,

```
Filesets in file system '<FS>':  
  
Attributes for fileset <Fileset>:  
=====
```

Status	Linked
Path	<Fileset_Linked_Path>
Id	236
Root inode	30932995
Parent Id	0
Created	Tue Mar 1 07:35:21 2020
Comment	
Inode space	59
Maximum number of inodes	100096
Allocated inodes	100096
Permission change flag	chmodAndSetacl
afm-associated	Yes
Target	nfs://<home>/<Fileset_Linked_Path>
Mode	single-writer
File Lookup Refresh Interval	30 (default)
File Open Refresh Interval	60 (default)
Dir Lookup Refresh Interval	60 (default)
Dir Open Refresh Interval	60 (default)
Async Delay	15 (default)
Last pSnapId	0
Display Home Snapshots	no
Number of Gateway Flush Threads	4
Prefetch Threshold	0 (default)
Eviction Enabled	yes (default)
I/O Flags	263168 (fastCreate)

Note: RPO or peer snapshot feature is not supported when the **afmFastCreate** parameter is set on an AFM or AFM-DR fileset.

```
# mmpsnap fs create -j fileset
```

The following error is displayed:

```
mmpsnap: [E] Peer snapshot creation failed. Error code 22.  
mmpsnap: Command failed. Examine previous error messages to determine cause.
```

Support of linking GPFS dependent fileset

Active File Management (AFM) supports to create and link a dependent fileset inside an independent fileset.

AFM can be enabled only on a GPFS independent fileset. However, the user can create a GPFS dependent fileset by using inode space of an AFM or AFM-DR fileset and link this dependent fileset under the AFM or

AFM-DR fileset for the data synchronization. All data and metadata operations that are generated inside the dependent fileset, which is linked inside AFM fileset, work seamlessly. The user can create data inside the AFM or AFM-DR fileset path and the linked dependent fileset path. AFM synchronizes the created data to the home or secondary fileset under the same path. AFM queues operations that come under the root path of the AFM fileset and the linked path of the dependent fileset and flushes the operations to the home. Also, during the prefetch, AFM can bring data from home dependent fileset path to the AFM fileset that has the dependent fileset that is linked at same path. The path location must be the same at the cache or primary and the home or secondary site.

All I/O requests that are generated under the linked dependent fileset use the same gateway node that is assigned to the AFM or AFM-DR fileset. The gateway node maintains the directory structure during the data synchronization and keeps the home or secondary in sync.

To define an AFM relationship between a cache or primary and a home or secondary, stop the specific AFM or AFM-DR fileset, and then link the dependent fileset at the same junction path under the AFM or AFM-DR fileset. That is, the relative path under the AFM or AFM-DR fileset at both sides must be the same so that all operations generated at one site sync to the appropriate path at the other site. When the dependent fileset is linked inside the AFM or AFM-DR fileset, the user must issue **Start** command to start the replication.

Prerequisite for linking dependent fileset under an AFM or AFM-DR fileset:

1. Set up a dependent fileset at both cache or primary and home or secondary sites and then start replication. Because AFM does not replicate dependent fileset information between cache or primary and home or secondary sites.
2. Ensure that the dependent fileset link path (junction path) is the same relative path to the AFM or AFM-DR fileset at both sites.

That is, the linking of dependent fileset inside AFM or AFM-DR fileset must be identical at both sites.

3. Create and link dependent fileset at a new location to perform failover or **changeSecondary** operation on the AFM or AFM-DR fileset, which has the linked dependent fileset.

Note: Ensure that the failover or change secondary site has an identical path to the cache or primary for the data synchronization before you perform the failover or **changeSecondary** operation.

4. Run the **mmafmconfig enable** command on each dependent fileset that is linked inside an independent fileset at the home/secondary.

To link a dependent fileset inside an AFM or AFM-DR fileset, do the following steps:

1. Create an AFM or AFM-DR fileset and set up an AFM relationship.
2. Create a dependent fileset by using inode space of the created AFM or AFM-DR fileset under which you want to link the dependent fileset at both sites.
3. Stop the specific AFM or AFM-DR fileset at the cache or primary.
4. Link the dependent fileset inside the AFM or AFM-DR fileset at the same relative path at both sites so that the directory structure is the same across the sites.
5. Start the AFM or AFM-DR fileset.

Note: To unlink the dependent fileset, an AFM or AFM-DR fileset must be stopped first.

Example

1. Create an AFM or AFM-DR fileset and set up an AFM relationship.

```
mmcrfileset <fs> <AFM/AFM-DR fileset> -p afmMode=<mode>,afmTarget=<AFMTarget> --inode-space  
New
```

```
mmlinkfileset <fs> <AFM/AFM-DR fileset> -J /< GPFS_PATH >/<AFM|AFM-DR link path>
```

For more information, see the topic *Administering AFM* in the *IBM Storage Scale: Administration Guide*.

2. Create a dependent fileset by using inode space of the created AFM or AFM-DR fileset of the created AFM or AFM-DR fileset.


```
mmcrfileset <fs> <dep-fileset> --inode-space <AFM/AFM-DR Fileset>
```

3. Stop the AFM or AFM-DR fileset at the cache.

```
mmafmctl <fs> stop -j <AFM/ADR Fileset>
```

For more information, see “Stop and start replication on a fileset” on page 118.

4. Link the dependent fileset at the relative junction path at both sites.

```
mmlinkfileset <fs> <dep-fileset> -J <GPFS_PATH/AFM|AFM-DR Path/Dependent Fileset>
```

5. Start the AFM or AFM-DR fileset at the cache or primary.

```
mmafmctl <fs> start -j <AFM/ADR Fileset>
```

For more information, see “Stop and start replication on a fileset” on page 118.

Note: This feature supports linking only the GPFS dependent fileset that is under the AFM or AFM-DR fileset root path. This feature does not allow linking the nested levels of the dependent fileset under the AFM or AFM-DR fileset.

Force flushing contents before Async Delay

Requests in the queue of a writable cache (SW/IW) can be flushed to the home before they get flushed automatically after Async Delay by using the following **mmafmctl flushPending** command.

```
mmafmctl Device flushPending [-j FilesetName [--list-file ListFile]] [-s LocalWorkDirectory]
```

For more information, see the **mmafmctl** command in *IBM Storage Scale: Command and Programming Reference Guide*.

Parallel data transfers

Parallel data transfer improves the AFM data transfer performance.

To transfer large files from a cache cluster to a home cluster by using the primary gateway, you can configure the cache cluster to use all gateways defined in the cluster. When NFS is used for AFM data transfers, multiple NFS servers are recommended on the home cluster to improve parallel data transfer performance through gateway-to-NFS server mappings that are configured by using the **mmafmconfig** command. Multiple gateways can share a single NFS server, which provides limited parallelism benefits such as split writes. However, a one-to-one mapping between a gateway and an NFS server delivers the best throughput. All NFS servers on the home cluster must export the home path using the same parameters.

In a cache cluster that uses NFS for AFM data transfer, you can map each gateway to a specific NFS server in the home cluster. This mapping replaces the NFS server name specified in the **afmTarget** parameter. An export server map allows you to define multiple NFS servers and associate them with specific AFM gateway nodes. The mapping can be modified without changing the **afmTarget** parameter of a fileset. However, a fileset relink or a file system remount is required for the updated mapping to take effect.

Use the **mmafmconfig** command to define, display, update, and delete gateway-to-NFS server mappings.

To define multiple NFS servers for an **AFMTarget** parameter and use parallel data transfers:

1. Define a mapping by using the **mmafmconfig** command.
2. Use the mapping as the **AFMTarget** parameter for one or more filesets.
3. Update parallel read and write thresholds, in chunk size, as required.

Creating AFM filesets by using NFS gateway mapping

The following example demonstrates how to configure gateway node mapping for an NFS target and then create AFM filesets by using the defined mappings.

- Four AFM cache gateway nodes are available:

- hs22n18
- hs22n19
- hs22n20
- hs22n21

- Two home NFS servers are used:

- js22n01 (192.168.200.11)
- js22n02 (192.168.200.12)

Gateway nodes are mapped to specific NFS servers to enable controlled parallel data transfer.

1. Define gateway node mappings. Create mappings that associate cache gateway nodes with home NFS servers.

- a. Create mapping1.

```
mmafmconfig add mapping1 --export-map js22n01/hs22n18,js22n02/hs22n19
```

A sample output is as follows:

```
mmafmconfig: Command successfully completed
mmafmconfig: Propagating the cluster configuration data to all affected nodes.
This is an asynchronous process.
```

- b. Create mapping2.

```
mmafmconfig add mapping2 --export-map js22n02/hs22n20,js22n01/hs22n21
```

A sample output is as follows:

```
mmafmconfig: Command successfully completed
mmafmconfig: Propagating the cluster configuration data to all affected nodes.
This is an asynchronous process.
```

2. Verify the mapping configuration.

```
mmafmconfig show
```

A sample output is as follows:

```
Map name:          mapping1
Export server map:  192.168.200.12/hs22n19.gpfs.net,
                   192.168.200.11/hs22n18.gpfs.net

Map name:          mapping2
Export server map:  192.168.200.11/hs22n20.gpfs.net,
                   192.168.200.12/hs22n21.gpfs.net
```

3. Create AFM filesets by using the mappings.

- a. Create a single-writer (SW) fileset.

```
mmcrfileset gpfs1 sw1 --inode-space new \
-p afmMode=sw,afmTarget=nfs://mapping1/gpfs/gpfs2/swhome
```

- b. Create a read-only (RO) fileset.

```
mmcrfileset gpfs1 ro1 --inode-space new \
-p afmMode=ro,afmTarget=nfs://mapping2/gpfs/gpfs2/swhome
```

Parallel read and write

Parallel reads and writes are effective only for files larger than the configured parallel threshold. The thresholds are defined by the following parameters:

- **afmParallelWriteThreshold**
- **afmParallelReadThreshold**

Parallel transfer applies to all file types except the following files:

- Reads on sparse files
- Files with partial file caching enabled

For these files, read operations are handled only by the primary gateway and are not split into chunks. The size of each transfer chunk is configured by using:

- **afmParallelWriteChunkSize**
- **afmParallelReadChunkSize**

These parameters control how data is split for parallel transfers across participating gateways.

AFM gateway node mapping

The following functions apply to an AFM gateway node mapping:

NSD protocol

When the NSD protocol is used and a fileset is created without a map, all gateway nodes participate in parallel data transfer.

NFS protocol

When the NFS protocol is used and multiple gateway nodes are mapped to the same NFS server:

- Only one gateway node performs read operations.
- Write operations are distributed across all mapped gateway nodes.

Gateway node mapping limitation

A gateway node can be mapped to only one NFS server.

Effect of map changes

Mapping changes take effect only after one of the following actions:

- The fileset is re-linked, or
- The file system is remounted.

Missing or mismatched mapping

If a map is not specified, or if the mapping does not match, parallel data transfer is not used. Normal data transfer is used instead.

Removing gateway designation

The gateway role can be removed from a node only if the node is not defined in any mapping configuration.

Note: If an AFM home cluster contains mixed architectures (x86 and PPC), parallel data transfer works only for nodes that belong to a single architecture. The architecture that initiates the data transfer first determines which node group is used.

Parallel data transfer and multiple remote mounts

This feature can be combined with the parallel data transfer by using multiple remote mounts feature to improve data transfer performance between an AFM cache and an AFM home cluster.

- Both features use the same AFM gateway node mapping defined by the **mmafmconfig** command.
- The features operate independently.
- You can enable either or both features depending on workload requirements.

Parallel data transfer using multiple remote mounts

Parallel data transfer by using the multiple remote mounts on the gateway node improves the AFM data transfer performance by sending unique data operations across multiple NFS server mounts.

A cache cluster can be configured for parallel data transfers by using multiple remote mounts when a target site has two or more NFS servers. This feature uses a mapping between the gateway nodes at a cache (primary) site and the multiple NFS serving nodes at a home (secondary) cluster by exporting the target path. When you use NFS as a backed protocol for AFM data transfers, multiple NFS servers can be used at the home cluster. All NFS servers at the home cluster must export the same target path by using the same parameters for a specific fileset.

The AFM parallel mounts feature uses the same mapping infrastructure as the parallel data transfer. The parallel data transfer feature is used to transfer files bigger than the read or write threshold values. However, when the AFM parallel mounts feature is enabled, NFS server nodes can send and receive data in parallel by using the available bandwidth.

In a cache cluster, by using NFS for AFM data transfer, each gateway node can be mapped to a specific NFS server at a home. Creating an export server map can be used to define more than one NFS server and map those NFS servers to specific AFM gateways.

The AFM parallel mounts feature can be enabled or disabled by setting the **afmParallelMounts** parameter in the **mmchconfig** command or the **mmchfileset** command.

This parameter is available at the global cluster level for the **mmchconfig** command and at the fileset level for the **mmchfileset** command.

Complete the following steps to configure and set up AFM parallel data transfer by using multiple remote mounts:

1. Enable the **afmParallelMounts** parameter at the cluster level.

Set up this tunable at the cluster level to enable the **afmParallelMounts** parameter for all filesets by using the following command:

```
mmchconfig afmParallelMounts=yes -i
```

2. Define a mapping by identifying gateway nodes and NFS servers.

To create a mapping, the user must identify the gateway node at the cache cluster and NFS servers at the home site. This mapping enables mounting remote export from two or more NFS servers at the gateway node and synchronizes unique data by using multiple remote mounts. The user can create an AFM or ADR fileset by using this mapping as the **AFMTarget** parameter.

3. Start data synchronization between the home and cache clusters, this synchronization uses multiple remote mounts at gateway nodes.

User can also enable this parameter for a specific AFM fileset. The user can set the **afmParallelMounts** parameter by issuing the following command:

```
mmchfileset <FS> -j <AFM/ADR fileset> -p afmParallelMounts=yes
```

Example

The following example shows a mapping for an NFS target. Two cache gateway nodes *cachegw_n1* and *cachegw_n2* are mapped to two home NFS servers *homenfs_n1* and *homenfs_n2* (192.168.200.11 and 192.168.200.12), and then SW filesets are created by using this mapping for the parallel data transfer by using multiple remote mounts.

1. Create a mapping between the NSF export server and the NFS server.

```
mmafmconfig add mapping1 --export-map homenfs_n1/cachegw_n1,homenfs_n2/cachegw_n2
mmafmconfig: Command successfully completed
mmafmconfig: Propagating the cluster configuration data to all affected nodes. This is an
asynchronous process.
```

The syntax is as follows:

```
mmafmconfig {add | update} MapName --export-map ExportServerMap
```

```
mmafmconfig show
Map name:          mapping1
Export server map:  192.168.200.12/cachegw_n2.gpfs.net,192.168.200.11/cachegw_n1.gpfs.net
```

2. Create filesets by using *mapping1* and link the fileset.

```
mmcrfileset gpfs1 sw1 --inode-space new -p afmMode=sw,afmTarget=nfs://mapping1/gpfs/gpfs2/
swhome
Fileset sw1 created with id 14 root inode 3670019.
```

3. Issue the following command to set parameter for a specific fileset:

```
mmchfileset gpfs1 sw1 -p afmParallelMounts=yes
Fileset sw1 changed.
```

```
mmlinkfileset fs1 sw1 -J /gpfs1/sw1
Fileset sw1 linked at /gpfs1/sw1
```

By using parallel remote mounts, all the data and metadata operations are uniquely replicated to the target site by queuing them to the unique remote mounts or channels. This method keeps the queue or operation processing completely within the primary gateway node for the fileset and involves multiple NFS servers that are defined in the mapping from the remote cluster. Primary gateway tracks operation that is played to different NFS mounts. This setup supports both write from the cache and read from the remote cluster.

This feature can be combined with the “Parallel data transfers” on page 77 feature to obtain better data transfer performance within an AFM cache and an AFM home. Both features use the same AFM gateway node mapping that is defined by using the **mmafmconfig** command. These features are independent of each other, and you can set these features by considering what is suited best for a workload.

Partial file caching

With partial file caching, the cache can fetch only the blocks that are read and not the entire file, thereby using network and local disk space more efficiently. This caching is useful when an application does not need to read the whole file. Partial file caching is enabled on an IBM Storage Scale block boundary.

Partial file caching is controlled by the **afmPrefetchThreshold** parameter that can be updated by using the **mmchfileset** command. The default value of this parameter is 0. Complete file caching and all blocks of a file are fetched after any three blocks are read by the cache and the file is marked as cached because of this value. This value is useful for sequentially accessed files that are read in their entirety, such as image files, home directories, and development environments.

The valid **afmPrefetchThreshold** values are in the range 1 – 100. This parameter value specifies the file size percentage that must be cached before the rest of the data blocks are automatically fetched into the cache. A large value is suitable for a file that is accessed partially.

An **afmPrefetchThreshold** value of 100 disables full file prefetching. This value caches only the data blocks that are read by the application. This value is useful for large random-access files that are either too large to fit in the cache or are never expected to be read in their entirety. When all data blocks are available in the cache, the file is marked as cached.

For sparse files, the percentage for prefetching is calculated as the ratio of the size of data blocks that is allocated in the cache and the total size of data blocks on the home. Holes in the home file are not considered in the calculation.

Writes on partially cached files

If a write is queued on a file that is partially cached, then the complete file is cached first. Only then the write is queued on the file. Appending to a partially cached file does not cache the whole file. In the LU

mode alone, the write inset or append on a file that is cached partially caches the whole file even if the prefetch threshold is set on the fileset.

Note: As partial file caching is not compatible with earlier versions, all nodes must be on GPFS 3.5.0.11 or later.

Prefetch

Prefetch fetches the file metadata (inode information) and data from home before an application requests the contents.

Prefetch is a feature that allows fetching the contents of a file into the cache before actual reads.

Prefetching files before an application starts can reduce the network delay when an application requests a file. Prefetch can be used to proactively manage WAN traffic patterns by moving files over the WAN during a period of low WAN usage.

Prefetch can be used to do the following tasks:

- Populate metadata
- Populate data
- View prefetch statistics

Use the following command to perform these activities:

```
mmafmcctl Device prefetch -j FilesetName [-s LocalWorkDirectory]
[--retry-failed-file-list|--enable-failed-file-list]
[{-directory LocalDirectoryPath [--skip-dir-list-file SkipDirListfile] |
--dir-list-file DirListfile [--policy]} [--nosubdirs]]
[{-list-file ListFile | --home-list-file HomeListFile} [--policy]]
[--home-inode-file PolicyListFile] --outband
[--home-fs-path HomeFilesystemPath] --delete --wait
[--metadata-only] [--gateway Node] [--empty-pttrash]
[--readdir-only] [--force] [--prefetch-threads nThreads]
```

For more information about the command, see *mmafmcctl* command in *IBM Storage Scale: Command and Programming Reference Guide*. If no options are given for prefetch, the statistics of the last prefetch command that is run on the fileset are displayed.

--metadata-only

Prefetches only the metadata and not the actual data. This option is useful in migration scenarios. This option requires the list of files whose metadata is to be populated. It must be combined with a list file option.

--list-file *ListFile*

The specified file contains a list of files and directories that need to be pre-populated, one entry per line. All files must have fully qualified path names. If the list of files or directories that need to be prefetched have file names with special characters, then a policy must be used to generate the listfile. This list file must be edited manually to remove all other entries except the file or directory names. The list of files can contain the following files and directories:

- Files with fully qualified names from cache
- Files with fully qualified names from home
- List of files from the home that are generated by using the policy. The file must not be edited.
- Directories with fully qualified names.

AFM fetches files and directories listed in the list of files:

- If an entry is for a file, AFM fetches the file.
- If an entry is for a directory, AFM recursively fetches all subdirectories and all files in the specified directory.

--enable-failed-file-list

Turns on generating a list of files that failed during prefetch operation at the gateway node. The list of files is saved as `.afm/.prefetchedfailed.list` under the fileset. Failures that occur during

processing are not logged in `.afm/.prefetchedfailed.list`. If you observe any errors during processing (before queuing), you might need to correct the errors and rerun prefetch.

--policy

Specifies that the `list-file` or `home-list-file` is generated by using a GPFS Policy by which sequences like `'\'` or `'\n'` are escaped as `'\\'` and `'\\n'`. If this option is specified, input file list is treated as already escaped. The sequences are unescaped first before queuing for prefetch operation.

Note: This option can be used only if you are specifying `list-file` or `home-list-file`.

--directory LocalDirectoryPath

Specifies path to the local directory from which you want to prefetch files. A list of all files in this directory and all its subdirectories is generated, and queued for prefetch. You can either specify `--directory` or `--dir-list-file` with **mmafmctl prefetch**. The `--policy` option can be used only with `--dir-list-file` and not with `--directory`.

For example,

```
mmafmctl fs1 prefetch -j fileset1 --dir-list-file /tmp/file1 --policy
```

The following example includes methods to name a directory for the `--directory` option, when the directory name contains special characters:

- When a directory name does not have terminal escape sequences, keep the absolute directory path within double quotation marks (" ").

```
mmafmctl fs2 prefetch -j roTestPrefetch_GPFS --directory "/gpfs/fs2/roTestPrefetch1/Dir_a\!  
h@j#k%l^k&78*9'\''"
```

A sample output is as follows:

```
mmafmctl: Performing prefetching of fileset: roTestPrefetch_GPFS  
mmafmctl(2020-04-13 02:35:39): Listing all files of directory "/gpfs/fs2/roTestPrefetch1/  
Dir_a\!h@j#k%l^k&78*9'\''"  
Queued          Failed          TotalData  
                25              0      (approx in Bytes)  
                25              0      131072000  
prefetch successfully queued at the gateway.  
(2020-04-13 02:35:41): Listed all files of directory "/gpfs/fs2/roTestPrefetch1/Dir_a\!  
h@j#k%l^k&78*9'\''"
```

- When the directory name has terminal escape sequences, do not keep the directory path within double quotation marks. The terminal auto-fills the escape sequences in the directory name when you press the <Tab> two times.

```
mmafmctl fs2 prefetch -j roTestPrefetch_GPFS --directory /gpfs/fs2/roTestPrefetch1/Dir_a\!  
\!h@j#k%l^k&78*9'\''\''/'
```

A sample output is as follows:

```
mmafmctl: Performing prefetching of fileset: roTestPrefetch_GPFS  
mmafmctl(2020-04-13 02:39:58): Listing all files of directory "/gpfs/fs2/roTestPrefetch1/  
Dir_a\!h@j#k%l^k&78*9'\''/'"  
Queued          Failed          TotalData  
                25              0      (approx in Bytes)  
                25              0      131072000  
prefetch successfully queued at the gateway.  
mmafmctl(2020-04-13 02:40:00): Listed all files of directory "/gpfs/fs2/roTestPrefetch1/  
Dir_a\!h@j#k%l^k&78*9'\''/'"
```

- When you press <Tab> two times to include the escape sequences in a directory name and keep the directory path within double quotation marks, the prefetch operation fails. The prefetch operation fails because the unescape of terminal escaped characters in the directory name is not performed.

```
mmafmctl fs2 prefetch -j roTestPrefetch_GPFS --directory "/gpfs/fs2/roTestPrefetch1/Dir_a\!  
\!h@j#k%l^k&78*9'\''\''/'"
```

A sample output is as follows:

```
mmafmctl: Performing prefetching of fileset: roTestPrefetch_GPFS
runPrepopSubcommand: Unexpected error from missing or incorrect prepop input path.
Return code: 1
mmafmctl: Command failed. Examine previous error messages to determine cause.
```

--skip-dir-list-file *SkipDirListfile*

Specifies a path to a file that contains a list of directories within an AFM fileset. These directories are not fetched during the prefetch operation. This option can only be used with the `--directory` option. When this option is enabled, the specified directories are not included in the prefetch process for the AFM fileset.

For example,

```
mmafmctl fs1 prefetch -j fileset1 --directory /gpfs/fs1/fileset1 --skip-dir-list-file /tmp/file1
```

Note: You cannot use this option with the `--dir-list-file` option.

--dir-list-file *DirListFile*

This parameter enables prefetching individual directories under AFM fileset. Input file specifies the unique path to a directory that you want to prefetch. AFM generates a list of files under the specified directory and subdirectories and queues it to the gateway Node. The input file can also be a policy-generated file for which you need to specify `--policy`

--nosubdirs

This option restricts the recursive behavior of `--directory` and `--dir-list-file` and prefetches only until the specified level of directory. If you specify this parameter, subdirectories under the directory are not prefetched. This parameter is optional and can be used only with `--directory`, `--dir-list-file`, and `--list-file`.

For example,

```
mmafmctl fs1 prefetch -j fileset1 --directory /gpfs/fs1/fileset1/dir1 --nosubdirs
```

```
mmafmctl fs1 prefetch -j fileset1 --dir-list-file /tmp/file1 --policy --nosubdirs
```

--retry-failed-file-list

Allows retrying prefetch of files that failed in the last prefetch operation. The list of files to retry is obtained from `.afm/.prefetchedfailed.list` under the fileset.

Note: To use this option, you must enable generating a list of failed files. Add **--enable-failed-file-list** to the command first.

--home-list-file *HomeListFile*

The specified file contains a list of files from home that need to be pre-populated, one file per line. All files must have fully qualified path names. If the list of files to be prefetched have file names with special characters, then a policy must be used to generate the `listfile`. A policy-generated file must be edited manually to remove all other entries except the file names. As of version 4.2.1, this option is deprecated. The **--list-file** option removes all other entries except the file names.

--home-inode-file *PolicyListFile*

The specified file contains the list of files from home that need to be pre-populated in the cache and this file is generated by using policy. This file must not be edited manually. This option is deprecated. The `-list-file` option removes all other entries except the file names.

--home-fs-path *HomeFileSystemPath*

Specifies the full path to the fileset at the home cluster and can be used along with **-list-file**. You must use this option, when in the NSD protocol the mount point on the gateway nodes of the **afmTarget** filesets does not match the mount point on the Home cluster. For example, the home file system is mounted on the home cluster at `/gpfs/homefs1`. The home file system is mounted on the cache by using NSD protocol at `/gpfs/remotefs1`.

For example,


```
mmafmctl gpfs1 prefetch -j cache1 -list-file /tmp/list.allfiles --home-fs-path /gpfs/
remotefs1
```

--wait

Completes the prefetch command after all data is downloaded to the AFM cache fileset. When you specify this option, the command waits until the data is queued and downloaded.

If you do not specify this option, the prefetch command returns after the data is queued, and the download operation runs in the background.

You can specify this option with all parameters when you run the prefetch command.

Wait parameter can also be specified along with -Y to displays the command output in a parseable format with a colon (:) as a field delimiter. Each column is described by a header.

Example:

```
mmafmctl fs1 prefetch -j afmR01 --directory /gpfs/fs1/afmR01 --wait
```

A sample output is as follows:

```
mmafmctl: Performing prefetching of fileset: afmR01
mmafmctl(2024-12-10 08:21:30): Listing all files of directory /gpfs/fs1/ afmR01
Queued (Processed) Failed TotalData
(approx in Bytes)
0 ( 202) 0 105381888
0 ( 303) 0 158334976
0 ( 606) 0 317194240
0 ( 1111) 0 581435392
0 ( 2121) 0 1050148864
0 ( 4141) 0 1050148864
0 ( 8282) 0 1050148864
2003 ( 12006) 0 1050148864

prefetch successfully queued at the gateway.

mmafmctl(2024-12-10 08:23:25): Listed all files of directory /gpfs/fs1/ afmR01
mmafmctl : Waiting for complete data prefetch.
Fileset Name Pending Failed Total Total data(Bytes)
Throughput(KB/s) FailedListPath
-----
afmR01 1616 0 2003 202899456
15228 -, -
afmR01 1335 0 2003 350224384
15228 -, -
afmR01 1070 0 2003 489160704
15228 -, -
afmR01 811 0 2003 624951296
15228 -, -
afmR01 549 0 2003 762314752
15228 -, -
afmR01 288 0 2003 899153920
15192 -, -
afmR01 24 0 2003 1037565952
15192 -, -
afmR01 0 0 2003 1050148864
0 -, -

mmafmctl: Prefetch completed successfully.
```

--readdir-only

Enables readdir operation on a dirty directory at the cache one last time and brings latest directory entries.

This option helps prefetching modified directory entries from the home, although the directory at the cache fileset was modified by the applications and AFM marked the dirty flag on the cache directory. This option overrides the dirty flag that is set when the data is modified at the local LU cache. In the LU mode, the dirty flag does not allow the readdir operation at the home and refreshes the directory file entries from the home.

This option helps in the migration process where new files were created at the home after the application was moved to the cache. The application already modified the directory and refresh

intervals were disabled. AFM queues `readdir` one last time on the cache directory and brings entries of the created files to the cache.

The **afmReadDirOnce** parameter must be set on an AFM fileset, and directory and files refresh intervals must be disabled.

For example,

1. To set **afmRefreshOnce** on an AFM fileset, issue the following command:

```
mmchfileset fs fileset -p afmRefreshOnce=yes
```

2. To check whether the **afmRefreshOnce** parameter value is set on an AFM fileset, issue the following command:

```
mmllsfileset fs fileset -L --afm
```

A sample output is as follows:

```
Filesets in file system '<fs>':
Attributes for fileset <fileset>:
=====
Status                               Linked
Path                                GPFS_PATH/fileset
Id                                  37
Root inode                           3145731
Parent Id                             0
Created                             Wed Mar  4 12:23:47 2020
Comment
Inode space                           6
Maximum number of inodes             100352
Allocated inodes                     100352
Permission change flag               chmodAndSetacl
afm-associated                       Yes
Target                               nfs://home/fileset
Mode                                 local-updates
File Lookup Refresh Interval         30 (default)
File Open Refresh Interval           30 (default)
Dir Lookup Refresh Interval          60 (default)
Dir Open Refresh Interval            60 (default)
Expiration Timeout                   disable (default)
Last pSnapId                         0
Display Home Snapshots               yes (default)
Number of Gateway Flush Threads      4
Prefetch Threshold                   0 (default)
Eviction Enabled                     yes (default)
IO Flags                             9216 (refreshOnce)
```

3. To run the prefetch operation for the `readdir` operation one last time, issue following command:

```
mmafmctl fs prefetch -j fileset --directory /fileset_path/directory --readdir-only
```

--force

Enables forcefully fetching data from the home during the migration process. This option overrides any set restrictions and helps to fetch the data forcefully to the cache. This option must be used only to forcefully fetch the data that was created after the migration process completion.

For example,

```
mmafmctl fs prefetch -j fileset --list-file listfile_path --force
```

--gateway Node

Allows selecting the gateway node that can be used to run the prefetch operation on a fileset, which is idle or less-utilized. This option helps to distribute the prefetch work on different gateway nodes and overrides the default gateway node, which is assigned to the fileset. It also helps to run different prefetch operations on different gateway nodes, which might belong to the same fileset or a different fileset.

For example,

```
mmafmctl fs prefetch -j fileset --list-file listfile_path --gateway Node2
```

--empty-pttrash

Cleans the .pttrash directory in an AFM fileset. Specify this option during the prefetch operation to delete all existing data in .pttrash directory. This option can be specified with all prefetch options.

In IBM Storage Scale AFM, the .pttrash directory is used to temporarily store files that are deleted or moved by clients. The .pttrash directory serves as a staging area for deleted or moved files. The files are stored in the .pttrash directory until the AFM policy is triggered to remove them, or until they are restored. This option allows potential data recovery or rollback operations before permanently deleting the files. Refer [Internal AFM directories](#) section for more information.

--prefetch-threads *nThreads*

Specifies the number of threads to be used for the prefetch operation. Valid values are 1 - 255. Default value is 4.

For example,

```
mmafmctl fs prefetch -j fileset --list-file listfile_path --prefetch-threads 6
```

Prefetch is an asynchronous process and the fileset can be used while prefetch is in progress. Prefetch completion can be monitored by using the **afmPrepopEnd** callback event or looking at **mmafmctl Device prefetch** command with no options.

Prefetch pulls the complete file contents from home (unless the **--metadata-only** flag is used), so the file is designated as cached when it is prefetched. Prefetch of partially cached files caches the complete file.

Prefetch can be run in parallel on multiple filesets, although only one prefetch job can run on a fileset.

While a file is getting prefetched, it is not evicted.

If parallel data transfer is configured, all gateways participate in the prefetch process.

If the file system unmounts during prefetch on the gateway, issue the prefetch again.

Prefetch can be triggered on inactive filesets.

Directories are also prefetched to the cache if specified in the prefetch file. If you specify a directory in the prefetch file and if that directory is empty, the empty directory is prefetched to cache. If the directory contains files or subdirectories, you must specify the names of the files or subdirectories that you want to prefetch. If you do not specify names of individual files or subdirectories inside a directory, that directory is prefetched without its contents.

If you run the prefetch command with data or metadata options, statistics like queued files, total files, failed files, total data (in bytes) is displayed.

For example,

```
mmafmctl FileSystem prefetch -j fileset --enable-failed-file-list --list-file /tmp/file-list
```

A sample output is as follows:

```
mmafmctl: Performing prefetching of fileset: <fileset>
Queued (Total) Failed TotalData (approx in Bytes)
0 (56324) 0 0
5 (56324) 2 1353559
56322 (56324) 2 14119335
```

Prefetch Recovery:

Note: This feature is disabled from IBM Storage Scale 5.0.2. If your cluster is running on an earlier version, prefetch recovery is possible.

If the primary gateway of a cache is changed while prefetch is running, prefetch is stopped. The next access to the fileset automatically retriggers the interrupted prefetch on the new primary gateway. The

list file used when prefetch was initiated must exist in a path that is accessible to all gateway nodes. Prefetch recovery on a single-writer fileset is triggered by a read on some file in the fileset. Prefetch recovery on a read-only, independent-writer, and local-update fileset is triggered by a lookup or readdir on the fileset. Prefetch recovery occurs on the new primary gateway and continues where it left off. It looks at which files did not complete prefetch and it rebuilds the prefetch queue. Examples of messages in the `mmfs.log` are as follows:

```
Wed Oct 1 13:59:22.780 2014: [I] AFM: Prefetch recovery started for the file system gpfs1
fileset iw1.
mmafmctl: Performing prefetching of fileset: iw1
Wed Oct 1 13:59:23 EDT 2014: mmafmctl: [I] Performing prefetching of fileset: iw1
Wed Oct 1 14:00:59.986 2014: [I] AFM: Starting 'queue' operation for fileset 'iw1' in
filesystem '/dev/gpfs1'.
Wed Oct 1 14:00:59.987 2014: [I] Command: tpscache /dev/gpfs1 1 iw1 0 257 42949 67295 0 0
1393371
Wed Oct 1 14:01:17.912 2014: [I] Command: successful tpscache /dev/gpfs1 1 iw1 0 257
4294967295 0 0 1393371
Wed Oct 1 14:01:17.946 2014: [I] AFM: Prefetch recovery completed for the filesystem gpfs1
fileset iw1. error 0
```

1. Metadata population by using prefetch:

```
mmafmctl fs1 getstate -j ro
```

A sample output is as follows:

Fileset Name	Fileset Target	Cache State	Gateway Node	Queue Length	Queue numExec
ro	nfs://c26c3apv1/gpfs/homefs1/dir3	Active	c26c4apv1	0	7

List Policy:
 RULE EXTERNAL LIST 'List' RULE 'List' LIST 'List' WHERE PATH_NAME LIKE '%'
 Run the policy at home:mmapplypolicy /gpfs/homefs1/dir3 -P px -f px.res -L 1 -N mount -I defer
 Policy creates a file which should be manually edited to retain only the file names. Thereafter
 this file is used at the cache to populate metadata.

```
mmafmctl fs1 prefetch -j ro --metadata-only --list-file=px.res.list.List
```

A sample output is as follows:

```
mmafmctl: Performing prefetching of fileset: ro
Queued      (Total)      Failed      TotalData
              (approx in Bytes)
          0      (2)          0              0
        100     (116)         5          1368093971
        116     (116)         5          1368093971

prefetch successfully queued at the gateway

Prefetch end can be monitored by using this event:
Thu May 21 06:49:34.748 2015: [I] Calling User Exit Script prepop: event afmPrepopEnd,
Async command prepop.sh.

The statistics of the last prefetch command can be viewed by running the following command:
mmafmctl fs1 prefetch -j ro
Fileset      Async Read      Async Read      Async Read      Async Read      Async Read
Name         (Pending)       (Failed)        (Already Cached) (Total)         (Data in Bytes)
-----
ro           0                1                0                7                0
```

2. Prefetch of data by giving list of files from home:

```
cat /listfile1
```

A sample output is as follows:

```
/gpfs/homefs1/dir3/file1
/gpfs/homefs1/dir3/dir1/file1
```

```
mmafmctl fs1 prefetch -j ro --list-file=/listfile1
```

A sample output is as follows:

```
mmafctl: Performing prefetching of fileset: ro
Queued (Total) Failed TotalData
      (2)      0      (approx in Bytes)
2      (2)      0      1368093971
```

```
mmafctl fs1 prefetch -j ro
```

A sample output is as follows:

Fileset Name	Async Read (Pending)	Async Read (Failed)	Async Read (Already Cached)	Async Read (Total)	Async Read (Data in Bytes)
ro	0	0	0	2	122880

3. Prefetch of data by using a list file, which is generated by using policy at home:

Inode file is created by using the policy at home, and must be used without editing manually.

```
List Policy:
RULE EXTERNAL LIST 'List' RULE 'List' LIST 'List' WHERE PATH_NAME LIKE '%'
```

For files with special characters, path names must be encoded with ESCAPE %.

```
RULE EXTERNAL LIST 'List' ESCAPE '%' RULE 'List' LIST 'List' WHERE PATH_NAME LIKE '%'
```

Run the policy at home:

```
mmapplypolicy /gpfs/homefs1/dir3 -P px -f px.res -L 1 -N mount -I defer
```

```
cat /lfile2
```

A sample output is as follows:

```
113289 65538 0 -- /gpfs/homefs1/dir3/file2
113292 65538 0 -- /gpfs/homefs1/dir3/dir1/file2
```

```
mmafctl fs1 prefetch -j ro -list-file=/lfile2
```

A sample output is as follows:

```
mmafctl: Performing prefetching of fileset: ro
mmafctl fs1 prefetch -j ro -list-file=/lfile2
mmafctl: Performing prefetching of fileset: ro
Queued (Total) Failed TotalData
      (2)      0      (approx in Bytes)
2      (2)      0      1368093971
```

4. Prefetch by using --home-fs-path option for a target with the NSD protocol:

```
mmafctl fs1 getstate -j ro2
```

A sample output is as follows:

Fileset Name	Fileset Target	Cache State	Gateway Node	Queue Length	Queue numExec
ro2	gpfs:///gpfs/remotefs1/dir3	Active	c26c4apv1	0	7

```
cat /lfile2
```

A sample output is as follows:

```
113289 65538 0 -- /gpfs/homefs1/dir3/file2
113292 65538 0 -- /gpfs/homefs1/dir3/dir1/file2
```

```
mmafmctl fs1 prefetch -j ro2 -list-file=/lfile2 --home-fs-path=/gpfs/homefs1/dir3
```

A sample output is as follows:

```
mmafmctl: Performing prefetching of fileset: ro2
```

Queued	(Total)	Failed	TotalData (approx in Bytes)
0	(2)	0	0
2	(2)	0	113292

```
mmafmctl fs1 prefetch -j ro2
```

A sample output is as follows:

Fileset Name	Async Read (Pending)	Async Read (Failed)	Async Read (Already Cached)	Async Read (Total)	Async Read (Data in Bytes)
ro2	0	0	0	2	122880

Peer snapshot -psnap

The peer snapshot function provides a snapshot at home and cache separately, ensuring application consistency on both sides (cache and home).

When you take a peer snapshot, it creates a snapshot in the cache fileset then queues a snapshot to be executed at the home. When all of the queued requests outstanding at the time the snap was performed in cache have been flushed to the home fileset so that the home data is consistent with cache, a snapshot of the corresponding home fileset is performed. The result is a pair of peer snapshots, one at the cache and one at the home. Both refer to the same copy.

Peer snapshots are created using the **mmpsnap** command on the cache. To create a peer snapshot, the cache fileset must be in the active state. The last successful snapshot is saved and can be viewed by running the **mm1ssnapshot** command at home or cache. Multiple outstanding peer snapshots can be queued on the gateway. Use the **mmpsnap** command to ensure that both the cache and home snapshots are removed. The **mmpsnap** command works only if the home cluster has run the **mmafmconfig enable ExportPath** command. The **mmpsnap** command can be used only in an SW cache. It cannot be used for RO, IW, or LU caches.

Note: Do not use the **mmde1snapshot** command to delete peer snapshots.

If the cache fileset is disconnected from the home fileset when the cache snapshot is created, the cache records that the peer snapshot on the home fileset has not been created. When connection is restored, it attempts to create the snapshot.

Peer snapshots are not allowed on an SW cache that uses the NSD protocol for communicating with home.

1. To create a fileset level snapshot in cache of a single-writer fileset called sw in file system fs1, run the following command:

```
mmpsnap fs1 create -j sw
```

The system displays the following output:

```
Writing dirty data to disk.
Quiescing all file system operations.
Writing dirty data to disk again.
Snapshot psnap-13882361812785649740-C0A80E85:4F44B305-59-12-03-01-02-27-28 created with id 8.
Snapshot psnap-13882361812785649740-C0A80E85:4F44B305-59-12-03-01-02-27-28 created at the
satellite.
Core snapshot has been queued.
```

2. To view the snapshot, run the following command:

```
mmlssnapshot fs1 -j sw
```

The system displays the following output:

```
Snapshots in file system fs1:
Directory SnapId Status Created Fileset
psnap-13882361812785649740-C0A80E85:4F44B305-59-12-03-01-02-27-28 8 Valid Thu Oct 27
02:27:29 2016 sw
```

3. To view the snapshot at home, run the following command at home.:

```
mmlssnapshot fs1
```

The system displays the following output:

```
Snapshots in file system fs1:
Directory SnapId Status Created Fileset
psnap-13882361812785649740-C0A80E85:4F44B305-59-12-03-01-02-27-28 8 Valid Thu Oct 27
02:23:16 2016
```

Gateway node failure and recovery

When the primary gateway of a fileset fails, another gateway node takes over the ownership of the fileset.

Gateway node failures are not catastrophic and do not result in the loss of data or the loss of the ability of AFM to communicate with the home cluster with updates and revalidations.

AFM internally stores all the information necessary to replay the updates made in the cache at the home cluster. When a gateway node fails, the in-memory queue is lost. The node rebuilt the queue in memory by taking over for the failed gateway. The process of rebuilding the queue is called recovery. As an administrator, ensure that you create sufficient disk space in `/var/mmfs/afm` for a smooth recovery. For the recovery of 1 million files or inodes, approximately 250 MB disk space is required in `/var/mmfs/afm`.

During recovery, outstanding cache updates are placed on the in-memory queue and the gateway starts processing the queue. AFM collects the pending operations by running a policy scan in the fileset. AFM uses the policy infrastructure in IBM Storage Scale to engage all the nodes that are mounting the file system. Pending requests are queued in a special queue that is called the priority queue. The priority queue is different from the normal queue where normal requests get queued. After the priority queue is flushed to home, the cache and home become synchronized and recovery is said to be completed and the cache returns to an Active state. In some cases, the files or directories that are deleted at cache might not be deleted at home. Therefore, the files or directories that are not deleted remain at the home. However, the recovery operations remain unaffected.

The beginning and end of the AFM recovery process can be monitored by using the **afmRecoveryStart** and **afmRecoveryEnd** callback events.

Recovery is used only for the single-writer mode filesets and the independent-writer mode filesets. It is triggered when the cache fileset attempts to move to the Active state, for example when the fileset is accessed for the first time after the failure.

Recovery can run in parallel on multiple filesets, although only one instance can run on a fileset at a time. The time taken to synchronize the contents of cache with home after the recovery of a gateway node depends on the number of files in the fileset and the number of outstanding changes since the last failure.

During the recovery or resync operation on an AFM fileset, complete data synchronization from the cache to the home might take some time. It depends on various factors such as the total amount of data, network, disk I/O. You can check the synchronization status by using the **--write-stats** and **--read-stats** options.

```
mmafmctl fs1 getstate -j filesetSW1 --write-stats
```

A sample output is as follows:

Fileset Name	Total Written Data (Bytes)	N/w Throughput (KB/s)	Total Pending Data to Write(Bytes)	Estimated Completion time

----- filesetSW1	----- 98359590	----- 68	----- 22620600	----- 5 (Min)

Note:

- If the **afmFastCreate** parameter value is set to yes or AFM to cloud object storage is enabled on a fileset, the --read-stats and --write-stats options show information such as N/w Throughput, Total Pending Data, Estimated Completion time only during the recovery or resync operation. During regular operations, the --read-stats or --write-stats option shows only Total Written Data.
- During recovery event, it might take some time for AFM to collect recovery data and queue operations to the AFM gateway node. The synchronization status is not shown until data is queued to the AFM gateway and the write operation is synchronized to the home.

In the multiple AFM caches environment, recovery is triggered after a failure. To limit the maximum number of AFM or AFM-DR filesets that can perform recovery at a time, set the **afmMaxParallelRecoveries** parameter. The recovery process is run on the number of filesets that you specify for **afmMaxParallelRecoveries**. The number of filesets that are specified in **afmMaxParallelRecoveries** are accessed for recovery. After recoveries are complete on these filesets, the next set of filesets is accessed for recovery. By default, **afmMaxParallelRecoveries** is set to 0, and the recovery process is run on all filesets. Specifying **afmMaxParallelRecoveries** restricts the number of recoveries, thus conserving hardware resources. For more information, see *Configuration parameters for AFM in IBM Storage Scale: Administration Guide*.

Peer snapshots that are created in cache and queued to home might get lost due to gateway node failure. These peer snapshots cannot be recovered through AFM recovery. For more information, see [“Peer snapshot -psnap” on page 90](#).

If any AFM-enabled fileset has no updates, the failure of a gateway node is harmless and application nodes do not experience delays. During recovery, application requests to all AFM filesets are momentarily blocked.

The following example indicates changes in the AFM fileset state during recovery:

```
node2:/gpfs/cache/fileset_SW mmafmctl fs1 getstate
Fileset Name Fileset Target Cache State Gateway Node Queue Length Queue numExec
-----
fileset_SW nfs://node4/gpfs/fshome/fset001 FlushOnly node2 0 0

node2:/gpfs/cache/fileset_SW mmafmctl fs1 getstate
Fileset Name Fileset Target Cache State Gateway Node Queue Length Queue numExec
-----
fileset_SW nfs://node4/gpfs/fshome/fset001 Recovery node2 0 0

node2:/gpfs/cache/fileset_SW mmafmctl fs1 getstate
Fileset Name Fileset Target Cache State Gateway Node Queue Length Queue numExec
-----
fileset_SW nfs://node4/gpfs/fshome/fset001 Active node2 0 3
```

For more information, see the *mmafmctl* command in *IBM Storage Scale: Command and Programming Reference Guide*.

An example of the messages in *mmfs.log* is as follows:

```
Thu Oct 27 15:28:15 CEST 2016: [N] AFM: Starting recovery for fileset 'fileset_SW' in fs 'fs1'
Thu Oct 27 15:28:15 CEST 2016: mmcommon afmrecovery invoked: device=fs1 filesetId=1...
Thu Oct 27 15:28:16 CEST 2016: [N] AFM: mmafmlocal /usr/lpp/mmfs/bin/mmappolypolicy...
Thu Oct 27 15:28:31.325 2016: [I] AFM: Detecting operations to be recovered...
Thu Oct 27 15:28:31.328 2016: [I] AFM: Found 2 update operations...
Thu Oct 27 15:28:31.331 2016: [I] AFM: Starting 'queue' operation for fileset 'fileset_SW' in
filesystem'fs1'.
Thu Oct 27 15:28:31.332 2016: [I] Command: tpspcache fs1 1 fileset_SW 0 3 1346604503 38 0 43
Thu Oct 27 15:28:31.375 2016: [I] Command: successful tpspcache fs1 1 fileset_SW 0 3 1346604503
38 0 43
Thu Oct 27 15:28:31 CEST 2016: [I] AFM: Finished queuing recovery operations for /gpfs/cache/
fileset_SW
```


Failures during recovery

Filesets can be in recovery state and might not complete recovery due to some conditions. The fileset might go to **Dropped** or **NeedsResync** state. This fileset state implies that recovery is failed.

The `mmfs.log` might contain the following lines: AFM: File system fs1 fileset adrSanity-160216-120202-KNFS-TC11-DRP encountered an error synchronizing with the remote cluster. Cannot synchronize with the remote cluster until AFM recovery is executed. remote error 28.

After the recovery fails, the next recovery is triggered after 120 seconds or when some operation is performed on the SW or IW fileset. After a successful recovery, modifications at the cache are synchronized with the home and the fileset state is 'Active'.

The following checks are needed from the administrator to ensure that the next recovery is successful:

1. Check the inode or block quota on cache and at home.
2. Ensure that home is accessible. Remount home file system and restart NFS at home.
3. Ensure that memory is not reached. If memory is reached, increase **afmHardMemThreshold**.
4. Check network connectivity with home.
5. If recovery keeps failing as eviction is triggered due to exceeding block quotas, increase block quotas or disable eviction to get recovery working.

Customizing a storage location to store the AFM recovery data

During the event of recovery, resync, or reconcile triggered on the AFM fileset, AFM stores recovery data under the default `/var/mmfs/afm` location. If the `/var` partition is full or not sufficient storage space, then AFM recovery might not complete.

To avoid the partition full issue, AFM can be configured to use a customized storage location. The recovery path can be configured for an AFM fileset by completing the following steps:

1. Set a cluster level config **afmRecoveryDir** by using the **mmchconfig** command to define the custom storage path for recovery or reconcile, resync, and then set individual fileset by using the **afmRecoveryUseFset=yes** using **mmchfileset** command.
2. After the **afmRecoveryDir** parameter is set at the cluster level, then any specific AFM fileset can be enabled to use the customized location by setting the **afmRecoveryUseFset** on the fileset. This parameter separates the requirement of enabling recovery for only those fileset that have huge data to process for synchronization during event of recovery or resync.

Consider the following points before you enable the **afmRecoveryDir** and **afmRecoveryDir** parameters to the custom storage location:

- If **afmRecoveryDir** parameter is not set at the cluster level, then AFM uses the default storage path for every fileset which is `/var/mmfs/afm`.
- If **afmRecoveryDir** parameter is set at the cluster level and **afmRecoveryUseFset** is set at the fileset level, then AFM uses the specified custom location for event of recovery, resync or reconcile.
- If **afmRecoveryDir** parameter is set at the cluster level and **afmRecoveryUseFset** is not set on the AFM fileset, then AFM recovery uses the `<fs_path>/<fileset_path>/<.ptrash/>` directory for recovery data for that fileset.

1. Specify a custom storage location for recovery at the cluster level that has sufficient storage space.

```
mmchconfig afmRecoveryDir=/bigstore -i
```

Valid value is a valid path name.

2. Enable existing individual AFM fileset to use custom location specified in the **afmRecoveryDir** parameter.

```
mmchfileset fs1 ADR-1 -p afmRecoveryUseFset=yes,....
```

The system output is as follows:

```
Fileset ADR-1 changed.
```

- a. Unset the **afmRecoveryUseFset** parameter on an AFM fileset.

```
mmchfileset fs1 ADR-1 -p afmRecoveryUseFset=no
```

The system output is as follows:

```
Fileset ADR-1 changed.
```

Valid values for **afmRecoveryUseFset** parameter is 'yes' and 'no'.

- b. When you are creating a AFM fileset to use custom location specified in the **afmRecoveryDir** parameter.

```
mmcrfileset fs1 ADR-1 -p afmRecoveryUseFset=yes,...
```

Valid value of the **afmRecoveryDir** parameter is a valid path name.

Cache eviction

Cache eviction is a feature where file data blocks in the cache are released when a fileset usage exceeds the fileset soft quota, and space is created for new files.

The process of releasing blocks is called eviction. However, file data is not evicted if the file data is dirty. A file whose outstanding changes are not flushed to home is a dirty file.

You can use automatic cache eviction or define your own policy to decide which file data needs to be evicted. To automatically enable cache eviction, define a fileset soft quota on the cache fileset. Eviction starts when fileset usage reaches the soft quota limit. A time lag might result when an eviction is triggered and the data is evicted. Tuning the soft and hard quota limits minimizes application delays because the data is being cached at a faster rate than it is being evicted.

Cache eviction based on inode quotas is not supported. Cached data from partially fetched files can be evicted manually by using **--file** parameter of **mmafmctl** command. When files are evicted from cache, all data blocks of those files are cleared. The files are then uncached. When a read operation is performed on the files the next time, the files fetch data from home.

Cache eviction is enabled by default on all AFM nodes and is controlled by the **afmEnableAutoEviction** parameter, and fileset block quota. Cache eviction can also be manually triggered by using the **mmafmctl evict** command. When a file is evicted, file data is removed from the cache, but the inode stays in the cache. Using eviction, you can build environments, where all objects from home are available but running in limited amount of space.

For example, a cache can be created in flash storage. File eviction opens a powerful method of providing small but high speed and low latency cache filesets to clusters.

Manual eviction can be done by using the **mmafmctl evict** command as follows:

```
mmafmctl Device evict -j FilesetName
      [--safe-limit SafeLimit] [--order {LRU | SIZE}]
      [--log-file LogFile] [--filter Attribute=Value ...]
      [--list-file ListFile] [--file FilePath] [--range]
```

For more information, see **mmafmctl** command in *IBM Storage Scale: Command and Programming Reference Guide*.

This option is applicable for RO/SW/IW/LU filesets. This command can be run manually or run in a script with a custom policy to implement a custom eviction policy. Options can be combined.

--safe-limit SafeLimit

This parameter is mandatory for the manual eviction, for order and filter attributes. It specifies target quota limit that is used as the low water mark for eviction in bytes – the value must be less than the

soft limit. This parameter can be used alone or can be combined with one of the following parameters (order or filter attributes). Specify the parameter in bytes.

--order *LRU* / *SIZE*

The order in which files are to be chosen for eviction: LRU - Least recently used files are to be evicted first. SIZE - Larger-sized files are to be evicted first.

LRU

AFM updates the file's atime when the specified file is read from the target/home even though the **atimeDeferredSeconds** parameter indicates otherwise. File's atime is not updated when the already cached file is read. This is to support auto eviction based on LRU algorithm. Reading the file from the target/home affects performance. So, atime is set after reading the file from target/home to make sure that the latest file read from target/home is evicted based on LRU.

--log-file *Log File*

The file where the eviction log is to be stored. By default no logs are generated.

--filter *Attribute=Value*

The attributes that you can use to control the way data is evicted from the cache. Valid attributes are: **FILENAME=File Name** - The name of a file to be evicted from the cache. This option uses an SQL-type search query. If the same file name exists in more than one directory, the cache evicts all the files with that name. The complete path to the file must not be given here. **MINFILESIZE=Size** - The minimum size of a file to evict from the cache. This value is compared to the number of blocks that are allocated to a file (KB_ALLOCATED), which might differ slightly from the file size. **MAXFILESIZE=Size** - The maximum size of a file to evict from the cache. This value is compared to the number of blocks that are allocated to a file (KB_ALLOCATED), which might differ slightly from the file size.

Possible combinations of options are:

- Only Safe limit
- Safe limit + LRU
- Safe limit + SIZE
- Safe limit + FILENAME
- Safe limit + MINFILESIZE
- Safe limit + MAXFILESIZE
- Safe limit + LRU + FILENAME
- Safe limit + LRU + MINFILESIZE
- Safe limit + LRU + MAXFILESIZE
- Safe limit + SIZE + FILENAME
- Safe limit + SIZE + MINFILESIZE
- Safe limit + SIZE + MAXFILESIZE

--list-file *ListFile*

The specified file contains a list of files to be evicted, one file per line. All files must have fully qualified path names. File system quotas need not be specified. If the list of files to be evicted have file names with special characters, then a policy must be used to generate the listfile. This policy output must be hand-edited to remove all other entries except the file names and can be passed to the evict command.

--file *FilePath*

The fully qualified name of the file to be evicted. File system quotas need not be specified.

enforceFilesetQuotaOnRoot – Enable to evict files created by root. By default, the files that are created by root are not evicted.

--range

Enables eviction of a range of specified data blocks from a file. When this option is used with the **--list-file** option, AFM calculates the offset and length data blocks of a file that is provided in

the `--list-file file`, and evicts only these data blocks from the file. With the range eviction, AFM provides fine granularity for eviction of specific data blocks from a file.

You can specify more than one files in the list of files. When **afmEvictRange** option is set on the fileset, auto-eviction and manual eviction (without range) evicts all the data blocks except first and last data blocks. By default the whole file is evicted. After this parameter enabled on an AFM fileset, a list of files contains offset and length that can be provided for eviction.

In the following example, 5G data blocks starting from 10G till 15G are evicted:

1. Enable the range byte eviction by setting the **afmEvictRange** parameter.

```
mmchfileset fs1 fset1 -p afmEvictRange=yes
```

2. Check whether the list-file is in the following format:

```
<Start_Offset> <length> <filepath>
```

3. List the list-file information.

```
ls -lsh /gpfs/fs1/ro1/bigfile1
```

A sample output is as follows:

```
19458 100G -rw-r--r-- 1 root root 100G Jul 16 23:35 /ibm/fs1/file1
```

4. Copy the file.

```
cat /tmp/rangefile
```

A sample output is as follows:

```
10G 15G /gpfs/fs1/ro1/bigfile1
```

5. Evict the range of data blocks from the file.

```
mmafmctl fs1 evict -j fileset --list-file /tmp/rangefile --range
```

AFM evicts data 5G data blocks starting from 10G till 15G, offset and length are aligned on block boundary. AFM evicts the range of data blocks from the file as per file system block size calculation.

1. Evicting by using the filter option:

```
node1:/gpfs/cache/fileset_IW mmlsfs fs1 -Q
flag value description
-----
Q user;group;fileset Quotas accounting enabled
none Quotas enforced
none Default quotas enabled
```

node1:/gpfs/cache/fileset_IW mmquotaon -v fs1

```
mmquotaon on fs1
```

node1:/gpfs/cache/fileset_IW mmlsfs fs1 -Q

```
flag value description
-----
Q user;group;fileset Quotas accounting enabled
user;group;fileset Quotas enforced
none Default quotas enabled
```

node1:/gpfs/cache/fileset_IW mmlsquota -j fileset_IW fs1

```
Block Limits | File Limits
Filesystem type KB quota limit in_doubt grace | files quota limit in_doubt grace
```

```
Remarks
fs1 FILESET 96 102400 2097152 0 none | 20 0 0 0 none
```

```
node1:/gpfs/cache/fileset_IW mmadmctl fs1 evict -j fileset_IW --filter
FILENAME='%work%'
```

```
mmadmctl: Run mmcheckquota command before running mmadmctl with evict option
mmadmctl: Command failed. Examine previous error messages to determine cause.
```

```
node1:/gpfs/cache/fileset_IW mmcheckquota fs1
```

```
fs1: Start quota check2 % complete on Thu Oct 27 09:19:41 2016
[...]
95 % complete on Thu Oct 27 09:20:03 2016
100 % complete on Thu Oct 27 09:20:03 2016
Finished scanning the inodes for fs1.
Merging results from scan.
```

```
node1:/gpfs/cache/fileset_IW mmadmlocal ls workfile1
```

```
-rw-r--r-- 1 root root 104857600 Oct 27 08:30 workfile1
```

```
node1:/gpfs/cache/fileset_IW mmadmctl fs1 evict -j fileset_IW --filter
FILENAME='%work%'
```

```
node1:/gpfs/cache/fileset_IW mmadmlocal ls workfile1
mmadmlocal: Command failed. Examine previous error messages to determine cause.
```

2. Manual eviction by using the **--list-file** option:

```
ls -lshi /gpfs/fs1/evictCache
total 6.0M
27858308 1.0M -rw-r--r--. 1 root root 1.0M Feb  5 02:07 file1M
27858307 2.0M -rw-r--r--. 1 root root 2.0M Feb  5 02:07 file2M
27858306 3.0M -rw-r--r--. 1 root root 3.0M Feb  5 02:07 file3M
```

```
echo "RULE EXTERNAL LIST 'HomePREPDAEMON' RULE 'ListLargeFiles' LIST
```

```
'HomePREPDAEMON' WHERE PATH_NAME LIKE '%" > /tmp/evictionPolicy.pol
```

```
mmapplypolicy /gpfs/fs1/evictCache -I defer -P /tmp/evictionPolicy.pol
-f /tmp/evictionList
```

```
#Edited list of files to be evicted
```

```
cat /tmp/evictionList.list.HomePREPDAEMON
```

```
27858306 605742886 0 --/gpfs/fs1/evictCache/file3M
```

```
mmadmctl fs1 evict -j evictCache --list-file /tmp/
evictionList.list.HomePREPDAEMON
```

```
Evicted (Total) Failed
1 (1) 0
tspccheevict: List of failures : /var/mmfs/tmp/evictionFailedList.mmadmctl.23349
mmadmctl: Command failed. Examine previous error messages to determine cause.
```

```
ls -lshi /gpfs/fs1/evictCache
```

```
total 3.0M
27858308 1.0M -rw-r--r--. 1 root root 1.0M Feb  5 02:07 file1M
27858307 2.0M -rw-r--r--. 1 root root 2.0M Feb  5 02:07 file2M
27858306   0 -rw-r--r--. 1 root root 3.0M Feb  5 02:07 file3M
```

3. Manual eviction by using the **--file** option:

```
ls -lshi /gpfs/fs1/evictCache
total 3.0M
27858308 1.0M -rw-r--r--. 1 root root 1.0M Feb  5 02:07 file1M
27858307 2.0M -rw-r--r--. 1 root root 2.0M Feb  5 02:07 file2M
27858306   0 -rw-r--r--. 1 root root 3.0M Feb  5 02:07 file3M
```

```
mmafmctl fs1 evict -j evictCache --file /gpfs/fs1/evictCache/file1M
```

```
ls -lshi /gpfs/fs1/evictCache/file1M
```

```
total 0
27858308 0 -rw-r--r--. 1 root root 1.0M Feb  5 02:07 file1M
```

4. Manual eviction of a partially cached file by using the **--file** option. In the following example, file1 is partially cached file of ro1 fileset. Out of 2MB, only 1MB data blocks are cached. You can use below steps to evict 1MB data blocks.

```
# ls -altrish /gpfs/fs1/ro1
```

```
total 2.0M
27858311 1.0M -rw-r--r--. 1 root root 2.0M Feb  5 02:07 file1
```

```
mmafmctl fs1 evict -j ro1 --file /gpfs/fs1/ro1/file1
```

```
# ls -altrish /gpfs/fs1/ro1/file1
```

```
total 0
27858311 0 -rw-r--r--. 1 root root 2.0M Feb  5 02:07 file1M
```

Note: If a file is evicted from an AFM fileset, the snapshot file of the evicted file contains zeros.

Operation with disconnected home

With a cache and home cluster that is separated by a wide area network (WAN), it might result in intermittent outages and possibly long-term disruptions.

If the primary gateway determines that the home cannot be accessed, the primary gateway waits until the interval time specified in the **afmDisconnectTimeout** parameter passes, and then changes the cache state to disconnected. This feature is available on all AFM filesets.

In a disconnected state, cached files are served to applications from the cache. Application requests for uncached data return an I/O error. All update operations from the cache complete, and return successfully to the application. These requests remain queued at the gateway until they can be flushed to home.

In a disconnected state, the home cannot be accessed for revalidation. Therefore, the latest updates from home are not available in the cache. Writes, which require revalidation with home might appear temporarily stuck if it is done after the Home fails, and before the cache moves to disconnected state. With revalidation stuck waiting on the unavailable home, the request is timed out because of the value set in the **afmDisconnectTimeout** parameter.

If cache filesets are based on the NSD protocol, when the home file system is not mounted on the gateway, the cache cluster puts the cache filesets into unmounted state. These cache filesets never enter the disconnected state.

If the remote cluster on the home cluster does not respond due to a deadlock, operations that require remote mount access, such as revalidation or reading uncached contents, stop responding until the remote mount becomes available again. This remote cluster response is true for AFM filesets that use the NSD protocol to connect to the home cluster. You can continue accessing cached contents without disruption by temporarily disabling all of the revalidation intervals until the remote mount is accessible again.

If a cache fileset is disconnected for an extended period, the number of file system updates might exceed the buffering capacity of the gateway nodes. In this situation, operations continue in the cache. When the connection to home is restored, AFM runs recovery and synchronizes its local updates to the home cluster.

AFM automatically detects when home is available and moves the cache into the Active state. The **afmHomeDisconnected** callback event and the **afmHomeConnected** callback event can be used to monitor when a cache changes state.

Fileset that is using a mapping target go into disconnected state if the NFS server of the primary gateway is unreachable, even when the NFS servers of all participating gateways are reachable.

Prefetch tasks that fail due to home disconnection, continue when home is available again.

The following example shows the number of read/write operations that were run while the home was in the disconnected mode.

ls -la

```
total 128
drwx----- 65535 root root 32768 Oct 13 16:50 .afm
-rw-r--r-- 1 root root 8 Oct 22 17:02 newfile2
drwx----- 65535 root root 32768 Oct 22 05:23 .pconflicts
drwx----- 65535 root root 32768 Oct 22 17:02 .ptrash
dr-xr-xr-x 2 root root 32768 Oct 13 16:20 .snapshots
```

for i in 1 2 3 4 5 ; do date > file\$i ; done

mmafmctl fs1 getstate

```
Fileset Name Fileset Target Cache State Gateway Node Queue Length Queue numExec
-----
fileset_IW nfs://node4/gpfs/fshome/fset002new Disconnected node2 27 102
```

turn on NFS on the home.

mmafmctl fs1 getstate

```
Fileset Name Fileset Target Cache State Gateway Node Queue Length Queue numExec
-----
fileset_IW nfs://node4/gpfs/fshome/fset002new Dirty node2 27 102
```

node1:/gpfs/cache/fileset_IW # mmafmctl fs1 getstate

```
Fileset Name Fileset Target Cache State Gateway Node Queue Length Queue numExec
-----
fileset_IW nfs://node4/gpfs/fshome/fset002new Dirty node2 27 102
```

date > file10 /gpfs/cache/fileset_IW/file10

mmafmctl fs1 getstate

```
Fileset Name Fileset Target Cache State Gateway Node Queue Length Queue numExec
-----
fileset_IW nfs://node4/gpfs/fshome/fset002new Active node2 0 134
```

Expiring a disconnected RO cache

Read-only (RO) filesets can be configured to expire cached data after the gateway nodes are in a disconnected state for a specified amount of time.

This feature provides control over how long a network outage between the cache and home can be tolerated before the data in the cache is considered stale. This control prevents access to old data, where old is defined by the amount of time that the WAN cache is out of synchronization with the data on the home site. Cached data is only available until the time period set in the **afmExpirationTimeout** parameter expires, at which point the cached data is considered 'expired' and cannot be read until the network is reconnected.

This feature is not available in other AFM modes.

RO cache filesets that are created by using the NSD protocol do not automatically go to the expired state, even if the expiration timeout is set for them. The RO cache filesets remain in the unmounted state until the home file system is unavailable. The administrator can manually expire the data in an RO cache before the expiration timeout is reached by using the **mmafmctl expire** command only if the expiration timeout is configured for the fileset. An RO cache that is created by using NSD protocol can also be manually expired. Use the following command:

```
mmafmctl Device {resync | expire | unexpire | stop | start} -j FilesetName
```

For more information, see *mmafmctl* command in *IBM Storage Scale: Command and Programming Reference Guide*.

1. Expiring an RO fileset -

```
mmafmctl fs1 expire -j ro1
```

mmafmctl fs1 getstate -j ro1

Name	Fileset	Target	Cache State	Gateway Node	Queue Length	Queue numExec
ro1		gpfs:///gpfs/remotefs1/dir1	Expired	c26c4apv1	0	4

2. Unexpiring an RO fileset -

```
mmafmctl fs1 unexpire -j ro1
```

#mmafmctl fs1 getstate -j ro1

Name	Fileset	Target	Cache State	Gateway Node	Queue Length	Queue numExec
ro1		gpfs:///gpfs/remotefs1/dir1	Active	c26c4apv1	0	4

The **afmFilesetExpired** callback event is triggered when a fileset moves into expired state. The **afmFilesetUnexpired** callback event is triggered when a fileset moves out of the expired state.

Viewing snapshots at home

All snapshots at home for RO and LU filesets can be viewed in the cache by setting the **afmShowHomeSnapshot** parameter to yes. This variable is not applicable for SW and IW filesets.

The home and cache can have different snapshot directories for clarity. If the **afmShowHomeSnapshot** parameter is changed after fileset creation, the change is reflected in the cache only after IBM Storage Scale is restarted, or the file system is remounted. However, if you change the **afmShowHomeSnapshot**

parameter value from yes to no during the lifetime of a fileset, it continues to show the home snapshots even after a restart, or file system remount.

For RO and LU filesets that use NFS, when a lookup is performed from the parent directory of `.snapshot` directory, the changes on the home are not reflected in the cache, as NFS cannot detect the changes in `.snapshot`. The **mtime/ctime** parameter of this directory is not modified on the home despite snapshots creates or deletes operations on the home. This is not true for RO and LU fileset that use GPFS backend as the latest snapshot directory gets reflected in cache.

AFM cannot detect a snapshot path on a home site that is not the IBM Storage Scale home site. To detect the snapshot path on this IBM Storage Scale or GPFS home when data is being migrated and prefetched, set the **afmShowHomeSnapshot** parameter. Because of this parameter, AFM detects the snapshot path on the home and does not pull snapshot files during migration.

You can set this parameter at a cluster level. Set this parameter if a home site is not the IBM Storage Scale site. On an IBM Storage Scale or GPFS home, you need not to set this parameter because AFM automatically detects the home snapshot path.

Changing home of AFM cache

AFM filesets continue to function though failures occurred on a home.

AFM filesets serve the cache applications with cached data. If the home is permanently lost, you can create a new home. An existing SW/IW AFM cache can make the following changes to the home. These changes are not supported on RO/LU filesets.

1. Replace the home with a new empty home where the new target is created by using the NFS or NSD mapping:

The administrator must run the **mmafmctl failover** command without **--target-only** option to point to the new home. The new home is expected to be empty. To ensure that extended attributes are synchronized, run the **mmafmconfig** command on the new home before you run the failover command. If the new target is a mapping, failover does not split data transfers and queues them as normal requests without parallel data transfer.

2. Replace any of the following on an existing home:

- a. Replace the communication protocol (NSD, NFS): The administrator must run the **mmafmctl failover** command without the **--target-only** option, by using the new target protocol.

Note: Only the protocol changes. The home path does not change.

- b. Enable or disable parallel data transfer by shifting between NFS and a mapping: The administrator must run the **mmafmctl failover** command without the **--target-only** option, by using the new target protocol or mapping.

Note: Only the protocol changes. The home path does not change.

- c. Replace either the IP address or the NFS server by using the same communication protocol and home path: The administrator must run the **mmafmctl failover** command with the **--target-only** option. The IP/NFS server must be on the same home cluster and must be of the same architecture as the old NFS server.

Note: Only the IP or NFS server changes.

During failover, ensure that the cache file system is mounted on all gateway nodes, and the new home file system is mounted on all the nodes in the home cluster.

Note: Failover does not use parallel data transfers.

When you create a new home, all cached data and metadata available in the cache are queued in the priority queue to the new home during the failover process. The failover process is not synchronous and completes in the background. The **afmManualResyncComplete** callback event is triggered when failover is complete. Resync does not split data transfers even if parallel data transfer is configured, and the target is a mapping.

If the failover is interrupted due to a gateway node failure or quorum loss, failover is restarted automatically when the cache fileset attempts to go to Active state.

When a cache has multiple IW filesets, the administrator must choose a primary IW cache and fail this cache over to a new empty home. All of the other IW cache filesets to the old home must be deleted and re-created. If another IW cache is failed over to the same home after it is failed over to another IW cache, all the data in that cache overwrites existing objects at home. For this reason, fail over to a non-empty home is discouraged.

When the failover function is likely to be used due to the failure of the old home, the admin must be cautious and disable automatic eviction. You need to ensure that the eviction does not free the cached data on the cache. The evicted data cannot be recovered if the old home is lost.

Each AFM fileset is independently managed and has a one-to-one relationship with a target, thus allowing different protocol backends to coexist on separate filesets in the same file system. However, AFM does not validate the target for correctness when a fileset is created. The user must specify a valid target. Do not use a target that belongs to the same file system as the AFM fileset. For more information, see **mmafmctl** command in *IBM Storage Scale: Command and Programming Reference Guide*.

The following example shows changing the target for an SW fileset. Consider a fileset_SW SW fileset of a fs1 file system that uses the nfs://node4/gpfs/fshome/fset001 home target. A failover is performed to a new target nfs://node4/gpfs/fshome/fset002new.

```
# mmlsfileset fs1 fileset_SW --afm
```

A sample output is as follows:

Name	Status	Path	afmTarget
fileset_SW	Linked	/gpfs/cache/fileset_SW	nfs://node4/gpfs/fshome/fset001

```
# mmafmctl fs1 getstate -j fileset_SW
```

A sample output is as follows:

Fileset Name	Fileset Target	Cache State	Gateway Node	Queue Length	Queue numExec
fileset_SW	nfs://node4/gpfs/fshome/fset001	Active	GatewayNode1	0	6

```
# mmafmctl fs1 failover -j fileset_SW --new-target nfs://node4/gpfs/fshome/fset002new
```

A sample output is as follows:

```
mmafmctl:Performing failover to nfs://node4/gpfs/fshome/fset002new
Fileset fileset_SW changed.
mmafmctl: Failover in progress. This may take while...
Check fileset state or register for callback to know the completion status.
```

```
# mmafmctl fs1 getstate -j fileset_SW
```

A sample output is as follows:

Fileset Name	Fileset Target	Cache State	Gateway Node	Queue Length	Queue numExec
fileset_SW	nfs://node4/gpfs/fshome/fset002new	NeedsResync	GatewayNode1	6	0

```
# mmafmctl fs1 getstate -j fileset_SW
```

A sample output is as follows:

Fileset Name	Fileset Target	Cache State	Gateway Node	Queue Length	Queue numExec
fileset_SW	nfs://node4/gpfs/fshome/fset002new	Recovery	GatewayNode1	6	0

```
# mmafmctl fs1 getstate -j fileset_SW
```

A sample output is as follows:

Fileset Name	Fileset Target	Cache State	Gateway Node	Queue Length	Queue numExec
fileset_SW	nfs://node4/gpfs/fshome/fset002new	Active	GatewayNode1	0 6	0

Note: After failover, the fileset changes to states such as NeedsResync or Recovery. Depending on the size or the elapsed time, the fileset remains in these transition states before the fileset turns into the Active state. The failover process is complete after the fileset is in Active state.

```
# mmlsfileset fs1 fileset_SW --afm
```

A sample output is as follows:

Fileset Name	Status	Path	afmTarget
fileset_SW	Linked	/gpfs/cache/fileset_SW	nfs://node4/gpfs/fshome/fset002new

Out-band Failover - You can choose to copy all cached data offline from the AFM cache to the new home with any tool that preserves modification time (**mtime**) with nanoseconds granularity. An example of such a tool is - rsync version 3.1.0 or later with protocol version 31. After the data is copied, you can run the **mmafmctl failover** command to compare **mtime** and **filesize** at home, and avoid queuing unnecessary data to home.

Resync on SW filesets

A conflict situation might arise due to an inadvertent write on an SW home. AFM runs a resynchronization automatically.

Writing data on an SW home is not allowed under normal circumstances. However, AFM does not enforce this condition. If an inadvertent write or corruption occurs on an SW home due to some error situation, it results in a conflict scenario for the single-writer fileset.

Note: Resync does not use parallel data transfers. Even if parallel data transfer is configured and the target is a mapping, the data transfers are not split.

If AFM detects inconsistencies during flushing, the fileset goes into NeedsResync state. When the fileset is in the NeedsResync state, AFM runs a resynchronization automatically during the next request to the primary gateway, and the inconsistencies are resolved.

If resync is not triggered automatically, the inconsistencies must be manually resolved by running the **resync** command. When you run a resync, the fileset must be in the Active or the NeedsResync state. Resync is not allowed on IW filesets as multiple cache filesets can update the same set of data. Use the following command - **mmafmctl Device {resync | expire | unexpire} -j FilesetName**. For more information, see *mmafmctl* command in *IBM Storage Scale: Command and Programming Reference Guide*.

```
mmafmctl fs1 resync -j sw1
```

Fileset Name	Fileset Target	Cache State	Gateway Node	Queue Length	Queue numExec
sw1	nfs://c26c3apv2/gpfs/homefs1/newdir1	Dirty	c26c2apv1	4067	10844

During a resync, all cached data and metadata are queued in the priority queue. The resynchronization is an asynchronous process and completes in the background. The `afmManualResyncComplete` callback event is triggered when the resynchronization is complete. When the complete priority queue is flushed, the fileset state changes to Active.

Evicted files are not synchronized to home.

A resync of partially cached files pushes only the cached data blocks to the new home. Uncached blocks are filled with null bytes.

Important: Do not use on an SW fileset, except when it is required to correct the home under such conflict scenarios.

Out-band Resync - You can choose to copy all cached data offline from the AFM cache to the new home with any tool that preserves modification time (**mtime**) with nanoseconds granularity. An example of such a tool is - rsync version 3.1.0 or later with protocol version 31. After the data is copied, you can run the **mmafmctl failover** command to compare **mtime** and **filesize** at home, and avoid queuing unnecessary data to home.

The Resync operation does not delete new files or directories that are created on the home nor new names that are given to files or directories at the home.

In some cases, a cache has pending changes such as delete and rename to replicate and the fileset recovery is triggered. In this case, the home might have old data and renamed files or directories are replicated to the home. Some extra files might exist at home because of the files or directories replication. The cache and the home have the latest copy of the files.

Note: When the **afmResyncVer2** parameter is enabled on an SW fileset or a primary fileset, the extra files are deleted. The resync operation is run on a cache fileset to ensure that the cache and the home are synchronized.

AFM synchronous read mount

Creates a separate mount path at the gateway node through which synchronous read calls are sent to the home cluster separately.

When read operation is performed on the AFM fileset, and the fileset is not available at cache, AFM fetches data from home to cache first, and then read operation is performed on the cached data blocks.

In few cases, where fetching data from home takes time, specially, in case where the file is migrated to the tape at the home and fetching this file at cache is slow. In cases of late response, EIO error is returned back to the application. This EIO error might cause queue to drop at the gateway node and AFM put the fileset into stale state.

Enable **afmSyncReadMount** creates a separate mount path at the gateway node through which synchronous read calls are sent to the home separately. This arrangement helps other operations that are running on the AFM fileset to synchronize with the home by using the default path. If due to any reason home is slow or busy, the read operation gets killed, and retried without affecting the other operations as this read operation is performed on the separate mount path.

Enabling **afmSyncReadMount** parameter might avoid dropping the whole queue at the gateway node. As a result, AFM avoids sending EIO error back to the application that performs the read operation and move AFM (SW/IW) fileset into stale state and avoid recovery in InfoSphere® MDM Server Workbench or IW fileset.

The **afmSyncReadMount** parameter can be enabled on the AFM Single-Writer (InfoSphere MDM Server Workbench), Independent-Writer (IW), Local-Update (LU), and Read-Only (RO) mode fileset by using **mmcrfileset** or **mmchfileset** command.

The valid values for the **afmSyncReadMount** are "yes" or "no".

To set or unset the **afmSyncReadMount** parameter, AFM fileset must be stopped or unlinked.

The **afmSyncReadMount** parameter can be enabled either at the time of fileset creation by using **mmcrfileset** or modify existing fileset by using **mmchfileset** command.

For more information, see *Stop and start replication on a fileset* in *IBM Storage Scale: Concepts, Planning, and Installation Guide*

Example

1. To create an AFM fileset and enable **afmSyncReadMount** on the fileset, issue the following command:

```
mmcrfileset fs1 afmRO -p afmMode=ro,afmTarget=home:/gpfs/fs1/target1,afmSyncReadMount=yes --
inode-space new
```

2. To update an existing AFM fileset and enable **afmSyncReadMount** on the fileset, AFM fileset must be stopped or unlinked. To stop the fileset, issue the following command:

```
mmafmctl fs1 stop -j afmRO
```

3. To set the **afmSyncReadMount** parameter, issue the following command:

```
mmchfileset fs1 afmRO -p afmSyncReadMount=yes
```

Note: Before you enable **afmSyncReadMount** parameter, ensure to upgrade the latest IBM Storage Scale cluster version and file system version. For more information, see *File system format changes between versions of IBM Storage Scale* in the *IBM Storage Scale: Administration Guide*.

4. After the parameter is set, fileset can be started or unlinked. To start the fileset, issue the following command:

```
mmafmctl fs1 start -j afmRO
```

AFM resync version 2

The AFM resync version 2 enhances the replication performance. This feature is suitable for a cluster that is employed for heavy stress and workloads. It uses an on-demand dependency resolution for queued messages, in case of resync and/or recovery is running. The message queuing performance is increased by using the runtime filtering or the dependency resolution. Thus, less memory is used on the gateway nodes and queued messages are replicated quickly. This feature also helps in the role reversal of AFM-DR filesets.

An AFM and AFM-DR primary fileset creates a file operation queue on the designated gateway nodes for a replication. These operations are run on the target or home cluster asynchronously based on the setting of the **afmAsyncDelay** parameter. For any dependent operations in the queue, this asynchronous delay helps to filter operations that need not to be sent to the target over an underlying protocol. This filtering of operations saves the bandwidth. For more information, see [“Asynchronous delay” on page 64](#).

When many incoming operations are generated because of heavy workloads, the queue length increases. On the gateway node, messages or operations are stored in the memory and processed for execution, filtering, or queued based on an operation type. For heavy workloads, the memory usage on nodes affects the system performance and the network throughput between cache and home filesets or primary and secondary filesets. When any gateway node fails, you must run demanding recovery and/or resync operations on a new gateway node. When the **afmResyncVer2** parameter is enabled queued messages are replicated based on the on-demand dependency resolution to the home or secondary whenever the queued messages are ready. The execution of queue to the home or secondary does not pause on certain dependent operations because the memory is available for more operations after the queues are replicated.

A single-writer (SW) cache fileset or an AFM-DR primary fileset can be configured for the AFM resync version 2 by setting the **afmResyncVer2** parameter in the **mmchfileset** command.

To set or unset the **afmResyncVer2** parameter on an AFM or AFM-DR fileset, you need to stop or unlink the fileset. For more information, see [“Stop and start replication on a fileset” on page 118](#).

AFM internally stores all the information necessary to replay the updates that are made in the cache to the home cluster. When a gateway node fails, the in-memory queue for any hosted fileset is lost. Any filesets that were hosted on the failed gateway nodes are transferred to another gateway nodes. The new gateway nodes rebuild queues in the memory. When a gateway node fails, the fileset that is hosted on the gateway node is transferred to another gateway node. This gateway node with the fileset builds the queue in the memory. This process is called recovery. During the recovery, outstanding cache updates are placed in the in-memory queue and the gateway processes the queue. AFM collects the pending operations by running a policy scan on the fileset. AFM uses the policy infrastructure in IBM Storage Scale to engage all the nodes that are mounting the file system to participate in the scan process. Pending requests, which are discovered by the recovery process, are queued in a special queue called the priority queue. At the same time, a normal queue is also created. The normal queue is used for new incoming operations to the fileset.

When the **afmResyncVer2** parameter is enabled, messages in the queue are checked for dependent operations. An example of dependent operations that need to be performed is as follows:

1. Create a file.
2. Change the file name.
3. Add some data the file.

However, there are other operations that can be run whenever they are available in the queue because they are not dependent on any other operations. The memory of a gateway node is saved where messages are queued.

Enable the AFM resync version 2 feature by using the following steps:

1. Verify the primary fileset information.

```
# mmlsfileset fs1 pri --afm -L
```

The sample output is as follows:

```
Filesets in file system 'fs1':

Attributes for fileset pri:
=====
Status                               Linked
Path                                /gpfs/fs1/pri
Id                                  1
Root inode                          524291
Parent Id                           0
Created                             Thu Feb 11 15:05:23 2021
Comment
Inode space                          1
Maximum number of inodes             100352
Allocated inodes                     100352
Permission change flag               chmodAndSetacl
afm-associated                        Yes
Target                               nfs://c7f2n06/gpfs/fs1/sec
Mode                                 primary
Async Delay                          15 (default)
Recovery Point Objective              disable (default)
Last pSnapId                         1
Number of Gateway Flush Threads      4
Primary Id                           2836795238842262449-C0A8693E60255310-1
IO Flags                             0x0 (default)
```

2. Stop the primary fileset.

```
# mmafmctl fs1 stop -j pri
```

3. Enable the AFM resync version 2 feature.

```
# mmchfileset fs1 pri -p afmResyncVer2=yes
```

4. Start the primary fileset.

```
# mmafmctl fs1 start -j pri
```

5. Verify the primary fileset information again.

```
# mmlsfileset fs1 pri --afm -L
```

The sample output is as follows:

```
Filesets in file system 'fs1':

Attributes for fileset pri:
=====
Status                               Linked
Path                                /gpfs/fs1/pri
Id                                  1
Root inode                          524291
Parent Id                           0
Created                             Thu Feb 11 15:05:23 2021
```

Comment	
Inode space	1
Maximum number of inodes	100352
Allocated inodes	100352
Permission change flag	chmodAndSetacl
afm-associated	Yes
Target	nfs://c7f2n06/gpfs/fs1/sec
Mode	primary
Async Delay	15 (default)
Recovery Point Objective	disable (default)
Last pSnapId	1
Number of Gateway Flush Threads	4
Primary Id	2836795238842262449-C0A8693E60255310-1
I/O Flags	0x10000 (afmResyncVer2)

AFM optimized recovery version v2

During the recovery event triggered on the AFM fileset, AFM determines all the operations that are pending and yet to synchronize with home. If rename and remove operations, AFM requires to read the home data.

If the AFM fileset contains large volume of data, recovery might take time to identify operations that are yet to synchronize and in pending state. To overcome this delay, AFM optimized recovery version v2 that can be enabled on the AFM fileset Single-Writer mode or AFM-DR primary mode fileset.

After the AFM optimized recovery version v2 is enabled, AFM avoids reading home directory, and uses GPFS policy to run at the home. This step helps in finding the relevant information to quickly identify the list of rename and recovery operations that needs to be synchronized and then queue them to the Gateway Node.

To run policy at home, AFM target must be running IBM Storage Scale GPFS filesystem. You can apply optimized AFM recovery version v2 on AFM Single-Writer and AFM-DR primary mode fileset only.

If newer AFM recovery version v2 observes any failure, AFM automatically uses the older version "v1" and try to complete the operation.

The AFM recovery version v2 can be enabled on the AFM Single-Writer or AFM-DR primary fileset. You can enable AFM recovery version v2 either at the time of fileset creation by using **mmcrfileset** command or by changing existing fileset by using **mmchfileset** command.

Before you enable the AFM recovery version "v2" on the existing fileset, the IBM Storage Scale cluster version and file system version must be upgraded to the latest release level.

Note: The **afmRecoveryVer2** parameter cannot be set on AFM Single Write mode or AFM DR primary mode fileset while running GPFS or NSD backend targets. This parameter is only supported on AFM fileset using NFS backend/target.

Example

1. To create an AFM fileset and enable recovery v2 on it, issue the following command:

```
mmcrfileset fs1 afmSW-recv2 -p afmMode=sw,afmTarget=home:/gpfs/fs1/
target1,afmRecoveryVer2=yes --inode-space newer
```

2. To update an existing AFM SW mode fileset and enable recovery v2 on it, AFM fileset must be stopped or unlinked. After the parameter is set, fileset can be started or unlinked.

```
mmafmctl fs1 stop -j afmSW-recv1
```

```
mmchfileset fs1 afmSW-recv1 -p afmRecoveryVer2=yes
```

```
mmafmctl fs1 start -j afmSW-recv1
```

For more information, see *File system format changes between versions of IBM Storage Scale* in the *IBM Storage Scale: Administration Guide*.

3. In IBM Storage Scale 5.1.6, new file systems are created at file system format level 29.00. To update a file system from an earlier format to format level 29.00, issue the following command:

```
mmchfs Device -V full
```

Where,

Device is the device name of the file system.

4. The "Enable AFM recovery version v2" feature of the IBM Storage Scale 5.1.6 requires a file system to be at format number 29.00 or later.

AFM recovery improvements for remove and rename operations

The AFM Recovery improvements for remove and rename operations are supported on AFM Single Writer and AFM DR primary mode where Home or Secondary is running IBM Storage Scale file system.

During a failure event, recovery is started by AFM to sync dirty data with home or secondary site. When recovery is triggered on an AFM fileset that contains large data, AFM performs various steps to synchronize the data between Cache and Home. This activity also includes identifying the remove and rename operations.

In case where the AFM fileset has large dirty data, this process might take some time to identify the rename and delete operations.

Finding rename and remove operations on AFM fileset data during recovery can be optimized by setting **afmRecoveryVer2** on the AFM Single Writer fileset or AFM DR primary fileset.

By enabling **afmRecoveryVer2**, AFM optimizes the process of finding IBM Storage Scale rename and remove operations by running policy at the home or secondary. Policy generates the list that is referred to identify the pending rename and remove operations at cache and AFM queues those operations to the Gateway Node.

Running policy scan at the home or secondary helps AFM to quickly identify the rename and recovery operations. This reduces time when directory has large files and various rename and remove operations that are pending for synchronization to the home.

After the **afmRecoveryVer2** is set on the AFM fileset, AFM tries to run recovery optimization. If optimized recovery cannot be performed, AFM fallback to the older recovery version and completes the recovery.

Before you enable **afmRecoveryVer2** parameter on the AFM fileset, cluster and file system must be upgraded to the latest release level.

Example

To set **afmRecoveryVer2** on the existing AFM fileset, stop the fileset first and then set the parameter. Do the following to perform this step:

```
mmafmctl fs1 stop -j SW1
```

```
mmchfileset fs1 SW1 -p afmRecoveryVer2=yes
```

```
mmafmctl fs1 start -j SW1
```

Using IBM Storage Protect for Space Management

IBM Storage Protect for Space Management can be used on the AFM filesets or on the home.

The following figure illustrates IBM Storage Protect for Space Management connected to home.




Figure 14. Sample setup of IBM Storage Protect for Space Management connected to home

A new file created at home becomes candidate for migration to the IBM Storage Protect server. When a migrated file at home is read from the cache, the file is recalled at home and served to the cache. If the cache does a write on a migrated file at home, the file is recalled at home and written to when the AFM queue is flushed. If multiple migrated files are written at the same time, the home issues recalls for all files at the same time. It is recommended that you exclude frequently changing files from IBM Storage Protect for Space Management migration process to avoid recalls.

The following figure illustrates IBM Storage Protect for Space Management connected to both home and cache.




Figure 15. IBM Storage Protect for Space Management connected to both home and cache

When using IBM Storage Protect for Space Management on an AFM fileset, the flag **AFMSKIPUNCACHEDFILES** must be set in the `dsm.sys` configuration file (IBM Storage Protect-related) to yes.

For example - **AFMSKIPUNCACHEDFILES** yes. This parameter should be used for read-write cache filesets. It prevents the migration of dirty and uncached files. If this flag is not set, it might result in long waiters or unexpected cache states and unexpected errors in terms of the IBM Storage Protect for Space Management migration processing. In the LU mode when this parameter is set and a file is made local because of updating data, migration of the file to tape might be prevented. For migration in the LU mode, do not unset this parameter. Migrated files cannot be evicted. File operations such as read, write, truncate, append or rename on migrated files recalls the files into the cache. When multiple migrated files are modified at the same time all recalls are submitted at the same time, IBM recommends that you exclude frequently changing files or frequently read files from IBM Storage Protect for Space Management migration process on both cache and home to avoid recalls.

It is recommended the following guidelines while using IBM Storage Scale AFM and IBM Storage Protect:

- Prevent cache eviction in combination with IBM Storage Protect on the same fileset. Both techniques have the same goal to reduce the space required in the fileset. The combination of both techniques unnecessarily increases the complexity of the environment.
- IBM Storage Scale snapshots and IBM Storage Protect have a limited compatibility. The deletion of a stub file that is reflected in a snapshot (the snapshot was generated before or after the file was migrated) causes the recall of the file data. The file data is stored in the snapshot so that it can be accessed later. Therefore, do not use snapshots for an AFM fileset (in home or cache) and in the file system hosting the AFM fileset, if you are using IBM Storage Protect for Space Management.
- When using IBM Storage Protect on home or cache be aware that access (read or write) to multiple migrated files at the same time causes bulk recalls. Access to multiple files can be caused by users such as when they copy an entire directory or by AFM when changed files are replicated to home where the previous versions are migrated. You can avoid these issues by using the optimized tape recall process, which requires a list of files to be recalled before processing.

When running IBM Storage Scale AFM and IBM Storage Protect backup operations, prevent cache eviction in combination with IBM Storage Protect backup on the same fileset, if possible. Evicted (uncached) files will be skipped from backup processing. This might lead to errors in terms of the versioning of files on the IBM Storage Protect server.

For detailed description about the setup and configuration of IBM Storage Protect for Space Management for AFM, see [Configuring IBM Spectrum Scale Active File Management](#).

Performing a planned maintenance by using the IW cache

You can perform a planned maintenance by using IW cache.

About this task

Let us assume that we have IW cache on Side A that is hosting applications. IW cache points to a home in Side B. Complete the following steps for a planned maintenance of Side A:

Important: Failback operations for Active File Management (AFM) independent writer (IW) cache mode are deprecated in IBM Storage Scale 6.0.1 and later. Deprecated indicates that this operation is still supported in the current release but is no longer recommended for use and might be removed in a future release.

Procedure

1. Stop all applications in IW cache (Side A).
2. Flush the pending queue on IW cache (Side A) to home (Side B) by using **mmafmctl Device flushPending -j FileSetName**.
3. Unlink IW cache (Side A).
4. Start applications at home (Side B).
5. Relink the cache after maintenance window.
6. Stop the applications at home (Side B).
7. Start the applications at cache (Side A).
8. Schedule the background prefetch if required (Side A).

The following figure illustrates IW cache (Side A) to home (Side B).




Figure 16. IW cache (Side A) to home (Side B)

Handling IW cache disaster

You can perform planned or unplanned maintenance by using the IW cache.

Performing an unplanned maintenance using IW cache

You can perform an unplanned maintenance by using the IW cache.

About this task

For an unplanned outage, some changes might not be synchronized with the home. IW can handle requests that are lost from the cache due to a disaster, while allowing application updates at the home to the same data.

The cache became unavailable because of the disaster at the IW cache. Therefore, the IW cache is down with pending updates (Side A). Steps to follow in an unplanned outage are:

Important: Failback operations for Active File Management (AFM) independent writer (IW) cache mode are deprecated in IBM Storage Scale 6.0.1 and later. Deprecated indicates that this operation is still supported in the current release but is no longer recommended for use and might be removed in a future release.

Procedure

1. Start applications at home (Side B). Record the time.
2. When the IW cache site (Side A) is ready to take over again, stop applications at the home site (Side B).

The following figure illustrates IW cache site (Side A) to home site (Side B).




Figure 17. IW cache site (Side A) and home site (Side B)

3. Link the cache site back (Side A) and ensure that the gateways are up. Do not access the fileset directories.

If the cache directories are accessed, recovery is triggered. Old contents from the cache, which were pending, are synchronized with the home at the end of recovery, over-writing latest application

changes at home. Perform the following steps so that the latest data is preserved and old or stale data is discarded. The data staleness is determined based on file `mtime` and failover time. The failover time is the time when applications were moved to the home, given as input to the **failback** command.

4. Run the following command from the cache site to synchronize the latest data, specifying the time when the home site became functional and including the complete time zone, in case home and cache are in different time zones:

```
mmafmctl FileSystem failback -j Fileset --start --failover-time  
'TimeIncludingTimezone'
```

The failback command resolves conflicts between pending updates in cache at the time of failover with the data changes at home. When the synchronization is in progress, the fileset state is `FailbackInProgress`. The fileset is read-only when the failback is in progress.

After failback completes, the fileset state is `FailbackCompleted`. The failback process resolves conflicts in the following way:

- a. Dirty data that was not synchronized is identified.
- b. If the changed files or directories were not modified at the home when the applications were connected to the home, the cache pushes the change to the home to avoid any data loss.
- c. If the files or directories are recently modified at the home, the cache discards the earlier updates from the cache. The next lookup on the cache brings the latest metadata from the home.

If the conflicted dirty data is a directory, it is moved to `.ptrash`. If the conflicted dirty data is a file, it is moved to `.pconflict`. The administrator must clean `.pconflict` directory regularly. If IW is converted to the other modes, `.ptrash` and `.pconflict` directories remain.

5. After achieving the `FailbackCompleted` state, run the following command to move the fileset from the Read-Only to the Active, where it is ready for use. If the command is not run successfully, run the command again.

```
mmafmctl FileSystem failback -j Fileset --stop
```

Note: If failback is not complete or if the fileset moves to `NeedFailback`, run the failback command again.

Results

The cache site is ready for use. All applications can start functioning from the cache site. New files that are created at home are reflected in the cache on the next access, based on the revalidation interval. Failback does not pull in data of uncached files from home, which needs to be done explicitly by the administrator by using the **mmafmctl prefetch** command. If failback might be used, the `ctime` of `RENAME` operations in the home file system must be updated. This is enabled by using **setCtimeOnFileRename** at home:

```
mmchconfig setCtimeOnFileRename=yes -i
```

Note: Failback does not work if revalidation is disabled in the IW cache site.

Using AFM with encryption

AFM supports file encryption. Encryption can be applied to AFM-managed filesets.

AFM home sites and cache sites can be enabled with encryption, independent of each other. The data is encrypted while at rest (on disk) and is decrypted on the way to the reader or application that is hosted on home and caches. However, AFM communication between home and cache is not encrypted.

With the data that is flowing between home and cache filesets not being encrypted by the adoption of file encryption. The communication between the clusters needs to be encrypted explicitly (if the privacy of the data over the network is a concern) by ensuring that a cipher list is configured. To ensure that the data is transmitted in the encrypted form, a cipher other than `AUTHONLY` must be adopted. `AES128-GCM-SHA256` is one of the recommended ciphers.

Run the **mmauth show** command to view the cipher lists used to communicate within the local and with the remote clusters. To ensure that all file content on disk and on the network is encrypted, configure file encryption at home and on the caches. Also, configure a cipher list on all the clusters, ensuring that ciphers are configured within and across clusters. Because of the file encryption, the data is transmitted in the encrypted form between NSD clients and servers (both directions). However, the file metadata or RPC headers are encrypted. Only the use of encrypted communications (cipher list) ensures that the entire message content gets encrypted.

If the NFS protocol is used for communication between the home and cache clusters, and privacy of the data over the network is a concern, then to encrypt the NFS transfers you need to set up an encrypted tunnel. If a cluster configured as cache includes an encrypted file system, then all nodes in the cluster, and especially the gateway nodes, require access to the Master Encryption Keys.

For encryption setup, see *Encryption* in the *IBM Storage Scale: Administration Guide*.

Using mmbackup

When the **mmbackup** command is run on a file system that has AFM filesets, only the cached data from the filesets is backed up.

The **mmbackup** command does not back up the uncached data from the AFM filesets.

In a file system that is IBM Storage Protect for Space Management managed, a backup operation in either the cache or home fileset skips migrated offline files that were not yet backed up. The backup of a file that is evicted from a cache causes the file to be read again from home to cache. To avoid reading the file again when eviction on the cache site is enabled, ensure that the backup occurs before the eviction.

Disabling AFM

An AFM fileset can be disabled and converted to a GPFS-independent fileset. However, GPFS-independent fileset cannot be converted to an AFM fileset.

Before you disable the AFM fileset, ensure that **afmShowHomeSnapshot** value is set to "no" on the specific AFM fileset. If **afmShowHomeSnapshot** is set to "yes" then you can modify the **afmShowHomeSnapshot** value by using **mmchfileset** command. After the value of **afmShowHomeSnapshot** is set to "no" you can proceed to disable the AFM fileset.

Issue the following command to set **afmShowHomeSnapshot** value to "no":

```
mmchfileset fs1 ro2 -p afmShowHomeSnapshot=no
```

The preceding command gives the output similar to the following as shown here:

```
mmlsfileset fs1 Fileset ro2 changed
```

To dissociate home and cache relationship of an AFM fileset, you can convert the AFM fileset to an independent fileset by using the **mmchfileset -p afmTarget=disable** command. Thereafter, the fileset is a GPFS or IBM Storage Scale-independent fileset and requests are not queued to home. An example is the case of NFS migration, where AFM fileset must be disabled after migration.

You must disable AFM filesets while the filesets are linked. In some cases, disabling of AFM fileset might fail with an error message.

Complete the following steps to disable an AFM fileset -

1. If the fileset is unlinked, run **mmlinkfileset gpfs0 FilesetName -J gpfs0/FilesetName** to link the fileset.
2. Run **mmchfileset gpfs0 FilesetName -p afmTarget=disable**.

The following warning message displays - Warning! Once disabled, AFM cannot be re-enabled on this fileset. Do you wish to continue? (yes/no)

Enter yes.

The following warning message displays - Warning! Fileset should be verified for uncached files and orphans. If already verified, then skip this step. Do you wish to verify same? (yes/no)

Enter yes.

- If error conditions are not found, the following message displays:

File system gpfs0 is mounted on *n* nodes or fileset *FilesetName* is not unlinked.

Unlink the fileset and run **mmchfileset gpfs0 FilesetName -p afmTarget=disable** again. Enter no to the warning message about verification of the fileset. The AFM fileset is disabled.

- In case of error conditions, disabling AFM fileset fails. See the following table for the error conditions with the corresponding messages.

Table 8. Conditions in which disabling AFM fileset fails, with corresponding messages	
Condition	Error message
Incomplete directories present	mmchfileset: 6027-2307 [E] Uncached files present, run prefetch first.
Uncached files present	Uncached files present, run prefetch first using policy output: /var/mmfs/tmp/cmdTmpDir.mmchfileset.18149/list-file.mmchfileset.18149
Orphans present	Orphans are present, run prefetch first using policy output: /var/mmfs/tmp/cmdTmpDir.mmchfileset.18149/list-file.mmchfileset.18149

You must run prefetch as suggested in the message to fix the error condition.

3. If uncached files or orphans are present, run prefetch by using the following command:

```
mmafmctl gpfs0 prefetch -j FilesetName --list-file /
var/mmfs/tmp/cmdTmpDir.mmchfileset.18149/list-file.mmchfileset.18149
```

4. If incomplete directories are present, generate a list file to prefetch the uncached files first. Use this generated list file that contains uncached file entries to run the command - **mmafmctl gpfs0 prefetch -j FilesetName --list-file list file path**. Here, *list file path* is the generated list file.
5. Ensure that all data in the AFM fileset is prefetched.
6. Run step “2” on page 114 again and complete disabling AFM fileset.

Disable AFM by using disable-online feature

Describes how to disable AFM by using disable-online feature.

An AFM fileset can be disabled by using the disable-online feature and converted to an independent fileset. Online disable feature can be used runtime and there is no need to unlink the fileset or take any downtime for applications. After the fileset is disabled, files and directories in that fileset will no longer sync with the target or remote cluster. The online disable feature is available only on AFM read only (RO) and Local updates (LU) mode filesets. After the fileset is AFM disabled by using online disable features, the fileset cannot be converted to an AFM fileset.

To dissociate home or target and cache relationship of an AFM fileset, you can convert the AFM fileset to an independent fileset by using **mmchfileset fsname filesename -p afmTarget=disable-online**. Thereafter, the fileset is an IBM Storage Scale independent fileset and file operations are not queued to home. An example is the case of NFS migration, where AFM fileset must be disabled after migration for local usage.

Note: Before you disable AFM by using disable-online feature ensure that all the files are cached (read completely) from the home or target fileset. Or, follow the interactive command to find the uncached files that needs to be prefetched.

The **mmchfileset** fsname filesetname -p afmTarget=disable-online command is interactive and provides an option to check whether files from home or remote fileset are cached into the AFM cache fileset. It is recommended that system administrators first verify the files that are ingested from home by providing “yes” input to verification of files and orphans prompt. After all the files are cached, the recommended step can be skipped and disable online can be run.

Example

Rotest1 is cache fileset

```
mmchfileset fs1 rotest1 --afm -L
```

```
Filesets in file system 'fs1':
Attributes for fileset rotest1:
=====
Status                               Linked
Path                                /gpfs/fs1/rotest1
Id                                  52
Root inode                           3407875
Parent Id                             0
Created                             Mon Jan 30 01:03:14 2023
Comment
Inode space                           26
Maximum number of inodes              100352
Allocated inodes                      100352
Permission change flag                chmodAndSetacl
afm-associated                        Yes
Permission inherit flag               inheritAclOnly
Target                               nfs://homeNode1/gpfs/fs1/rotest
Mode                                  read-only
File Lookup Refresh Interval          30 (default)
File Open Refresh Interval            60 (default)
Dir Lookup Refresh Interval           60 (default)
Dir Open Refresh Interval             60 (default)
Async Delay                           disable
Expiration Timeout                    disable (default)
Last pSnapId                          0
Display Home Snapshots                no (default)
Number of Gateway Flush Threads        4
Prefetch Threshold                    0 (default)
Eviction Enabled                       yes (default)
IO Flags                              0x0 (default)
IO Flags2                             0x0 (default)
```

Data in rotest1 fileset

```
ls -lR /gpfs/fs1/rotest1
```

```
/gpfs/fs1/rotest1:
total 16
drwxr-xr-x. 2 root root 8192 Jan 30 01:00 dir1
drwxr-xr-x. 2 root root 8192 Jan 30 01:00 dir2
-rw-r--r--. 1 root root    7 Jan 30 01:00 file1
-rw-r--r--. 1 root root    7 Jan 30 01:00 file2
-rw-r--r--. 1 root root    7 Jan 30 01:00 file3
-rw-r--r--. 1 root root    7 Jan 30 01:00 file4
-rw-r--r--. 1 root root    7 Jan 30 01:00 file5

/gpfs/fs1/rotest1/dir1:
total 0
-rw-r--r--. 1 root root 7 Jan 30 01:00 file1
-rw-r--r--. 1 root root 7 Jan 30 01:00 file2
-rw-r--r--. 1 root root 7 Jan 30 01:00 file3

/gpfs/fs1/rotest1/dir2:
total 0
-rw-r--r--. 1 root root 7 Jan 30 01:00 file1
```



```
-rw-r--r--. 1 root root 7 Jan 30 01:00 file2
-rw-r--r--. 1 root root 7 Jan 30 01:00 file3
```

Reading files as below (These files are cached)

```
cat /gpfs/fs1/rotest1/file* > /dev/null
cat /gpfs/fs1/rotest1/dir1/file* > /dev/null
```

Note: The files are not read from (as an example) /gpfs/fs1/rotest/dir2. These files are uncached.

Perform disable-online

```
mmchfileset fs1 rotest1 -p afmTarget=disable-online
```

```
Warning! Once disabled, AFM cannot be re-enabled on this fileset. Do you wish to continue?
(yes/no) yes
Warning! Fileset should be verified for uncached files and orphans. If already verified,
then skip this step.
Do you wish to verify same? (yes/no) yes
Verifying if all the data is cached. This may take a while...
mmchfileset: [E] Uncached files present, run prefetch first
Files list file: /var/mmfs/tmp/cmdTmpDir.mmchfileset.1152157/list-file.mmchfileset.1152157
```

Verification of cached files shows the files (uncached) in the list file

```
cat /var/mmfs/tmp/cmdTmpDir.mmchfileset.1152157/list-file.mmchfileset.1152157
%2Fgpfs%2Ffs1%2Frotest1%2Fdir2%2Ffile2
%2Fgpfs%2Ffs1%2Frotest1%2Fdir2%2Ffile1
%2Fgpfs%2Ffs1%2Frotest1%2Fdir2%2Ffile3
```

Use the list that is provided by the above command for prefetching as below

```
mmafmctl fs1 prefetch -j rotest1 --list-file /var/mmfs/tmp/cmdTmpDir.mmchfileset.1152157/
list-file.mmchfileset.1152157
```

```
mmafmctl: Performing prefetching of fileset: rotest1
Queued      ( Processed/      Total)      Failed      TotalData
              3          (      3/          3)          0      (approx in Bytes)
                                   21

prefetch successfully queued at the gateway.
```

```
mmchfileset fs1 rotest1 -p afmTarget=disable-online
Warning! Once disabled, AFM cannot be re-enabled on this fileset. Do you wish to continue?
(yes/no) yes
Warning! Fileset should be verified for uncached files and orphans. If already verified,
then skip this step.
Do you wish to verify same? (yes/no) yes
Verifying if all the data is cached. This may take a while...
All the data is cached, rerun the command to disable AFM.
mmchfileset: Command failed. Examine previous error messages to determine cause.
```

Above command make sure that all the data is cached. Proceed to disable without verification

```
mmchfileset fs1 rotest1 -p afmTarget=disable-online
```

```
Warning! Once disabled, AFM cannot be re-enabled on this fileset. Do you wish to continue?
(yes/no) yes
Warning! Fileset should be verified for uncached files and orphans. If already verified,
then skip this step.
Do you wish to verify same? (yes/no) no
Fileset rotest1 changed.
```

After fileset is disabled, it will no longer associate to AFM target or home fileset

```
mmllsfileset fs1 rotest1 --afm -L
```

```
Filesets in file system 'fs1':

Attributes for fileset rotest1:
=====
Status                               Linked
Path                                /gpfs/fs1/rotest1
Id                                  52
Root inode                          3407875
Parent Id                           0
Created                             Mon Jan 30 01:03:14 2023
Comment
Inode space                          26
Maximum number of inodes             100352
Allocated inodes                     100352
Permission change flag               chmodAndSetacl
afm-associated                       No
Permission inherit flag               inheritAclOnly
```

Stop and start replication on a fileset

The stop and start operations can be used during the downtime planning of the replication activities on an AFM or AFM-DR fileset. When you stop the replication on an AFM fileset, the fileset state changes to 'Stopped' and the replication between both sites stops until you perform the 'start' operation on the stopped fileset. These operations are useful to control the replication during the downtime and manage the bandwidth efficiently.

When the fileset state is 'Stopped', any data modification at the cache or the home is not synchronized. That is, the changed data at the cache does not reflect in filesets at the home, and the home data is not available at the configured cache.

After the downtime, you need to perform the 'start' operation to start the replication activities on the fileset. The start operation first triggers the 'recovery' phase on the fileset, which determines the local changes and prepares the changed list for the replication. This list is sent to the assigned gateway node to synchronize the data with the home.

During the recovery phase, all pending operations are synchronized by using the Priority Queue on the gateway node. AFM or AFM-DR replicates a full file again if the file has any changes during the recovery phase because AFM does not track changed data at the file level. After the recovery phase is complete and all changes are synchronized to the home, the AFM fileset state changes to 'Active'.

If AFM-DR primary fileset is not synchronized to the secondary fileset and the primary fileset is in 'PrimInitInProg' during the maintenance, AFM queues all files again and synchronizes the queued files to the AFM-DR secondary. When the primary is synchronized with the secondary and 'psnap0' is created at both sites, AFM changes the primary fileset state and leaves the 'PrimInitInProg' state.

- To stop replication of an AFM fileset, issue the following command:

```
mmafmctl <fs> stop -j <fileset>
```

You can check the fileset state by using the following command:

```
mmafmctl <fs> getstate -j <fileset>
```

A sample output is as follows:

Fileset Name	Fileset Target	Cache State	Gateway Node	Queue Length	Queue numExec
afmFileset	nfs://p7fbn09/gpfs/fs0/swfileset3	Stopped			

The state of the fileset is now 'Stopped'.

- To start replication on a specific fileset, issue the following command:

```
mmafmctl <fs> start -j <fileset>
```

If you perform other operations like stop or flushpending on a fileset that is in the 'Stopped' state, the system displays an error message as follows:

```
Operation not permitted
mmafmctl: [E] Operation is not permitted as state of the fileset is Stopped.
mmafmctl: Command failed. Examine previous error messages to determine cause.
```

A similar message like this error message is displayed, when you perform the start operation on a fileset that is not in the 'Stopped' state.

AFM IPv6 Support

Internet Protocol Version 6 (IPv6) is the protocol that is designed by the IETF to replace Internet Protocol Version 4 (IPv4). Active File Management (AFM) fileset now supports connectivity and administration over IPv6.

AFM, AFM-based asynchronous disaster recovery (AFM DR), and AFM to cloud object storage fileset supports connectivity and administration over IPv6. The IBM Storage Scale cluster can be configured to use its daemon and admin communication over IPv6 addresses.

AFM gateway nodes and mapping support IPv6 address as well. After the IBM Storage Scale cluster is set up for using IPv6 as its network, filesets can be created by using these IPv6 addresses. Also, the filesets can be converted to use IPv6 addresses by using **mmafmctl failover -target-only** option and AFM DR filesets can be converted to use IPv6 by using **mmafmctl ChangeSecondary --target-only** option.

To set up IPv6 support for IBM Storage Scale cluster, see *Enabling a cluster for IPv6* in the *IBM Storage Scale: Administration Guide*.

To change IP addresses or hostnames of cluster nodes, see *Changing IP addresses or host names of cluster nodes* in the *IBM Storage Scale: Administration Guide*.

Important: The IPv6 addresses must be mentioned in the square brackets when the IPv6 is configured with the AFM-related commands.

Examples

1. Setting a remote cluster AFM target as IPv6.

```
mmcrfileset fs2 drp-fset-2 --inode-space=new -p afmtarget=[2001:192::321:5eff:febf:a69b]:/
gpfs/fs1/drs-fset-2 -p afmmode=primary
```

The output looks similar to the following as shown:

```
Fileset drp-fset-2 created with id 258 root inode 63963139.
Primary Id (afmPrimaryId) 134157187695697640-C0A876075E27E51D-258
```

2. Verifying fileset status.

```
mmafmctl fs1 getstate
```

The output looks similar to the following as shown:

Fileset Name	Fileset Target	Cache State	Gateway
Node Queue Length Queue numExec			
-----	-----	-----	
obj1	https://s3.amazonaws.com:443/mufileset	Active	
c7f2n04	0 143119		
ip1	nfs://[2001:192::310:18ff:fec6:ecd8]/fileset	Active	
c7f2n03	0 1167977		
ip2	nfs://[2001:192::310:18ff:fec6:ecd8]/fileset1	Active	
c7f2n03	0 524752		

Note: To verify whether the hostname resolves to an IPv6 addresses, use **mmcmi host <hostname>** command.

Example of changing the target address from IPv4 to IPv6 by using failover

Use the following details for changing the target address from IPv4 to IPv6 by using failover.

After the IBM Storage Scale cluster is enabled with IPv6 addresses, following steps can be used to failover existing filesets that are connected to the remote cluster over IPv4 to IPv6 addresses.

Note: Remote IBM Storage Scale cluster is also enabled with IPv6 addresses.

Cache cluster

```
mmiscluster
```

```
GPFS cluster information
```

```
=====
```

```
GPFS cluster name:      pry.Node2
GPFS cluster id:        2836795238850206208
GPFS UID domain:        pry.Node2
Remote shell command:   /usr/bin/ssh
Remote file copy command: /usr/bin/scp
Repository type:        CCR
```

Node	Daemon node name	IP address	Admin node name	Designation
1	Node2	2001:192::42f2:e9ff:fe0a:5bfc	Node2	quorum-manager
2	Node3	2001:192::42f2:e9ff:fe0a:5094	Node3	quorum-gateway
3	Node4	2001:192::42f2:e9ff:fe0a:619c	Node4	quorum-manager-gateway
4	Node5	2001:192::42f2:e9ff:fe0a:5794	Node5	quorum

Note: Here, the AFM cache cluster and AFM remote or home cluster is converted to IPv6.

Existing connection over IPv4

1. `mmafmctl fs1 getstate -j ip1`

Fileset Name	Fileset Target	Cache State	Gateway Node	Queue Length	Queue numExec
ip1	nfs://192.168.38.101/fileset	Active	Node3	0	57588

2. Performing failover by using **mmafmctl failover --target-only** option.

```
mmafmctl fs1 failover -j ip1 --new-target [2001:192::210:18ff:fec6:ecd8]:/fileset --target-only
```

```
mmafmctl: Performing failover to nfs://[2001:192::210:18ff:fec6:ecd8]:/fileset
Fileset ip1 changed.
```

Note: **--target-only** option changes the Ips of the target. When the **--target-only** option is not provided, it synchronizes all the data from Cache to target.

3. Create a file in a fileset and make sure that the relation becomes active.

```
echo filedata > /gpfs/fs1/ip1/file1
```

```
mmafmctl fs1 getstate -j ip1
```

Fileset Name	Fileset Target	Cache State	Gateway Node	Queue Length	Queue numExec
ip1	nfs://[2001:192::210:18ff:fec6:ecd8]:/fileset	Active	Node3	0	3

Creation of mapping for parallel data transfer by using IPV6 address and hostname

To enable parallel data transfer on IPV6, both home clusters nodes and cache cluster nodes must be running on IPV6. A Gateway Node that runs IPV6 must be able to communicate to the individual home export node to synchronize data across the clusters.

Mapping of home cluster nodes and gateway nodes at cache that are running on IPV6 can be mapped together either by using IPV6 address or the hostname.

Examples of creating mapping by using IPV6 address and hostname

1. Mapping creation by using IPV6 address

As you start creating mapping by using IPV6 addresses, IPv6 address must be written inside a square bracket '['].

The following example demonstrates how to create mapping by using IPV6 address:

```
mmafmconfig add map1 --  
export-map [2001:192::42f2:e9ff:feaf:6aa8]/[2001:192::42f2:e9ff:feaf:85a0],  
[2001:192::42f2:e9ff:feaf:6c30]/[2001:192::d756:f6d2:2710:330f]
```

System output:

```
mmafmconfig: Command successfully completed  
mmafmconfig: Propagating the cluster configuration data to all affected nodes  
This is an asynchronous process.
```

2. Mapping creation by using hostname

Creation of mapping by using hostname is exactly similar to IPv4 mapping creation.

The following example demonstrates how to create mapping by using the hostname:

```
mmafmconfig add mapx1 --export-map c80f3m5n06/c80f2m5n02,c80f3m5n05/c80f5m5n14
```

System output:

```
mmafmconfig: Command successfully completed  
mmafmconfig: Propagating the cluster configuration data to all affected nodes  
This is an asynchronous process.
```

In the preceding command example, c80f3m5n06 and c80f3m5n05 are export nodes at the home cluster and c80f2m5n02 and c80f5m5n14 are Gateway nodes at the cache cluster.

In the preceding examples, all the hostname resolution is pointing to the IPv6 address.

Note:

- Mix of IPv6 and IPv4 mapping is not supported.
- For more information about Parallel data transfer, see [Parallel data transfers](#).

After the mapping is created, user can create an AFM fileset by using mapping. The following command demonstrates how a user can create an AFM fileset by using mapping:

```
mmcrfileset <FS> <fset_name> -p afmMode=<AFM Mode>,afmTarget=<protocol>://<Mapping>/  
<remoteFS_Path>/<Target> --inode-space new
```

AFM fast remove for large directories

The AFM fast remove is an option for performance optimization. This option deletes large directory trees in Active File Management (AFM) filesets faster to reduce the delete operation at a home cluster.

The AFM fast remove option is an AFM feature that changes how directory delete operations are handled when a parent directory is deleted at a home cluster. AFM handles eligible child directory delete operations locally at the cache rather than replicating them to the home cluster.

During directory deletion, AFM evaluates each queued delete operation. If the parent directory no longer exists at the home cluster, AFM treats the remaining child directories as no longer relevant to home-side consistency. Those child directories are then removed locally at the cache.

At the home cluster, the directory is moved into a designated AFM trash location, `.afm/.afmtrash` directory, and is later cleaned up by internal AFM processes.

Components

Cache cluster

The system where directory deletes can be completed locally when fast remove conditions are met.

Home cluster

The authoritative location that tracks parent directory removal and manages deferred cleanup.

AFM trash directory

A temporary holding area at the home cluster for directories pending final cleanup. The trash directory is at `.afm/.afmtrash` within the fileset.

AFM control logic

Internal mechanisms that schedule and perform deferred directory deletion. The control logic automatically runs cleanup operations to remove content from the trash directory.

Use cases

The AFM fast remove feature is designed for the following scenarios:

- Removal of large directory trees with many nested subdirectories.
- Environments where delete performance affects application availability.
- Workloads that require minimum delete-related network traffic.
- Load reduction on the home cluster during large cleanup operations.

Limitations and considerations

Consider the following limitations when using the AFM fast remove feature:

- The AFM fast remove feature applies only when the parent directory is deleted at the home cluster.
- The feature is supported for single-writer and primary AFM filesets only.
- The fileset is not an AFM to cloud object storage fileset because the AFM fast remove option is not supported on this fileset.
- If automated cleanup cannot complete, directories can remain temporarily in the AFM trash location.

AFM cache synchronization with outband updates

AFM out-of-band cache synchronization maintains data consistency between an Active File Management (AFM) cache fileset and a home or target fileset without dependence on user or application access.

In distributed file system environments, reliance on access-driven cache population can result in inconsistent cache state and variable access latency. AFM out-of-band cache synchronization addresses such conditions through background evaluation and alignment of cache contents with a home or target fileset.

The mechanism supports predictable cache readiness in environments where data update frequency exceeds access-based refresh patterns.

AFM scans the cache fileset, stores scan state information, and compares the current scan with the previous scan. The comparison identifies new, modified, and deleted files and directories. The system combines new and changed data into a prefetch set and pulls the required data from the home fileset into the cache. This process runs independently of applications, users, and administrator monitoring.

The out-of-band updates operations

- **Namespace scan:** Collection of file and directory structure within the cache fileset. All files and directories are scanned.
- **State record:** Preservation of snapshot data for comparison. The state is stored in the `/.afm/.outbandupdate/` directory.
- **Delta analysis:** Identification of differences between current and previous snapshots.
 - New files and directories (`.newFiles`)
 - Modified files and directories (`.changedFiles`)
 - Deleted files and directories (`.deletedFiles`)
- **Prefetch selection:** Identification of entries that require cache population or update. The `.prefetchList` contains new and changed entries.
- **Prefetch operation:** Operation independent of user or application access paths. Updated content from home are pulled in to the cache.

Use case

- Frequent data updates at the home cluster.
- Scheduled workloads that require cache readiness before execution.
- Access patterns that do not refresh cache contents naturally.
- Read-dominant workloads with strict latency requirements.
- Operational environments that require continuous cache consistency.

Limitations and considerations

- Support applies only to specific AFM cache modes and excludes Local Update (LU) configurations.
- Data transfer uses network bandwidth and cache storage resources.
- Initial cache population requires greater resource usage than incremental execution cycles.
- Cache eviction does not trigger automatic repopulation if a subsequent change does not occur.

Relationship with on-demand cache

On-demand AFM cache population relies on application access to trigger data transfer, which can introduce unpredictable latency. Out-of-band cache synchronization shifts refresh activity to background execution and supports predictable cache availability.

Administrative value

AFM out-of-band cache synchronization enables deterministic cache behavior in distributed storage environments. Separation of cache refresh from application access reduces first-access latency and improves performance consistency for critical workloads.

Writable snapshots in the AFM local update mode

Writable snapshots provide read and write access to point-in-time data copies for safe testing, development, and analysis workflow without modifying the original file system. An isolated writable layer by using local update mode fileset is created over read-only snapshots.

Snapshots and writable snapshots

A global snapshot preserves the contents of an entire file system at a single point in time. A fileset snapshot preserves the contents of an independent fileset at a single point in time. Therefore, these types of snapshots are read-only; you cannot modify the preserved data.

A writable snapshot extends snapshot capabilities by adding a writable layer over a read-only snapshot. In IBM Storage Scale, the writable layer is implemented as a cache fileset that uses AFM in local update (LU) mode and references the snapshot path as the fileset data source. You can read, create, modify, and delete files from a point-in-time copy without changing the original file system or snapshot.

For more information about fileset snapshots, see the *Fileset-level snapshots* section in *IBM Storage Scale: Administration Guide*.

Writable snapshot implementation

You can configure a cache fileset that uses AFM LU mode as a writable snapshot by setting home target of the fileset to a snapshot path. The home target parameter specifies the snapshot path that the fileset uses as a data source.

The home file system contains the primary workload and source data. You can create a snapshot from the entire home file system or an independent fileset in the file system. You then create the writable snapshot as a cache fileset that uses AFM LU mode in a separate file system that acts as the cache file system.

The cache file system can be in the same IBM Storage Scale cluster as the home file system or in a different cluster. The cache fileset connects to the snapshot path by using the GPFS protocol (including remote mount for access across clusters) or NFS protocol.

Notes:


- If the cache and home file systems are in the same cluster, use GPFS as the protocol for the cache fileset.
- If you use NFS as the AFM protocol, you must export the snapshot path as a read-only NFS export.

Data flow and file behavior

After you create a writable snapshot, all files and directories appear in the fileset when applications run lookup operations or when metadata prefetch operations occur. Data blocks remain in the snapshot path until read operations require the data blocks.

- When an application reads a file or when prefetch operations retrieve files, AFM retrieves the file from the snapshot and stores the file in the cache fileset.
- When an application creates a file or modifies an existing file, the system stores the new or changed data only in the fileset.
- When an application deletes a file in the fileset, the deletion applies only to the writable snapshot. The system deletes newly created or modified files only in the fileset, not in the source snapshot.

Writable snapshot components



Home file system

The file system that contains the primary workload and source data.

Snapshot path

The path to the global snapshot or fileset snapshot that the fileset uses as a data source.

Cache file system

The separate file system that hosts the cache fileset, which acts as the writable snapshot.

Cache fileset

The fileset that uses AFM LU mode to present the snapshot contents as writable data and stores all local changes.

Protocol

The connection method from the cache fileset to the snapshot path. Use GPFS (including remote mount for configurations across clusters) when both file systems are in IBM Storage Scale clusters. Use NFS when the snapshot path is exported as a read-only export.

Use cases

- Clone an entire file system or directory tree for development, testing, or analysis without changing production data.
- Process files that existed when a snapshot was created without changing the source data.
- Provide multiple isolated writable views of the same snapshot for different users or workloads.

Considerations

- Snapshots remain read only. The cache fileset provides write access without changing snapshot behavior.
- Ensure that the cache file system has sufficient capacity to store new and modified data.
- If you delete the source snapshot, data that was not yet retrieved into the cache fileset becomes inaccessible.
- Changes that you make in the writable snapshot are isolated and do not propagate to the original file system or snapshot.

For more information about the writable snapshots configuration, see the *Configuring writable snapshots* section in *IBM Storage Scale: Administration Guide*.

The failover mapping for AFM NFS targets

The failover mapping provides NFS resiliency for Active File Management (AFM) filesets that use mapped NFS export targets by selecting another available export server when the current server becomes unavailable.

The failover mapping is an AFM feature that maintains data synchronization between the cache and home clusters when a mapped home NFS export server becomes unavailable. This applies only to AFM filesets that use a target map for exported NFS targets.

Without the failover mapping, AFM can detect that a mapped NFS export server is unavailable, but AFM continues to depend on the failed server assignment. As a result, synchronization can stop until a working export server is selected through configuration changes or recovery actions.

When the **afmFailoverMap** parameter is enabled at the fileset level, and the **afmParallelMounts** parameter is enabled at the cluster level, AFM evaluates the NFS export servers that are listed in the target map and uses the next available server for the fileset. This behavior reduces disruption when a home export server fails and synchronization is continues through another mapped export server.

The target map defines the relationship between gateway nodes in the cache cluster and NFS export servers in the home cluster. AFM uses that mapping to determine which alternate export servers are available for the fileset.

If another mapped export server is available, AFM switches data synchronization to that server. If no mapped export server is available, AFM moves the fileset to the unmounted state.

The valid values for the **afmFailoverMap** parameter are yes and no. The default value is no.

To enable the failover mapping, see the *Enabling the failover mapping* section in *IBM Storage Scale: Administration Guide*.

Asynchronous auto prefetch for AFM NFS-based filesets

Asynchronous auto prefetch synchronizes created, modified, or deleted data from the home cluster to the AFM cache fileset at specified intervals. This automatic synchronization provides quick insights to AI applications, reduces costs, improves performance, enhances security, and streamlines operations.

Asynchronous auto prefetch synchronizes data changes from the home cluster to an AFM cache fileset at regular intervals. When you set the **afmAsyncPrefetchInterval** parameter, the system detects file and directory creations, modifications, and deletions on the home cluster and replicates them to the cache. This feature is available for filesets in independent writer (IW) and read-only (RO) modes. For cloud object storage filesets, see [“Asynchronous auto prefetch” on page 181](#).

Configuration parameters

The **afmAsyncPrefetchInterval** parameter controls the synchronization interval. Set this parameter by using the **mmchfileset** command for existing filesets or the **mmcrfileset** command during the creation of a fileset.

- The default value is disabled.
- The recommended value is 30 minutes.
- The valid range is 0 to 2147483647 minutes.

Watch folder integration

When you enable clustered watch folder for an AFM fileset in the RO or IW mode, the system generates event notifications for file and directory operations. The clustered watch folder detects file system activities such as creation, deletion, modification at the AFM fileset level and streams notifications to an external message queue. For more information, see [“Clustered watch folder” on page 256](#).

Watch folder events are generated at the cache cluster when the asynchronous prefetch process downloads changed metadata and data from the home cluster. For each change on the home cluster, the cache fileset creates or updates the corresponding file or directory. Deletions on the home cluster trigger corresponding deletions on the cache.

Example: Setting an asynchronous prefetch interval for existing filesets

```
mmchfileset fs1 afmR01 -p afmAsyncPrefetchInterval=60
```

A sample output is as follows:

```
Fileset afmR01 changed.
```

This command sets the asynchronous prefetch interval to 60 minutes for the *afmR01* fileset in the *fs1* file system.

Example: Setting an asynchronous prefetch interval during the creation of a fileset

```
mmcrfileset fs1 afmR01 -p afmtarget=home1:/gpfs/fs1/export1 \  
--inode-space=new --inode-limit=10000 \  
-p afmmode=ro,afmFileOpenRefreshInterval=disable,\  
afmFileLookupRefreshInterval=disable,\  
afmDirOpenRefreshInterval=disable,\  
afmDirLookupRefreshInterval=disable,\  
afmAsyncPrefetchInterval=60
```

A sample output is as follows:

```
Fileset afmR01 created with id 5 root inode 2097155.
```

This command creates a new AFM read-only fileset with an asynchronous prefetch interval set to 60 minutes. The fileset is configured with:

- The home1:/gpfs/fs1/export1 target path.
- AFM mode is read-only (RO).
- All refresh intervals are disabled.
- Asynchronous prefetch interval set to 60 minutes.

Example: Verifying asynchronous prefetch interval settings

```
mmlsfileset fs1 afmR01 --afm -L
```

A sample output is as follows:

```
Filesets in file system 'fs1':

Attributes for fileset afmR01:
=====
Status                               Linked
Path                                 /gpfs/fs1/afmR01
Id                                   5
Root inode                           2097155
Parent Id                             0
Created                               Wed May 13 18:54:23 2026
Comment
Inode space                           4
Maximum number of inodes              10000
Allocated inodes                      10000
Permission change flag                chmodAndSetacl
afm-associated                         Yes
Target                                home1:/gpfs/fs1/export1
Mode                                  read-only
File Lookup Refresh Interval          disable
File Open Refresh Interval            disable
Dir Lookup Refresh Interval           disable
Dir Open Refresh Interval              disable
Async Delay                           15 (default)
Async Prefetch Interval                60
Last pSnapId                          0
Display Home Snapshots                no
Number of Gateway Flush Threads       4
Prefetch Threshold                    0 (default)
Eviction Enabled                       yes (default)
IO Flags                              0x0 (default)
```

Note: The asynchronous prefetch interval is set to 60 minutes. The system synchronizes changes from the home cluster every 60 minutes.

For more information about configuration, see [Configuring asynchronous auto prefetch for AFM filesets with NFS targets](#)

Benefits

Asynchronous auto prefetch provides the following benefits:

Reduced latency

Applications access current data from the local cache instead of the remote home cluster.

Lower network usage

Only changed data transfers between clusters to reduce bandwidth consumption.

Simplified operations

Synchronization occurs automatically at the configured interval without manual data transfers.

Consistent data access

Multiple applications use the same synchronized cache to maintain consistent data across workloads.

Requirements and limitations

- Available only for filesets in IW and RO modes.
- The **afmAsyncPrefetchInterval** parameter is disabled by default. Set a value to enable synchronization.

AFM limitations

Active File Management (AFM) limitations are as follows:

- If you change an AFM SW/IW home export from Kernel NFS or CNFS to the new IBM Storage Scale CES in 4.1.1 or later, AFM does not recognize the home. The best way to manage this situation is to run the **mmafmctl failover** command. See the IBM Storage Scale documentation for more details.
- AFM RO/LU home export cannot be changed from Kernel NFS or CNFS to the new IBM Storage Scale CES in 4.1.1. or later.
- Customers who enabled DMAPI/IBM Storage Protect for Space Management in the AFM home cluster must be running at RHEL 6.3, or later, or SLES 11 SP2, or later, on the cache cluster. There is a bug that is identified in earlier levels of RHEL, which requires that the exported path at home is excluded from DMAPI/IBM Storage Protect for Space Management. For more information, see [IBM Support](#).
- Code upgrade limitations:
 - If there are AFM filesets that use the Kernel NFS backend to communicate with a GPFS home that is running V4.1 or earlier, upgrading home to IBM Spectrum Scale V4.1.1 or later causes the AFM filesets to disable synchronization to home with a message as follows:

```
GPFS: 6027-3218 Change in home export detected. Synchronization with home
is suspended until the problem is resolved
```

Note: This upgrade affects only AFM filesets that were originally created with a GPFS home that runs V4.1 or earlier. This upgrade does not affect AFM filesets that were created with a home that runs IBM Storage Scale V4.1.1 or later. This issue is fixed in 4.1.1.4 and 4.2.0.1 available at Fix Central ([IBM Fix central](#)).

- AFM filesets do not support the gpfs API – **gpfs_prealloc**.
- Because of a known issue with AFM SMB integration, the cache cluster might experience queued messages on gateway nodes. As a result, SMB clients might experience a performance impact on data transfers.
- Do not change an *fsid* of the NFS export path on the home site.

Related concepts

“General guidelines and recommendations for AFM SW and IW mode fileset” on page 454

Use the IBM guidelines and the recommendations to configure and maintain AFM Single Writer (SW) and Individual Writer (IW) mode fileset.

AFM and AFM DR limitations

The common limitations for both AFM and AFM DR functions include:

- The following table shows the OS and architecture support matrix and supported AFM function:

Table 9. Supported function and supported OS/architecture			
OS and architecture	Fileset	Node roles	Comments
Linux on Power® and x86	SW/IW/RO/LU primary/secondary	Gateway Application Manager	
Linux on Z	SW/IW/RO/LU primary/secondary	Gateway Application Manager	
AIX on Power	SW/IW/RO/LU primary	Gateway Application	AIX cannot be an AFM home site or an AFM-DR secondary site.

- An AFM home site that is not an IBM Storage Scale cluster:
 - An AFM cache site can be configured to connect with any home site that provides access to the home data through the NFSv3 protocol. The home site is not required to be an IBM Storage Scale cluster.
 - In this scenario, the **mmafmconfig** command cannot not be issued on the AFM home site. Therefore, the AFM cache site cannot cache or update extended attributes (EAs), access control lists (ACLs), file sparseness, or AFM psnaps.
 - The IBM Storage Scale support team helps customers to investigate problems in this scenario but only to address issues with the AFM cache site that runs IBM Storage Scale. Any problems at the AFM home site, or its file system, must be addressed by the customer.
 - For AFM-DR configurations, both the primary and secondary sites must be IBM Storage Scale clusters.
- An AFM home site that runs IBM Storage Scale on AIX:
 - The **mmafmconfig** command is not supported on AIX. Therefore, the AFM cache site cannot cache or update extended attributes (EAs), access control lists (ACLs), file sparseness, or AFM psnaps.
 - AFM does not support AIX nodes as an exporter of NFS. That is, AFM does not support NFS mounting exports that are reached through a system that runs on AIX.
 - Because of the previous limitations, AIX nodes cannot act as an AFM home site or as an AFM-DR secondary site.
 - AIX is not supported as a gateway node for AFM fileset modes.
 - AIX is not supported as an AFM home site or AFM-DR secondary site.
- The **mmclone** command is not supported on AFM cache and AFM DR primary filesets. Clones that are created at home for AFM filesets are treated as separate files in the cache.
- Quality of service is not tested on AFM and AFM Async DR filesets.
- Connecting an AFM-DR fileset to IBM Storage Protect for Space Management is not supported.
- The NSD protocol for communication does not support AFM and AFM DR filesets on Linux on Z.
- Connecting AFM DR fileset to IBM Spectrum Archive is not supported.
- AFM and Async DR is not supported on clusters that have Windows nodes.
- Cascading relationships with AFM caches and AFM primary filesets are not tested.
- Fileset snapshot restore is not supported for AFM and AFM DR filesets.
- Files in an AFM or AFM DR fileset, which are shared by using SMB, an error message E_ACCESS might appear if SMB holds locks on the files. When the locks are released by SMB, AFM replicates again and processes the files from the AFM or AFM DR fileset.
- The **dm_write_invis()** function is not supported on AFM and AFM DR filesets.
- The **dm_read_invis()** function is not supported on AFM and AFM DR filesets, if a file is not cached.

- The File Audit logging feature is not supported on AFM or AFM DR filesets.
- **Immutability** and **appendOnly** features are not supported on AFM filesets.
- AFM and AFM DR filesets do not replicate security, system, or trusted extended attributes.
- The **--iam-mode** option is not supported on AFM filesets.
- AFM and AFM DR do not support DM Punch Hole method (GPFS API '**dm_punch_hole**').
- To enable the Kerberos security for an AFM-DR fileset, the AFM-DR secondary must be given the read/write access.
- If a file is evicted from an AFM fileset, the snapshot file of the evicted file contains zeros.
- Eviction is not supported on an AFM-DR primary fileset and an AFM-DR secondary fileset.
- The following file attributes are maintained locally in the cache but these are not replicated:
 - Control attributes
 - Direct I/O
 - Replication factors
 - Fileset quotas
 - Storage pool flags
 - Special file types such as FIFO, socket, block, character device
- Hard links can be created on the cache. Hard links on the home are displayed as individual files in the cache during the lookup. To prefetch hard links from home, run the **--metadata-only** option as the first operation.
- File locking across the cache and the home is not supported.
- User extended attributes, ACLs, and sparse files supported only on the home where the **mmafmconfig** command is run.
- Parallel data transfer is supported only in Linux-only clusters. If a home is a mix of architectures (x86 and ppc), parallel data transfer works only for the set of nodes that belong to any one architecture, depending on which architecture serves the data transfer first.
- If encryption is configured on a home site that is running on a file system, which AFM uses as a GPFS backend target (multi-cluster remote mount), of an IBM Storage Scale cluster, ensure that the cache cluster is also configured the same way as the home cluster. Because of this configuration, AFM can access the files on the target file system for the replication.
- Online replication of a dependent fileset is not supported on an AFM DR fileset with a GPFS/multi cluster backend.

Related concepts

[“General guidelines and recommendations for AFM SW and IW mode fileset” on page 454](#)

Use the IBM guidelines and the recommendations to configure and maintain AFM Single Writer (SW) and Individual Writer (IW) mode fileset.

AFM-based asynchronous disaster recovery (AFM DR)

The following topics introduce you to AFM-based Asynchronous Disaster Recovery (AFM DR).

The following figure illustrates the AFM disaster recovery process.




Figure 18. Asynchronous disaster recovery

Introduction

Active File Management-based asynchronous disaster recovery (AFM DR) is a filesset-level replication disaster-recovery capability.

The disaster recovery solution includes the following features:

- Providing business stability during a disaster
- Restoring business stability after the disaster is repaired.
- Enduring multiple disasters
- Minimizing data loss because of a disaster

AFM-based asynchronous disaster recovery is an AFM-based filesset-level replication disaster-recovery capability that augments the overall business recovery solution. This capability is a one-to-one active-passive model and is represented by two sites: primary and secondary.

The primary site is a read/write filesset where the applications are currently running and has read/write access to the data. The secondary site is read-only. All the data from the primary site is asynchronously synchronized with the secondary site. The primary and secondary sites can be independently created in storage and network configuration. After the sites are created, you can establish a relationship between the two filesets. The primary site is available for the applications even when communication or secondary fails. When the connection with the secondary site is restored, the primary site detects the restored connection and asynchronously updates the secondary site.

The following data is replicated from the primary site to the secondary site:

- File-user data
- Metadata including the user-extended attributes except the inode number and a time
- Hard links
- Renames

The following file system and fileset-related attributes from the primary site are not replicated to the secondary:

- User, group, and fileset quotas
- Replication factors
- Dependent filesets

AFM DR can be enabled on GPFS-independent filesets only.

Note: An independent fileset that has dependent filesets cannot be converted into an AFM DR fileset.

A consistent view of the data in the primary fileset can be propagated to the secondary fileset by using fileset-based snapshots (psnaps). Recovery Point Objective (RPO) defines the frequency of these snapshots and can send alerts through events when it is unable to achieve the set RPO. RPO is disabled by default. The minimum time that you can set as RPO is 60 minutes. AFM-based asynchronous DR can reconfigure the old primary site or establish a new primary site and synchronize it with the current primary site.

Individual files in the AFM DR filesets can be compressed. Compressing files saves disk space. For more information, see *File compression* in the *IBM Storage Scale: Administration Guide*.

Snapshot data migration is also supported. For more information, see *ILM for snapshots* in the *IBM Storage Scale: Administration Guide*.

When a disaster occurs on the primary site, the secondary site can be failed over to become the primary site. When required, the filesets of the secondary site can be restored to the state of the last consistent RPO snapshot. Applications can be moved or failed over to the acting primary site. This application movement helps to ensure stability with minimal downtime and minimal data loss. This makes it possible for applications to eventually be failed back to the primary site as soon as the (new) primary is on the same level as the acting primary.

Note: AFM DR does not offer any feature to check consistency of files across primary and secondary sites. However, you can use any third-party utility to check that consistency after files are replicated.

You can simultaneously configure a site for continuous replication of IBM Storage Scale data along with AFM DR site. With IBM Storage Scale continuous replication, you can achieve a near disaster recovery and a far disaster recovery with AFM DR site.

Recovery time objective (RTO)

Recovery time objective (RTO) is the time that is taken to convert a secondary fileset into an active primary fileset and move applications to that site after a disaster.

After a disaster at the primary site, the secondary site becomes the acting primary by running the **mmafmctl failoverToSecondary** command. The **--norestore** option is set by default. The secondary fileset is converted to acting primary but avoids restoration of data from the last RPO snapshot and maintains the existing data. As an administrator, you can run the **mmafmctl failoverToSecondary** command with the **--restore** option to clean the existing data at the secondary fileset and restore all data from the last RPO snapshot.

RTO can be estimated if the administrator knows the volume of data that is generated in an RPO interval. The volume of data that must be restored from the last snapshot depends on the number of files that changed in an RPO interval. Because a snapshot restore requires a complete scan of the snapshot to determine which files changed. The failover time is one of the components that an administrator must use to calculate the RTO for the application.

Modes and concepts

AFM DR uses the same underlying infrastructure as AFM. AFM DR is characterized by two modes: the fileset in the primary cluster uses the primary mode and the fileset in the secondary cluster uses the secondary mode.

AFM DR is supported over both NFS v3 and GPFS protocol. The primary fileset is owned by the primary gateway, which communicates with the NFS server on the secondary side. The primary-secondary relationship is strictly one-to-one.

AFM revalidation does not apply to primary filesets. All files are always cached because primary is the only writer and secondary is in the read-only mode.

You can convert the SW/IW relationship to a DR relationship. However, you cannot convert a DR relationship to an SW/IW relationship.

AFM-based Asynchronous Disaster Recovery features

The following sections describe the AFM-DR features.

The following AFM features are offered on AFM DR filesets:

- Force flushing contents before Async Delay
- Parallel data transfers
- Peer snapshot - psnap
- Gateway node failure and recovery
- Operation with disconnected secondary
- Using IBM Storage Protect for Space Management
- Disabling AFM DR
- Using AFM DR with encryption
- Stop and start replication on a fileset

You can use **mmbackup** command to back up all files from primary, as all files are in a cached state on the primary fileset. Similar to AFM filesets, IBM Storage Protect can be connected to primary or secondary, or both sides. When IBM Storage Protect is connected to the primary side, set AFMSKIPUNCACHEDFILES yes in dsm.sys file. AFM features such as revalidation, eviction, prefetch, partial file caching, expiration, resynchronization, failover, and showing home snapshots are not offered on AFM DR filesets.

Immutability and appendOnly for AFM-DR filesets

By enabling an integrated archive manager (IAM) mode on an Active File Management Disaster Recovery (AFM-DR) primary and secondary fileset, you can set the Immutability and appendOnly flags on a file or a directory in an AFM-DR fileset.

You cannot change or rename an immutable file. On a file or directory with the appendOnly flag enabled, you can do append operations, but you cannot do delete, modify, or rename operations.

After you set these flags on a file or a directory in an AFM-DR primary fileset, AFM automatically synchronizes these metadata flags on the secondary fileset. You can set the Immutability and appendOnly flags independently. If both Immutability and appendOnly flags are set on a file, the Immutability flag takes effect.

To enable the Immutability and appendOnly flags, you need to enable an IAM mode on an AFM-DR primary and secondary fileset. All IAM modes are supported. The behavior of the Immutability and appendOnly flags is the same for a GPFS file and a GPFS directory. To configure an IAM mode, see *Enabling integrated archive manager (IAM) modes on AFM-DR filesets* in the *IBM Storage Scale: Administration Guide*.

For more information about these flags, see *Immutability and appendOnly features* in the *IBM Storage Scale: Administration Guide*.

Note: You must do the following upgrade for the Immutability and appendOnly flags:

- Upgrade all nodes in a cluster to the latest release level.
- Upgrade the file system version to the latest version.
- Upgrade the cluster minRelease version to the latest version.

If an AFM-DR primary fileset and an AFM-DR secondary fileset has a linked dependent fileset, ensure that the same IAM mode, which is set on the primary fileset and the secondary fileset, is set on all linked dependent filesets.

If an AFM-DR primary and an AFM-DR secondary fileset has one or more linked dependent filesets and you want to set the compliant mode, do the following steps:

1. Set the compliant IAM mode on the primary fileset and all linked dependent filesets.
2. Set the compliant IAM mode on the secondary fileset and all linked dependent filesets.

Note: To enable IAM mode on a dependent fileset, you need not to run the **mmafmctl stop** command or the **mmafmctl start** command. These commands are needed only to set an IAM mode on an AFM-DR primary fileset.

RPO snapshots

Recovery point objective (RPO) snapshots are peer snapshots that are taken at the same time on the primary and the secondary sides. RPO is disabled by default, and the minimum value you can set is 60 minutes and increment in multiples of 5 minutes. You can update the **afmRPO** parameter while you create the primary to change the interval, or after you created the fileset by using the **mmchfileset -p** command with the **afmRPO** parameter.

An example of the command -

```
mmchfileset <FileSystemName> <filesetname> -p afmRPO=<time_in_minutes>
```

The appropriate RPO intervals for each application setup are determined by the following factors:

- The rate of data change
- The fileset parameters such as **afmAsyncDelay** and **afmNumFlushThreads**
- The network bandwidth between the two sites

Each RPO interval triggers a snapshot at fileset level on both the primary and the secondary sides. Snapshot quiescing is performed at the individual fileset level, not at the entire file system level. Quiescing a fileset is a data-transfer intensive action. As further optimization, RPO snapshots are only taken for primary filesets whose data or metadata is modified since the last RPO period.

The **afmAsyncDelay** parameter specifies the minimum time that the primary site must wait before flushing the pending data that needs to be transferred to the secondary site. An RPO request in queue flushes the queue before you create a snapshot on the secondary site. The general guideline for the **afmAsyncDelay** parameter is to set the asynchronous delay less than the RPO interval.

The **afmRPOMiss** event occurs when RPO is missed because of a network delay or failure to create the RPO snapshot on the secondary site. If RPOs are missed, AFM waits for the expiration of the next RPO and during next RPO, a message is logged in **mmfs.log** and a new RPO snapshot is created. Failed RPOs are queued on the gateway again, and are run again at the secondary site until they succeed. At any point in time, two RPO snapshots that are not **psnap0** are retained on both sides. If RPOs are not played at the secondary due to some reason and primary does not get acknowledgment for the create from the secondary, the RPO is also deleted from primary. To improve the performance of more than one fileset taking RPOs at the same time, the RPO snapshots are batched together. In this process, the RPOs of few filesets might get slightly adjusted to use the batching facility.

Normal RPO deletions from the primary site are performed and queued when the new RPO snapshots are received. While there is every attempt to ensure that the delete operations are successful on the secondary site, there can be some extra snapshots as some delete operations might not be successful. In such cases, you must manually delete the extra snapshots on the secondary site by using

mmde1snapshot -p. Apart from automatic RPO snapshots, you can create user snapshots by running the **mmpsnap create** command with the **-rpo** option. These snapshots are intermediate between scheduled RPOs and are not deleted during the RPO snapshot delete process.

When the primary site fails, you can roll back the last consistent snapshot present on the secondary site while converting it to the current primary site. However, it is recommended to convert the secondary as-is into the acting primary because in most cases the data in the secondary is always greater than or equal to the last RPO snapshot. When applications are moved to the current primary fileset after the failover operation, RPO snapshots are not taken because no secondary is associated with the current primary. Therefore, RPO is not applicable to the current primary site. The gateway node mapped to serve a fileset is called the primary gateway of the fileset. The primary gateway acts as the owner of the fileset. If the primary gateway fails, another gateway node takes over the fileset as primary gateway. The recovery is triggered on the new primary gateway while taking over the fileset. If the next RPO snapshot is due at the time when recovery gets triggered, a new RPO snapshot is taken and used for recovery. Whenever RPO snapshots are used for recovery, the normal cleanup process of the older RPO snapshots does not take place. Therefore, some extra RPO snapshots are temporarily displayed on the primary and secondary sites and are cleaned up on subsequent RPO intervals. No regular RPO snapshots are taken on a fileset that is running recovery until the recovery process is complete.

All user-created and system-created RPO snapshots except the **psnap0** belonging to an active primary-secondary must be deleted before the primary-secondary is deleted. You can change the RPO interval by running **mmchfileset**. The time for the next RPO snapshot (**Tnext**) is calculated by adding the RPO interval that you set (**Irpo**) to the time at which the previous snapshot occurred (**Tprev**): **Tnext = Tprev + Irpo**. If no file in the RPO snapshot changes during the interval, AFM does not take the snapshot, and the time for the next RPO snapshot is calculated as **Tnext = Tnext + Irpo**.

Note: The RPO interval is not added to the time at which you set the RPO interval, but to the time of the last RPO snapshot. For example, the previous snapshot was at 9:00 AM and the RPO interval is 720 minutes. If you change the RPO interval to 780 minutes during the 12 hours from 9:00 AM, the next RPO snapshot occurs after 780 minutes from 9:00 AM.

The RPO management deletes the oldest snapshot when it creates a new snapshot but it never deletes the **psnap0**. The **mmpsnap** command can be run to delete **psnap0**. By deleting **psnap0**, storage utilization improves because data blocks used by **psnap0** are not held down. These data blocks can be significant over a period of time. Also, deletion of **psnap0** can improve the performance of subsequent creation and deletion of RPO snapshots. However, you must delete **psnap0** only when other RPO snapshots are present in the primary and secondary filesets to handle disaster recovery. If **psnap0** is deleted without any other RPO snapshots, data is lost. Also, **psnap0** can be deleted from secondary in cases like failing over to the secondary, and a new primary being set up.

How to decide if you need RPO

Failover

RPO should be enabled if you need to go back in time after a failover by restoring the data on the secondary from the last RPO snapshot. But this is NOT recommended, as the data in the secondary is always greater than or equal to last RPO snapshot. Hence in most cases, RPO is not needed for failover.

Failback

If the primary comes back after a temporary failure, you would only need to copy the data that has changed in the secondary to the primary. For this, RPO function needs to be enabled. However, the RPO can be a large value, such as 24 hours. This value will dictate how much data is copied back to primary from the secondary during a failback. The higher the value, the larger the amount of data copied.

Factors for deciding the RPO interval

RPO interval dictates how often snapshots are created on the primary and secondary. The entire file system must be quiesced before taking a snapshot, even if the snapshot is for a single fileset in a massive file system.



CAUTION: Setting RPO intervals on primary fileset has significant implications on the performance of the fileset.

- Classify your filesets based on how critical the data is - the more critical the data the lower the async delay.
- To calculate this approximately, review the frequency/pattern of data generation at the primary and considering the number of GW nodes, network latency and other parameters, predict the time taken for this data to get synchronized with the secondary. The RPO must be set to this time.
- Do not set many filesets to the same RPO interval.

Role reversal

When a planned or unplanned failover occurs, the role reversal process reverses a secondary site to an acting primary site and the old primary site to the secondary site. The role reversal is recommended to handle the failover.

When a primary site fails, applications can move to a secondary site after the secondary site is promoted to the primary role. The secondary site acts as a primary site for the application. After reversal of role, the secondary site can continue to act as a primary site until the old primary site is back. When the failed site (old primary site) is restored, it takes over the role of the secondary site.

In the traditional “Failback to primary” method, when the failed primary site is restored, workloads are transferred from the acting primary site (old secondary site) to the restored primary site. This method is recommended when the rate of change at the primary is not high, and workloads can be easily transferred between two sites.

In the role reversal method, the secondary site permanently acts as a primary site. After the primary site is restored, it is designated as a secondary site. The role reversal method is recommended when the rate of change at the primary is high, and the filesets are huge. If you adopt the role reversal method, ensure that you remove extra files and old snapshots from the new secondary site.

Important: For a planned or unplanned failure, it is recommended to use the role reversal method instead of failover and failback methods. The failover and failback methods will be deprecated soon.

To reverse the role, do the following steps:

1. When you plan the role reversal, ensure that the primary and secondary sites are in synchronization.
2. Issue the following command on the secondary site after the primary site failure:

```
mmafmcctl Device failoverToSecondary -j FilesetName [--norestore |--restore ]
```

3. After the secondary site is promoted to the primary site, move applications on the primary site.
4. After the old primary site is restored, prepare the old primary site to promote to the secondary site.
5. If the role reversal is unplanned, ensure that the primary site has dirty files. These files were created or modified but were not replicated to the secondary site because of failures.
 - Do not delete the dirty files that were available at both sites but were not in sync because of the primary site failure. After the role reversal, the acting primary site overwrites these files on the secondary site (the old primary site) when the site is restored and configured.
 - Dirty files that were newly created at the old primary site but not replicated to the secondary site because of the primary site failure. These files are extra files at the old primary (secondary) after the role reversal.
 - These files are overwritten if the same named files are created at the acting primary site and replicated to the secondary (old primary) site.
 - If you choose again the old primary site as a primary site, then these created files are replicated to the secondary site.
6. To prepare the old primary site, do the following steps:
 - a. Unlink the fileset by issuing the following command:

```
# mmunlinkfileset fs fileset -f
```

- b. Disable the role of fileset at the primary site by running the following command:

```
# mmchfileset fs fileset -p afmTarget=disable
```

- c. Link back the fileset by issuing the following command:

```
# mmlinkfileset fs fileset -J path_to_fs/fileset
```

- d. In case of an unplanned role reversal, some changes are not yet replicated from an old primary to a new primary site. However, before the replication, the primary that have an active queue fails. In this case, run the **failoverToSecondary** parameter as in step “2” on page 136. If you want to keep the old primary, which is the reversed secondary site, to the same RPO snapshot level, run the **mmrestorefs** command to revert the fileset to the RPO snapshot level for the application consistency.

```
# mmrestorefs fs latestRPO -j fileset
```

Note: If the old primary is not restored, it does not cause the data inconsistency at a file system level or at a fileset level. But, it might cause data inconsistency for the application. For example, a file on an old primary with new data changes that are not replicated to an old secondary. If before this replication completes, there is a failure at the old primary and this data does not replicate to the old secondary, and the old primary is not reverted to the RPO snapshot level during the role reversal, the both sites are synchronized because of the following scenarios:

- Some data of the file is replicated to the old secondary, but some data is still being replicated. In this case, the reverse relationship establishment overwrites the file from the reversed primary to the reversed secondary, and both the sites are synchronized.
- The file does not exist on the old secondary. It is still in a queue on the old primary when it failed. In this case, one of the following things might happen:
 - In the reverse relationship, the reversed primary has a file with the same name. Because of the same name the new file from the reversed primary is overwritten to the reverse secondary, and both the sites are synchronized.
 - The reverse relationship does not affect the new file. Therefore, the reversed primary does not know whether the new file is on the reversed secondary. For such a new file, when the role is reversed second time, the original primary becomes primary again. The relationship establishment synchronizes the new file to the original secondary and both filesets data becomes consistent at the fileset level and on a snapshot.

7. Get the primary ID of the fileset at the acting primary site by issuing the following command:

```
# mmafmctl fs getPrimaryId -j fileset
```

8. Convert the fileset of the old primary site to the fileset of the secondary site by issuing the following command at the old primary:

```
# mmafmctl fs convertToSecondary -j fileset --primaryid PrimaryID
```

9. Prepare to export this secondary fileset path for use by the acting primary fileset.

10. Issue the following command at the acting primary to build the relationship:

```
# mmafmctl fs changeSecondary -j fileset --new-target oldPrimary:filesetpath --inband
```

11. Remove old snapshots from the old primary because the old snapshots contain old data.

The primary site (the old secondary site) is now the primary site and the old primary site is the secondary site. The roles of the primary site and the secondary site are interchanged and the relationship is restored.

Note: The role reversal method is used if filesets are large, and a high rate of change at the primary site.

Failover to the secondary site

When the primary site stops functioning, all applications must move to the secondary site to ensure continuity. Run the following command on the secondary to convert the secondary site into acting primary site: **mmafmcctl <FileSystemName> failoverToSecondary -j <FilesetName> [--norestore | --restore]**.

You can use the **--restore** option to restore the data from latest snapshot or the **--norestore** option to keep the all existing data during the failover process on the secondary site. The **--norestore** option is the default option. After the failover, the secondary site acts as a primary site and hosts applications.

The customer must ensure that the NFS export used for the secondary site is not removed so that after failback to the primary, the original primary site can communicate with the secondary site, which is the current primary. Also, the RPO snapshots existing on the current primary must not be deleted because the RPO snapshots are used during the failback process when the original primary site starts functioning again.

RPO snapshots are temporarily disabled on the active primary site. The new data is not replicated to any other site until a new primary is set up, or the original primary is reinstalled, and the current primary is converted back to the secondary site.

Failing back to the old primary site

The old primary site must be restored when it is repaired.

About this task

Important: For a planned or unplanned failure, it is recommended to use the role reversal method instead of failover and failback methods. For more information about the role reversal, see [“Role reversal” on page 136](#).

Complete the following steps to restore the old primary site when it is repaired and it is back online after the disaster:

Procedure

1. Issue the following command on the old primary:

```
# mmafmcctl Device failbackToPrimary -j FilesetName { --start | --stop }[--force]
```

The **--start** option restores the primary to the contents from the last RPO on the primary before the disaster. With the **--start** option, the primary is in the read-only mode. This mode avoids accidental corruption until the failback process is completed. After the old primary site starts functioning again, all RPOs that were present before the disaster can be accessed. If a common RPO snapshot `psnap0` is not present, the old primary site can be converted to a normal GPFS fileset. To set up the primary site, see the steps in [“Failing back to the new primary site” on page 139](#).

If the **--start** option that is run on the fileset is unsuccessful, next the **--start failbackToPrimary** option might not be allowed. You can use the **--force** option to start failback again.

While the primary is coming back, as an administrator ensures that I/O does not occur on the primary fileset before the start of the **failback --start** process.

2. Issue the following command on the old primary:

```
mmafmcctl Device { applyUpdates | getPrimaryId } -j FilesetName
```

This command applies differences that are created by applications on the old primary site as the current primary site took over the applications.

All the differences can be brought over in a single or multiple iterations. For minimizing the application downtime, this command can be run repeatedly to synchronize the contents of the original primary site

with the current primary site. When the contents on both the sites are as close as possible or have minimal differences, applications must take a downtime and this command must be run one last time. applyUpdates might fail with an error during instances when the acting primary is overloaded. In such cases, the command needs to be run again.

3. On the old primary, complete the failback process by running **mmafmctl** with the **failbackToPrimary --stop** option.

With this command, the fileset is in the read/write mode. The primary site is ready for starting the applications. If the **--stop** option of the failback does not complete due to errors and you cannot stop the failback, it can be forced to stop with the **--force** option.

4. Convert the current primary site back to the secondary site, and set the primary ID. Unlink the acting primary site, change it to secondary by issuing the following command on the acting primary (secondary):

```
# mmchfileset device fileset -p afmMode=secondary -p afmPrimaryID=primaryid)
```

NFS can be restarted on the secondary site to ensure that the secondary export is accessible to the primary site. The primary and secondary sites are connected back as before the primary disaster and all data from the primary is played on the secondary site. Regular RPO also resumes on the primary site.

Failing back to the new primary site

You can configure a new primary if the old primary site stopped functioning.

About this task

Important: For a planned or unplanned failure, it is recommended to use the role reversal method instead of failover and failback methods. For more information about the role reversal, see [“Role reversal” on page 136](#).

Complete the following steps to configure a new primary site:

Procedure

1. Create a GPFS fileset on the new primary site. This fileset is configured as the new primary.
2. Create the latest snapshot from the acting primary by running the **mmafmctl** command with the **replacePrimary** option.
A new psnap0 with all the latest contents is created on the current primary server. This snapshot is used in the subsequent steps in this task.
3. Copy the contents from the snapshot to the new primary site by using **scp** or other means outside of AFM.
4. The snapshot that is created from the acting primary is an argument to the **--secondary-snapname** option. Issue the following command to convert the fileset on the primary site:

```
mmafmctl Device convertToPrimary -j FilesetName  
[--afmtarget Target { --inband | --secondary-snapname SnapshotName}]  
[ --check-metadata | --nocheck-metadata ][--rpo RPO] [-s LocalWorkDirectory]
```

The primary ID is generated and a psnap0 with the same name is created on the primary site. After the ID generation and the snapshot creation, complete all the steps in [“Failing back to the old primary site” on page 138](#).

Changing the secondary site

This topic lists the steps to follow when the secondary site stops functioning.

About this task

Complete the following steps if the secondary site stops functioning:

Procedure

1. Create a GPFS independent fileset on the new secondary site.
2. Copy the data on the primary site to the new GPFS fileset by using **ftp** or **scp**. This copying data is known as outband trucking.
3. Run the following command to establish a new secondary site for the primary site:

```
mmafmctl Device changeSecondary-j FilesetName  
--new-target NewAfmTarget [ --target-only | --inband ]  
[-s LocalWorkDirectory]
```

The **--inband** option is used depending on the method that is being used to truck data. If you want to change the mount path or the IP address in the target path, you can use the **--target-only** option. The new NFS server must be in the same home cluster and must be of the same architecture as the existing NFS server in the target path. The **--target-only** option must not be used if the home is changed completely, or if you want to change the protocol while you are using the same home cluster. While you are using the same secondary path, if you want to change only the target protocol from GPFS to NFS and vice versa, you can use **mmafmctl changeSecondary** command.

4. Convert the GPFS fileset on the new secondary site to a secondary and set the primary ID by running the **mmchfileset** command or the **mmafmctl** command with **convertToSecondary** option.

Ensure that the NFS export on the secondary site is accessible on all the gateway nodes on the primary cluster if NFS is used for defining AFM target on the primary site. If the GPFS protocol is used for the target, the secondary file system must be mounted on all gateway nodes on the primary site.

After the primary and secondary sites are linked, the `psnap0` queued from the primary fileset is played on the secondary fileset. The new secondary site is now linked to this primary site. Ensure that the primary file system is mounted on all gateway nodes and new secondary file system is mounted on all the nodes in the secondary cluster.

Using the primary fileset to make changes to the secondary fileset

Procedure

1. Replace the secondary with a new empty secondary where the new target uses NFS/NSD/mapping.

The administrator must run the **mmafmctl changeSecondary** with **--inband** option to point to the new secondary. The new secondary is expected to be empty. If the new target is a mapping, **changeSecondary** does not split I/Os and queues them as normal requests without parallel data transfer.

2. Replace any of the following on an existing secondary:

- Replace the communication protocol (NSD/NFS):

The administrator must run the **mmafmctl changeSecondary** with **--inband** option by using the new target protocol.

Note: Only the protocol changes and not the secondary path.

- Enable or disable parallel data transfer PIO by shifting between NFS/NSD and a mapping:

The administrator must run the **mmafmctl changeSecondary** with **--inband** option, by using the new target protocol or mapping.

Note: Only the protocol changes and not the secondary path.

- Replace only the IP address or NFS server by using the same communication protocol and secondary path:

The administrator must run the **mmafmctl changeSecondary** command with the **-target-only** option.

Note: Only the IP or NFS server changes. The IP/NFS server must be on the same secondary cluster and must be of the same architecture as the old NFS server.

Online creation and linking of a dependent fileset in AFM DR

On an independent fileset, AFM is always enabled for granularity and use of features such as quotas and eviction. However, you can create a dependent fileset by using inode space of an AFM-DR fileset and link this dependent fileset to an AFM-DR primary fileset for the data synchronization.

Online creation and linking of a dependent fileset ensures that the dependent fileset with the same name is also created on the secondary for data movement. All data and metadata operations that are generated in the dependent fileset, which is linked to a primary fileset, work seamlessly. Applications can create data on the AFM-DR fileset path and the linked dependent fileset path. AFM synchronizes the created data to the secondary fileset.

On an AFM-DR primary fileset when a dependent fileset is created and linked, a dependent fileset is created and linked to the dependent fileset on the secondary cluster. Similarly, when a fileset is unlinked and deleted from a primary fileset, it is unlinked and deleted from the AFM-DR secondary fileset. All I/Os generated in the linked dependent fileset use the same gateway node that is assigned to the AFM-DR fileset.

Example

1. Create a primary fileset.

```
mmcrfileset fs1 primary1 -p afmMode=primary --inode-space=new -p afmTarget=SecondaryNode1:/gpfs/fs1/secondary1
```

A sample output is as follows:

```
Fileset primary1 created with id 288 root inode 27787267.  
Primary Id (afmPrimaryId) 2836795238918359130-3F69A8C064BE2942-288
```

2. On a secondary cluster, create a secondary fileset by using the primary ID.

```
mmcrfileset fs1 secondary1 -p afmMode=secondary -p  
afmPrimaryID=2836795238918359130-3F69A8C064BE2942-288 --inode-space new
```

A sample output is as follows:

```
Fileset secondary1 created with id 209 root inode 6291459.
```

3. On a secondary cluster, link the secondary fileset and add its NFS export.

```
mmlinkfileset fs1 secondary1 -J /gpfs/fs1/secondary1
```

A sample output is as follows:

```
Fileset secondary1 linked at /gpfs/fs1/secondary1
```

4. Create an NFS export entry for the linked secondary fileset path.

```
/gpfs/fs1/primary1  
<GatewayIPs>(rw,nohide,insecure,no_subtree_check,sync,no_root_squash,fsid=1)
```

5. Export this fileset path.

```
systemctl restart nfs-server
```

6. On the primary cluster link the primary fileset

```
mmmlinkfileset fs1 primary1 -J /gpfs/fs1/primary1
```

A sample output is as follows:

```
Fileset primary1 linked at /gpfs/fs1/primary1
First snapshot name is psnap0-rpo-3F69A8C064BE2942-288
Flushing dirty data for snapshot primary1:psnap0-rpo-3F69A8C064BE2942-288...
Quiescing all operations in the inode space 53.
Snapshot primary1:psnap0-rpo-3F69A8C064BE2942-288 created with id 497.
```

7. On the primary cluster create and link a dependent fileset by using primary1 inode space.

```
mmcrfileset fs1 primary1dep1 --inode-space=primary1
```

A sample output is as follows:

```
Fileset primary1dep1 created with id 290 root inode 27811842.
```

```
mmcrfileset fs1 primary1dep2 --inode-space=primary1
```

A sample output is as follows:

```
Fileset primary1dep2 created with id 291 root inode 27811843.
```

```
mmmlinkfileset fs1 primary1dep1 -J /gpfs/fs1/primary1/primary1dep1
```

A sample output is as follows:

```
Fileset primary1dep1 linked at /gpfs/fs1/primary1/primary1dep1
```

```
mmmlinkfileset fs1 primary1dep2 -J /gpfs/fs1/primary1/primary1dep2
```

A sample output is as follows:

```
Fileset primary1dep2 linked at /gpfs/fs1/primary1/primary1dep2
```

8. On the secondary cluster check that the filesets are created and linked.

```
mmllsfileset fs1
```

A sample output is as follows:

Filesets in file system 'fs1':		
Name	Status	Path
root	Linked	/
gpfs/fs1		
secondary1	Linked	/gpfs/fs1/secondary1
primary1dep1	Linked	/gpfs/fs1/secondary1/primary1dep1
primary1dep2	Linked	/gpfs/fs1/secondary1/primary1dep2

9. On the primary fileset create some data and let it replicate on the secondary fileset.

```
dd if=/dev/urandom of=/gpfs/fs1/primary1/primary1dep1/file1 count=4 bs=256K
```

A sample output is as follows:

```
4+0 records in
4+0 records out
1048576 bytes (1.0 MB, 1.0 MiB) copied, 0.00924786 s, 113 MB/s
```

```
dd if=/dev/urandom of=/gpfs/fs1/primary1/primary1dep1/file2 count=4 bs=256K
```

A sample output is as follows:

```
4+0 records in
4+0 records out
1048576 bytes (1.0 MB, 1.0 MiB) copied, 0.0077314 s, 136 MB/s
```

```
dd if=/dev/urandom of=/gpfs/fs1/primary1/primary1dep2/file1 count=4 bs=256K
```

A sample output is as follows:

```
4+0 records in
4+0 records out
1048576 bytes (1.0 MB, 1.0 MiB) copied, 0.0085112 s, 123 MB/s
```

```
dd if=/dev/urandom of=/gpfs/fs1/primary1/primary1dep2/file2 count=4 bs=256K
```

A sample output is as follows:

```
4+0 records in
4+0 records out
1048576 bytes (1.0 MB, 1.0 MiB) copied, 0.0086434 s, 121 MB/s
```

```
ls -R /gpfs/fs1/primary1
```

A sample output is as follows:

```
/gpfs/fs1/primary1:
primary1dep1 primary1dep2

/gpfs/fs1/primary1/primary1dep1:
file1 file2

/gpfs/fs1/primary1/primary1dep2:
file1 file2
```

10. After data is replicated, verify the data is present on the secondary.

```
ls -R /gpfs/fs1/secondary1
```

A sample output is as follows:

```
/gpfs/fs1/secondary1:
primary1dep1 primary1dep2

/gpfs/fs1/secondary1/primary1dep1:
file1 file2

/gpfs/fs1/secondary1/primary1dep2:
file1 file2
```

11. On the primary fileset unlink and delete the dependent fileset.

```
mmunlinkfileset fs1 primary1dep1 -f
```

A sample output is as follows:

```
Fileset primary1dep1 unlinked.
```

```
mmdeleteset fs1 primary1dep1 -f
```

A sample output is as follows:

```
Checking fileset ...
Checking fileset complete.
Deleting user files ...
100.00 % complete on Fri Oct 20 09:36:28 2023 ( 100352 inodes with total 392 MB data
processed)
Deleting fileset ...
Fileset primary1dep1 deleted.
```

```
ls -R /gpfs/fs1/primary1
```

A sample output is as follows:

```
PrimaryNode1 1]
/gpfs/fs1/primary1:
primary1dep2

/gpfs/fs1/primary1/primary1dep2:
file1 file2
```

12. On the secondary fileset verify that the fileset is deleted.

```
mmlsfileset fs1
```

A sample output is as follows:

```
Filesets in file system 'fs1':
Name Status Path
root Linked /gpfs/fs1
secondary1 Linked /gpfs/fs1/secondary1
primary1dep2 Linked /gpfs/fs1/secondary1/primary1dep2

SecondaryNode1 0] ls -R /gpfs/fs1/secondary1
/gpfs/fs1/secondary1:
primary1dep2

/gpfs/fs1/secondary1/primary1dep2:
file1 file2
```

Note: When recovery is running because of some failures, dependent filesets must not be unlinked.

AFM DR limitations

AFM DR limitations include:

- If you change an AFM DR Secondary export from Kernel NFS or CNFS to the new IBM Spectrum Scale CES in 4.1.1 or later, AFM does not recognize the secondary. The best way to manage this situation is to run the **mmafmctl changeSecondary** command. See the IBM Storage Scale documentation for more details.
- DR administration by using the **mmafmctl** command is not supported when run from the AIX OS.
- AFM ADR (primary/secondary filesets) is not supported on an FPO enabled system.
- AFM features such as revalidation, eviction, prefetch, show home snapshots, partial file caching, expire and unexpire, mmafmctl resync and failover, and IW failback are not applicable for AFM DR filesets.
- Administrative operations such as fileset conversion to primary, failover, and failback can be run at the same time on multiple filesets. However, it is not recommended because commands might fail with a "busy" error. The commands might fail in cases where the system runs out of the resources that are needed to perform multiple operations at the same time. If this error occurs, rerun the failed command after the conflicting command completes.

AFM and AFM DR limitations

The common limitations for both AFM and AFM DR functions include:

- The following table shows the OS and architecture support matrix and supported AFM function:

Table 10. Supported function and supported OS/architecture			
OS and architecture	Fileset	Node roles	Comments
Linux on Power and x86	SW/IW/RO/LU primary/secondary	Gateway Application Manager	

Table 10. Supported function and supported OS/architecture (continued)

OS and architecture	Fileset	Node roles	Comments
Linux on Z	SW/IW/RO/LU primary/secondary	Gateway Application Manager	
AIX on Power	SW/IW/RO/LU primary	Gateway Application	AIX cannot be an AFM home site or an AFM-DR secondary site.

- An AFM home site that is not an IBM Storage Scale cluster:
 - An AFM cache site can be configured to connect with any home site that provides access to the home data through the NFSv3 protocol. The home site is not required to be an IBM Storage Scale cluster.
 - In this scenario, the **mmafmconfig** command cannot not be issued on the AFM home site. Therefore, the AFM cache site cannot cache or update extended attributes (EAs), access control lists (ACLs), file sparseness, or AFM psnaps.
 - The IBM Storage Scale support team helps customers to investigate problems in this scenario but only to address issues with the AFM cache site that runs IBM Storage Scale. Any problems at the AFM home site, or its file system, must be addressed by the customer.
 - For AFM-DR configurations, both the primary and secondary sites must be IBM Storage Scale clusters.
- An AFM home site that runs IBM Storage Scale on AIX:
 - The **mmafmconfig** command is not supported on AIX. Therefore, the AFM cache site cannot cache or update extended attributes (EAs), access control lists (ACLs), file sparseness, or AFM psnaps.
 - AFM does not support AIX nodes as an exporter of NFS. That is, AFM does not support NFS mounting exports that are reached through a system that runs on AIX.
 - Because of the previous limitations, AIX nodes cannot act as an AFM home site or as an AFM-DR secondary site.
 - AIX is not supported as a gateway node for AFM fileset modes.
 - AIX is not supported as an AFM home site or AFM-DR secondary site.
- The **mmclone** command is not supported on AFM cache and AFM DR primary filesets. Clones that are created at home for AFM filesets are treated as separate files in the cache.
- Quality of service is not tested on AFM and AFM Async DR filesets.
- Connecting an AFM-DR fileset to IBM Storage Protect for Space Management is not supported.
- The NSD protocol for communication does not support AFM and AFM DR filesets on Linux on Z.
- Connecting AFM DR fileset to IBM Spectrum Archive is not supported.
- AFM and Async DR is not supported on clusters that have Windows nodes.
- Cascading relationships with AFM caches and AFM primary filesets are not tested.
- Fileset snapshot restore is not supported for AFM and AFM DR filesets.
- Files in an AFM or AFM DR fileset, which are shared by using SMB, an error message E_ACCESS might appear if SMB holds locks on the files. When the locks are released by SMB, AFM replicates again and processes the files from the AFM or AFM DR fileset.
- The `dm_write_invis()` function is not supported on AFM and AFM DR filesets.
- The `dm_read_invis()` function is not supported on AFM and AFM DR filesets, if a file is not cached.
- The File Audit logging feature is not supported on AFM or AFM DR filesets.
- **Immutability** and **appendOnly** features are not supported on AFM filesets.
- AFM and AFM DR filesets do not replicate security, system, or trusted extended attributes.

- The **--iam-mode** option is not supported on AFM filesets.
- AFM and AFM DR do not support DM Punch Hole method (GPFS API '**dm_punch_hole**').
- To enable the Kerberos security for an AFM-DR fileset, the AFM-DR secondary must be given the read/write access.
- If a file is evicted from an AFM fileset, the snapshot file of the evicted file contains zeros.
- Eviction is not supported on an AFM-DR primary fileset and an AFM-DR secondary fileset.
- The following file attributes are maintained locally in the cache but these are not replicated:
 - Control attributes
 - Direct I/O
 - Replication factors
 - Fileset quotas
 - Storage pool flags
 - Special file types such as FIFO, socket, block, character device
- Hard links can be created on the cache. Hard links on the home are displayed as individual files in the cache during the lookup. To prefetch hard links from home, run the **--metadata-only** option as the first operation.
- File locking across the cache and the home is not supported.
- User extended attributes, ACLs, and sparse files supported only on the home where the **mmafmconfig** command is run.
- Parallel data transfer is supported only in Linux-only clusters. If a home is a mix of architectures (x86 and ppc), parallel data transfer works only for the set of nodes that belong to any one architecture, depending on which architecture serves the data transfer first.
- If encryption is configured on a home site that is running on a file system, which AFM uses as a GPFS backend target (multi-cluster remote mount), of an IBM Storage Scale cluster, ensure that the cache cluster is also configured the same way as the home cluster. Because of this configuration, AFM can access the files on the target file system for the replication.
- Online replication of a dependent fileset is not supported on an AFM DR fileset with a GPFS/multi cluster backend.

Related concepts

[“General guidelines and recommendations for AFM SW and IW mode fileset” on page 454](#)

Use the IBM guidelines and the recommendations to configure and maintain AFM Single Writer (SW) and Individual Writer (IW) mode fileset.

AFM DR deployment considerations and best practices

For AFM DR deployment, it is important to know the characteristics of AFM DR including advantages and limitations, its deployment method, and best practices.

AFM DR is a file-level asynchronous disaster recovery solution. Replication of data to a secondary cluster is asynchronous, which means that the data that is created on the primary cluster is copied to the secondary cluster with delay. The amount of delay is determined by the network bandwidth between the sites and some IBM Storage Scale tuning parameters.

AFM DR copies only the data blocks that changed on the primary cluster to the secondary cluster, which makes it an efficient replication mechanism.

Since data replication is asynchronous by using AFM DR to recover data after a primary cluster failure might cause some data loss. The amount of data that is lost depends on the amount of data that is in the primary queue when a failure occurs. Therefore, AFM DR is not the right disaster recovery solution if your application cannot tolerate possible data loss. If this loss is not acceptable, customers must implement a synchronous disaster recovery mechanism.

Characteristics of AFM DR

The following sections give a brief description of the various characteristics of AFM DR.

Independent filesets in AFM DR

The primary and secondary fileset clusters are independent filesets.

Features

- The disaster protection of data happens at an independent fileset level.
- A single file system can contain multiple AFM DR relationships, one per fileset.

Advantages

- Fileset-level replication allows a high level of granularity and flexibility to protect data.
- An RPO can be specified for each fileset, which allows to set it based on the criticality of the data.
- AsyncDelay can be set on a fileset level to control rate of replication.

Limitations

- Can be defined only on an independent fileset.
- A single AFM DR relationship cannot be created at the file system level.

One-on-one relationships in AFM DR

AFM DR creates a one-to-one relationship between the primary and the secondary filesets.

Advantages

- RPO intervals can be defined per AFM DR relationship.
- AsyncDelay can be specified per AFM DR relationship.
- Quotas can be specified for the primary and secondary filesets.

Limitations

You can create only two copies of the data.

Active-passive relationships in AFM DR

AFM DR creates an active-passive relationship between the primary and the secondary clusters.

Features

- Only the primary can write to the secondary fileset. The secondary is read-only for all other users.
- You can take a snapshot of the secondary fileset and use the secondary as a source to create backups or additional copies.

Advantages

The secondary fileset is protected, so no other program or application other than the primary can write to the secondary.

Limitations

- All operations are allowed on the primary fileset, but only snapshot operations are allowed on the secondary fileset.

- The secondary can be converted into the primary even when the primary is writing to the secondary. There is no protection against an administrator doing a failover operation on the secondary, even if the DR relationship is healthy.
- The administrator must ensure that the primary cannot accidentally write to secondary after a failover.

NFSv3 versus NSD exports in AFM DR

The primary accesses the secondary fileset by using either NFSv3 or NSD protocol.

Features of NFSv3

NFSv3 is a stateless protocol, which is resilient on low-bandwidth and high-latency networks.

Advantages of NFSv3

It is recommended to use NFSv3 because NFSv3 is more tolerant when deployed on an unstable or high-latency network.

Note: It is recommended to use the NSD protocol only if the NFS protocol does not meet the expected performance requirements.

Limitations of NFSv3

For parallel I/O to work by using NFSv3, create a map between the gateway nodes in the primary to the NFS servers in the secondary. For more information about parallel I/O configuration, see *Configuration parameters for AFM* in the *IBM Storage Scale: Administration Guide*.

Features of NSD

The NSD protocol is sensitive to packet drops and network latency.

Advantages of NSD

NSD provides high performance to LAN-based installations.

Limitations of NSD

If the secondary cluster is inaccessible, the NSD mount hangs, which in turn might cause the application to hang on the primary cluster.

Trucking features in AFM DR

AFM DR trucking feature can be used to synchronize data between the primary and secondary filesets when the DR relationship is established for the first time.

Features

- AFM DR requires the same data between the primary and secondary before a DR relationship is established by using the peer-to-peer snapshot (PSNAP) mechanism. If the data on the primary and secondary are the same, the peer-to-peer snapshot mechanism creates a special snapshot called PSNAP0. The PSNAP0 is created on both the primary fileset and the secondary fileset.
- Two trucking methods are supported by AFM DR:
 - Inband trucking: AFM DR copies the data from the primary to the secondary using the DR transport protocol (NFSv3 or NSD). After the copy is complete, the primary and secondary have the exact same data and PSNAP0 is created.
 - Outband trucking: It is the user's responsibility to copy all the data from the primary and secondary by using any mechanism that they choose. For example, rsync with **mtime** preserved. After all the data is copied, the user can use the AFM DR commands to establish the DR relationship that creates the PSNAP0.

Note: Both methods use the **--inband** option of the **mmafmctl** command while setting up the relationship between the primary and secondary sites.

For more information about the two trucking methods, see *Converting GPFS fileset to AFM DR* in *IBM Storage Scale: Administration Guide*. See the following use cases for recommendation on when to use inband versus outband trucking.

AFM DR trucking use cases

Create the primary and secondary filesets from scratch

The primary and secondary clusters are empty and AFM DR establishes the relationship by creating PSNAP0 that is empty.

Trucking method recommendation: Use the inband trucking method for this scenario because the clusters have no data to copy and create PSNAP0.

Create a primary and secondary fileset from scratch, where primary has data and secondary does not have data

In this case, either the inband or the outband mechanism can be used to copy the data to the secondary cluster.

- **Inband:**

By using the inband method, the existing data in the primary is copied by AFM DR to the secondary. To copy this data a work queue is created on the gateway that contains all of the files that need to be copied and file system operations that need to be performed on the secondary. Then, the gateway processes this queue as quickly as possible by copying data to the secondary. When all of the work, which was outstanding when the command was run, is flushed to the secondary, a PSNAP0 is created. The trucking process ends when a snapshot is created.

It is recommended to start applications on the primary after the trucking process is complete. However, if applications are started while the trucking process is in progress, the gateway node queues the data generated by the application and the requests from the trucking process at the same time. The gateway node queue buffer is filled with many requests, which causes the queue memory overflow. In such a case, the queue is dropped and the gateway node goes into the recovery mode. This can result in a cycle where the primary cluster always lags behind the secondary cluster.

The potential solutions are as follows:

- Add bandwidth to handle the initial synchronization workload. This bandwidth can involve tuning – for example, increasing the number of flush threads to increase the gateway throughput, additional hardware or network bandwidth.
- Keep the application from writing to the primary until the trucking process is complete and PSNAP0 is created.
- Increase the gateway node memory allocation to prevent memory overflows.

- **Outband:**

In this case, a complete copy of the data exists before the primary-secondary relationship is created. The data can be copied to the secondary by using any method that includes a restore from a backup or a rsync process. When an identical copy of the data exists at the secondary, AFM DR can establish the DR relationship. The primary and secondary need to be identical when the relationship is established, so during this time the data must not change in the primary. After the PSNAP0 is created, applications can start by using the AFM DR filesets.

Note: Both methods use the **--inband** option of **mmafmctl** command while setting up the relationship between the primary and secondary sites. For data that is already copied to the secondary site, **mtime** of files at the primary and secondary site is checked. If **mtime** values of both files match, data is not copied again and existing data on the secondary site is used. If **mtime** values of both files do

not match, existing data on the secondary site is discarded and data from the primary site is written to the secondary site.

The primary fileset is empty and the secondary fileset has data

Trucking method recommendation: In such a situation, the recommended mechanism is outband trucking. The administrator copies data from the secondary to the primary and then uses the AFM DR conversion mechanism to establish the DR relationship.

Data exists on the 'to-be' primary and the secondary filesets

Ensure that the data is identical on the primary and the secondary before the AFM DR conversion is performed.

Trucking method recommendation: This can be done by using only the outband trucking. AFM DR checks whether the data is the same when the relationship is established.

Failover and failback

Failover

After a failure of the primary, applications can fail over to the secondary fileset. Performing a failover operation on the secondary fileset converts it to the primary and makes it writable. AFM DR failover has two modes: “As is” and “Revert to the latest RPO snapshot on the secondary”:

- As is:

The “As is” data corresponds to the last byte of data that was copied to the secondary cluster before the primary cluster failed. This is always the latest data available on the secondary cluster. It is the recommended option and also the default behavior.

- Revert to the latest RPO snapshot on the secondary fileset:

Restoring data from the last RPO snapshot results in data being restored from the past. As a result, this process is not recommended in AFM DR as it might lead to data loss. For example, although the RPO interval was set to 24 hours, the failover occurred at the 23rd hour, then restoring from the last snapshot (taken 24 hours ago) can result all the data that is created in the last 23 hours being lost. This failover option can take extra time as the data from the snapshot must be restored completely before the application can start by using the fileset.

Failback

There are two options in AFM DR to failback to the original primary:

- Failback Option 1: New primary

Create the primary from scratch because all the data in the primary was lost. To create a primary from scratch, all the data needs to be copied to the original primary from the active primary by using whichever mechanism the administrator chooses.

- Fail back Option 2: Temporary loss of primary

A temporary failure of the primary that results in a failover of the application to the secondary. To reestablish the original primary:

A common snapshot between the primary and secondary is needed because AFM DR uses the last RPO snapshot to calculate exactly what data changed on the secondary since the failover occurred. The following steps give a simplified illustration of that a scenario:

1. A PSNAP was created between the primary and secondary at 8 AM. At 8 AM, the snapshot in the primary and secondary have the exact same data.
2. At 9 AM, the primary fails.
3. At 9 AM, the applications failover to secondary.

4. Primary is back up at 10 AM that is 2 hours later.
5. Restore the T0 snapshot on the primary fileset.
6. Restore the 8 AM snapshot on the primary fileset.
7. Copy the changed data that is calculated in step 6 from the secondary fileset to the primary fileset.
Now, the primary is in sync with the secondary fileset.

Note: For more information about the failback, see *Failback procedures* in the *IBM Storage Scale: Administration Guide*.

For a temporary primary failure, the PSNAP (peer-to-peer snapshot) mechanism can be used to periodically take synchronized snapshot between the primary and the secondary. The snapshots can be scheduled at fairly large intervals, for example, one or two times a day. Only the data created in the secondary since the last PSNAP needs to be copied back to the primary during a failback.

Note:

The amount of data copied back to the primary depends on the following:

- How long the primary is unavailable.
- The amount of data that is created since the primary failure.
- How long the failure was from the last snapshot. For example, if the RPO snapshot is captured every 24 hours and the primary fails at the 23rd hour, then changes for the last 23 hours need to be recovered.

Deployment considerations for AFM DR

See the following considerations to understand whether AFM DR fulfills your production requirements for disaster recovery:

Amount of data created per hour or per day

With data replication the amount of data created in a day determines the network bandwidth and the gateway node design, and dictates whether the replication rates can be supported by AFM DR. You can determine the requirements by looking at:

- Network bandwidth

The network must have the bandwidth to accommodate a transfer rate equivalent to the rate data is generated. You should consider how you want to handle high traffic events like fail-back where I/O requirements can be greatly increased for a short period of time.

- Gateway nodes(s)

A gateway tracks all the changes and replicated the data from the primary to the secondary. As the number of filesets increase you can increase the number of gateway nodes in the primary cluster to increase throughput. The number of gateway nodes needed depends on:

- The number of primary filesets.
- The rate of data changes generated by each fileset.
- The bandwidth of the gateway hardware and the network connections between clusters.
- The method in which the filesets are distributed across gateway nodes. For example, it is possible that all the heavily loaded filesets are allocated to the same gateway node.

Note: There is no manual way of controlling the fileset allocation to a gateway node. Currently, the method in which AFM DR allocates fileset to the gateway node is ad hoc and creates another challenge for the user.

The number of filesets on the primary cluster

The number of filesets impact the following factors:

- You might need to increase the number of gateway nodes for even distribution of the workload.

- The AFM DR RPO mechanism creates and deletes a snapshot after every RPO interval. As the number of filesets increase, creating the RPO snapshots simultaneously for all the filesets might cause a significant load on the system. The RPO can be disabled if there is no need to failback after a temporary primary failure. Alternatively, setting a large RPO interval can relieve some of this pressure however, once the RPO is enabled then all of the above points need to be considered carefully.

Note: A gateway node failure and recovery causes redistribution of the fileset workload, which in turn might cause changes in the performance of the replication characteristics.

Best practices for AFM DR

Minimizing data loss during primary failure

AFM DR is asynchronous, so there is always a chance data can be lost in the event of a failure. The amount of data lost is affected by:

- The network bandwidth between the primary and secondary.
- The performance of the gateway which depends on IBM Storage Scale tuning, the amount of memory and CPU available to gateway node and the number of filesets allocated to each gateway node. If the gateway node is overloaded, it can result in replication rate reduction and reduced network bandwidth utilization.
- The ability of the primary and secondary to read and write the data to disk.

The AFM DR replication should be tuned to minimize the data loss in case of primary failure by keeping-up with the data creation rate. Some workloads occur in bursts, in this case you can design for average data transfer rates though keep in mind the time to sync is much greater right after a burst of changes.

Note: The AFM DR replication rate is independent of the RPO interval. The RPO interval does not affect the data loss sustained during primary failure.

Generating notification for failed replication

There is no automated notification mechanism in AFM DR to monitor the replication rates that are falling behind as long as they are within the RPO. However, a script can be written to periodically test the gateway node to see how fast the message queue is being processed. This provides an estimate of the replication rate sustained by the gateway node.

To monitor the RPO you can use the AFMRPOMISS callback event. This event is triggered if the RPO snapshot is not taken at the set interval. The event indicates something is wrong within the system, and can be used as a trigger to start an analysis of what needs to be rectified within the system to bring it back to optimal performance.

Using tuning parameters to improve performance

There are several AFM DR tuning parameters that can be used to tune performance. For more information on tuning parameters, see *Configuration parameters for AFM And AFM-DR* in the *IBM Storage Scale: Administration Guide*.

AFM DR use case

This use case describes the AFM-based Asynchronous Disaster Recovery (AFM DR) process when NFS and SMB clients are mounted from the primary IBM Storage Scale cluster.

Overview of AFM DR when NFS and SMB clients are mounted from the primary cluster

During scheduled maintenance on the primary cluster or a failure at the primary site cluster, those NFS, and SMB clients that mount exports from the primary cluster experience a brief and temporary outage, when the `mmafmctl failover` commands are run. After the `mmafmctl failover` commands

complete, the NFS and SMB clients can mount from the new acting primary cluster (formerly the secondary cluster).

Note: To ensure a successful failover, system administrators at each site need to create identical export definitions on the secondary cluster. It is important to keep the definitions up to date when changes occur on the primary cluster.

After you create initial AFM DR relationships, you can create NFS and SMB exports so that applications can access the IBM Storage Scale data via protocols.

Use case 1: The path to the data is identical on both the primary and secondary clusters

1. On both the primary and the secondary clusters, enter the same path and client (-c) definitions as shown in the following example:

```
mmnfs export add /ibm/gpfs0/export1 -c
9.11.102.0/24(Access_Type=R0,Squash=root_squash;
9.11.136.0/23(Access_Type=R0,Squash=root_squash;
10.18.40.0/22(Access_Type=RW,Squash=root_squash;
10.18.48.0/22(Access_Type=RW,Squash=root_squash;
```

2. On the NFS clients, run the following commands.

- Issue the **showmount -e** command and the IP address of the primary cluster as shown in the following example:

```
showmount -e 10.18.96.100
Export list for 10.18.96.100:
/ibm/gpfs0/export1
9.11.102.0,9.11.136.0,10.18.40.0,10.18.48.0
```

- Issue the **showmount -e** command and the IP address of the secondary cluster as shown in the following example:

```
showmount -e 10.18.50.30
Export list for 10.18.50.30:
/ibm/gpfs0/export1
9.11.102.0,9.11.136.0,10.18.40.0,10.18.48.0
```

3. On the NFS clients, issue the following command to mount from the primary CES IP addresses as shown in the following example:

```
mount -o hard, sync
10/18.96.100:/ibm/gpfs0/export0 /mnt_pnt
```

4. If a failure occurs on the primary cluster, initiate failover.
5. On the NFS clients, issue the following commands to unmount from the primary CES IP addresses and then mount from the secondary CES IP addresses as shown in the following example:

```
umount -f /mnt_pnt
mount -o hard, sync
10.18.50.30:/ibm/gpfs0/export0 /mnt_pnt
```

Use case 2: The path to the data is different on the primary and secondary cluster

1. Create the export with the same client (-c) definitions.

- Issue the following command and the IP address of the primary cluster:

```
mmnfs export add /ibm/gpfs0/export1 -c
9.11.102.0/24(Access_Type=R0,Squash=root_squash;
9.11.136.0/23(Access_Type=R0,Squash=root_squash;
10.18.40.0/22(Access_Type=RW,Squash=root_squash;
10.18.48.0/22(Access_Type=RW,Squash=root_squash;
```

- Issue the following command and the IP address of the secondary cluster:

```
mmnfs export add /ibm/gpfs3/export9 -c
9.11.102.0/24(Access_Type=R0,Squash=root_squash;
9.11.136.0/23(Access_Type=R0,Squash=root_squash;
10.18.40.0/22(Access_Type=RW,Squash=root_squash;
10.18.48.0/22(Access_Type=RW,_quash=root_squash;
```

2. On the NFS clients, run the following commands.

- Issue the **showmount -e** command and the IP address of the primary cluster:

```
showmount -e 10.18.96.100
Export list for 10.18.96.100:
/ibm/gpfs0/export1
9.11.102.0,9.11.136.0,10.18.40.0,10.18.48.0
```

- Issue the **showmount -e** command and the IP address of the secondary cluster:

```
showmount -e 10.18.50.30
Export list for 10.18.50.30:
/ibm/gpfs3/export9
9.11.102.0,9.11.136.0,10.18.40.0,10.18.48.0
```

3. On the NFS clients, issue the following command to mount from the primary CES IP addresses:

```
mount -o hard, sync
10/18.96.100:/ibm/gpfs0/export0 /mnt_pnt
```

4. On the NFS clients, issue the following commands to unmount from the primary CES IP addresses and then mount from the secondary CES IP addresses:

```
umount -f /mnt_pnt
mount -o hard, sync
10.18.50.30:/ibm/gpfs3/export9 /mnt_pnt
```

AFM to cloud object storage

The AFM to cloud object storage feature enables placement of files or objects in an IBM Storage Scale cluster to a cloud object storage.

Cloud object services such as Amazon S3, IBM Cloud® Object Storage, Seagate Lyve Cloud, and Microsoft Azure Blob storage offer industry-leading scalability, data availability, security, and performance. AFM to cloud object storage supports Amazon S3, IBM Cloud Object Storage, Seagate Lyve Cloud, Microsoft Azure Blob, and Google Cloud Platform. Microsoft Azure Blob storage is supported by using two features.

- AFM to cloud object storage direct Microsoft Azure Blob support: AFM to cloud object storage fileset connects to Azure containers without any intermediary application.
- By deploying MinIO as an S3 gateway between AFM to cloud object storage and Microsoft Azure Blob storage.
- AFM to cloud object storage supports S3 endpoints as backend also.

The AFM to cloud object storage allows associating an IBM Storage Scale fileset with a cloud object storage. Customers use a cloud object storage to run workloads such as mobile applications, backup and restore, enterprise applications, and big data analytics, file server. These workloads can be cached on AFM to cloud object storage filesets for faster computation and synchronize back to the cloud object storage server.

The front-end for object applications is an AFM to cloud object storage fileset with the data exchange between the fileset and cloud object storage buckets through the AFM to cloud object storage in the background by providing high performance for the object applications. Object applications can also span across AFM to cloud object storage filesets and on a cloud object storage. Both the fileset and the cloud object storage can be used as a backup of important data.

The following figure illustrates the AFM to cloud object storage:




Figure 19. AFM to cloud object storage

The AFM to cloud object storage on an IBM Storage Scale fileset becomes an extension of cloud object storage buckets for high-performance or used objects. Depending upon the modes of AFM to cloud object storage fileset configurations, objects required for applications such as AI and big data analytics can be downloaded, worked upon, and can be uploaded to a cloud object storage. The objects that are created by applications can be synchronized to the objects on a cloud object storage asynchronously. An AFM to cloud object storage fileset can cache only metadata or both metadata and data.

The AFM to cloud object storage also allows data center administrators to free the IBM Storage Scale storage capacity by moving less useful data to the cloud storage. This feature reduces capital and operational expenditures. The AFM-based cache eviction feature can be used to improve the storage capacity manually and by using policies. For more information about the AFM cache eviction, see [“Cache eviction”](#) on page 94.

The AFM to cloud object storage uses the same underlying infrastructure as AFM. For more information, see [“Active File Management”](#) on page 54.

An AFM to cloud object storage fileset supports setting an access control list (ACL) on a file of up to 2 KB size. When an ACL is assigned on a file of more than 2 KB size, the file is discarded and only file data is synchronized with the bucket.

The AFM to cloud object storage is available on all IBM Storage Scale editions.

AFM to cloud object storage by using the Microsoft Azure Blob storage

AFM to cloud object storage fileset can be configured to use Microsoft Azure Blob storage by deploying MinIO. MinIO runs as an S3 gateway, and it converts S3 API to/from Azure Blob native. All communication between AFM to cloud object storage fileset and the bucket at the Microsoft Azure Blob storage goes through S3 gateway. All the operations, modes, and functions are seamlessly supported with S3 gateway that communicates all occurrences to Microsoft Azure Blob.

The following figure illustrates the AFM to cloud object storage by using Microsoft Azure Blob as backend through MinIO S3 gateway:




Figure 20. AFM to cloud object storage by using Azure Blob

AFM to cloud object storage operation modes

AFM to cloud object storage filesets can be created for all AFM modes: single writer (SW), independent writer (IW), local update (LU), read only (RO), or manual update (MU).

AFM to cloud object storage synchronizes operations between the cache fileset and the cloud object storage bucket by using `http` or `https` protocol. Both push and pull of files or objects are supported based on the mode of the fileset.

All operations from the AFM to cloud object storage gateway are queued in the memory on the primary gateway and then all operations are pushed to the cloud object storage asynchronously based on the **`afmAsyncDelay`** interval.

Fileset operation modes

AFM to cloud object storage filesets can be created in similar ways for the different AFM modes: single writer (SW), independent writer (IW), local update (LU), read only (RO), or manual update (MU).

Functioning of all AFM to cloud object storage fileset modes are similar with AFM to cloud object storage except the local-update (LU) mode, where AFM to cloud object storage LU mode fileset can upload data back to the cloud object storage target.

Unlike AFM LU-mode, where data cannot be pushed to the home, AFM to cloud object storage LU-mode fileset can upload data to the target bucket of the cloud object storage server by using the **mmafmcctl upload** command. The upload operation is a manual operation. You can use this operation to run applications in the LU-mode cache and synchronize the data to the bucket.

Single writer (SW)

In this mode, only an AFM to cloud object storage fileset does all the writing and this fileset does not check a cloud object storage for object updates. The AFM to cloud object storage fileset does not enforce writing or modifying an object on a cloud object storage. After an SW-mode AFM to cloud object storage fileset is created, if a cloud object storage has pre-existing data, this data can be prefetched by using the download option from the **mmafmcctl** command.

Independent writer (IW)

This mode allows multiple AFM to cloud object storage filesets to point to the same cloud object storage bucket. Multiple AFM to cloud object storage filesets can be on the same IBM Storage Scale cluster or on a different cluster and points to a cloud object storage. There is no synchronous locking between clusters while updating objects on a cloud object storage. Each AFM to cloud object storage fileset reads from a cloud object storage and makes updates to the cloud object storage independently based on the revalidation intervals and asynchronous delay.

This mode is used to access different objects from each IW AFM to cloud object storage site, for example, unique users at each site that update objects in their cloud object storage bucket. This mode allows multiple AFM to cloud object storage clusters to modify the same set of objects. Only advanced users must do this modification because there is no locking or ordering between updates. Updates are propagated to a cloud object storage in an asynchronous manner and can be delayed because of network disconnections. Therefore, conflicting updates from multiple AFM to cloud object storage sites can cause the data on the cloud object storage site to be undetermined.

Read-only (RO)

In this mode, data in the AFM to cloud object storage fileset is read only. You cannot create or modify files in an RO-mode AFM to cloud object storage fileset. If an application is using an object while it is being re-created, that is, deleted and re-created with the same name on a cloud object storage, the object gets re-created in AFM to cloud object storage fileset. The **mmafmcctl download** command can be used to download the objects on priority or preference.

Local updates (LU)

This mode is similar to the RO mode, although you can create and modify objects in an AFM to cloud object storage fileset. Updates in the AFM to cloud object storage fileset are considered local to the AFM to cloud object storage fileset and get decoupled from the corresponding object on a cloud object storage. Local updates are never pushed back to a cloud object storage.

When an object is downloaded and modified in an AFM to cloud object storage LU-mode fileset, the modified data does not synchronize back with the bucket on the cloud object storage server even if the object on the bucket is modified. This behavior is similar to the AFM LU-mode fileset. For more information about AFM modes, see [“Operations with AFM modes” on page 64](#).

Behaviors with local objects

In AFM, LU-mode objects have one of the following states:

Uncached

Objects on a cloud object storage are shown as uncached. For these objects, only metadata is copied into an AFM to cloud object storage fileset. The object does not reside on the AFM to

cloud object storage fileset, but only on the cloud object storage. Changes on the cloud object storage are reflected in the AFM to cloud object storage fileset.

Cached

If an uncached object is read in an AFM to cloud object storage fileset or pre-fetched, the state of the object changes to cached. In the cached state, all changes to the object on a cloud object storage are reflected in the AFM to cloud object storage fileset. The object resides on the AFM to cloud object storage fileset.

When symlinks are created on an AFM to cloud object storage cache fileset, the relationship between the symlink file and the target file is maintained and the files are uploaded to the bucket.

The symlinks relationship is maintained only if they are created in an AFM to cloud object storage fileset. If the symlinks relationship is not created in the AFM to cloud object fileset, it is not maintained.

Local

Object data or metadata that is modified on an AFM to cloud object storage fileset becomes local to the AFM to cloud object storage fileset. The cached objects relationship to the object in a cloud object storage is broken. Changes on the cloud object storage are not reflected in the AFM to cloud object storage fileset anymore and object changes are not copied to the cloud object storage.

Note: Objects can be downloaded from a cloud object storage and uploaded back without affecting the LU-mode. When the object is downloaded, it is synced-in to the object on the cloud object storage.

Manual updates (MU)

The manual update (MU) mode is only available with AFM to cloud object storage filesets. This mode supports manual replication of the files or objects by using ILM policies or user-provided object list. MU mode is supported on AFM to cloud object storage backend. The IBM Storage Scale independent fileset can be converted to MU mode fileset or MU mode fileset can be created as a new relationship to cloud object storage.

MU mode fileset provides the flexibility to upload and download files or objects to and from cloud object storage after you finalize the set of objects to upload or download. Unlike other AFM to cloud object storage object fileset modes, MU mode depends on manual intervention from administrators to upload and download the data to be in sync. As administrators, you can also automate upload and download by using ILM policies to search specific files or objects to upload or download.

MU mode semantics:

- Applications or users create, update, or delete data in the MU mode fileset.

The changes that are done in the MU mode are not automatically synchronized to cloud object storage buckets. An administrator must manually run the upload or reconcile command as required to complete the synchronization. Until the synchronization is completed, the data lying in the MU mode is dirty data and does not get replicated. After the synchronization is done, data copy is available both in cache and in cloud object storage.

- Either all the files or specific files can be determined and uploaded by using the **mmafmcscctl** upload command. Or, policies can help determine these specific files.
- When the files are added to cloud object storage or the data is changed on cloud objects storage, these files can be downloaded by using the **mmafmcscctl download** command.

Important: The MU mode does not recognize the files added to cloud object storage or modified data automatically, to download these files or data administrators can create object list that contains the names of the files from cloud object storage.

- While files are being modified in the cloud object storage bucket (target), the modified metadata or data can be fetch to the cache to synchronize both copies. Run **mmafmcscctl download** with **--metadata** or **--data** to synchronize the copies as needed.

- All the changes to files are local to the fileset except remove operations, which can be made automatically replicated by using the `afmMUAutoRemove` configuration option.

Deleting files from an MU mode fileset

Files on an MU mode fileset take inode and space from fileset. Files can be deleted from the MU mode fileset but the file inodes are reclaimed from the fileset inode space only when **mmafmcscctl delete** commands are issued. The MU mode has two delete options backed by **mmafmcscctl delete** command. The files can be deleted locally by using **posix** commands. For example, **rm**. The MU mode does not reclaim inodes immediately after the deletion until the delete operation is executed to the COS target.

The **mmafmcscctl delete** commands – from cache option looks for deleted files in the MU fileset and reclaim the file inodes, whereas –from-target option looks for the deleted files in MU fileset, reclaims the inodes and queue the delete operation to COS so that files deleted from MU fileset gets deleted on COS as well.

Note: After the delete --from-cache is performed, the inodes are reclaimed from MU mode fileset. After execution of this command if the delete --from-target is executed, the deletes are not queued on the cloud object storage.

Example : Deleting and reclaiming inodes from an MU mode fileset

1. Create an MU mode fileset.

```
node1 1] mmafmcscconfig fs1 mufilesetdemo --endpoint http://s3.us-east.cloud-object-
storage.appdomain.cloud --bucket demobucketx --object-fs --mode mu
node1 1]
node1 1]
node1 1] mmafmctl fs1 getState
Fileset Name      Fileset Target      Cache State      Gateway
Node      Queue Length      Queue numExec
-----
mufilesetdemo      http://s3.us-east.cloud-object-storage.appdomain.cloud:80/demobucketx Inactive
node1 1]
```

2. Check the used inodes.

```
node1 24:16 1] mmlsfileset fs1 mufilesetdemo -i
Collecting fileset usage information ...
Filesets in file system 'fs1':
Name      Status      Path      InodeSpace
MaxInodes AllocInodes UsedInodes
mufilesetdemo Linked      /gpfs/fs1/mufilesetdemo
3          100352      100352      3
node1 1]
```

3. Create five files.

```
node1 1] echo "filedata" > /gpfs/fs1/mufilesetdemo/file1
node1 1] echo "filedata" > /gpfs/fs1/mufilesetdemo/file2
node1 1] echo "filedata" > /gpfs/fs1/mufilesetdemo/file3
node1 1] echo "filedata" > /gpfs/fs1/mufilesetdemo/file4
node1 1] echo "filedata" > /gpfs/fs1/mufilesetdemo/file5
```

4. Check inode stats : Here inode numbers are increased by five.

```
node1 1] mmlsfileset fs1 mufilesetdemo -i
Collecting fileset usage information ...
Filesets in file system 'fs1':
Name      Status      Path      InodeSpace      MaxInodes
AllocInodes UsedInodes
mufilesetdemo Linked      /gpfs/fs1/mufilesetdemo      3          100352
100352      8
node1 1]
```

5. Create an object list with all files and upload.

```
node1 1] cat /root/objectlist
/gpfs/fs1/mufilesetdemo/file1
/gpfs/fs1/mufilesetdemo/file2
/gpfs/fs1/mufilesetdemo/file3
/gpfs/fs1/mufilesetdemo/file4
/gpfs/fs1/mufilesetdemo/file5
node1 1]
node1 1]
node1 1] mmafmcscctl fs1 mufilesetdemo /gpfs/fs1/mufilesetdemo upload --object-list /root/
```

```

objectlist
  Queued      (Total)      Failed      TotalData
                        (approx in Bytes)
      5        (5)         0         45
Object Upload successfully queued at the gateway.
node1 1]

```

6. Remove all the files from the MU fileset.

```

node1 1] rm -rf /gpfs/fs1/mufilesetdemo/file*
node1 1]

```

7. Check the inodes stats : inodes are not yet reclaimed from MU fileset.

```

node1 1] mmlsfileset fs1 mufilesetdemo -i
Collecting fileset usage information ...
Filesets in file system 'fs1':
Name              Status      Path              InodeSpace      MaxInodes
AllocInodes      UsedInodes
mufilesetdemo    Linked    /gpfs/fs1/mufilesetdemo    3              100352
100352           8
node1 1]

```

8. Issue the following command to reclaim the inodes from MU mode fileset:

```

node1 1] mmafmcosctl fs1 mufilesetdemo /gpfs/fs1/mufilesetdemo/ delete --from-cache

```

9. Check the inode stats. In the following example, the used inodes are reclaimed.

```

node1 1] mmlsfileset fs1 mufilesetdemo -i
Collecting fileset usage information ...
Filesets in file system 'fs1':
Name              Status      Path              InodeSpace      MaxInodes
AllocInodes      UsedInodes
mufilesetdemo    Linked    /gpfs/fs1/mufilesetdemo    3              100352
100352           3
node1 1]

```

Note: The **mmafmcscctl** command with option **delete --from-cache**, does not queue deletes to cloud object storage.

Example: Deleting files, reclaiming inode, and queueing deletes to COS

1. Check MU mode fileset inode stats.

```

Node1 1] mmlsfileset fs1 mufilesetdemo -i
Collecting fileset usage information ...
Filesets in file system 'fs1':
Name              Status      Path              InodeSpace      MaxInodes      AllocInodes
UsedInodes
mufilesetdemo    Linked    /gpfs/fs1/mufilesetdemo    3              100352      100352      3
Node1 1]

```

2. Create five files in fileset1.

```

Node1 1] echo "filedata" > /gpfs/fs1/mufilesetdemo/file1
Node1 1] echo "filedata" > /gpfs/fs1/mufilesetdemo/file2
Node1 1] echo "filedata" > /gpfs/fs1/mufilesetdemo/file3
Node1 1] echo "filedata" > /gpfs/fs1/mufilesetdemo/file4
Node1 1] echo "filedata" > /gpfs/fs1/mufilesetdemo/file5
Node1 1]

```

3. Check the inodes stats:Used inodes increased by five.

```

Node1 1] mmlsfileset fs1 mufilesetdemo -i
Collecting fileset usage information ...
Filesets in file system 'fs1':
Name              Status      Path              InodeSpace      MaxInodes      AllocInodes
UsedInodes
mufilesetdemo    Linked    /gpfs/fs1/mufilesetdemo    3              100352      100352      8
Node1 1]

```

4. Create an object list and upload all the files to cloud object storage.

```

Node1 1]
Node1 1] mmafmcosctl fs1 mufilesetdemo /gpfs/fs1/mufilesetdemo upload --object-list /root/
objectlist
  Queued      (Total)      Failed      TotalData
                        (approx in Bytes)

```

```

5          (5)          0          45
Object Upload successfully queued at the gateway.
Node1 1]

```

5. Remove the files from the MU mode fileset.

```

Node1 1] rm -rf /gpfs/fs1/mufilesetdemo/file*
Node1 1]

```

6. Check the inode usage : Inodes are not reclaimed after delete.

```

Node1 1] mmlsfileset fs1 mufilesetdemo -i
Collecting fileset usage information ...
Filesets in file system 'fs1':
Name          Status Path          InodeSpace  MaxInodes  AllocInodes  UsedInodes
mufilesetdemo Linked /gpfs/fs1/mufilesetdemo 3          100352     100352      8
Node1 1]

```

7. Issue **mmafmcscctl** delete command with **-from-target** option.

```

Node1 1] mmafmcscctl fs1 mufilesetdemo /gpfs/fs1/mufilesetdemo/ delete --from-target

```

8. Check that the remove operations are queued and executed to cos.

```

Node1 1] mmafmcctl fs1 getstate
Fileset Name  Fileset Target  Cache State  Gateway Node
Queue Length  Queue numExec
-----
mufilesetdemo http://s3.us-east.cloud-object-storage.appdomain.cloud:80/demobucketx Dirty         c7f2n04
5          22

Node1 1] mmafmcctl fs1 getstate
Fileset Name  Fileset Target  Cache State  Gateway Node
Queue Length  Queue numExec
-----
mufilesetdemo http://s3.us-east.cloud-object-storage.appdomain.cloud:80/demobucketx Active        c7f2n04
0          27
Node1 1]

```

9. Check that the inodes in MU fileset are reclaimed.

```

Node1 1] mmlsfileset fs1 mufilesetdemo -i
Collecting fileset usage information ...
Filesets in file system 'fs1':
Name          Status Path          InodeSpace  MaxInodes  AllocInodes  UsedInodes
mufilesetdemo Linked /gpfs/fs1/mufilesetdemo 3          100352     100352      3
Node1 1]

```

MU mode autoremoval

To keep the deletes in sync with cloud object storage, the autoremoval feature can be configured on the MU fileset by using **mmchfileset** command. When the parameter **afmMUAutoRemove** is set to yes then removes are automatically queued to the cloud object storage.

For setting the **afmMUAutoRemove** parameter option first, complete the following steps:

1. Stop the fileset by using the **mmafmcctl** stop command.
2. Set the parameter on the fileset by using the **mmchfileset <fs> <fileset> -p afmMuAutoRemove=yes** command.
3. Again start the fileset by using the **mmafmcctl** start command.

Note: When the **afmMuAutoRemove** option is configured, the inodes are reclaimed from the MU mode fileset automatically.

Use case scenarios for the MU mode:

- The MU mode supports scenarios where the data changes at AFM to cloud object storage fileset is not needed to be in sync with the target cloud object storage. Applications can perform data modification at fileset level and whenever it is necessary to sync the data to the cloud object storage that can be synced.
- Similarly, application running on cloud can modify the data on cloud and when it is required to be in sync with AFM to cloud object storage MU mode fileset, the data can be downloaded. This helps in bandwidth use as data can be uploaded or downloaded at specific times. For example,

application works on files in the MU fileset and uploaded to COS at later time in the day when the load on the network is less.

Renaming files from an MU mode fileset

Files renamed in an AFM to cloud object storage fileset that is in MU mode do not get renamed on the cloud object storage target instantly, no matter the value of `afmMUAutoRemove`. System administrators can run **reconcile** command or **upload** command to synchronize with the cloud object storage targets.

Object operation modes

AFM to cloud object storage supports **ObjectFS** and **ObjectOnly** behaviors.

These two behaviors are defined during the creation of a fileset by using the `mmafmcosconfig --object-fs` command. If you specify the `--object-fs` option while you create an AFM to cloud object storage fileset, the **ObjectFS** mode behavior is enabled. If you skip the `--object-fs` option, an AFM to cloud object storage fileset is created with the **ObjectOnly** mode. The behavior is determined based on the workload requirement. When a mode is enabled, you cannot change it later.

objectFS mode

In the **objectFS** mode, an AFM to cloud object storage fileset is synchronized with a cloud object storage. AFM to cloud object storage RO, LU, and IW modes of filesets synchronize metadata to and from a cloud object storage. It includes operations such as `readdir` and `lookups` to synchronize the data. Objects are downloaded when they are read on demand or worked upon by an application that is running on an AFM to cloud object storage fileset. The `objectfs` enabled AFM to cloud object storage fileset behaves like an AFM fileset. For SW and IW modes-enabled AFM to cloud object storage fileset, the AFM to cloud object storage uploads files as objects to the object storage server. For an AFM RO, LU, and IW modes-enabled AFM to cloud object storage fileset, AFM automatically synchronizes objects from the cloud object storage to the AFM to cloud object storage fileset as files. Enable this parameter if the AFM to cloud object storage fileset behaves like an AFM modes fileset.

ObjectOnly mode

If the `object-fs` parameter is not defined, the **ObjectOnly** parameter will get enabled. This behavior is set on the AFM to cloud object storage filesets by default.

With the **ObjectOnly** mode, refresh of an AFM to cloud object storage fileset (AFM RO, LU, and IW mode fileset) with a cloud object storage will not be on-demand or frequent. You need to manually download data or metadata from the cloud object storage to the AFM to cloud object storage fileset, meanwhile data transfer from the AFM to cloud object storage fileset to the cloud object storage works automatically without manual intervention.

Enable this parameter on an AFM to cloud object storage fileset to avoid frequent trips and reduce the network contention by selectively performing the download or upload operations.

Multi-site replication

By using the multi-site replication, you can connect an AFM to cloud object storage fileset to multiple cloud object storage buckets by using AFM gateway nodes. Each gateway node connects to a bucket on a respective cloud object storage from the same or different cloud object storage provider. This connection enables placement of files or objects in an IBM Storage Scale cluster to multiple cloud object storage buckets simultaneously.

By enabling IBM Storage Scale multi-site replication (by using AFM), you can move data across multiple object storage providers. This feature provides high availability, fault tolerance, reduced latency, parallel access, and migration between clouds. Also, you can move datasets within regions for AI and hybrid cloud environments, which allow you to access data on a cheaper storage and where compute is available.

You can enable this feature on IBM Storage Scale Advanced Edition, IBM Storage Scale Data Management Edition, or IBM Storage Scale Erasure Code Edition.

Multi-site replication features

Disaster recovery and high availability

The multi-site replication ensures that stored data in an IBM Storage Scale cluster is simultaneously replicated in buckets of multiple object storage services. When a disaster or data center outage occurs, you can seamlessly switch to other site to minimize downtime and maintain business continuity.

Data redundancy and resilience

The replication of data to buckets on any cloud storage services enhances redundancy and resilience. When any cloud storage provider experiences an issue, data remains accessible from other sites, which improves system reliability.

Cross-cloud redundancy and high availability

The multi-site replication ensures that stored data in an IBM Storage Scale cluster is simultaneously replicated to multiple object storage providers. By using the multi-site replication, you can create redundancy on multiple cloud storage services. Because of redundancy, you can mitigate risks that are associated with potential issues or outages in a specific cloud platform.

Global collaboration and data distribution

Organizations with a global presence can use multi-site replication to distribute data across multiple regions. This feature facilitates collaboration among geographically dispersed teams and makes data available locally for better performance.

Efficient cloud migration and hybrid deployments

The multi-site replication feature can be valuable during cloud migration or hybrid cloud deployments. It seamlessly replicates and distributes data between on-premises storage and cloud storage, facilitates a smooth transition.

To enable this feature, see the *Configuring the multi-site replication of AFM to cloud object storage* section in *IBM Storage Scale: Administration Guide*.

AFM to cloud object storage parallel read data transfer

Parallel read data transfer improves the overall data read transfer performance of an AFM to cloud object storage fileset by using multiple gateway nodes. When this feature is enabled, objects that are stored on a cloud object storage can be fetched in an IBM Storage Scale cluster fileset by using multiple gateway nodes.

To help the primary gateway read large files from the cloud object storage, an AFM cache cluster can be configured to use multiple gateways. On a cloud object storage, the same endpoints can point to a bucket or can set up in the distributed mode, that is, multiple servers by using same storage paths. Cloud object storage servers also provide multi-threaded multi-part download to achieve a maximum read performance.

In an AFM cache, gateway nodes can be mapped to endpoints of cloud object storage services. An export map replaces the `--endpoint` server name in the **mmafmcconfig** command to set up parallel reads. An export map can be changed without modifying the `--endpoint` option for a fileset. However, the fileset must be relinked or the file system must be remounted for the mapping to take effect. For more information about how to define, display, delete, and update mappings, see the *mmafmconfig* command in *IBM Storage Scale: Command and Programming Reference Guide*.

To define and enable the parallel data read transfer for the AFM to cloud object storage, complete the following steps:

1. Define the mapping by using the **mmafmconfig** command.

A gateway can be defined manually by setting the **afmGateway** parameter. For more information, see *mmchfileset* command, in *IBM Storage Scale: Command and Programming Reference Guide*.

2. Set up the keys for the target buckets by using the **mmafmckeys** command.
3. Use the mapping as an `--endpoint` when you set the AFM to cloud object storage relationship by using the **mmafmcconfig** command.

4. Define the parallel read parameters. For more information about the parameters, see *Configuration parameters for AFM and AFM-DR in IBM Storage Scale: Administration Guide*.

All gateway nodes other than the primary gateway that is defined in a mapping are called participating or helper gateway nodes. The primary gateway of an AFM to cloud object storage fileset communicates with each of the participating gateway nodes, depending on their availability. When parallel data read transfer is configured, a single data transfer request is split into multiple chunks. Chunks are sent across to the participating gateway nodes in parallel for the transfer from a cloud object storage by using the respective mapping. The primary gateway processes the replies from all the participating gateway nodes, handles all data transfer failures, and coordinates activities until all data transfer is completed. If any participating gateway node fails, the primary gateway attempts to retry the failed task on the next available gateway and generates an error message in the mmfs log.

Note: For parallel data transfers, the gateway nodes in the mapping must have the same architecture and Linux operating system.

Parallel reads are effective on files or objects with sizes larger than files or objects that are specified by the parallel threshold. The threshold is defined by using the **afmParallelReadThreshold** parameter and is true for all types of objects and files.

Use the **afmParallelReadChunkSize** parameter to configure the size of each chunk. This parameter defines the minimum chunk size of the read that needs to be distributed among the gateway nodes during parallel reads. A zero value disables the parallel reads across multiple gateways.

Example

1. Create an export map by using multiple gateway nodes.

```
mmafmconfig add map1_cos --export-map lb1.ait.cleversafelabs.com/
c7f2n03,lb1.ait.cleversafelabs.com/c7f2n04
--no-server-resolution
```

A sample output is as follows:

```
mmafmconfig: Command successfully completed
mmafmconfig: Propagating the cluster configuration data to all affected nodes. This is
an asynchronous process.
```

Note: When the `--no-server-resolution` option is specified, AFM skips the DNS resolution of a specified hostname and creates the mapping. The specified hostname in the mapping must not be replaced with an IP address. Ensure that the specified hostname is resolvable when the mapping is in operation. This resolution is important if an endpoint resolves to multiple addresses and does not bind to a single IP address.

```
mmafmconfig show map1_cos
Map name: map1_cos
Export server map: lb1.ait.cleversafelabs.com/c7f2n03,lb1.ait.cleversafelabs.com/c7f2n04
```

2. Set up an access key and a secret key for the bucket by using the export map.

```
mmafmcoskeys cosbucket1:map1_cos set AccessKey SecretKey
```

3. Create an AFM to cloud object storage relation by using the **mmafmcosconfig** command.

```
mmafmcosconfig fs1 cosbucket1 --endpoint http://map1_cos --object-fs --bucket
cosbucket1 --debug --mode iw
afmobjfs=fs1 fileset=cosbucket1
bucket=cosbucket1 newbucket= objectfs=yes dir=
policy= tmpdir= tmpfile= cleanup=yes mode=iw
xattr=no ssl=no autoRemove=no acls=no gcs=no vhb=
bucketName=cosbucket1 region= serverName=map1_cos cacheFsType=http
Linkpath=/gpfs/fs1/cosbucket1 target=http://map1_cos/cosbucket1
map=http://map1_cos cacheHost=map1_cos
endpoint=lb1.ait.cleversafelabs.com cachePort=
endpoint=lb1.ait.cleversafelabs.com ENDPOINT=--endpoint http://map1_cos
XOPT= -p afmParallelWriteChunkSize=0 -p afmParallelReadChunkSize=0
```


Here,

cosbucket1

An existing bucket with objects presents on the cloud object storage for reading.

--new-bucket

This option can be used to create a new bucket.

4. Tune the parallel transfer thresholds for parallel reads. The **afmParallelReadThreshold** parameter value is 1 GB and the **afmParallelReadChunkSize** parameter value is 512 MB.

```
mmchfileset fs1 cosbucket1 -p afmParallelReadThreshold=1024
```

A sample output is as follows:

```
Fileset cosbucket1 changed.
```

```
mmchfileset fs1 cosbucket1 -p afmParallelReadChunkSize=536870912
```

A sample output is as follows:

```
Fileset cosbucket1 changed.
```

Note: Set these parameters before you do any operation on the fileset to get these values in effect. If any operations are already done on the fileset, use the **mmafmct1 stop** command, and then set these parameters and start the fileset.

5. Check that the values are set on the fileset.

```
mmlsfileset fs1 cosbucket1 --afm -L
```

A sample output is as follows:

```
Filesets in file system 'fs1':

Attributes for fileset cosbucket1:
=====
Status                               Linked
Path                                /gpfs/fs1/cosbucket1
Id                                   5
Root inode                           1572867
Parent Id                             0
Created                               Thu Aug 19 15:24:39 2021
Comment
Inode space                           3
Maximum number of inodes              100352
Allocated inodes                      100352
Permission change flag                chmodAndSetacl
afm-associated                         Yes
Target                               http://map1_cos:80/cosbucket1
Mode                                  independent-writer
File Lookup Refresh Interval          120
File Open Refresh Interval            120
Dir Lookup Refresh Interval           120
Dir Open Refresh Interval              120
Async Delay                           15 (default)
Last pSnapId                          0
Display Home Snapshots                no
Parallel Read Chunk Size              536870912
Parallel Read Threshold                1024
Number of Gateway Flush Threads        8
Prefetch Threshold                    0 (default)
Eviction Enabled                       yes (default)
Parallel Write Chunk Size              0
IO Flags                              0x0 (default)
```

6. List and read the object in the cache.

```
ls -lash /gpfs/fs1/cosbucket1/object1
0 -rwxrwx--- 1 root root 2.0G Aug 19 2021 object1
```

7. Get all content of the object.

```
# cat /gpfs/fs1/cosbucket1/object1 > /dev/null
```

8. List and read again the object in the cache.

```
ls -lash /gpfs/fs1/cosbucket1/object1
2.0G -rwxrwx--- 1 root root 2.0G Aug 19 2021 object1
```

9. Check the fileset status.

```
mmafmctl fs1 getstate
```

A sample output is as follows:

Fileset Name	Fileset Target	Cache State	Gateway Node	Queue Length	Queue numExec
cosbucket1	http://map1_cos:80/cosbucket1	Active	c7f2n03	0	10

10. To view the queue on the primary gateway and the helper gateway, use the following command:

```
mmfsadm saferdump afm cosbucket1
```

A sample output is as follows:

```
Output snip -
Normal Queue: (listed by execution order) (state: Active)
  0 Read [1572868.1572868] inflight (4194304 @ 12582912) tid 3454523
  2 ReadSplit [1572868.1572868] inflight (486539264 @ 1610612736) tid 3401451
```

Here,

ReadSplit

Operations on gateways that are defined in the mapping can be seen.

For more information, see the *mmfsadm command* in *IBM Storage Scale: Problem Determination Guide*.

Connectivity to cloud object storage

Each AFM to cloud object storage fileset in a cluster is served by one of the nodes that is designated as a gateway in the cluster.

The gateway node that is mapped to a fileset is called the primary gateway of the fileset. The primary gateway acts as the owner of the AFM fileset and communicates with the cloud object storage by using the **mmafmtransfer** command. An IBM Storage Scale cluster that has many AFM to cloud object storage filesets can have multiple gateway nodes. These gateway nodes connect to the cloud object storage by using HTTP or HTTPS protocols.

Target backend

AFM to cloud object storage can communicate with different cloud object storage backend. Some cloud object storage uses different API other than S3 standard API for the following reasons:

- To enable an AFM fileset to synchronize objects with different cloud to object storage backend such as Microsoft Azure Blob, Google GCP or AWS VHB style.
- To enable AFM to synchronize with specific cloud object storage backend, AFM to cloud object fileset must be created with a specified parameter.

The options are as follows:

--azure

Providing this option during the creation of an AFM to cloud object storage fileset by using the **mmafmcconfig** command enables AFM to use Microsoft Azure Blob storage specific API.

--gcs

Providing this option during the creation of an AFM to cloud object storage fileset by using the **mmafmcconfig** command enables AFM to use Google cloud object storage specific API.

--vhb

Provides support of S3 style virtual hosting of bucket such as URL to be used during the creation of an AFM to cloud object storage fileset. For VHB, the format of endpoint is such as `https://bucket-name.s3.region-code.amazonaws.com`. AFM requires a region to be added to the URL and separated it by using '@' symbol.

An example is as follows:

1. Add the key by issuing the following command:

```
mmafmckeys regbkt1**:**ap-south-1@regbkt1.s3.amazonaws.comset akey1 skey1
```

2. Verify that the key was added successfully by issuing the following command:

```
mmafmckeys regbkt1:ap-south-1@regbkt1.s3.amazonaws.com get
```

A sample output is as follows:

```
regbkt1:ap-south-1@regbkt1.s3.amazonaws.com=COS:akey1:skey1
```

3. Create an AFM to cloud object storage fileset by issuing the following command:

```
**#**mmafmcconfig fs1 iw1 --endpoint https://ap-south-1@regbkt1.s3.amazonaws.com --object-fs --xattr --bucket regbkt1 --mode iw --acls --vhb --directory-object
```

Note: To enable AFM to cloud object storage for standard other S3 cloud object storage backends, above options are not needed. Default option implicitly enables AFM to use Standard S3 APIs.

Underlying protocol

Communication between IBM Storage Scale and a cloud object storage can be enabled by using HTTP or HTTPS.

The AFM to cloud object storage also enables sharing of data across AFM to cloud object storage clusters and a cloud object storage, even if the networks are unreliable or have high latency by using the built-in queuing optimization techniques of AFM. For more information, see [“Active File Management” on page 54](#).

Note:

Within a single IBM Storage Scale cluster where AFM to cloud object storage is enabled, application nodes comply with POSIX semantics. The file or object locking across cache and a cloud object storage is not supported.

When you perform operations on AFM to cloud object storage filesets, ensure that the operations are supported on a cloud object storage over the chosen protocol. Because the operations that are performed from an AFM to cloud object storage fileset are replayed on the cloud object storage as normal object operations. Ensure that the bucket restrictions and limitations from a cloud object storage provider are followed.

The cache state of an AFM to cloud object storage is similar to an AFM or AFM-DR fileset. You can check the state by using the **mmafmctl getstate** command. This command is used to check the status of synchronization of the specified fileset. For example, if a bucket is not accessible on the cloud object storage server, the fileset is moved to the Unmounted state and the following log error is displayed:

```
2020-08-27_10:32:01.571-0400: [E] AFM: COS access of 192.168.1.1://bucket1 failed with error 13 for file system fs1 fileset ID 9. Caching will be disabled and the access will be tried again after 300 seconds, on the next request to the gateway.
```

When the bucket is accessible, AFM to cloud object storage tries the synchronization again. Similarly, other states become valid with the AFM to cloud object storage fileset.

For more information about AFM fileset states, see *Monitoring AFM and AFM DR* in the *IBM Storage Scale: Problem Determination Guide*.

Eviction in AFM to cloud object storage

AFM to cloud object storage filesets use a storage space management process that is called eviction. There are two types of eviction methods, data eviction and metadata eviction.

Eviction is used to make space for new objects and files on an AFM to cloud object storage fileset. Eviction releases object or file data blocks (space) or its metadata (inode) in the fileset, if the fileset usage exceeds the fileset quota. This feature is useful when the storage that is provided by IBM Storage Scale is less than the cloud object storage. The quotas can be defined when you set the AFM to cloud object storage relationship by using the **mmafmcconfig** command with **--quota-files NumberOfFiles** and **--quota-blocks NumberOfBlocks** parameters.

The process of releasing blocks or inodes is called eviction. However, file data or metadata is not evicted if the file data is dirty. A file whose outstanding changes are not flushed to the home is called a dirty file. When files or objects are evicted from the cache, all data blocks of those files are cleared. The files or objects are uncached. When a read operation is performed on the files or objects, the files or objects fetch data from a cloud object storage. Similarly, any metadata operation such as lookup or ls populate the inodes that were evicted from the AFM to cloud object storage fileset.

While provisioning to control the number of inodes that are used by a cache or home fileset, you can specify a limit on the number of inodes for that fileset during the fileset creation or change the limits by using the **mmchfileset** command with the **--inode-limit** option. For more information, see *mmchfileset* command in the *IBM Storage Scale: Command and Programming Reference Guide*.

For more information about the data and metadata eviction, see *Evicting files or objects data and Evicting files or objects metadata* in the *IBM Storage Scale: Administration Guide*.

Data eviction from an AFM to cloud object storage fileset after uploading

AFM to cloud object storage fileset employs high-performance caching with its extension bucket located on the cloud object storage. A caching fileset is planned with faster and optimal storage and the data created in the cache can be sent across to the respective cloud object storage backends.

With ever demanding workloads and performance centric cache applications, the storage needs to be optimized at the fileset level. This feature optimizes the storage at the AFM to cloud object storage fileset level by evicting the data of files or objects after these files and objects are successfully uploaded on the cloud object storage buckets. The data eviction of uploaded files and objects frees space in the fileset. This available space helps applications to create new or other files and objects and work on them without hampering performance.

By using **mmafmcctl upload --evict-after-upload** command all files or files specified by **--object-list** can be uploaded, and then evicted from the cache for the space management.

This feature supports the manual updates (MU) mode because of the enabled reconcile feature. By using reconcile, a policy can be defined to reconcile files with AFM to cloud object storage fileset and its respective cloud bucket. The files curated by a user-defined policy or a default policy are uploaded to the bucket, and then evicted from the cache. See the **mmafmcctl reconcile --evict-after-upload** command in the *IBM Storage Scale: Command and Programming Reference Guide* for more information.

When an application or administrators require the data of an evicted file from the cache, it can be fetched on demand or can be downloaded by using the **mmafmcctl download** command.

Advantages of the data eviction after uploading

Optimized storage utilization

By offloading data to the typically lower-cost cloud object storage backend, AFM maximizes the efficiency of the local cache. This prevents unnecessary storage costs and ensures optimal use of high-performance storage.

Enhanced performance

A well-managed cache is crucial for maintaining high application performance. The eviction of data to cloud object storage frees up cache space for new, frequently accessed data, preventing performance degradation due to cache misses.

Cloud object storage reduction

Depending on storage pricing models, storing data in cloud object storage might be more cost-effective than retaining it in the local cache. Eviction can help reduce overall storage expenses.

Use cases

Quick Data workload: For temporary files, intermediate results, or data with a short access window, eviction after upload is particularly beneficial. It prevents cache getting denser and ensures optimal performance for frequently accessed data.

Long-term Data Retention: Data that needs to be retained for extended periods can be migrated to the more durable and cost-effective cloud object storage. This approach extends data retention while optimizing cache usage.

Workload Optimization: By selectively evicting data based on access patterns or data age, organizations can optimize cache utilization and workload distribution.

Note: The renamed file is not evicted immediately until data is modified.

Example

1. Create two filesets in the single writer (SW) mode and the manual update (MU) mode.

```
mmafmcconfig fs1 fileset1 --endpoint http://s3.us-east-1.amazonaws.com --object-fs --xattr --new-bucket fileset1bucket --mode sw --acls --directory-object
mmafmcconfig fs1 fileset2 --endpoint http://s3.us-east-1.amazonaws.com --object-fs --xattr --bucket fileset2bucket --mode mu --acls --directory-object
```

where:

fileset1

Is created in the SW mode and this fileset connects to *fileset1bucket* buckets on the on the cloud object storage.

fileset2

Is created in the MU mode and this fileset connects *fileset2bucket* buckets on the cloud object storage.

2. Create data.

```
for a in `seq 3`; do dd if=/dev/urandom of=/gpfs/fs1/fileset1/file$a count=40 bs=256k ; done
for a in `seq 3`; do dd if=/dev/urandom of=/gpfs/fs1/fileset2/file$a count=40 bs=256k ; done
```

3. List data with space usage.

```
ls -sh /gpfs/fs1/fileset1
```

A sample output is as follows:

```
total 30M
10M file1 10M file2 10M file3

Node1] ls -sh /gpfs/fs1/fileset2
total 30M
10M file1 10M file2 10M file3
```

4. Evict the data after uploading from the respective filesets.

```
mmafmcctl fs1 fileset1 /gpfs/fs1/fileset1 upload --all --evict-after-upload
```

A sample output is as follows:

```
Queued Failed TotalData
(approx in Bytes)
3 0 31457280
Object upload successfully queued at the gateway.
```

5. Reconcile the data after uploading to the respective filesets.

```
mmafmcscctl fs1 fileset2 /gpfs/fs1/fileset2/ reconcile --evict-after-upload
```

A sample output is as follows:

```
Dirty file list : /var/mmfs/afm/fs1-7/recovery/policylist.data.list.dirtyFiles
Threads in use: 1
Queued (Total) Failed TotalData
(approx in Bytes)
3 (3) 0 31457280
Object Upload successfully queued at the gateway.
```

6. Verify that objects are stored on the cloud object storage.

```
amazonconsole] /root/mc ls aws/fileset1bucket
```

A sample output is as follows:

```
[2024-07-26 05:49:38 EDT] 10MiB file1
[2024-07-26 05:49:41 EDT] 10MiB file2
[2024-07-26 05:49:43 EDT] 10MiB file3
amazonconsole] ls aws/fileset2bucket
[2024-07-26 05:54:41 EDT] 10MiB file1
[2024-07-26 05:54:42 EDT] 10MiB file2
[2024-07-26 05:54:42 EDT] 10MiB file3
```

7. Check data usage after the data eviction.

```
ls -sh /gpfs/fs1/fileset1
```

A sample output is as follows:

```
total 0
0 file1 0 file2 0 file3
Node1] ls -sh /gpfs/fs1/fileset2
total 0
0 file1 0 file2 0 file3
```

Audit messages support for the AFM to cloud object storage

The audit messages support provides better troubleshooting and improves the cluster security by sending audit message to the configured logs.

The AFM to cloud object storage feature is enabled with the support of audit message logging for the following commands:

- **mmafmcaccess**
- **mmafmcconfig**
- **mmafmcscctl**
- **mmafmckeys**

IBM Storage Scale can send an audit message to the syslog and the GPFS log whenever these commands are invoked.

For more information about how to configure, see *Audit messages for cluster configuration changes* in the *IBM Storage Scale: Problem Determination Guide*.

Partial file or object caching for AFM to cloud object storage

When the partial file or object caching feature is enabled, an application fetches only required data from an object or a file on a cloud object storage. Therefore, network and local disk space are used more

efficiently. If an application does not need to read an entire object or file, this feature must be enabled. This feature is enabled on a block boundary.

The partial caching of an object or a file is controlled by setting the **afmPrefetchThreshold** parameter value. The default value of this parameter is 0. Because of this value, an entire file is cached and all blocks of a file are fetched after any three blocks are read by the cache and the file is marked as cached. This value is useful for sequentially accessed files that are completely read such as image files. To configure this parameter, see *mmchfileset command in IBM Storage Scale: Command and Programming Reference Guide*.

Valid values of the **afmPrefetchThreshold** parameter are in the range 1 – 100. This parameter value specifies the file size percentage that must be cached before all other the data blocks are automatically fetched into the cache. A higher value is suitable for a file that is accessed partially. When the **afmPrefetchThreshold** value is set to 100, it disables prefetching an entire file. This value caches only the data blocks that are read by an application. Also, large random-access files that do not fit in the cache are not read completely. When all data blocks are available in the cache, the file is marked as cached.

Note:

- The download or the prefetch feature fetches full objects from a cloud object storage though they are partially cached previously. Prefetch of partially cached objects pulls the entire objects to the cache.
- Failover of partially cached objects pushes only the cached data blocks to the target. Uncached blocks are filled with null bytes.
- Any writes queued on a partially cached object fetch the entire file, even if a prefetch threshold limit is set on the object.
- Object uses default 16M as chunk size and download in multiple of 16M data.

If a write operation is queued on a file that is partially cached, the entire file is cached first, and then the write operation is queued on the file. Appending to a partially cached file does not cache the entire file. Only in the LU mode, the write inset or append on a file that is cached partially caches the entire file even if the prefetch threshold is set on the fileset.

Example

1. The *objectbucket* bucket is present on the cloud object storage with two objects as follows:

```
Name      : object100M
Date       : 2021-05-21 08:26:11 EDT
Size       : 100 MiB
ETag       : 0073370fd78dd34e5043d13ba515b5a2
Type       : file
Metadata   :
  Content-Type: application/octet-stream
```

```
Name      : object1G
Date       : 2021-05-21 08:26:11 EDT
Size       : 1000 MiB
ETag       : b7fabalddde52a27fb925858102db50b-8
Type       : file
Metadata   :
  Content-Type: application/octet-stream
```

2. Set keys on an IBM Storage Scale AFM cache cluster.

```
# mmafmcoskeys objectbucket:192.168.118.121 set sdjnlSDLknsf3093mkey1
skdjfnrkfjnergergnwegwrgulKjfu12
```

3. Create an AFM to cloud object storage relationship.

```
# mmafmcosconfig fs1 objectfileset --endpoint http://192.168.118.121 --bucket objectbucket
--mode iw --object-fs --debug
afmobjfs=fs1 fileset=objectfileset
bucket=objectbucket newbucket= objectfs=yes dir=
policy= tmpdir= tmpfile= cleanup=no mode=iw
xattr=no ssl=no acls=no gcs=no vhb=
```

```
bucketName=objectbucket region= serverName=192.168.118.121 cacheFsType=http
Linkpath=/gpfs/fs1/objectfilesset target=http://192.168.118.121/objectbucket
map=http://192.168.118.121 cacheHost=192.168.118.121
endpoint=192.168.118.121 ENDPOINT=--endpoint http://192.168.118.121
XOPT= -p afmParallelWriteChunkSize=0 -p afmParallelReadChunkSize=0
```

4. Verify the fileset information.

```
# mmlsfilesset fs1 objectfilesset --afm -L
```

A sample output is as follows:

```
Filesets in file system 'fs1':
Attributes for fileset objectfilesset:
=====
Status                               Linked
Path                                /gpfs/fs1/objectfilesset
Id                                   7
Root inode                           2097155
Parent Id                             0
Created                             Fri May 21 08:31:56 2021
Comment
Inode space                           4
Maximum number of inodes              100352
Allocated inodes                      100352
Permission change flag                chmodAndSetacl
afm-associated                         Yes
Target                               http://192.168.118.121:80/objectbucket
Mode                                 independent-writer
File Lookup Refresh Interval          120
File Open Refresh Interval            120
Dir Lookup Refresh Interval           120
Dir Open Refresh Interval             120
Async Delay                           15 (default)
Last pSnapId                          0
Display Home Snapshots                no
Parallel Read Chunk Size              0
Number of Gateway Flush Threads       4
Prefetch Threshold                    0 (default)
Eviction Enabled                      yes (default)
Parallel Write Chunk Size             0
IO Flags                              0x0 (default)
```

Note: Prefetch threshold limit is set to 0 by default.

5. Stop the fileset.

```
# mmafmctl fs1 stop -j objectfilesset
```

6. Change the prefetch threshold limit.

```
# mmchfilesset fs1 objectfilesset -p afmprefetchthreshold=100
```

A sample output is as follows:

```
Fileset objectfilesset changed.
```

7. Start the fileset.

```
# mmafmctl fs1 start -j objectfilesset
```

8. Verify the prefetch threshold limit.

```
# mmlsfilesset fs1 objectfilesset --afm -L
```

A sample output is as follows:

```
Filesets in file system 'fs1':
Attributes for fileset objectfilesset:
=====
Status                               Linked
Path                                /gpfs/fs1/objectfilesset
```



```

Id 7
Root inode 2097155
Parent Id 0
Created Fri May 21 08:31:56 2021
Comment
Inode space 4
Maximum number of inodes 100352
Allocated inodes 100352
Permission change flag chmodAndSetacl
afm-associated Yes
Target http://192.168.118.121:80/objectbucket
Mode independent-writer
File Lookup Refresh Interval 120
File Open Refresh Interval 120
Dir Lookup Refresh Interval 120
Dir Open Refresh Interval 120
Async Delay 15 (default)
Last pSnapId 0
Display Home Snapshots no
Parallel Read Chunk Size 0
Number of Gateway Flush Threads 4
Prefetch Threshold 100
Eviction Enabled yes (default)
Parallel Write Chunk Size 0
IO Flags 0x0 (default)

```

Note: The prefetch threshold is set to 100. That means only the read blocks are cached.

9. To check the objects on a cloud object storage, issue the **ls** command.

```
# ls -lash /gpfs/fs1/objectfilesset/
```

A sample output is as follows:

```

total 259K
512 drwxrws--- 5 root root 4.0K May 21 08:34 .
256K drwxrwxrwx 9 root root 256K May 21 08:31 ..
0 -rwxrwxrwx 1 root root 100M May 21 2021 object100M
0 -rwxrwxrwx 1 root root 1000M May 21 2021 object1G

```

10. Read the object partially by using the **dd** command.

```
# dd if=/gpfs/fs1/objectfilesset/object100M bs=4M count=10 > /dev/urandom
```

```

10+0 records in
10+0 records out
41943040 bytes (42 MB, 40 MiB) copied, 1.43027 s, 29.3 MB/s

```

```
# dd if=/gpfs/fs1/objectfilesset/object1G bs=4M count=100 > /dev/urandom
```

```

100+0 records in
100+0 records out
419430400 bytes (419 MB, 400 MiB) copied, 14.129 s, 29.7 MB/s

```

11. Check the disk usage of these objects. They are the same as the data read by application or the **dd** command.

```
# du -h /gpfs/fs1/objectfilesset/object1G
```

```
400M /gpfs/fs1/objectfilesset/object1G
```

```
# du -h /gpfs/fs1/objectfilesset/object100M
```

```
40M /gpfs/fs1/objectfilesset/object100M
```

AFM to cloud object storage directory object support

With directory object support, you can set access control lists (ACLs) and extended attributes on empty and non-empty directories, both. These empty or non-empty directories are transferred to the home along with the files as normal operations.

AFM to cloud object storage supports directory objects. All directories with and without objects can now be synchronized to the cloud object storage. Now, you can set extended attributes and ACLs on the directories with the directory object support.

Enabling directory object support for AFM to cloud object storage

Following are the two ways to enable the directory object support for AFM to cloud object storage:

1. When you start creating a new AFM to cloud object storage relationship, use the `--directory-object` option to enable the directory object support.

```
mmafmcconfig fs1 fileset1 --endpoint http://region@endpoint --object-fs --xattr --acls
--bucket bucket1 --mode ro --directory-object
```

2. For old AFM to cloud object storage file sets that are created in old version of the IBM Storage Scale, use the **mmchfileset** command after you stop the file set and then start the file set.

```
node1] mmlsfileset fs1 fileset1 -L --afm
Filesets in file system 'fs1':

Attributes for fileset fileset1:
=====
Status                               Linked
Path                                /gpfs/fs1/fileset1
Id                                  42
Root inode                          18350083
Parent Id                           0
Created                             Mon Apr 11 05:01:34 2022
Comment
Inode space                          19
Maximum number of inodes             100352
Allocated inodes                     100352
Permission change flag               chmodAndSetacl
afm-associated                       Yes
Permission inherit flag               inheritAclOnly
Target                               https://s3.amazonaws.com:443/fileset1
Mode                                 single-writer
File Lookup Refresh Interval         120
File Open Refresh Interval           120
Dir Lookup Refresh Interval          120
Dir Open Refresh Interval            120
Async Delay                          1 (default)
Last pSnapId                         0
Display Home Snapshots               no
Parallel Read Chunk Size              0
Number of Gateway Flush Threads       8
Prefetch Threshold                   0 (default)
Eviction Enabled                     yes (default)
Parallel Write Chunk Size             0
IO Flags                             0x48080000
(afmObjectXattr,afmObjectACL,afmObjectFastReaddir)
IO Flags2                            0x0
```

```
Node1] mmafmctl fs1 stop -j fileset1
```

```
Node1] mmchfileset fs1 fileset1 -p afmObjectDirectoryObj=yes
```

```
Node1] mmlsfileset fs1 fileset1 -L --afm
Filesets in file system 'fs1':

Attributes for fileset fileset1:
=====
Status                               Linked
Path                                /gpfs/fs1/fileset1
Id                                  42
```

```

Root inode                18350083
Parent Id                 0
Created                  Mon Apr 11 05:01:34 2022
Comment
Inode space              19
Maximum number of inodes 100352
Allocated inodes         100352
Permission change flag   chmodAndSetacl
afm-associated           Yes
Permission inherit flag  inheritAclOnly
Target                   https://s3.amazonaws.com:443/fileset1
Mode                     single-writer
File Lookup Refresh Interval 120
File Open Refresh Interval 120
Dir Lookup Refresh Interval 120
Dir Open Refresh Interval 120
Async Delay              1 (default)
Last pSnapId             0
Display Home Snapshots  no
Parallel Read Chunk Size 0
Number of Gateway Flush Threads 8
Prefetch Threshold       0 (default)
Eviction Enabled         yes (default)
Parallel Write Chunk Size 0
IO Flags                  0x48280000
                        (afmObjectXattr,afmObjectDirectoryObj,afmObjectACL,afmObjectFastReaddir)
IO Flags2                 0x0
Node1] mmafmctl fs1 start -j fileset1

```

AFM to cloud object storage support for more than 2 K metadata

AFM to cloud object storage supports setting more than 2 K metadata on the fileset objects.

The S3 protocol has limitation of 2 K total object metadata. This limitation restricts true synchronization of metadata to and from cloud object storage where objects are set with more than 2 K metadata. AFM achieves this objective by splitting the extended metadata objects and keeping them as separate files in `.afm` directory that is created in the cloud object storage bucket. This method is useful when fileset is used for managing large numbers of users and groups or extended attributes and access controlled lists (ACLs) that is set through distributed file systems such as NFS.

When more than 2 K metadata (xattr and ACL size exceeds 1800 bytes. As 200 bytes are reserved for other AFM metadata attributes) is set on the object, its metadata is separated in the file that is stored in the `.afm` directory. The `.afm` directory is available on the cloud object storage, with the same path as the original object.

This feature is enabled by default on AFM to cloud object storage filesets.

Renaming an object with more than 2 K metadata deletes the old metadata object that is stored in `.afm` directory on cloud object storage. This process creates a new object that is kept with the renamed object name.

Similarly, deleting an object with more than 2 K metadata deletes the object in `.afm` directory on cloud object storage.

Note: Setting more than 2 K metadata on the AFM to cloud object storage fileset root (also known as the fileset path) is not supported.

Examples

```
mmafmctl fs1 getstate
```

Fileset Name	Fileset Target	Cache State	Gateway Node	Queue Length	Queue numExec
fileset1	https://s3.amazonaws.com:443/fileset1	Active	c7f2n04	0	12

More than 2 K metadata is set on object object1 as shown in the following sample.

```
mmlsattr -L -d /gpfs/fs1/fileset1/object1
```


A policy rule is an SQL-like statement that directs **mmafmcscctl** reconcile command to upload data to the cloud object storage based on criteria defined in the policy file.

Note: Based on policy, the **mmafmcscctl reconcile** command does not upload a directory with objects that do not meet the policy criteria to cloud objects storage.

A policy rule specifies one or more conditions. When the conditions are true, the specific rule applies. Conditions can be specified by SQL expressions, which can include SQL functions, variables, and file attributes.

Few available file attributes are shown in the following list:

- File name or extension
- File size
- User ID and group ID
- Date and time when the file was last accessed
- Date and time when the file was last modified

Steps to policy-based upload for the manual updates mode

1. According to the data created in the manual updates fileset, system administrator defines a policy. For example, upload all the objects or files with the specific names. For more information, see *Creating a policy* in the *IBM Storage Scale: Administration Guide*.
2. As a system administrator, you can either install this policy permanently by using **mmafmcscctl -add-policy** command or run it when needed, by using **mmafmcscctl -policy** command.
3. When the policy is installed permanently by using **-add-policy**, **mmafmcscctl** reconcile command uses this policy to run when **mafmcscctl Device FilesetName path reconcile** command is run without any option.
4. As a system administrator, you can remove the installed policy by using **mmafmcscctl Device FilesetName Path reconcile -remove-policy** command.
5. As a system administrator, you can also run a policy right away by using **mmafmcscctl Device FilesetName Path reconcile -policy** command. Here, the policy is not stored internally.
6. As a system administrator, you can view the installed policy by using **--list-policy** option.

Note: Reconcile command also performs cleanup of files when it is run along with the policy based uploads. Files that are deleted from manual update mode fileset are queued for deletes on the Cloud Object Storage and their inodes are reclaimed.

Policy creation

As a system administrator you can define a policy for variety of matching options, here is the policy that uploads all files and directories with matching name "IBM".

Important: When policy is created, make sure that the LIST option must have prerequisite names that are, "dirtyFiles" for files and "dirtyDirs" for directories.

Example

Manual updates mode fileset

```
Node1] mmlsfileset fs1 mufileset --afm -L
```

Filesets in file system 'fs1':

```
Attributes for fileset mufileset:
=====
Status                               Linked
Path                                /gpfs/fs1/mufileset
Id                                   125
Root inode                          102236163
```

```

Parent Id                                0
Created                                Wed Apr 20 03:54:03 2022
Comment
Inode space                            99
Maximum number of inodes                100352
Allocated inodes                       100352
Permission change flag                  chmodAndSetacl
afm-associated                          Yes
Permission inherit flag                 inheritAclOnly
Target                                 https://s3.amazonaws.com:443/mufilesset
Mode                                    manual-updates
File Lookup Refresh Interval            120
File Open Refresh Interval              120
Dir Lookup Refresh Interval             120
Dir Open Refresh Interval               120
Async Delay                             disable
Last pSnapId                            0
Display Home Snapshots                  no
Parallel Read Chunk Size                 0
Number of Gateway Flush Threads          8
Prefetch Threshold                      0 (default)
Eviction Enabled                        yes (default)
IO Flags                                0x8280000
(afmObjectXattr,afmObjectDirectoryObj,afmObjectACL)
IO Flags2                                0x0
Node1]

```

Files present in the manual-updates mode fileset

```
Node1] ls -l /gpfs/fs1/mufilesset/
```

```

total 0
-rw-r--r-- 1 root root 0 Apr 29 08:42 extrafile1
-rw-r--r-- 1 root root 0 Apr 29 08:42 extrafile2
-rw-r--r-- 1 root root 0 Apr 29 08:42 extrafile3
-rw-r--r-- 1 root root 0 Apr 29 08:42 extrafile4
-rw-r--r-- 1 root root 0 Apr 29 08:42 extrafile5
-rw-r--r-- 1 root root 0 Apr 29 08:41 file1
-rw-r--r-- 1 root root 0 Apr 29 08:41 file2
-rw-r--r-- 1 root root 0 Apr 29 08:41 file3
-rw-r--r-- 1 root root 0 Apr 29 08:41 file4
-rw-r--r-- 1 root root 0 Apr 29 08:41 file5
-rw-r--r-- 1 root root 0 Apr 29 08:42 IBM1
-rw-r--r-- 1 root root 0 Apr 29 08:42 IBM2
-rw-r--r-- 1 root root 0 Apr 29 08:42 oneIBM
-rw-r--r-- 1 root root 0 Apr 29 08:42 twoIBM

Policy to transfer files which have IBM in their names respectively
Node1] cat policyfile1
RULE EXTERNAL LIST 'dirtyFiles'
RULE 'dirtyFilesRule' LIST 'dirtyFiles'
WHERE NAME LIKE '%IBM%' AND
PATH_NAME NOT LIKE '$filesetPath/.pconflicts/%' AND
PATH_NAME NOT LIKE '$filesetPath/.afm/%' AND
PATH_NAME NOT LIKE '$filesetPath/.ptrash/%'

```

Issue the following command on Node1]:

```
mmafmcctl fs1 mufilesset /gpfs/fs1/mufilesset reconcile --policy policyfile1
```

```

Dirty file list : /var/mmfs/afm/fs1-125/recovery/policylist.data.list.dirtyFiles
  Queued      (Total)      Failed      TotalData
                        (approx in Bytes)
        4          (4)          0              0

Object Upload successfully queued at the gateway.
Could not find any deleted files.

```

```
Node1] mmafmctl fs1 getstate
```

Fileset Name	Fileset Target	Cache State	Gateway Node	Queue
Length	Queue numExec			
-----	-----	-----	-----	
-----	-----			

mufileset 4	https://s3.amazonaws.com:443/mufileset 691	Active	c7f2n04

File match in name as IBM are transferred to cloud objects (seen by using cloud object shell):

```
[2022-04-29 08:45:37 EDT] 0B STANDARD IBM1
[2022-04-29 08:45:37 EDT] 0B STANDARD IBM2
[2022-04-29 08:45:37 EDT] 0B STANDARD oneIBM
[2022-04-29 08:45:37 EDT] 0B STANDARD twoIBM
Node1]
```

Support of Google Cloud Platform for AFM to cloud object storage

With Google Cloud Platform (GCP) support, data can be synchronized between AFM, S3 object to Google Cloud Platform storage.

AFM to cloud object storage files can be configured to use Google Cloud Storage (GCS) Platforms backend object storage server by using **-gcs** parameter during creation of the AFM fileset. After a fileset is created by using **-gcs** option, AFM access GCS backend storage server by using supported APIs. Here, Google cloud storage APIs are different than the traditional S3 API.

After a fileset is configured with GCS, the fileset can only be used for GCS backend server. AFM internally determines which API to use for communicating to the target cloud object storage server for all the communications.

All the modes and operations are supported by the GCP object storage server backend.

Note:

- While you use Google Cloud Storage platform server as backend cloud object storage server, AFM is unable to create new bucket. User with an administrator privilege must create bucket at the object storage server before user creates AFM to cloud object storage fileset. This preceding limitation only exists with the GCS backend and not with other supported cloud object storage server.
- Failover of AFM to cloud object storage fileset that is configured by using **-gcs** option cannot be failover to another backend.

Symbolic links

Active File Management (AFM) to cloud object storage supports symbolic links, which are also called symlinks.

When symlinks are created on an AFM to cloud object storage cache fileset, the relationship between the symlink file and the target file is maintained and the files are uploaded to the bucket. The symlinks relationship is maintained only if they are created in an AFM to cloud object storage fileset. If the symlinks relationship is not created in the AFM to cloud object fileset, it is not maintained.

The eviction of symlinks is not supported. AFM to cloud object storage does not follow the symlink to evict an original file. The eviction might throw an error while processing a symlink file and the next file is processed.

Replication by using the manual updates mode of the AFM to cloud object storage

IBM Storage Scale provides a powerful solution for managing large-scale file systems across multiple nodes and sites. The replication at the file system level by using the AFM to cloud object storage provides a reliable and efficient solution for replicating or tiering data from a file system to a cloud-based object storage system.

The IBM Storage Scale file system level replication by using the AFM to cloud object storage in the manual updates mode gives you control over the replication process, and you can replicate data from an IBM Storage Scale file system to a cloud object storage manually and automatically by using reconcile policies. For more about information reconciliation, see the *Reconciling files between IBM Storage Scale file system and cloud storage tier* in the *IBM Storage Scale: Administration Guide*.

This feature supports the following use cases:

- **Backup and recovery:** By replicating data from a file system to a cloud object storage, you can create a backup copy of critical data. In the manual update mode, you have the flexibility to schedule replication based on your backup strategy. This ensures that your data is protected and can be restored in the event of a disaster.
- **Data archival and retention:** Storing data for long-term archival and compliance purposes is essential for many organizations. By using the manual updates mode of the AFM to cloud object storage of the replication, you can periodically replicate data to a cloud object storage by ensuring its preservation and adherence to data retention policies. This reduces storage cost on the primary file system while maintaining data accessibility when needed.
- **Migration to a cloud storage:** The replication by using the manual updates mode of the AFM to cloud object storage is valuable for migrating data from on-premises file systems to a cloud storage. You can tier the data incrementally, validating each step before you initiate the replication process. This approach allows a smooth transition and minimizes impact on ongoing operations.
- **Data analysis and processing:** When you analyze data or process tasks, it is often beneficial to replicate relevant data to a cloud storage. By using the manual updates mode of the AFM to cloud object storage for the replication, you can choose time to replicate data. It ensures that the most up-to-date information is available for processing, without any impact on running operations.
- **Hybrid cloud storage:** The replication by using the manual updates mode of the AFM to cloud object storage replication is ideal for hybrid cloud storage, where you have a combination of on-premises file systems and cloud-based storage. When this feature is enabled, you can synchronize data between the two environments based on your requirement, and it ensures that data is consistent across locations.

For more information about the manual updates mode, see [“Fileset operation modes” on page 157](#).

For more information the replication by using the manual updates mode, see the following sections:

- [“The limitations of the replication at the file system level by using the manual updates mode” on page 185](#)
- *Configuring the replication at the file system level by using the manual updates mode of the AFM to cloud object storage* section in the *IBM Storage Scale: Administration Guide*
- *Evicting data or objects by using the manual updates mode of the AFM to cloud object storage* section in the *IBM Storage Scale: Administration Guide*

Parameters for a new file system

Set the parameters that are listed in the following table after a file system is created.

Table 11. Parameters for the replication of a new file system by using the AFM to cloud object storage manual updates mode	
Parameter	Description
afmObjectXattr	Specifies user-extended attributes to be synchronized with a cloud object storage. If this option is skipped, the AFM to cloud object storage does not synchronize user-extended attributes with the cloud object storage.
afmObjectDirectoryObj	AFM to cloud object storage supports directory objects. All directories with and without objects can now be synchronized to the cloud object storage. You can now set extended attributes on directories with the directory object support.
afmObjectACL	Enables the cache to synchronize ACL to the cloud object storage server. If not specified, AFM will not synchronize ACLs to cloud object storage. Default is disabled ACL synchronization.

Table 11. Parameters for the replication of a new file system by using the AFM to cloud object storage manual updates mode (continued)

Parameter	Description
afmObjectFastReaddir	Improves the read-dir performance of an AFM to cloud object storage fileset by skipping synchronization of extended attributes and ACLs of an object between a cloud object storage server and a cache fileset. When this option is set, deleted objects on the cloud object storage server will not be reflected immediately on the cache fileset. This option shall not be used with the afmObjectXattr option.
afmObjectSSL	Specifies SSL certificate verification. This option is valid only with HTTPS. Default value of this parameter is disabled.
afmObjectGCS	This parameter is used for Google cloud object storage backend. Enabling this parameter makes AFM to cloud object storage compatible with the list operations used by the Google cloud object.
For more information about these parameters, see the <i>mmchfileset</i> command in the <i>IBM Storage Scale: Command and Programming Reference Guide</i> .	

Microsoft Azure Blob support for the AFM to cloud object storage

AFM to cloud object storage feature enables placement of files or objects in an IBM Storage Scale cluster to a cloud object storage, which supports Microsoft Azure Blob storage from IBM Storage Scale 5.1.9. AFM to cloud object storage allows associating an IBM Storage Scale fileset directly with Microsoft Azure Blob containers that are created under the Microsoft Azure Blob storage accounts.

Application interface uses an AFM fileset of a preferable mode that connects to a Microsoft Azure Blob storage container. In the background, data is seamlessly transferred through AFM between these filesets and Microsoft Azure Blob containers, which ensures high-performance and robust support for object applications. Object applications can extend their reach across both AFM-connected cloud object storage filesets and cloud object storage itself. Additionally, both the fileset and cloud object storage offer dependable backup solutions for critical data.

To enable this feature, see the **Configuring an AFM to cloud object storage fileset with Microsoft Azure Blob** section in the *IBM Storage Scale: Administration Guide*.

Asynchronous auto prefetch

AFM can periodically fetch changed data asynchronously from a cloud object storage bucket to an AFM to cloud object storage fileset to synchronize the cache automatically. During the set prefetch interval, the created, modified, or deleted objects are synchronized to the cache. This synchronization helps AI applications with quick insights, reduced costs, improved performance, enhanced security, and streamlined operations.

Watch folder for the asynchronous auto prefetch

AFM supports all the clustered watch folder events to monitor filesets for file accesses events. With the clustered watch folder, you can monitor file operations across clusters by using a centralized tool that has scalability and resiliency built-in. The clustered watch folder captures file system activities, such as removal, deletion, or modification, at the IBM Storage Scale AFM fileset level and generates and streams notifications to topics in an external sink that you can manage. For more information, refer to [“Clustered watch folder”](#) on page 256.

You can enable an AFM to cloud object storage fileset for watch folder events to generate all the events fetched from the cloud object storage bucket to the cache. Watch folder events are generated on the cache when AFM downloads changed metadata and data to the cache.

For every created object in the cloud object storage target bucket, AFM fileset fetches and creates a file or directory with the same name on the cache. Similarly, after any object modification or object deletion, AFM synchronizes automatically to the cache. If a file is deleted from a cloud object storage bucket, then a fileset in the AFM RO, IW mode moves the deleted file in to the `.ptrash` directory under the AFM fileset.

Parameters for the asynchronous auto prefetch

The asynchronous auto prefetch feature enables automatic and periodic synchronization of created, modified, or deleted objects from a cloud object storage to an AFM to cloud object storage fileset cache. This feature is only available for filesets in independent writer (IW) and read-only (RO) modes. Use the following parameters to enable the asynchronous auto prefetch:

- Use the **--async-prefetch-interval** parameter to set the interval. Valid values are in the range 1-2147483647. By default, this option is disabled. You can enable this option and set a 30-minute interval.
- Use the **--fast-readdir2** parameter to enhance object synchronization, particularly for deleted objects. However, this parameter does not fetch or synchronize ACLs and xattrs.
- Use **--async-prefetch-interval** and **--fast-readdir2** parameters during the fileset creation for new filesets.
- Use the following commands for the existing filesets:

```
mmchfileset afmAsyncPrefetchInterval=<value>
```


```
mmchfileset afmObjectFastReaddir2=yes
```

Note: Both parameters must be enabled together for the asynchronous auto prefetch feature to function.

To configure the asynchronous auto prefetch, see the *Configuring the asynchronous auto prefetch* in the *IBM Storage Scale: Administration Guide*.

Asynchronous notifications and Amazon SQS integration

The asynchronous notifications operation notifies AFM about any changes in a cloud object storage. For example, when objects are added, modified, or deleted in a specific bucket.



By using asynchronous notifications, AFM can periodically fetch changed data from a cloud object storage bucket to an AFM to cloud object storage fileset to synchronize all changes with the AFM cloud object storage mode filesets.

This asynchronous notifications operation is different from the asynchronous auto prefetch operation. For more information about Asynchronous auto prefetch, see “Asynchronous auto prefetch” on page 181. By using the asynchronous notifications operation, AFM detects object changes in a bucket and pulls only updated objects without polling. AFM identifies changes by querying to the Amazon S3 Simple Queue Service (SQS), which is configured on a bucket or an account.

Important:

- For asynchronous notifications the **AmazonSQSFullAccess** permission must be granted to an account.
- The asynchronous notifications operation is available only for Amazon S3 storage.
- AFM downloads data and metadata of changed objects on an SQS-enabled fileset.

After the following configuration, AFM internally queries SQS to create a list of changed objects and synchronizes those updates to the AFM RO mode fileset:

- Creation of an AFM to cloud object storage mode fileset by using the **mmchconfig --sqs** command.
- The SQS configuration on existing read-only (RO) or independent-writer (IW) mode filesets by using the **mmchconfig afmObjectSQS=yes**.
- Creation of an RO mode fileset by using the **mmafmcconfig --sqs**, or update an RO mode fileset by using **mmchfileset fs1 ro1 -p afmObjectSQS=yes**.

For more information about configuration, see the *Configuring SQS-based asynchronous notifications* section in the *IBM Storage Scale: Administration Guide*.

You can also set up an IW or SW mode AFM fileset to write local changes to the cloud object bucket and an RO mode AFM fileset where SQS is enabled. These filesets are connected to the same cloud object bucket. The IW mode synchronizes local updates to the bucket and the RO mode automatically fetches all objects changes through SQS. AFM creates a list of changed data and fetch it to the cache.

When all updates are synchronized to the AFM RO mode fileset, watch folder events are triggered to monitor clusters. For more information, see the *Configuring asynchronous notifications with Amazon SQS for watch folder events* section in the *IBM Storage Scale: Administration Guide*.

Glacier storage class support for the AFM to cloud object storage

Many S3 providers offer different S3 storage classes that you can choose based on performance, data access, resiliency, and cost requirements of workloads. S3 storage classes help to classify the nature of a workload, cost, and different access pattern.

For more information about AWS S3 Glacier, refer to [Amazon S3 Storage Classes](https://aws.amazon.com/s3/storage-classes) (<https://aws.amazon.com/s3/storage-classes>).

For more information about IBM Storage Deep Archive on Diamondback, refer to <https://www.ibm.com/products/blog/the-s3-glacier-storage-classes-bring-the-benefits-of-tape-into-the-mainstream>.

AFM to cloud object storage mode filesets support S3 standard and S3 Glacier storage classes as a target. A storage class must be specified as a target when you create an AFM to cloud object storage mode fileset. If storage class is not specified, AFM uses the default S3 standard. All AFM to cloud object storage fileset modes supports Glacier Object Storage Server as a target.

You must specify the Glacier storage class when you create an AFM to cloud object storage mode fileset:

```
mmafmcconfig --storage-class GLACIER
```

For a Glacier target, AFM is tested with IBM Storage Deep Archive tape server as a Glacier backend that is supported for AFM to cloud object storage mode filesets.

While you use a Glacier storage class as a target in AFM to cloud object storage mode fileset, data downloading takes time. Because data fetching from a Glacier server (tape) takes time to retrieve the data. Hence, re-try download after the data is available from the Glacier storage.

Note: The Glacier storage class does not support the directory object feature. Therefore, when you create an AFM to cloud object storage mode fileset, create the fileset by using `--storage-class GLACIER`. That is, `--directory-object` parameter must not be used with the **mmafmcconfig** command. This parameter is not supported for the Glacier backend.

AFM to cloud object storage limitations

The AFM to cloud object storage limitations are as follows:

- Hard links are supported only in the Local-Update (LU) mode. The creation of hard links fails with permission denied or E_PERM error on other modes.
- The primary resources of object servers are buckets and objects.
- AFM can synchronize empty directories to the cloud object server only if an AFM to cloud object fileset is created by using the **mmafmcconfig --directory-object** command.
- The ChangeTimes operation is supported on the cache filesets. However, this operation does not replicate to the target cloud object server.
- To replicate the chmod and chown metadata operations to a cloud object storage, you need to use the **mmafmcconfig** command with the `--xattr` option.
- Renaming a non-empty directory is only supported for filesets in LU mode. When the fileset is not in LU mode and you attempt to rename a non-empty directory, the operation fails with the E_NOTEMPTY error. Empty directories and local directories can be renamed.
- Parallel data transfer for write operations is not supported on an AFM to cloud object storage enabled fileset. Objects are not synchronized by splitting chunks across gateway nodes but objects are queued on different gateway nodes based on the mapping.
- Deletion of objects on the cloud object storage synchronizes the deletion of files on the cache. But empty directories might remain on the cache.
- Support of file and directory naming convention is in accordance with the cloud object storage server supported guidelines. Some cloud object storage servers do not support special characters. The file or directory name must not contain special characters.
- When **Fastreaddir** and **FastReaddir2**, **symlink** are enabled, AFM to cloud object storage mode files are pulled as regular files from the cloud object storage target bucket.
- File paths that have more characters than maximum characters limitation are not supported.
- You can run AFM to cloud object storage commands such as **mmafmckeys**, **mmafmcconfig**, **mmafmcctl**, and **mmafmcaccess** only from a Linux node.
- **Immutability** and **appendOnly** features are not supported on AFM to cloud object storage filesets.
- The `--iam-mode` option is not supported on AFM to cloud object storage filesets.
- Reading AFM uncached files from the snapshot results in the error 5 (EIO). To avoid the error, set the **afmSnapUncachedRead** option on the fileset. This configuration returns zeros for the evicted blocks instead of error 5.
- A dependent fileset linking is not supported in an AFM to cloud object storage fileset.
- In some cases, an AFM cache has pending changes such as delete and rename to replicate, and the fileset recovery is triggered. In this case, the cloud object storage might have old data and renamed files or directories are replicated. Because of the files or directories replication, some extra files might exist on the cloud object storage. The AFM cache and the cloud object storage have the latest copy of the files.
- AFM to cloud object storage does not maintain sparseness of files when uploading the files to cloud object storage.
- Based on the user defined policy, **reconcile** command does not upload a directory with objects that do not meet the policy criteria to cloud objects storage.

- Renaming of an uncached file is not supported on an AFM to cloud object storage MU-mode fileset. To rename a file, you need to download this file data first and then rename it.
- While creating samba exports of an AFM to cloud object storage enabled fileset, Samba Oplocks needs to be disabled.
- AFM to cloud object storage does not delete empty directories on the cloud bucket.
- Adding a new node to the cluster might not synchronize the AFM to cloud object storage keys on the node. To synchronize the AFM to cloud object storage keys on the newly added node, you can run the following command on the new node:

```
/usr/lpp/mmfs/bin/mmccrnv _afmcos refresh
```

- AFM does not support storing an object list file in the AFM fileset that is used for uploading or downloading. To run upload or download by using the object list file, you can create an object list file outside the AFM fileset linked-path.
- If a file is evicted from an AFM to cloud object storage fileset, you cannot rename the evicted file from the gateway node. You can rename the evicted file from a non-gateway node.
- Failover from an AFM fileset that uses an NFS or NSD backend to an AFM to cloud object storage mode fileset is not supported. Disable the AFM NFS or NSD backend fileset before conversion to AFM to cloud object storage MU mode. Failover from an AFM to cloud object storage mode fileset to NFS or NSD backend targets is also not supported.

ObjectOnly mode AFM to cloud object storage fileset support:

Upload or write and download or read operations are replicated to a target object server. The operations such as delete and rename are not synchronized between a cache and an object server.

- If data on a cache is deleted or renamed, the data is not synchronized to an object server.
- If an object on an object storage server is deleted or renamed, and the object is already replicated to a cache, the object is not synchronized to the cache.

The limitations of the replication at the file system level by using the manual updates mode

- Automatic eviction is not supported. Manual eviction can be run by using the **mmafmcostl** command.
- After the file system level replication is AFM disabled (**mmchfileset fsname root -p afmtarget=disable**), AFM cannot be enabled on it. Refer the [Disabling AFM](#) section for more details.
- After the replication is enabled at the file system level, AFM filesets must not be created and linked in this file system.
- The replication process does not migrate any file system-specific parameters from an old system to a new system. For example,
 - Quotas
 - Snapshots
 - File system-level tuning parameters
 - Policies
 - Fileset definitions
 - Encryption keys
- AFM filesets that are linked in a file system must be disabled before you convert the file system to AFM to cloud object storage replication.
- For an AFM to cloud object storage fileset in the MU mode, when directories are deleted the inodes are not reclaimed automatically. You can run the **mmafmcostl delete --from-target** or **reconcile** command to delete these inodes and sync the data.

Introduction to IBM Storage Scale on AWS

IBM Storage Scale addresses the needs of the applications for which performance or performance-to-capacity ratio demands cannot be met by traditional scale-up storage systems.

Note: In IBM Storage Scale, the term Node is typically used to refer to any running instance of an operating system. The nodes deployed in this Amazon Machine Images (AMIs) are all Amazon Elastic Compute Cloud (Amazon EC2) instances, so the term instance is used in the place of node.

AWS Services

The following section gives a brief introduction to the various AWS services that are available.

AWS CloudFormation

AWS CloudFormation is used to create and manage a collection of related AWS resources, and provision and update them in an orderly and predictable way. You use a template to describe all the AWS resources that are needed, while the AWS CloudFormation creates and configures the resources, and figures out dependencies. For more information, see [AWS CloudFormation Documentation](#).

Amazon VPC

The Amazon VPC service is used to provision a private, isolated section of the AWS Cloud, where you can start AWS services and other resources in a virtual network that you define. You have complete control over your virtual networking environment, including selection of your own IP address range, subnet creation, and configuration of route tables and network gateways. For more information, see [Amazon Virtual Private Cloud Documentation](#).

Amazon EC2

The Amazon EC2 service is used to start virtual machine instances with various operating systems. You can choose from the existing AMIs or import your own virtual machine images. For more information, see [Amazon Elastic Compute Cloud Documentation](#).

Amazon Auto Scaling

The Amazon Auto Scaling service helps maintain high availability and manage capacity by automatically increasing or decreasing the EC2 instance fleet. You can use auto scaling to run your fleet at optimal utilization by increasing instance capacity during demand spikes and decreasing capacity during downtimes. For more information, see [AWS Auto Scaling Documentation](#).

Amazon EBS

Amazon Elastic Block Store (Amazon EBS) provides persistent block storage volumes for use with Amazon EC2 instances in the AWS Cloud. For more information, see [Amazon Elastic Block Store Documentation](#).

Amazon S3

Amazon Simple Storage Service (Amazon S3) is storage for Internet. You can use Amazon S3 to store and retrieve any amount of data at any time, from anywhere on the web. You can accomplish these tasks using the simple and intuitive web interface of the AWS Management Console. For more information, see [Amazon Simple Storage Service Documentation](#).

Amazon CloudWatch

Amazon CloudWatch is a monitoring service for AWS Cloud resources and the applications you run on AWS. You can use CloudWatch to collect and track metrics, collect and monitor log files, set alarms, and automatically react to changes in your AWS resources. For more information, see [Amazon CloudWatch Documentation](#).

Amazon IAM

AWS Identity and Access Management (IAM) enables you to securely control access to the AWS services and resources for your users. With IAM, you can manage users, security credentials such as access keys, and permissions that control which AWS resources users can access, from a central location. For more information, see [AWS Identity and Access Management Documentation](#).

AWS Lambda function

AWS Lambda is a stateless compute service that allows you run code without provisioning or managing servers. For more information, see [AWS Lambda documentation](#).

AWS Systems Manager

AWS Systems Manager is a service that is used to view and control AWS infrastructure. The Systems Manager enables the viewing of operational data from other AWS services, and the automation of operational tasks. In addition, the Systems Manager can also be used to store and retrieve configuration data. For more information, see [AWS Systems Manager](#).

Regions and Availability Zones

IBM Storage Scale deployed on the AWS cloud uses regions to place or replicate resources and data in multiple Availability Zones. If the `multiple availability zone` template is selected, this architecture deploys the IBM Storage Scale cluster nodes across three Availability Zones within an AWS region.

Note: A flow represents a single, point-to-point network connection. For more information, see [The Floodgates Are Open – Increased Network Bandwidth for EC2 Instances](#).

The supported regions are included in the following list:

- us-east-1
- us-east-2
- us-west-1
- us-west-2
- ca-central-1
- eu-central-1
- eu-west-1
- eu-west-2
- eu-west-3
- ap-northeast-1
- ap-northeast-2
- ap-southeast-1
- ap-southeast-2
- ap-south-1
- sa-east-1

When you deploy the `multi availability zone` template configuration for an IBM Storage Scale cluster with replication, each element of the data and the metadata is replicated in a separate availability zone. The replication is done to avoid the loss of data when hardware failures occur in one of the Availability Zones. By default, instances are distributed evenly between the Availability Zones. Therefore, each availability zone has one private subnet. However, these zones remain a single or logical cluster.

- If you choose to deploy a `single availability zone` configuration, all the server and compute instances are deployed into a single availability zone, and IBM Storage Scale creates only one copy of each data element.
- If you choose to deploy a `multi availability zone` configuration, all the server and compute instances are deployed into multiple Availability Zones, and IBM Storage Scale creates two copies of each data element.

IBM Storage Scale instance types and operating systems

The IBM Storage Scale deployment on AWS supports a large selection of EC2 instance types for the IBM Storage Scale cluster nodes. The instance type must be chosen after a detailed evaluation of the IBM Storage Scale performance and resource requirements, and also the resource requirements for any application workloads that might be running on these nodes. For more information, see [Amazon EC2 Instance Types](#).

IBM Storage Scale on AWS BYOL 1.5.0 offers support for the following Operating Systems versions:

- Red Hat® Enterprise Linux (RHEL) versions 9.4

Customers must specify an appropriate AMI running a supported operating system. Red Hat Enterprise Linux AMIs are available from the AWS Marketplace and the AWS Community. Additionally, customers can choose to create their own custom AMI. For more information, see *Creating custom AMI* in *IBM Storage Scale: Administration Guide*.

Important: The customer is responsible to ensure that they have the necessary support agreements in place for the Red Hat Enterprise Linux Operating Systems AMIs they specify to facilitate resolution with Red Hat support for any issues.

IBM Storage Scale usage restrictions

This version does not support the following features of IBM Storage Scale:

- Protocol support, including the use of Cluster Export Services (CES) nodes and protocol access such as Network File System (NFS), Object, and Server Message Block (SMB)
- Transparent Cloud Tiering (TCT)
- Compression
- Encryption
- Data Management API (DMAPI) support, including Hierarchical Storage Management (HSM) to tape
- Hadoop Distributed File System (HDFS) connector support
- Call home
- Multi-cluster support: Exporting an IBM Storage Scale file system from one IBM Storage Scale cluster to another IBM Storage Scale
- User name space management and quota management
- Snapshots and clones

Additional limitations include the following:

- Using EBS volume encryption for IBM Storage Scale file systems is not supported.
- The archiving and restoring of IBM Storage Scale data through the use of AWS services is not supported.

Introduction to system health and troubleshooting

IBM Storage Scale comes with several functions to monitor and maintain the health of a system.

System health

In IBM Storage Scale, system health monitoring is performed by the `Sysmonitor` daemon. The `Sysmonitor` daemon monitors all critical aspects of the entire cluster, where IBM Storage Scale is used, to ensure that potential issues are detected as soon as possible. `Sysmonitor` daemon does hundreds of checks on the relevant cluster nodes and raises RAS events. Based on these event checks, it informs the user whether anything is not working as expected and also provides guidance to solve existing problems.

You can configure IBM Storage Scale to raise events when certain thresholds are reached. As soon as one of the metric values exceeds or drops beyond a threshold limit, the `Sysmonitor` daemon receives an event notification from the monitoring process. The `Sysmonitor` daemon then generates a log event and updates the health status of the corresponding component. For more information about event type and health status, see *Event type and monitoring status for system health* in the *IBM Storage Scale: Problem Determination Guide*.

Every component has certain events that are defined for it. The `mmhealth node eventlog` command gives an overview of the happenings across all components on the local node that is sorted by time. For more information, see *System health monitoring use cases* in the *IBM Storage Scale: Problem Determination Guide*. The user can also create and raise custom health events in IBM Storage Scale.

For more information, see *Creating, raising, and finding custom defined events* in the *IBM Storage Scale: Problem Determination Guide*.

The **mmhealth node show** command displays the results of the health monitoring of the node and its services, which run in the background. The role of a node in monitoring determines the components that need to be monitored. For many IBM Storage Scale components, separate categories exist in the output of the **mmhealth node show** command. For example, the typical components that are presented on a node are GPFS, network, file system, disk, and Perfmon. For a complete list of supported components, see *Monitoring the health of a node* in the *IBM Storage Scale: Problem Determination Guide*.

For these services, the **mmhealth node show <service>** command displays the results of the health monitoring, aggregated health state of a service, and recent active events for this service. This view also can be used to get more details about the particular health states of a service subcomponent by using the **--verbose** option. For more information about node role and functions, see *Monitoring the health of a node* in the *IBM Storage Scale: Problem Determination Guide*.

You can also use the **mmhealth cluster show** command to see an overview of the health monitoring for the complete cluster.

You can use the **mmhealth** command to do the following tasks:

- View the health of a node or cluster.
- View current events, and get tips for a better system configuration.
- View details of any raised event.
- Browse event history.
- Manage performance thresholds.
- Configure monitoring intervals.

For more information, see *mmhealth command* in the *IBM Storage Scale: Command and Programming Reference Guide*.

Protocol monitoring

You can monitor system health, query events, and perform maintenance and troubleshooting tasks that are related to CES by using the **mmces** command. If a CES node is unable to export information by using the configured protocols, then its CES IP address is reassigned to another node. The reassignment is done so that when a single node goes down, the availability is not impacted. For more information, see *CES configuration issues* in the *IBM Storage Scale: Problem Determination Guide*.

You can use the **mmces** command to manage protocol addresses, services, node state, logging level, and balancing the load. For more information about **mmces** command, see *mmces command* in the *IBM Storage Scale: Command and Programming Reference Guide*.

You can use the **mmprotocoltrace** command to collect trace information for debugging system problems or performance issues. For more information, see *mmprotocoltrace command* in *IBM Storage Scale: Command and Programming Reference Guide*.

Troubleshooting

Despite all the functions that are meant to maintain your system's health, you might still face some issues with your storage clusters. To start the troubleshooting process, collect details of the issues reported in the system.

IBM Storage Scale provides the following options for collecting details:

- Logs
- Dumps
- Traces
- Diagnostic data collection through CLI.

Note: For more information, see *CLI commands for collecting issue details* in the *IBM Storage Scale: Problem Determination Guide*.

- Diagnostic data collection through GUI

For more information, see *Troubleshooting* in the *IBM Storage Scale: Problem Determination Guide*.

To diagnose the cause of an issue, it might be necessary to gather some extra information from the cluster. This information can then be used to determine the root cause of an issue. Collection of debugging information, such as configuration files and logs can be gathered by using the **gpfs.snap** command. This command gathers data about GPFS, operating system information, and information for each of the protocols. For more information, see *gpfs.snap command* in the *IBM Storage Scale: Command and Programming Reference Guide*.

For more high-level analysis of an issue, you can also use the tracing feature. Tracing is logging at a high level. For more information, see *Collecting details of issues by using logs, dumps, and traces* in the *IBM Storage Scale: Problem Determination Guide*.

Introduction to performance monitoring

IBM Storage Scale provides functionality to monitor and maintain the performance of a cluster and its nodes.

IBM Storage Scale comes with a scalable performance monitoring solution built out of sensors and proxies, which reads the individual data, and a collector database that stores it for later use. Sensors, collectors, and proxies, along with performance metrics form the basic components of the Performance Monitoring tool. The performance metrics are counters that monitor the working of a system on which IBM Storage Scale is deployed. The tool can be configured to have several sensors that collect performance data from the nodes in a cluster, and scales up to many nodes. In IBM Storage Scale, there are more than 50 performance sensors that track more than a thousand performance metrics, and collect capacity and usage information. For more information, see *Using the performance monitoring tool* in the *IBM Storage Scale: Problem Determination Guide*.

The sensors first receive the data for one point in time and then parse the data into a format that is understood by the collector. The sensors then send the data directly to the collector. Queries are used by the customer or other applications to see and further use the time series data. A single collector can easily support up to 150 sensor nodes. The Performance Monitoring tool can be configured with multiple collectors to increase scalability and fault-tolerance. This latter configuration is referred to as federation. In a multi-collector federated configuration, collectors must be aware of their peers to work properly. All collectors that are part of the federation must be specified in the peer configuration option in the collector's configuration file.

You can use the **mmperfmon** command to configure the Performance Monitoring tool and its components to set up metrics, run queries, and compare node metrics. For more information, see *mmperfmon command* in the *IBM Storage Scale: Command and Programming Reference Guide*.

The Performance Monitoring tool monitors the system on a per-component basis. A list of all supported components and metrics can be found in the *List of performance metrics* in the *IBM Storage Scale: Problem Determination Guide*.

A file system or disk can have several items of this kind that are called *entities* in the system. The monitoring of such a component is broken down to the monitoring of the individual file systems or disks.

All the information that is collected by the Performance Monitoring tool can be viewed by using one of the following options:

IBM Storage Scale GUI

A fully customizable dashboard that provides hard-wired performance charts in detail views.

mmperfmon query

A CLI command to query performance data.

Grafana bridge

An open source monitoring dashboard.

REST API to query performance data

The APIs to query and chart the performance data by user-defined dashboard software.

All the performance metrics data that is collected by the Performance Monitoring tool can be monitored by using threshold monitoring. Threshold monitoring is a service that helps the user to identify performance issues by defining threshold rules for selected performance metrics. An IBM Storage Scale user can set user-defined thresholds on any performance metric. For more information, see *Threshold Monitoring for system health* in the *IBM Storage Scale: Problem Determination Guide*.

While troubleshooting, you can also find more detailed information about the performance monitoring by viewing the perfmon tool logs. The performance monitoring tool logs can be found in the `/var/log/zimon` directory on each node that is configured for performance monitoring. For more information, see *Performance monitoring tools logs* in the *IBM Storage Scale: Problem Determination Guide*.

Data protection and disaster recovery in IBM Storage Scale

The IBM Storage Scale installation must be protected against data loss to ensure continuity of operation after a malfunction.

Data loss can be prevented by protecting the four types of key data:

- Cluster configuration data
- File system configuration data
- File system contents (user data, metadata, configuration)
- Protocol configuration data

Cluster configuration data is the administrative data that associates the nodes, addresses, networks, and software installation on each node. The system administrators must save the following configuration information:

- The output of the **mmlscluster** command to ensure that reconstruction of this data is possible if needed.
- CCR backup file or the `mmsdrifs` file depending on type of repository. For more information, see the *mmsdrrestore* command in *IBM Storage Scale: Command and Programming Reference Guide*.
- Snapshots.

File system configuration data consists of a wide variety of information involving all the file systems in the cluster. To protect this data, it is essential to use the **mmbbackupconfig** command for each file system and save the output files for future use. This configuration data details which disks are associated as NSD components of which file systems, how much storage is in use, the filesets defined, the quotas defined, and other valuable configuration data that describes the file system structure. These do not include the file data in the user files. User file data is the lowest level of information and most frequently changing contents that need protection.

IBM Storage Scale has built in data protection processes that allows the users not only to back up and restore valuable data, but also recover data that could be potentially lost or corrupted by their own actions. For example, unintentional deletion or overwriting of a user's file.

Data backup options in IBM Storage Scale

This section describes the various options available in IBM Storage Scale to back up data.

Backing up the data

IBM Storage Scale creates a second copy of the data within the system as backup. These secondary copies are usually point-in-time copy of the original data, and can be created on a daily, monthly, or weekly basis. In a situation where the actual data is corrupted, the system uses the backed-up copy to restore data to a previous point in time.

You can use the **mmbackup** command to back up the user data from a GPFS file system or independent fileset to an IBM Storage Protect server or servers. The **mmbackup** command can be used only to back up file systems owned by the local cluster. For more information on how to back up your data, see *Backup your data* in the *IBM Storage Scale: Problem Determination Guide* and *mmbackup command* in the *IBM Storage Scale: Command and Programming Reference Guide*.

Fileset and backup

IBM Storage Scale supports backup from independent filesets in addition to backup of the whole file system. For more information on fileset backup, see *Filesets and backup* in the *IBM Storage Scale: Administration Guide*. If you are planning to use IBM Storage Protect to back up IBM Storage Scale file systems, see [“Backup considerations for using IBM Storage Protect” on page 362](#)

Data restore options in IBM Storage Scale

This section describes the various options available in IBM Storage Scale to restore data.

Restoring data

The system can restore the data if it becomes corrupted due to user errors, hardware failure, and loss or even if the data gets corrupted due to bugs in other application software. For more information on how to restore system data, see *Restoring data and system configuration* in the *IBM Storage Scale: Problem Determination Guide*.

Data mirroring in IBM Storage Scale

In IBM Storage Scale you can copy data from one location to a secondary storage location, thus creating a mirror image of the original. This replication is called data mirroring.

In data mirroring the data is copied in real time, so the information stored in the secondary copy is always an exact replica of the data in the primary copy. Data mirroring is useful in the speedy recovery of critical data after a disaster. Data mirroring can be implemented locally or offsite at a different location. For more information on data mirroring and replication, see *Data mirroring and replication* in the *IBM Storage Scale: Administration Guide*.

Protecting file data using snapshots

A snapshot of an entire file system or of an independent fileset can be created to preserve the contents of the file system or the independent fileset at a single point in time.

A fileset snapshot is a snapshot of the entire inode space. Any snapshot of an independent fileset also includes any dependent filesets contained within that independent fileset. You cannot create a snapshot of a dependent fileset. For more information on creating and maintaining snapshots, see *Creating and maintaining snapshots of file systems* in the *IBM Storage Scale: Administration Guide*.

Snapshots can be used in environments where multiple recovery points are necessary. Ensure that the file system is mounted before restoring the snapshot. For information about mounting a file system, see *Mounting a file system* in *IBM Storage Scale: Administration Guide*. For more information on restoring data using snapshots, see *Restoring a file system from a snapshot* in the *IBM Storage Scale: Administration Guide*.

Note:

For object protocol, the best method to ensure consistency in the event that data has to be restored from a snapshot, is to configure the `cesSharedRoot` directory to be in the same file system as the object filesets.

Introduction to Scale Out Backup and Restore (SOBAR)

Scale Out Backup and Restore (SOBAR) is a data protection mechanism used specifically for disaster recovery situations.

SOBAR is used to back up and restore GPFS files that are being managed by IBM Storage Protect for Space Management. For more information on SOBAR, see *Scale Out Backup and Restore (SOBAR)* in the *IBM Storage Scale: Administration Guide*.

Commands for data protection and recovery in IBM Storage Scale

The following section lists the commands used for data protection and recovery in IBM Storage Scale.

mmbackup

Used to back up the user data from a GPFS file system or independent fileset to an IBM Storage Protect server or servers. For more information on **mmbackup**, see *mmbackup command* in the *IBM Storage Scale: Command and Programming Reference Guide*.

mmbackupconfig

Used to collect basic file system configuration information. The configuration information backed up by this command includes block size, replication factors, storage pool layout, filesets and junction points, quota information, and a number of other file system attributes. This command forms a part of the SOBAR utilities along with the **mmrestoreconfig**, **mmimgbackup** and **mmimgrestore** commands. For more information on **mmbackupconfig**, see *mmbackupconfig command* in the *IBM Storage Scale: Command and Programming Reference Guide*.

mmfsck

Used to detect and repair conditions that can cause problems in a file system. The **mmfsck** command operates in two modes: online or offline. The online mode can run while the file system is mounted. It detects and repairs lost blocks (blocks that are allocated but do not belong to any file) and corruptions in the block allocation map. The offline mode can run only when the file system is unmounted, but it can detect and fix problems that the online mode cannot. In general, you do not need to run the command in offline mode unless you are directed to by the IBM Support Center. For more information see the topic *mmfsck command* in the *IBM Storage Scale: Command and Programming Reference Guide*.

mmimgbackup

Used to perform a backup of a single GPFS file system metadata image. This command forms a part of the SOBAR utilities along with the **mmbackupconfig**, **mmrestoreconfig** and **mmimgrestore** commands. Make sure that you run the **mmbackupconfig** command before you run the **mmimgbackup** command. For more information on **mmimgbackup**, see *mmimgbackup command* in the *IBM Storage Scale: Command and Programming Reference Guide*.

mmimgrestore

Used to restore a single GPFS file system metadata image. This command forms a part of the SOBAR utilities along with the **mmbackupconfig**, **mmrestoreconfig** and **mmimgbackup** commands. The **mmrestoreconfig** command must be run before running the **mmimgrestore** command. For more information on **mmimgrestore**, see *mmimgrestore command* in the *IBM Storage Scale: Command and Programming Reference Guide*.

mmrestoreconfig

Used to restore file system configuration information. The **mmrestoreconfig** command allows you to either query or restore, or both query and restore the output file generated by the **mmbackupconfig** command. This command forms a part of the SOBAR utilities along with the **mmbackupconfig**, **mmimgbackup** and **mmimgrestore** commands. For more information on **mmrestoreconfig**, see *mmrestoreconfig command* in the *IBM Storage Scale: Command and Programming Reference Guide*.

mmrestorefs

Used to restore user data and attribute files to a file system or an independent fileset by using those of the specified snapshot. For more information on **mmrestorefs**, see *mmrestorefs command* in the *IBM Storage Scale: Command and Programming Reference Guide*.

mmsdrrestore

Restores the latest GPFS system files on the specified nodes. When the CCR committed directories of all the quorum nodes have corrupted or missing files, the **--ccr-repair** parameter repairs the CCR committed directories of all the quorum nodes. For more information see *mmsdrrestore command* in the *IBM Storage Scale: Command and Programming Reference Guide*.

IBM Storage Scale GUI

You can configure and manage various features that are available with the IBM Storage Scale system by using the IBM Storage Scale management GUI.

You can complete the following important tasks through the IBM Storage Scale management GUI:

- Monitoring the performance of the system based on various aspects.
- Monitoring system health.
- Creating and managing file systems.
- Creating filesets and snapshots.
- Managing NFS and SMB data exports.
- Defining default, user, group, and fileset quotas.
- Creating and manage node classes.
- Monitoring the capacity details at various levels, such as file system, pools, filesets, users, and user groups.
- Monitoring and managing various services that are available in the system.
- Monitoring remote clusters.
- Configuring call home.
- Configuring and monitoring thresholds.
- Creating and managing user-defined node classes.
- Configuring user authentication for NFS and SMB users.
- Creating and managing S3 buckets and user accounts.
- Creating and managing AFM filesets
- Configuring fileset level multitenancy

The following table provides an overview of the features that are associated with each GUI page.

Table 12. Features associated with IBM Storage Scale GUI pages	
GUI Page	Function
Home	Provides an overall summary of the IBM Storage Scale system configuration, status of its components, and services that are hosted on it.
Monitoring > Dashboards	Provides a dashboard to display the predefined performance charts and the customized charts that are marked as favorite charts in the Monitoring > Statistics page.
Monitoring > Events	Monitor and troubleshoot the issues that are reported in the system.
Monitoring > Tips	Monitor events of type "tips". The tip events give recommendations to the user to avoid certain issues that might occur in the future.
Monitoring > Statistics	View performance and capacity data in customizable charts.

Table 12. Features associated with IBM Storage Scale GUI pages (continued)

GUI Page	Function
Monitoring > Thresholds	Configure various thresholds based on the performance monitoring metrics. You can also monitor the threshold rules and the events that are associated with each rule.
Monitoring > Command Audit Log	Logs the commands that are issued and tasks that are completed by the users and administrators. These logs can also be used to troubleshoot issues that are reported in the system.
Monitoring > Event Notifications	Configure event notifications to notify administrators when events occur in the system.
Multi-Tenancy	Configure fileset-level multitenancy. This enables multiple tenants to share a single filesystem while ensuring logical isolation through dedicated filesets.
Nodes	Monitor the performance and health status of the nodes that are configured in the cluster. You can also create and manage user-defined node classes from the Nodes page.
Cluster > Network	Monitor the performance and health status of various types of networks and network adapters.
Cluster > Remote Connections	Request access to the GUI node of the remote cluster to establish communication between the local and remote clusters. You can also use the Remote Connections page to grant access to the remote clusters when an access request is received.
Files > File Systems	Create, view, and manage file systems.
Files > Filesets	Create, view, and manage filesets.
Files > Snapshots	Create snapshots or snapshot rules to automate snapshot creation and retention.
Files > User Capacity	Monitor the capacity details of users and user groups.
Files > Quotas	Create and manage default, user, group, and fileset quotas at the file system level.
Files > Information Lifecycle	Create, manage, and delete policies that manage automated tiered data storage.
Files > Transparent Cloud Tiering	Provides both a summarized and attribute-wise details of the Transparent Cloud Tiering service, which is integrated with the IBM Storage Scale system. The Transparent Cloud Tiering page is displayed if the feature is enabled on the cluster.

Table 12. Features associated with IBM Storage Scale GUI pages (continued)

GUI Page	Function
Files > Active File Management	Helps to monitor the performance, health status, and configuration aspects of AFM, AFM DR, and gateway nodes. You can see the Active File Management page in the GUI that is part of the AFM cache cluster with at least one AFM-enabled fileset.
Files > File System ACL	Define ACL templates and apply ACLs on files and directories.
Files > Clustered Watch Folder	Provides the list of the clustered watch folders that are configured in the system. You can also view the details of each watch folder and the events that are raised against each record from the detailed view.
Files > Fileset tiering	Helps to run AFM policy to tier MU filesets for a specific file system.
Storage > Pools	Helps to monitor the performance, health status, and configuration aspects of all available pools in the IBM Storage Scale cluster.
Storage > NSDs	Helps to monitor the performance, health status, and configuration aspects of all the network shared disks (NSD) that are available in the IBM Storage Scale cluster.
Protocols > NFS Exports	Create and manage NFS exports. Protocols pages are displayed in the GUI only when the protocol feature is enabled on the cluster.
Protocols > SMB Shares	Create and manage SMB shares. Protocols pages are displayed in the GUI only when the protocol feature is enabled on the cluster.
Protocols > S3	Create and manage S3 buckets and user accounts.
Services	Monitor and manage various services that are hosted on the system.
Support > Diagnostic Data	Download diagnostic data to troubleshoot the issues that are reported in the system.
Support > Call Home	Configure call home feature that automatically notifies the IBM Support personnel about the issues that are reported in the system. You can also manually upload diagnostic data files and associate them with a PMR through the GUI.

Note:

- After installing the system and GUI package, you need to create the first GUI user to log in to the GUI. This user can create other GUI administrative users to perform system management and monitoring tasks. When you launch the GUI for the first time after the installation, the GUI welcome page provides options to create the first GUI user from the command line prompt by using the `/usr/lpp/mmfs/gui/cli/mkuser <user_name> -g SecurityAdmin` command.
- Creation of 50k filesets through REST API is supported.

Assistance for understanding the features associated with a GUI page

The following three levels of assistance are available for the GUI users:

Hover help

A brief description of a feature that is associated with a field. It appears when you hover the mouse over the tiny question mark that is placed next to the field label. Hover help is available only for the important and complex fields.

Context-sensitive help

A detailed explanation of the features that are associated with the page. The context-sensitive help files are available in the help menu, which is placed on the header bar of the GUI page.

IBM Docs

The third-level of information helps users to find entire details of the product. You can also access the IBM Storage Scale documentation from the help menu, which is placed on the header bar of the GUI page.

Related concepts

[“Manually installing IBM Storage Scale management GUI” on page 497](#)

The management GUI provides an easy way for the users to configure, manage, and monitor the IBM Storage Scale system.

[“Manually upgrading the IBM Storage Scale management GUI” on page 708](#)

You can upgrade the IBM Storage Scale management GUI to the latest version to get the latest features. You can upgrade one GUI node at a time without shutting down IBM Storage Scale on other nodes to ensure high availability.

IBM Storage Scale management API

With the IBM Storage Scale management API, you can develop scripts to automate labor-intensive cluster management tasks. These APIs provide a useful way to integrate and use the IBM Storage Scale system.

The IBM Storage Scale management API is a REST-style interface for managing IBM Storage Scale cluster resources. It runs on HTTPS and uses JSON syntax to frame data inside HTTP requests and responses. For more information, see [“Functional overview” on page 198](#).

The IBM Storage Scale management API implementation is based on GUI stack. A GUI server is managing and processing the API requests and commands. The following figure illustrates the architecture of the API.




Figure 21. IBM Storage Scale management API architecture



Warning: Replacing remote shell (SSH) and file copy (SCP) commands with IBM Storage Scale native REST API administration daemon–based scripts may affect the behavior of certain IBM Storage Scale management APIs.

The IBM Storage Scale management API has the following features:

- Reuses the GUI deployment's backend infrastructure, which makes the introduction of new API commands easier.
- Uses the same role-based access feature that is available to authenticate and authorize the GUI users. No additional configuration is necessary for the API users. However, access to individual REST APIs might be defined at the user group-level with the help of access control lists (ACLs). Users belonging to the specific user group can access the REST APIs based on the privileges that are defined within those ACLs. Every ACL is distinct for a user group. That is, user privileges for individual REST APIs vary for different user groups.
- Makes deployment easier as the GUI installation takes care of the basic deployment.
- Fixes scalability issues and introduces new features such as, filter and field parameters, and it supports paging.
- Highly scalable and can support large clusters with up to 1000 nodes.
- The APIs are driven by the same WebSphere® server and object cache that is used by the IBM Storage Scale GUI.

Installing and configuring IBM Storage Scale management API

For information about how to configure the IBM Storage Scale system to start using the APIs, see [Chapter 9, “Installing and configuring IBM Storage Scale management API,”](#) on page 633.

IBM Storage Scale management API command reference

For information about the API endpoints, see the topic *IBM Storage Scale management API commands* in the *IBM Storage Scale: Command and Programming Reference Guide*.

You can also access the documentation corresponding to each API endpoint from the GUI itself. The API documentation is available in the GUI at: <https://<IP address or host name of API server>:<port>/api/explorer/>. For example, <https://scalegui.ibm.com:443/api/explorer>.

The IBM Storage Scale management API features

The following topics describe the API features in detail.

Functional overview

The IBM Storage Scale management API is a REST-style API that provides interoperability between a client and server over a network. These APIs allow authenticated users to manage day-to-day storage tasks.

The following list provides the salient features of the REST-style APIs:

- Resource-based
- Stateless
- Client/server
- Cacheable
- Layered system

REST is a resource-based service system in which requests are made to the resource's Universal resources identifier (URI). The request invokes a response from the resource in the JSON format. An IBM Storage Scale management API *resource* is a collection of all the objects of the same type in a cluster, such as all the filesets. A *resource element* is an instance of a resource, such as a particular fileset named fileset1. In a REST API request, a resource or a resource followed by a resource element appears as the last term in the request, except for any trailing HTTPS parameters. The following example request ends

with a resource, `filesets`. It indicates that this operation affects all the filesets that are part of a file system named `gpfs0`:

```
https://198.51.100.1:443/scalemgmt/v2/filesystems/gpfs0/filesets
```

In contrast, the following example ends with a resource that is followed by a resource element `filesets/fileset1`. It indicates an operation that affects this particular fileset, such as creating or deleting the fileset:

```
https://198.51.100.1:443/scalemgmt/v2/filesystems/gpfs0/filesets/fileset1
```

The kind of operations that can be performed on the resources or a resource element are directed by the HTTP methods, such as GET, POST, PUT, DELETE. In some cases, the parameters in the HTTPS request also direct the operations. The following list provides the meanings of the basic HTTP methods that are used in the requests:

- GET - Reads a specific resource or a collection of resources and provides the details as the response.
- PUT - Updates a specific resource or a collection of resources.
- DELETE - Removes or deletes a specific resource.
- POST - Creates a resource.

API requests

API requests include a URL and in some cases HTTP parameters or JSON data.

URL for an API request

The URL for an API request has the following format:

```
https://IP address or host name:port/scalemgmt/v2/resource_type/[resource_element]
```

where

IP address or host name:port

Specifies the IP address or hostname and the port of the IBM Storage Scale management API server, which is IBM Storage Scale GUI server.

/scalemgmt/v2

Specifies the path of the API server. This path is not configurable; it must be `/scalemgmt/v2`.

resource_type

Specifies a resource type. The following list specifies the supported resource types.

```
ces/addresses
ces/services
cluster
config
filesets
filesystems
info
jobs
nfs/export
nodes
nodes/health
nodeclasses
quotas
smb/shares
snapshots
perfmon
```

```
/filesystems/{filesystemName}/filesets/{filesetName}/psnaps  
thresholds
```

Note: The full list of resource types can be accessed from the REST API explorer that is available at <https://<GUI server IP>:<port number>/api/explorer/>.

resource_element

Specifies the name of a particular resource, such as the name of a file system, a fileset, a node, or a snapshot. Optional.

As described earlier, an API request that ends with a resource type indicates a request that affects all the instances of that resource in the cluster. For example, the following HTTP request, sent with a GET method, returns information about all the nodes in the cluster:

```
'https://198.51.100.1:443/scalemgmt/v2/filesystems/nodes
```

In contrast, an API request that ends with a resource element that follows the resource type, indicates a request that affects only the specified resource element. For example, the following HTTP request, sent with a GET method, returns information about the specified node, node1.

```
'https://198.51.100.1:443/scalemgmt/v2/filesystems/nodes/node1
```

Data in API requests

Some types of API requests require one or more input parameters that follow the resource or resource element. For example, sent with a GET method, the following request returns information about all the filesets in the specified file system:

Table 13. Getting information about filesets	
Request method	GET
URL	<code>https://198.51.100.1:443/scalemgmt/v2/filesystems/gpfs0/ filesets/fs1</code>
JSON data	None.

Also, some API requests require input parameters to be included in the body of the HTTP request in JSON format. For example, the following request creates a new snapshot in a file system. The JSON data specifies the name of the file system and the name of the new fileset:

Table 14. Creating a snapshot	
Request method	POST
URL	<code>https://198.51.100.1:443/scalemgmt/v2/filesystems/gpfs0/ snapshots</code>
JSON data	<pre>{ "snapshotName": "snap1" }</pre>

Request headers

The following headers are needed for requests.

```
'content-type: application/json' 'accept: application/json'
```

Fields parameter

The field parameter can be used to specify the fields of the returned object by adding the query parameter "?fields" to a GET URL. For example, ?fields=field1,field2,parentField1.field1.

You can use certain keywords to get the expected results in the response. The following list provides the basic keywords that can be used with the fields parameter:

'none:'

Returns only the fields that are mandatory to uniquely identify an object. For example, filesystemName and filesetName for a fileset. Use this keyword for queries that return more than one object. It indicates that only the keys that are needed to retrieve the full object are returned.

'all:'

Returns all fields. Use this keyword in queries to a single object.

Apart from these keywords, you can also use the following options to customize the response by using the fields parameter:

- Use the field name "." to query fields that are wrapped in a parent object.

For example, use ?fields=config.iamMode to query for "iamMode" in the following object.

```
"config" : {  
  "iamMode" : "off"  
}
```

- Use parent name to query for all fields of a parent. For example, ?fields=config,afm includes all fields of the parent objects "config" and "afm"
- Multiple fields can be separated by using ",". For example, ?fields=config.iamMode,status.id

Examples

The following examples display the use of various keywords in the fields parameter.

1. curl -k -u admin:admin001 -X GET --header 'accept:application/json' 'https://198.51.100.1:443/scalemgmt/v2/filesystems/gpfs0/filesets?fields=:all:' . This GET call returns the following response:

```
{  
  "status": {  
    "code": "200",  
    "message": "..."  
  },  
  "paging": {  
    "next": "https://localhost:443/scalemgmt/v2/filesystems/gpfs0/filesets?lastId=1001"  
  },  
  "filesets": [  
    {  
      "config": {  
        "filesetName": "myFset1",  
        "filesystemName": "gpfs0",  
        "path": "/mnt/gpfs0/myFset1",  
        "inodeSpace": "0",  
        "maxNumInodes": "4096",  
        "permissionChangeMode": "chmodAndSetAcl",  
        "comment": "Comment1",  
        "iamMode": "off",  
        "oid": "158",  
        "id": "5",  
        "status": "Linked",  
        "parentId": "1",  
        "created": "2016-12-13 13.59.15",  
        "isInodeSpaceOwner": "1",  
        "inodeSpaceMask": "0",  
        "snapId": "0",  
        "rootInode": "131075"  
      },  
      "afm": {  
        "afmTarget": "nfs://10.0.100.11/gpfs/afmHomeFs/afmHomeFileset2",  
        "afmAsyncDelay": "0",  
        "afmDirLookupRefreshInterval": "60",  
        "afmDirOpenRefreshInterval": "60",  
      }  
    }  
  ]  
}
```

```

        "afmEnableAutoEviction": "false",
        "afmExpirationTimeout": "100",
        "afmFileLookupRefreshInterval": "30",
        "afmMode": "read-only",
        "afmNumFlushThreads": "4",
        "afmParallelReadChunkSize": "134217728",
        "afmParallelReadThreshold": "1024",
        "afmParallelWriteChunkSize": "0",
        "afmParallelWriteThreshold": "0",
        "afmPrefetchThreshold": "0",
        "afmPrimaryID": "string",
        "afmRPO": "0",
        "afmShowHomeSnapshots": "false"
    }
}
]
}

```

2. GET /filesystems/gpfs0/filesets - The default keyword : none : is used in the GET request and this call returns only the key fields as shown in the following example.

```

{
  "filesets": [
    {
      "config": {
        "filesetName": "(fs1)",
        "filesystemName": "gpfs0"
      },
      "links": {
        "self": "https://198.51.100.1:8191/scalemgmt/v2/filesystems/gpfs0/filesets/fs1"
      }
    },
    {
      "config": {
        "filesetName": "afmFset1",
        "filesystemName": "gpfs0"
      },
      "links": {
        "self": "https://198.51.100.1:8191/scalemgmt/v2/filesystems/gpfs0/filesets/afmFset1"
      }
    }
  ],
  "status": {
    "code": 0,
    "message": ""
  }
}

```

3. GET /filesystems/gpfs0/filesets?fields=state.rootInode,state.id - Uses the field name "." for querying fields that are wrapped in a parent object. This GET call return the following response.

```

{
  "filesets": [
    {
      "config": {
        "filesetName": "(fs1)",
        "filesystemName": "gpfs0"
      },
      "links": {
        "self": "https://198.51.100.1:8191/scalemgmt/v2/filesystems/gpfs0/filesets/fs1"
      },
      "state": {
        "id": 1,
        "rootInode": "latest",
      }
    },
    {
      "config": {
        "filesetName": "afmFset1",
        "filesystemName": "gpfs0"
      },
      "links": {
        "self": "https://198.51.100.1:8191/scalemgmt/v2/filesystems/gpfs0/filesets/afmFset1"
      },
      "state": {
        "id": 2,
        "rootInode": 50692
      }
    }
  ]
}

```

```

    ]
    "status": {
      "code": 0,
      "message": ""
    }
  }
}

```

Filter parameter

Use the filter parameter to filter the retrieved objects. For example, all filesets with ID less than 2 can be retrieved by using `?filter=status.id<2`. The fields that are used in the filter parameter are always added to the result set and therefore they do not have to be specified by using the **field** parameter. Filters can be concatenated by using `and`.

The following list provides the supported operators for different field types:

String

`"=", "!="` using a regular expression as the value is supported. For example, `?filter=config.filesetName='.*fs.*'`.

Numeric

`"=", "!="`, `"<`, `">`. For example, `?filter=status.id<2`.

Boolean

`"=", "!="` by using `"true"` or `"false"` as the value. For example, `?filter=config.isIndependent=false`.

The following examples show how to use filter parameter in a request:

```

curl -k -u admin:admin001 -X GET -H content-type:application/json
'https://198.51.100.1:443/scalemgmt/v2/filesystems/gpfs0/filesets?filter=status.rootInode>3'
{
  "filesets" : [ {
    "config" : {
      "filesetName" : "fs1",
      "filesystemName" : "gpfs0"
    },
    "status" : {
      "rootInode" : 10752
    }
  } ],
  "status" : {
    "code" : 0,
    "message" : "List of filesets for filesystem"
  }
}

```

Retrieve all file systems with encryption="yes":

```

curl -H content-type:application/json -K -u admin:admin001 -X GET
'https://198.51.100.1:443/scalemgmt/v2/filesystems?filter=encryption=yes'
(
  "filesystems" : [ (
    "encryption" : "yes"
    "filesystemName" : "objfs"
  ) ],
  "status" : {
    "code" : 200,
    "message" : "List of all file systems"
  }
}

```

API responses

API responses provide a response code from the IBM Storage Scale management API server and a return code and a non-mandatory return message from IBM Storage Scale.

Return codes and messages

A return code and an optional return message from IBM Storage Scale are stored in JSON format in a status structure at the end of any returned data. A return code of 0 indicates that IBM Storage Scale processed the request successfully:

```
"status": {  
  "message": "",  
  "code": 200  
},
```

A response code from the IBM Storage Scale management API server describes the result of the API request. The following table lists the response codes.

Table 15. Response codes. A four-column table that describes the response codes.			
Description	Code	Client retry	Return code
Success	20x	None	200 - Success (GET or PUT). 201 - New resource created (POST). 202 - Request accepted (long-running command). 204 - No content (DELETE).
Client error	4xx	No, the command continues to fail.	400 - Invalid request (format error in request data) 401 - Unauthorized request (wrong credentials). 403 - Resource not accessible. 404 - Resource not found (wrong URL). 405 - Method that is not supported for this resource.
Server error	5xx	Yes, in most cases	500 - Internal server error, retry (system problem) 503 - Service not available (the server is busy with other requests or is down)

Data in REST API responses

IBM Storage Scale management API responses always include some data in JSON format. The status structure with a return code and with a return message is always present. The return message is not mandatory and might not be displayed with the code. See the status structure at the end of the following JSON example. Also, GET methods return the requested information in JSON format. The following example is from a response to a request to GET information about CES addresses. In JSON, the parenthesis [] indicate an array. The following example contains an array that is named cesaddresses that contains two elements, each of which describes one CES address.

```
{  
  "cesaddresses": [  
    {  
      "cesNode": 1,  
      "attributes": "",  
      "cesAddress": "198.51.100.10",  
      "cesGroup": "",  
    },  
    {  
      "cesNode": 2,  
      "attributes": "",  
      "cesAddress": "198.51.100.14",  
      "cesGroup": "",  
    }  
  ]  
}
```



```

    }
  ]
  "status": {
    "message": "",
    "code": 200
  }
}

```

Response headers

The same header types appear in all responses from the REST API server. The content type must be application/json. The following set of response headers is typical:

```

{
  HTTP/1.1 200 OK
  X-Frame-Options: SAMEORIGIN
  Content-Type: application/json
  Content-Language: en-US
  Content-Length: 210
  Set-Cookie:
    LtPaToken2=yyYwMkAHcegg0e74hZhsqf5iCU9r2n19QBG09nH7uPm3Mt/vpQfVEHuhwRIWKfq1fi1t8EVn6sZJx7
    +6EpVvUqxqs9PdmIXX28DzU/wwQFxFGIMA5a2AuNAFYmZ7lFWqYEEhWq5tx0oQMBHQOL7AbkyR6TUq+F1wzvuZnTe1
  cwu0AYmwZr6WrWdLDj8ZMJ22k95s2PmLbmNMsuSEeSbUFmc1nZdYueieRBgL7QgokS9014lX2gN8YSwWxx1jFzCwsed
  MsdYvhawLhYJNA3Ik0FnFVJpeopLP4EQtcSdMPXzpX+AQHn/0XQdd6iaWfHppt;
  Path=/; HttpOnly
  Set-Cookie: JSESSIONID=0000SRpFCu8e03i3WXhDyJS2qnn:8bf7532d-fb68-4d4b-90fc-4aa4e3deefd5;
  Path=/; Secure; HttpOnly
  Date: Thu, 16 Feb 2017 15:14:52 GMT
  Expires: Thu, 01 Dec 1994 16:00:00 GMT
  Cache-Control: no-cache="set-cookie, set-cookie2"
}

```

Paging

Paging happens on the response data if more than 1000 objects are returned by the query to avoid server overload. Every response object contains an object that is called "paging", which contains a link that can then be used to retrieve the next set of objects, if available.

The URL to retrieve the next page can then be found under "paging" → "next" as shown in the following example:

```

curl -k -u admin:admin001 -X GET -H content-type:application/json
"https://198.51.100.1:443/scalemgmt/v2/filesystems/gpfs0/filesets?
filter=filesetName='myFset1.*'"
{
  "filesets" : [ {
    "config" : {
      "filesetName" : "myFset134",
      "filesystemName" : "gpfs0"
    },
    "status" : {}
  }, {
    "config" : {
      "filesetName" : "myFset133",
      "filesystemName" : "gpfs0"
    },
    "status" : {}
  } ],
  "paging": {
    "next" : "https://198.51.100.1:443/scalemgmt/v2/filesystems/gpfs0/filesets?
filter=filesetName='myFset1.*'"
  },
  "status" : {
    "code" : 200,
    "message" : "List of filesets for filesystem"
  }
}

```

Asynchronous jobs

All POST, PUT, and DELETE requests always run asynchronously, and a job object is returned in the response data. The returned job ID can then be used to retrieve the status of a job by using the GET /scalemgmt/v2/jobs/{jobId} endpoint.

The following example shows how to get information about the job 12345.

Request URL:

```
curl -k -u admin:admin001 -X GET --header 'accept:application/json'
'https://198.51.100.1:443/scalemgmt/v2/jobs/12345'
```

Response data:

```
{
  "status": {
    "code": "200",
    "message": "...",
  },
  "paging": {
    "next": "https://localhost:443/scalemgmt/v2/filesystems/gpfs0/filesets?lastId=1001"
  },
  "jobs": [
    {
      "result": {
        "commands": ["'mmcrfileset gpfs0 restfs1001'", "..."],
        "progress": ["'(2/3) Linking fileset'"],
        "exitCode": "0",
        "stderr": ["'EFSSG0740C There are not enough resources available to create
a new independent file set.'", "..."],
        "stdout": ["'EFSSG4172I The file set {0} must be independent.'", "..."]
      },
      "request": {
        "type": "GET",
        "url": "/scalemgmt/v2/filesystems/gpfs0/filesets",
        "data": {
          "config": {
            "filesetName": "restfs1001",
            "owner": "root",
            "path": "/mnt/gpfs0/rest1001",
            "permissions": "555"
          }
        }
      },
      "jobId": "12345",
      "submitted": "2016-11-14 10.35.56",
      "completed": "2016-11-14 10.35.56",
      "status": "COMPLETED"
    }
  ]
}
```

Note: In the JSON data that is returned, the return code indicates whether the command is successful. The response code 200 indicates that the command successfully retrieved the information. Error code 400 represents an invalid request and 500 represents an internal server error.

The following endpoints are available to monitor or delete a job:

- GET /jobs: Get details of asynchronous jobs.
- GET /jobs/{jobId}: Get details of a specific asynchronous job.
- DELETE /jobs/{jobId}: Cancel an asynchronous job that is running.

Accessing the IBM Storage Scale REST API endpoint details through Swagger and API explorer

You can access the details of the API endpoints through the following three options:

- Swagger editor.
- API explorer option that is available on the GUI node.
- IBM Storage Scale Documentation.

The Swagger option is widely used because you can view the descriptions, run the API endpoints with the required parameters, and generate client code.

Accessing API details by using Swagger

The Swagger editor uses the API description file, which can be either in YAML or JSON format. Such a description file describes the API formally in a machine-readable format with all the attributes, operations and endpoints. You can get the API description from the GUI node.

The following steps help you to access the API details and try out the endpoints by using the Swagger editor:

1. Retrieve the description file for the IBM Storage Scale management API from the following location: `https://[host name or IP address of the GUI server]:443/api/docs?root=/scalegmt/v2/`. The following screen appears.




Figure 22. Option to download API description files in JSON format from the GUI node

2. Click **Save** to save the description file in the JSON format.

Note: If you are using Chrome or Microsoft Edge, you must right-click the browser page and select **Save as** from the menu that appears to save your file in the JSON format.

3. Go to <http://editor.swagger.io/> to open the Swagger editor in the web browser.
4. In the **File** menu, select **Import file** and then browse and select the JSON file that you downloaded.

You can now see the content of the description file on the left side and the generated documentation on the right side as shown in the following figure:




Figure 23. REST API documentation opened in Swagger editor

The left side of the view provides the code. The right side of the view displays all the endpoints with their input and output parameters. It also provides the option to try out the code by using Curl. The following topics provide details about the request parameters, response, and status of the request:

- [“API requests” on page 199](#)
- [“API responses” on page 204](#)

The following steps explain how to work with the endpoints in the Swagger editor:

1. Open Swagger editor and import the description file as explained before. You can see the content of the description file on the left side and the generated documentation on the right side.
2. Click **Authorize** and enter your user credentials that you use to access the GUI node.
3. Click the endpoint that you want to explore or try out, from the list of endpoints. The endpoint details appear. Each endpoint provides the details such as the list of request parameters, option to try out the parameters, and response data that includes the requested information and endpoint execution status. The following figure shows the options that are available in the GET `/scalemgmt/v2/ces/services/{service}` endpoint.

GET

/scalemgmt/v2/ces/services/{service}

Get detailed information about a CES Service

🔒

Returns detailed information for a CES Service

Parameters

Cancel

Name	Description
service * required string (path)	Service name to query smb

Execute

Responses

Response content type

application/json

Code	Description
200	successful operation Example Value Model <pre>{ "status": { "code": 200, "message": "...", }, "cesservices": { "protocolStates": [{</pre>

Figure 24. Example for options that are available in the Swagger editor for each endpoint

4. Type the name of the service for which you need the details, in the Parameters section. In this example, SMB is selected as the target service.
5. Click **Execute** to view the details of the specified service.

The **Responses** section displays the following details:

- Curl command that is run in the background. For example, `curl -X GET "http://198.51.100.12:1000/scalemgmt/v2/ces/services/smb" -H "accept: application/json" -H "authorization: Basic YWRtaW46YWRtaW4wMDE="`
- Request URL. In this case, it is `http://198.51.100.12:1000/scalemgmt/v2/ces/services/smb`
- Response data: The response data includes the status of the request and the requested data in a GET request.

Generating client code from Swagger

To generate the client code from the Swagger editor, click **Generate Client** and select the format in which you need to generate the client code. You can use it to develop the client application that uses the IBM Storage Scale API endpoints.

Accessing API details from GUI node by using the API explorer

To display a swagger-like documentation, you can access the following page in your browser:

```
https://[guiHost]:443/api/explorer/
```

Where, [guiHost] is the host name or IP address of the IBM Storage Scale GUI server.

The options that are available to view the endpoint details or try it out, are the same as the ones that are available in the Swagger editor. When you use the API explorer that is available on the GUI node, you do not need to separately import the JSON file that contains the API endpoint descriptions.

Accessing API documentation through the IBM Storage Scale Documentation

To access the details of each endpoint from the IBM Storage Scale Documentation, see *IBM Storage Scale management API endpoints* in *IBM Storage Scale: Command and Programming Reference Guide*.

List of IBM Storage Scale management API commands

The following table lists the currently supported operations for each type of resource and resource element. The first and second columns list the resource type and the operation. The third column lists the resource and resource element terms that appear in the HTTPS request. The braces {} are a notation for indicating a resource element in text and do not appear in the actual HTTPS request. The fourth column lists the HTTP method that is called to do the operation. A footnote indicates a request that requires HTTP parameters.

Note: You can access the documentation corresponding to each API command from the GUI node. The API documentation is available in the GUI at: `https://<IP address or host name of GUI server>:<port>/api/explorer/`. For example, `https://scalegui.ibm.com:443/api/explorer`.

Table 16. Operations supported for resources and resource elements in API endpoints			
Resource type	Operation	Resource/{element}	Method
CES addresses	Get information about CES addresses.	/ces/addresses	GET
	Get information about a CES address.	/ces/addresses/{cesAddress}	GET
CES services	Get information about CES services.	ces/services	GET
	Get information about a CES service.	/ces/services/{service}	GET
Cluster	Get information about the current cluster.	/cluster	GET
GPFS snap	List the GPFS snap files containing the diagnostic data that is collected by the gpfs . snap command.	/diagnostic/snap	GET
	Create a GPFS snap file.	/diagnostic/snap	POST
	Download the specific snap file that comprises the diagnostic data to troubleshoot issues in the system.	/diagnostic/snap/{snapPath}	GET
	Upload a snap file to a PMR.	/diagnostic/snap/{snapPath}/pmr/{pmrID}	PUT
	Delete the snap file that contains the diagnostic data that is collected by the gpfs . snap command.	diagnostic/snap/{snapPath}	DELETE
CLI audit log	Gets details of the CLI commands that are issued on the cluster.	/cliauditlog	GET
Config	Get the current configuration attributes.	/config	GET

Table 16. Operations supported for resources and resource elements in API endpoints (continued)

Resource type	Operation	Resource/{element}	Method
File systems	Get information about file systems in the cluster.	/filesystems	GET
	Get information about a file system.	/filesystems/{filesystemName}	GET
	Resume a file system.	filesystems/{filesystemName}/resume	PUT
	Suspend a file system.	/filesystems/{filesystemName}/suspend	PUT
	Unmount a file system.	/filesystems/{filesystemName}/unmount	PUT
ACLs	Get access control list of a file or directory in a file system.	/filesystems/{filesystemName}/acl/{path}	GET
	Set an access control list for a file or directory in a file system.	/filesystems/{filesystemName}/acl/{path}	PUT
	Get the list of all REST API access control lists (ACL).	/access/acls	GET
	Get the list of the REST API access control lists (ACL) defined for a user group.	/access/acls/{userGroup}	GET
	Delete all the REST API access control lists (ACL) that is defined for a user group.	/access/acls/{userGroup}	DELETE
	Delete a specified REST API access control lists (ACL) entry that is defined for a user group.	/access/acls/{userGroup}/entry/{entryID}	DELETE
	Copy a REST API ACL entry from an existing group and overwrite the entries that are defined for the specified group.	/access/acls/{userGroup}	POST
	Add a REST API access control list (ACL) entry for a user group.	/access/acls/{userGroup}	PUT
File audit logging	Enable or disable file audit logging for a file system.	/filesystems/{filesystemName}/audit	PUT
Storage Pool	Get information about a particular storage pool.	/filesystems/{filesystemName}/pools/{poolName}	GET
	Get the list of storage pools in a specific file system.	/filesystems/{filesystemName}/pools	GET
Disks	Get details about disks that are part of a file system.	/filesystems/{filesystemName}/disks	GET
	Get details about a disk that is part of a file system.	/filesystems/{filesystemName}/disks/{diskName}	GET
Filesets	Get information about filesets in a file system.	/filesystems/{filesystemName}/filesets	GET
	Get information about a fileset.	/filesystems/{filesystemName}/filesets/{filesetName}	GET
	Create a fileset.	/filesystems/{filesystemName}/filesets	POST
	Modify a fileset.	/filesystems/{filesystemName}/filesets/{filesetName}	PUT
	Delete a fileset.	/filesystems/{filesystemName}/filesets/{filesetName}	DELETE
	Link an existing fileset.	/filesystems/{filesystemName}/filesets/{filesetName}/link	POST
	Unlink a fileset.	/filesystems/{filesystemName}/filesets/{filesetName}/link	DELETE
	Control AFM for a fileset.	/filesystems/{filesystemName}/filesets/{filesetName}/afmctl	POST

Table 16. Operations supported for resources and resource elements in API endpoints (continued)

Resource type	Operation	Resource/{element}	Method
Directory	Creates a directory on a GPFS file system.	/filesystems/{filesystemName}/directory/{path}	POST
	Remove a directory from a GPFS file system.	/filesystems/{filesystemName}/directory/{path}	DELETE
	Copy a directory on a file system.	/filesystems/{filesystemName}/directoryCopy/{sourcePath}	PUT
	Create a directory on a fileset.	/filesystems/{filesystemName}/filesets/{filesetName}/directory/{path}	POST
	Remove a directory from a fileset.	/filesystems/{filesystemName}/filesets/{filesetName}/directory/{path}	DELETE
	Copy a directory on a fileset.	/filesystems/{filesystemName}/filesets/{filesetName}/directoryCopy/{sourcePath}	PUT
Quotas	Get information about quota defined at the fileset level.	/filesystems/{filesystemName}/filesets/{filesetName}/quotas	GET
	Set quota limits at the fileset level.	/filesystems/{filesystemName}/filesets/{filesetName}/quotas	POST
	Get information about quota defined at the file system level.	/filesystems/{filesystemName}/quotas	GET
	Set quota limits at the file system level.	/filesystems/{filesystemName}/quotas	POST
	List quota defaults in a cluster.	/filesystems/{filesystemName}/filesets/{filesetName}/quotadefaults	GET
	Set the quota defaults USR, GRP, or FILESET for a file system.	/filesystems/{filesystemName}/filesets/{filesetName}/quotadefaults	POST
	Enable/disable the quota defaults for a fileset.	/filesystems/{filesystemName}/filesets/{filesetName}/quotadefaults	PUT
	List quota defaults in the cluster.	/filesystems/{filesystemName}/quotagracedefaults	GET
	Set the default quota for user, group, and fileset of a file system.	/filesystems/{filesystemName}/quotagracedefaults	POST
	Enable or disable the quota management for a file system.	/filesystems/{filesystemName}/quotamanagement	PUT
Cloud Object Store	Create a Cloud Object Store fileset.	/filesystems/{filesystemName}/filesets/cos	POST
	Create a Cloud Object Store-related directory in a fileset.	/filesystems/{filesystemName}/filesets/{filesetName}/cos/directory	POST
	Download files from Cloud Object Store.	/filesystems/{filesystemName}/filesets/{filesetName}/cos/download	POST
	Evict files from object store.	/filesystems/{filesystemName}/filesets/{filesetName}/cos/evict	POST
	Upload files to the object store.	/filesystems/{filesystemName}/filesets/{filesetName}/cos/upload	POST
Owner	Get details of the owner of the file or directory in a file system.	/filesystems/{filesystemName}/owner/{path}	GET
	Set owner of the file or directory in a file system.	/filesystems/{filesystemName}/owner/{path}	PUT
Policies	Lists the policies that are applied on the file system.	/filesystems/{filesystemName}/policies	GET
	Sets a policy for a file system.	/filesystems/{filesystemName}/policies	PUT
Information	Get information about the IBM Storage Scale management API.	/info	GET
Jobs	Get details about asynchronous jobs.	/jobs	GET
	Get details about an asynchronous job.	/jobs/{jobID}	GET
	Cancel a specific asynchronous job.	/jobs/{jobID}	DELETE

Table 16. Operations supported for resources and resource elements in API endpoints (continued)

Resource type	Operation	Resource/{element}	Method
NFS exports	Get information about NFS exports configured in the system.	/nfs/exports	GET
	Get information about an NFS export configured in the system.	/nfs/exports/{exportPath}	GET
	Create an NFS export.	/nfs/exports	POST
	Modify an existing NFS export.	/nfs/exports/{exportPath}	PUT
	Delete an NFS export.	/nfs/exports/{exportPath}	DELETE
Nodes	Get information about nodes in the cluster.	/nodes	GET
	Get information about a node.	/nodes/{nodeName}	GET
	Add one or more nodes in the cluster.	/nodes	POST
	Delete one or more nodes from the cluster.	/nodes/{name}	DELETE
	Change node designation for a node or node class.	/nodes/{name}	PUT
Node mapping	List all AFM node mappings.	/nodes/afm/mapping	GET
	Create AFM node mapping.	/nodes/afm/mapping	POST
	Delete existing mapping.	/nodes/afm/mapping/{mappingName}	DELETE
	Gets the details of a specific map.	/nodes/afm/mapping/{mappingName}	GET
	Change an existing map.	/nodes/afm/mapping/{mappingName}	PUT
Node health	Get details about system health events that are reported on a node.	/nodes/{name}/health/events	GET
	Get details about system health states of the events that are reported in a node.	/nodes/{name}/health/states	GET
Node class	Get details of the node classes.	/nodeclasses	GET
	Create a node class.	/nodeclasses	POST
	Delete a specific node class.	/nodeclasses/{nodeclassName}	DELETE
	Get details of a specific node class.	/nodeclasses/{nodeclassName}	GET
	Make changed to a specific node class.	/nodeclasses/{nodeclassName}	PUT
Symlink	Remove a symlink from a GPFS file system.	/filesystems/{filesystemName}/symlink/{path}	DELETE
	Create a symlink for a path from a GPFS file system.	/filesystems/{filesystemName}/symlink/{linkPath}	POST
Services	Gets the details of the services that are configured on a node.	/nodes/{name}/services	GET
	Gets the details of a specific service that is configured on a node.	/nodes/{name}/services/{serviceName}	GET
	Starts or stops a service on a node or node class.	/nodes/{name}/services/{serviceName}	PUT
NSDs	Gets information about NSDs that are part of the system.	/nsds	GET
	Gets information about an NSD.	/nsds/{nsdName}	GET

Table 16. Operations supported for resources and resource elements in API endpoints (continued)

Resource type	Operation	Resource/{element}	Method
Snapshots	Get information about fileset snapshots.	/filesystems/{filesystemName}/filesets/{filesetName}/snapshots	GET
	Get information about a fileset snapshot.	/filesystems/{filesystemName}/filesets/{filesetName}/snapshots/{snapshotName}	GET
	Get information about file system snapshots.	/filesystems/{filesystemName}/snapshots	GET
	Get information about a file system snapshot.	/filesystems/{filesystemName}/snapshots/{snapshotName}	GET
	Create a fileset snapshot.	/filesystems/{filesystemName}/filesets/{filesetName}/snapshots	POST
	Create a file system snapshot.	/filesystems/{filesystemName}/snapshots	POST
	Delete a fileset snapshot.	/filesystems/{filesystemName}/filesets/{filesetName}/snapshots	DELETE
	Delete a file system snapshot.	/filesystems/{filesystemName}/snapshots	DELETE
	Copy a file system snapshot.	/filesystems/{filesystemName}/filesets/{filesetName}/snapshotCopy/{snapshotName}	PUT
	Copy a directory from a path relative to a snapshot of the file system.	/filesystems/{filesystemName}/filesets/{filesetName}/snapshotCopy/{snapshotName}/path/{source Path}	PUT
AFM peer snapshots	Create a peer snapshot.	/filesystems/{filesystemName}/filesets/{filesetName}/psnaps	POST
	Delete a peer snapshot.	/filesystems/{filesystemName}/filesets/{filesetName}/psnaps/{snapshotName}	DELETE

Table 16. Operations supported for resources and resource elements in API endpoints (continued)

Resource type	Operation	Resource/{element}	Method
AFM	List AFM state in a file system.	/filesystems/{filesystemName}/afm/state	GET
	Upload an AFM file.	afm/upload	POST
	Delete an uploaded file.	afm/upload/{filename}	DELETE
	Get all local bucket definition.	bucket	GET
	Set or change keys for a bucket.	bucket/keys	PUT
	Delete local bucket definition with keys.	bucket/keys/{bucketName}	DELETE
	Get listing of fileset.	filesystems/{filesystemName}/filesets	GET
	Create a new fileset.	filesystems/{filesystemName}/filesets	POST
	Create a new cloud object storage related fileset	filesystems/{filesystemName}/filesets/cos	POST
	Delete a fileset.	filesystems/{filesystemName}/filesets/{filesetName}	DELETE
	Get detailed information about a fileset.	filesystems/{filesystemName}/filesets/{filesetName}	GET
	Change an existing fileset	filesystems/{filesystemName}/filesets/{filesetName}	PUT
	The AFM control functions	filesystems/{filesystemName}/filesets/{filesetName}/afmctl	POST
	Specifies the delete operation that is performed on an AFM to cloud object storage fileset.	filesystems/{filesystemName}/filesets/{filesetName}/cos	DELETE
	Create a new cloud object storage related directory in a fileset.	filesystems/{filesystemName}/filesets/{filesetName}/cos/directory	POST
	Download files from an object store.	filesystems/{filesystemName}/filesets/{filesetName}/cos/download	POST
	Shows download progress for the previously run or currently running download task executed by using the mmafmcscctl download command.	filesystems/{filesystemName}/filesets/{filesetName}/cos/downloadStats	POST
	Evict files from an object store.	filesystems/{filesystemName}/filesets/{filesetName}/cos/evict	POST
	Specifies the remove operation by using user defined policy.	filesystems/{filesystemName}/filesets/{filesetName}/cos/policy	DELETE
	Specifies the list operation by using user defined policy.	filesystems/{filesystemName}/filesets/{filesetName}/cos/policy	GET
	Specifies the add operation by using user defined policy.	filesystems/{filesystemName}/filesets/{filesetName}/cos/policy	POST
	Specifies the reconcile update operation by using user defined policy.	filesystems/{filesystemName}/filesets/{filesetName}/cos/policy	PUT
	Upload files to an object Store.	filesystems/{filesystemName}/filesets/{filesetName}/cos/upload	POST
SMB shares	Get information about SMB shares.	smb/shares	GET
	Get information about an SMB share.	/smb/shares/{shareName}	GET
	Create an SMB share.	/smb/shares	POST
	Modify an existing SMB share.	/smb/shares/{shareName}	PUT
	Delete an SMB share.	/smb/shares/{shareName}	DELETE
	Delete complete access control list of an SMB share.	/smb/shares/{shareName}/acl	DELETE
	Get access control list of share.	/scalemgmt/v2/smb/shares/{shareName}/acl	GET
	Delete an entry from the access control list of a SMB share.	/smb/shares/{shareName}/acl/{name}	DELETE
	Get access control list entry of a specific user/group/system of share.	/smb/shares/{shareName}/acl/{name}	GET
	Adds an entry to the access control list of an SMB share.	/smb/shares/{shareName}/acl/{name}	PUT

Table 16. Operations supported for resources and resource elements in API endpoints (continued)

Resource type	Operation	Resource/{element}	Method
Performance monitoring	Get performance data.	/perfmon/data	GET
	Gets the sensor configuration.	/perfmon/sensor	GET
	Gets details about a specific sensor configuration.	/perfmon/sensor/{sensorName}	GET
	Modifies the sensor configuration.	/perfmon/sensor/{sensorName}	PUT
	Run a custom query for performance data.	/perfmon/stats	GET
Thresholds	Get list of threshold rules.	/thresholds	GET
	Create a threshold rule.	/thresholds	POST
	Delete an existing threshold rule.	/thresholds/{name}	DELETE
	Get details of a threshold rule.	thresholds/{name}	GET
For more information about the commands, see <i>REST API commands</i> in the <i>IBM Storage Scale: Command and Programming Reference Guide</i> .			

IBM Storage Scale native REST API

The IBM Storage Scale native REST API is a new feature that provides native REST APIs for secure and remote administration of IBM Storage Scale.

With this native REST API feature, you can manage the IBM Storage Scale cluster through a new daemon that runs on each IBM Storage Scale node. This feature replaces the administrative operations that were previously done with mm-commands. It also eliminates several limitations of the mm-command design, including the dependency on SSH, the requirement for privileged users, and the need to issue commands locally on the node where IBM Storage Scale runs.

The native REST API feature includes the following functions for IBM Storage Scale:

- New control plane daemon (administration daemon)
- New CLI (**scalectl**)
- Remote administration
- Role-based access control (RBAC)
- Allow non-root administration
- No dependency on SSH

New control plane daemon

A new administration daemon (**scaleadmd**) is deployed on the IBM Storage Scale nodes to use native REST API functions. The purpose of this daemon is to serve the native REST API interface, process the administrative requests from users, and handle all associated authentication and authorization controls for each request.

Important: When the administration daemon starts or restarts, there might be a brief delay before it can serve requests. As a result, the first CLI or REST API request might be rejected if the daemon startup routines are still completing. This behavior is similar to the delay that occurs before the I/O daemon can serve all requests.

The administration daemon restart can introduce this delay in the following situations:

- Immediately after installation of the native REST API rpm (the administration daemon starts at the end of the installation process).
- After nodes are added to the cluster (the remote node administration daemon internal services restart on the added nodes).
- After you manual start or restart the administration daemon by using the **systemd** or a similar command
- After node deletion (the administration daemon restarts on the node that is deleted).

A common error status that is returned is **UNAVAILABLE**, but other errors might also occur. If a request fails after the administration daemon restarts under these conditions, wait briefly (a few seconds) and retry the command. Repeat this process as needed.

Remote administration

With the native REST API feature, users can now issue admin operations from outside the cluster. The new **scalectl** CLI can communicate with the REST API through HTTPS, which enables commands to be run from outside the IBM Storage Scale node.

Role-based access control (RBAC)

The native REST API includes RBAC, allowing the security administrator to grant granular permissions for a user. A user can be allowed or denied to run operations at the resource level. For example, a user can be granted permission to operate on a specific file system or fileset only.

Allow non-root administration

Running the existing administrative mm-commands requires elevated privileges. The native REST API feature does not require the elevated privileges. RBAC policies can determine which users have the authority to run certain administrative operations. The user does not require the elevated privileges on the system to request administrative actions to take place.

No dependency on SSH

With the native REST API feature, operations that are done by the API use secure connections between administration daemons for remote code execution, eliminating the need for SSH.

High-level architecture

The administration daemon authenticates and authorizes user requests, either granting or denying permission based on the RBAC policy rules at the time of the request. Every request goes through the authentication and authorization middleware.

When you run the **scalectl**, it connects to the administration daemon through HTTPS (when the **scalectl --url <host> <port>** option is used) or through UNIX domain socket (UDS) by default.

The **scalectl** binary is a lightweight client of the administration daemon. It does minimal user input checks and processing. The administration daemon handles all request processing, authentication, and authorization.




Figure 25. High-level architecture

The following figure depicts the default **scalectl** flow, which requires running the command on the API node itself:




Figure 26. Default **scalectl** flow

The following figure depicts the **scalectl** flow with **--url** option (available to run from outside the cluster):




Figure 27. **scalectl** flow with **--url** option

REST API overview

The administration daemon supports a REST API that allows applications to communicate with it directly, without the use of **scalectl** command-line interface. For more information about the native REST API endpoints, including request and response payloads, see the OpenAPI specification. You can access this OpenAPI specification at <https://x.x.x.x:46443/openapi/>, where x.x.x.x is the public IP address of the API server node. Ensure that the trailing slash (/) is included in the URL.

Non-root administration daemon

An administration daemon (**scaleadm**) is deployed on the IBM Storage Scale nodes to use IBM Storage Scale native REST API functions. The purpose of this daemon is to serve the IBM Storage Scale native REST API interface, process the administrative requests from users, and handle all associated authentication and authorization controls for each request.

When you install the **gpfs.scaleapi** package, a new user and group are created. If a user or group with the same name exists, the installation fails. The installation creates the following items: The home directory for the user is at **/var/mmfs/scaleadm**.

- **User:**scaleapiadm
- **Group:**scaleapiadm
- **Home directory:**/var/mmfs/scaleadm

Installation, upgrade, and support operations require root access on all cluster nodes. You must maintain this access and enable or disable it based on your operational and security requirements.

Some IBM Storage Scale operations require root privileges. The `/usr/lpp/mmfs/bin/scaleadmexec` binary uses Linux capabilities to elevate privileges during execution. The administration daemon invokes this binary for operations such as direct disk access, PAM authentication, SSH or SCP replacement calls (such as **`scaleadmremoteexecute`** or **`scaleadmremotetransfer`**), among others.

Note: Make sure any output directory that is specified by a command has the correct permissions for the administration daemon to access it.

Important: When the administration daemon starts or restarts, there might be a brief delay before it can serve requests. As a result, the first CLI or REST API request might be rejected if the daemon startup routines are still completing. This behavior is similar to the delay that occurs before the I/O daemon can serve all requests.

The administration daemon restart can introduce this delay in the following situations:

- Immediately after installation of the native REST API rpm (the administration daemon starts at the end of the installation process).
- After nodes are added to the cluster (the remote node administration daemon internal services restart on the added nodes).
- After you manual start or restart the administration daemon by using the **`systemd`** or a similar command
- After node deletion (the administration daemon restarts on the node that is deleted).

A common error status that is returned is **UNAVAILABLE**, but other errors might also occur. If a request fails after the administration daemon restarts under these conditions, wait briefly (a few seconds) and retry the command. Repeat this process as needed.

API nodes for IBM Storage Scale native REST API

The IBM Storage Scale native REST API feature includes a new node class called API nodes. The native REST API is available only through the API nodes, which expose the full suite of APIs provided by this feature. API nodes also support local commands through the **`scalectl`** CLI without the **`--url`** option. By default, API nodes are quorum nodes.

API nodes manage and replicate metadata that is required for the native REST API, including long-running operations (LRO) data. These nodes can ensure high availability of the API. Because quorum is required for IBM Storage Scale to function, quorum nodes are designated as API nodes by default. The following figure illustrates the architecture of a cluster with three quorum or API nodes:

API Nodes




Figure 28. High-level architecture of a cluster with API nodes

The native REST API server does not start on non-API nodes. These nodes expose only a subset of native REST API functions through **scalectl** commands, which are limited to managing local configurations, such as node identities.

Only the following **scalectl** commands are supported locally on non-API nodes:

- **scalectl node config**
- **scalectl cluster**
- **scalectl config**

To use additional native REST API functions from a non-API nodes, issue the **scalectl** command with the **--url <FQDN/IP>:<port>** flag. The FQDN or IP must point to an API node, and the request is sent over HTTPS.

Note:

- If the native REST API feature is enabled and the quorum designation of a node is changed, restart **scaleadmd** service on nodes where the quorum designation was changed.

Authentication for IBM Storage Scale native REST API

When the IBM Storage Scale native REST API feature is deployed, authentication is enabled by default.

Authentication

The following authentication methods are supported for IBM Storage Scale native REST API:

- Pluggable authentication modules (PAM)
- Lightweight directory access protocol (LDAP)
- OpenID connect (OIDC)
- Active directory (AD) (through LDAP queries)

- Certificate-based authentication

Note: After the **authentication_type** and its associated configuration are changed, restart the **scaledmd** service for the new authentication type to take effect.

PAM

PAM authentication is supported by using the **pam_unix** module. It is the default authentication method for IBM Storage Scale native REST API. The **pam_unix** module validates a username and password against the system password database (/etc/passwd and /etc/shadow). The user who issues commands through the **scalectl** or the REST API must exist on the nodes where the commands are issued. IBM Storage Scale does not create these users. Cluster administrators must create these users, and ensure that these usernames do not overlap across cluster nodes.

When you use the **scalectl** command locally on a node, the administration daemon gets the peer client information from the unix domain socket (UDS). It uses the UID of the client to find the username for authorization. Because the user is already logged in locally, no additional authentication is done at this stage. The client username is then used for authorization as defined in the RBAC rules.

When you use the **scalectl** command with the **--url** flag or through third-party clients such as **curl**, authentication is sent through the HTTP header basic authentication scheme. The **scalectl** command requires a **.netrc** file in the home directory of the user profile. Use the **NETRC** environmental variable to specify a different path. **scalectl** does not attempt to resolve IP addresses or hostnames that are defined in the **.netrc** file. The IP address, hostname or fully qualified domain name (FQDN) must match the value that is provided in the **--url** flag. To match any machine, use the *default* keyword in the **.netrc** file. The same **.netrc** file can be used with **curl** by using the **-n** or **--netrc** flags. This file generates the same basic authentication header.

A sample **.netrc** file can include the following entries:

- To match a specific host defined by the **--url** flag, use one or more machine entries:

```
machine <host> login <user> password <password>
```

- To match any host defined by the **--url** flag, use a default entry. Include only one default entry, and place it after all machine entries:

```
default login <user> password <password>
```

The **basic** authorization scheme is used to pass the encoded username and password to the PAM for validation. This scheme is base64-encoded but not encrypted. However, HTTPS is used to connect to the administration daemon. Therefore, the username and password are encrypted in transit. The username and password are never cached or stored locally.

OIDC

OIDC authentication uses your organization identity provider to authenticate with IBM Storage Scale native REST API. Modify the administration daemon configuration to switch the authentication type to **oidc-userinfo** by using **scalectl** or REST APIs. After you modify the configuration, the subconfiguration that is defined under **rest_server.oidc** is used. Depending on your environment, the following options are supported:

```
rest_server.oidc.callback_url
rest_server.oidc.authorization_url
rest_server.oidc.token_url
rest_server.oidc.userinfo_url
rest_server.oidc.discovery_url
rest_server.oidc.client_id
rest_server.oidc.client_secret
rest_server.oidc.scope
```

Provide the **callback_url** in the subconfiguration and to your identity provider. If your identity provider supports the **discovery_url**, use this option instead of providing individual URLs for **authorization_url**, **token_url**, and **userinfo_url**. A specific URL provided for **authorization_url**, **token_url**, or **userinfo_url** takes precedence over any URL obtained through the **discovery_url** endpoint. The **OAUTH_CLIENT_SECRET** environment variable takes

precedence over the **client_secret** configuration option. By default, *openid* is always included in the **scopes**. If required, add additional **scopes**, including *email* and *profile*.

Note: Restart **scaleadmd** for the changes to take effect.

To generate an OIDC token from your identity provider, complete the following steps:

1. Open a browser and go to the following URL, replacing <FQDN> with your organization's fully qualified domain name (FQDN):

```
https://<FQDN>:46443/login
```

2. Enter your credentials to log in.
3. After successful authentication, the token is displayed in the browser in the following format:

```
{ "token": "<your_token>" }
```

For future requests, use the **--bearer** option with the **scalectl** command. If this option is set, **scalectl** requires the token to be specified in the **OIDC_TOKEN** environment variable.

The authorization bearer header scheme is used to pass the token for validation. The connection to **scaleadmd** uses HTTPS, ensuring that the token is encrypted in transit.

When **scaleadmd** authenticates a request, it forwards the token to the identity provider for validation. To reduce network traffic, validation results are cached for 10 seconds. If a token is revoked or expires, it might allow additional requests for up to 10 seconds due to caching.

LDAP

Similar to the PAM authentication, LDAP requires the `.netrc` file to contain the username and password that can be validated by the LDAP server. Modify the administration daemon configuration to switch the authentication type to **ldap** by using **scalectl** or REST API.

After you modify the configuration, the subconfiguration that is defined under **rest_server.ldap** is used. Depending on your environment, the following options are supported:

```
rest_server.ldap.url
rest_server.ldap.basedn
rest_server.ldap.filter
rest_server.ldap.binddn
rest_server.ldap.bind_password
```

The **filter** option must return a single result. If multiple results are returned, an error occurs. A bind is attempted by using the provided **binddn** and **bind_password**. If they are not set, an unauthenticated bind is used. The **filter** option replaces the **%s** placeholder with the username to be authenticated. The DN and password from the returned result are used to authenticate another bind request. If the bind is successful, authentication is completed and authorization proceeds with the provided username.

You can comment out or remove any configuration options that are not in use.

Note: Restart **scaleadmd** for the changes to take effect.

AD (through LDAP queries)

If your environment requires AD, enable the AD authentication through the LDAP configuration. LDAP can query AD servers, allowing the same LDAP configuration to be used for AD authentication.

Similar to the PAM and LDAP authentication, AD authentication requires the `.netrc` file to contain the username and password for validation by the AD server. Modify the administration daemon configuration to switch the authentication type to **ldap** by using the **scalectl** command or REST API. After you modify the configuration, the subconfiguration that is defined under **rest_server.ldap** is used. Depending on your environment, the following options are supported:

```
rest_server.ldap.url = "ldap://test1.xxx.stgxxx.ibm.com"
rest_server.ldap.basedn = "cn=users,dc=storage3test,dc=tuc,dc=stglabs,dc=ibm,dc=com"
rest_server.ldap.filter = "(sAMAccountName=%s)"
```

```
rest_server.ldap.binddn = "ADuser1@test1.xxx.stgxxx.ibm.com"
rest_server.ldap.bind_password = "&lt;password&gt;"
```

Note: Restart **scaledmd** for the changes to take effect.

Certificate-based authentication

Certificate-based authentication allows the use of certificates in the IBM Storage Scale native REST API. Modify the administration daemon configuration to switch the authentication type to **cert** by using **scalectl** or REST APIs.

After you modify the configuration, the subconfiguration that is defined under **rest_server.cert** is used. Depending on your environment, the following options are supported:

```
rest_server.cert.cachain
rest_server.cert.allowed_cns
rest_server.cert.allowed_cns_regex
```

- **cachain:** Specifies the certificate authority (CA) certificate that signs the client certificates. If this field is not specified, the cluster node identity **cachain** is used.
- **allowed_cns:** Specifies a list of common names (CNs) that are allowed. The CN from the authenticating certificate must exactly match one of the CNs in this list.
- **allowed_cns_regex:** Uses a regular expression to match the CN of the authenticating certificate. For example, "local...\$" matches localhost but not localhost1.
- **allowed_cns** takes precedence over **allowed_cns_regex**. If neither is set, then it allows any CN if the certificate was signed by the **cachain**.

Note: If **allowed_cns** and **allowed_cns_regex** are both set, **allowed_cns** takes precedence. If both the **allowed_cns** and **allowed_cns_regex** are not set, any CN is allowed if the certificate was signed by the **cachain**.

After certificate-based authentication is enabled, use the **--cert** and **--key** flags with **scalectl** to issue requests to the administration daemon. The CN from the authenticated certificate is extracted and used as the username for authorization.

Role-based access control

Role-based access control (RBAC) is a security feature that allows a security administrator to grant granular permissions to a user. A user can be allowed or denied permission to run operations at the resource level. For example, a user can be granted permission to operate only on a specific file system or fileset. It is enabled when the IBM Storage Scale native REST API feature is deployed.

IBM Storage Scale includes one predefined user, root, and one predefined role, **SecurityAdmin**. By default, only the root user can run the initial **scalectl** commands to configure node identities and create the cluster. To allow a non-privileged user to perform the initial setup, run the `/usr/lpp/mmfs/bin/scaleuseraddrbac` script to allow the user to create the node identities and create the cluster. This script grants the specified user the **SecurityAdmin** role, which allows full access to all API endpoints. The script must be run before the cluster is created. The assigned user retains **SecurityAdmin** privileges unless explicitly removed.

Domains provide logical groupings of memberships, permissions, and resources that are required for IBM Storage Scale components to function correctly. If needed, you can assign additional users to the predefined **SecurityAdmin** role that is included with IBM Storage Scale. If no domain is specified in **scalectl --domain** or the REST header **X-StorageScaleDomain**, the **StorageScaleDomain** is used by default. The default **StorageScaleDomain** cannot be deleted.

The following core components implement RBAC:

Domain

A logical grouping of resources, users or roles, and actions. No built-in restrictions on which resources can be included in a domain.

Resources

A resource is represented as a URL endpoint. A resource can be a file system, fileset, disk, or NSD. A wildcard (*) is used to match any resource. For example, **/scalemgmt/v3/filesystems/fs0/filesets/*** matches any filesets in the *fs0* file system for the specified action.

Actions

Actions are the various operations that can be done on a resource. Not all actions apply to every resource. Supported actions are *create*, *delete*, *get*, *list*, *update*, *link*, *unlink*, *mount*, *unmount*, and *cani* or *impersonate*.

User

The person who makes the request.

Effect

Defines whether access is allowed or denied. The rule evaluation follows these conditions:

- By default, set to **Deny**.
- The system checks for an allow rule that matches the request.
- If an explicit deny rule exists, it overrides an allow rule.
- To grant permission, there must be an explicit allow rule without a conflicting deny rule.
- **Deny** rules are useful when using wildcards (*) to allow all actions except those explicitly denied.

Membership

The relationship between a user and their role within a domain.

Permission

The relationship between a role, action, effect, and resource within a domain.

Domain filtering in the RBAC

RBAC uses the domain specified in a request to determine whether a user has permission to access a resource. If a user has permission in a domain, they can create resources in that domain. No additional domain checks occur during creation. After a resource is created, only users with permission in the same domain can perform actions such as **get**, **list**, **update**, or **delete** on the resources. This behavior is known as domain filtering.

For example,

- If a user has delete permissions for wildcard **Delete /scalemgmt/v3/filesystems/*** resources in *domain1*, they can delete only the file system in *domain1* and not in *domain2*. To delete a file system in *domain2*, the user must have explicit delete permissions for **/scalemgmt/v3/filesystems/*** in *domain2*.

When a user creates a resource, that resource is associated with the domain in which the user was authorized. To perform actions such as **get**, **list**, **update**, or **delete**, the user must have permission for the endpoint in the same domain as the resource.

When a user lists resources, they see only those associated with their domain. If an administrator grants a user wildcard **GET /scalemgmt/v3/filesystems** permissions, the user can view only the file systems in their domain. This filtering enforcement allows administrators to grant broad permissions while restricting access to resources from other domains. This also applies when creating sub-resources that depend on others.

For example,

- If *userA* in *domainA* creates NSDs, those NSDs belong to *domainA*. If *userB* in *domainB* tries to create a file system using those NSDs, they appear as non-existent because they are filtered from view. Similarly, if *userB* lists NSDs, they see only the NSDs in *domainB*.

The following operations are exception to this domain filtering:

Creation of fileset

If a user has permission to create a fileset in a specific file system, the creation is allowed, even if the file system belongs to another domain. This exception is allowed because the fileset creation is a common and requiring a filesystem per domain is often impractical.

Clear ID of NSDs

If a user has permission to clear ID of NSDs (**DELETE /scalemgmt/v3/troubleshooting/nsds/clearId**), they can affect all devices. This operation is highly privileged, intended for recovery scenarios only. Use this resources only if the NSD was successfully deleted but failed to clear the ID from the disk. Since the NSD no longer exists, domain filtering is skipped. The administration daemon ensures that the NSD does not exist before allowing this command, so it cannot be run against defined NSDs.

You can manage domains and policy rules by using the **scalectl** command or REST API. The default policy file allows basic RBAC rules. If the user has permission for an endpoint, they can run the **scalectl authorization** commands or use authorization API endpoints to create or manage the domains. For more information, see *scale authorization command or Authorization endpoints* section in the *IBM Storage Scale: Command and Programming Reference Guide*.

Long-running operations

The IBM Storage Scale native REST API feature supports long-running operations (LRO). An LRO is an internal task that represents business logic to process user-requested actions that take a significant amount of time to complete. For example, deleting NSDs is an LRO. When an LRO is submitted, the LRO manager handles its lifecycle, including accepting, running, and monitoring operations.

An LRO can be in one of the following states during its lifecycle:

- Accepted: The LRO manager acknowledged receiving an LRO.
- Initializing: The LRO manager started the operation, and it is preparing to run.
- Running: The LRO is processing the IBM Storage Scale native REST API user-requested action.
- Failed: The LRO encountered an error that forced it to stop execution.
- Interrupted: The LRO was unexpectedly terminated.
- Canceled: The LRO was stopped due to a cancellation request.
- Done: The LRO successfully completed the user-requested action.
- Stale: The LRO is unresponsive and might be deadlocked.

Note: The LRO status represents its lifecycle but does not necessarily indicate the outcome of the operation. If an LRO stops due to an error, its status is *Failed*. If an LRO completes but encounters noncritical errors that do not stop execution, its status is *Done*. The **result** field in the operation resource contains a response if the LRO completes successfully or an error if execution stops due to an error.

LRO behavior in REST API and CLI

When an LRO completes, REST APIs return HTTP 200 (OK). If the request has been accepted but the operation is still in progress, the API returns HTTP 202 (Accepted).

For some LIST-type LROs that are invoked by using the **scalectl** CLI with pagination enabled and JSON output requested (for example, **scalectl fileset list fs1 --json**), the CLI might internally issue multiple LIST requests to retrieve all pages of data. Returning individual LRO responses for each request is not practical in this scenario. Instead, the CLI aggregates the results from multiple requests and returns a list of the corresponding objects rather than individual LRO responses.

To receive an individual response wrapped in an LRO object for each LIST request when using the CLI, pagination must be disabled.

Note: When you do not use the **--json** flag with the **scalectl** CLI on LRO-type endpoint methods, the request waits for the full LRO operation to reach a completion status. When you use the **--json** flag with the **scalectl** CLI, the request returns quickly. You must check the status manually by using the **scalectl operation get <operation_id>** CLI until the job reaches a completion state.

When **scalectl** processes long-running operations (LROs), it polls the returned operation name until the request is complete, if the request is not completed immediately. This behavior can have role-based access control (RBAC) implications. For an LRO, users might require permissions to query LROs in addition to the permissions to create the resource. These permissions are necessary for **scalectl** to retrieve the final output during the polling phase. If a user does not have the required RBAC permissions to query LROs, **scalectl** returns a **PERMISSION_DENIED** error. This error indicates that the user does not have the necessary permissions on the `/scalemgmt/v3/operations/` resource. For more information about the role-based access control, see *Role-based access control* section in the *IBM Storage Scale: Concepts, Planning, and Installation Guide*.

Managing long-running operations

Users can interact with LRO through the LRO manager, which provides the endpoints in both the REST API and CLI. These endpoints provide an interface to display the execution status and results of LROs. They also allow for the cleanup of historical LROs that are no longer needed. The following endpoints allow to interact with the LRO manager:

LIST

The **LIST** operation retrieves all the LROs managed by the LRO manager. For more information about the operations CLI and API, see **scalectl operation list** command or operations API endpoints in the *IBM Storage Scale: Command and Programming Reference Guide*.

GET

The **GET** operation retrieves information about a single LRO that is managed by the LRO manager. For more information about the operations CLI and API, see **scalectl operation get** command or operations API endpoints in the *IBM Storage Scale: Command and Programming Reference Guide*.

OUTPUT

The **OUTPUT** operation retrieves any console output that is generated by an LRO job. LRO might generate console output during their execution. For more information about the operations CLI and API, see **scalectl operation output** command or operations API endpoints in the *IBM Storage Scale: Command and Programming Reference Guide*.

CANCEL

The **CANCEL** operation submits a request to the LRO manager to stop an LRO job. The job status changes to *Canceled* once it stops. For more information about the operations CLI and API, see **scalectl operation cancel** command or operations API endpoints in the *IBM Storage Scale: Command and Programming Reference Guide*.

You can cancel the following LRO endpoints:

File system endpoints

- `filesystems: POST`

File system disk endpoints

- `filesystems/{filesystem}/disks: POST`
- `filesystems/{filesystem}/disks/{disk_name}: DELETE`
- `filesystems/{filesystem}/disks:batchAdd: POST`
- `filesystems/{filesystem}/disks:batchDelete: POST`

Snapshot endpoints

- `filesystems/{filesystem}/filesets/snapshots:batchCreate: POST`
- `filesystems/{filesystem}/filesets/snapshots:batchDelete: POST`
- `filesystems/{filesystem}/filesets/{fileset}/snapshots: GET`
- `filesystems/{filesystem}/filesets/{fileset}/snapshots: POST`
- `filesystems/{filesystem}/filesets/{fileset}/snapshots/{snapshot_name}: DELETE`
- `filesystems/{filesystem}/snapshots: GET`

- filesystems/{filesystem}/snapshots: POST
- filesystems/{filesystem}/snapshots/{snapshot_name}: DELETE
- filesystems/{filesystem}/snapshots:batchDelete: POST

DELETE

The **DELETE** operation removes a non-running job from the LRO manager. This action permanently deletes all LRO information and metadata. Deleting an active or running job is not allowed. A successful delete action returns an empty JSON object. If an error occurs, the JSON response contains error codes and messages. For more information about the operations CLI and API, see **scalectl operation delete** command or operations API endpoints in the *IBM Storage Scale: Command and Programming Reference Guide*.

Target nodes field in IBM Storage Scale native REST API

In the IBM Storage Scale native REST API feature, the target nodes are used to execute requests on one or more cluster nodes.

The administration daemon exposes APIs that require actions on a subset of cluster nodes specified by the user. Traditionally, you can accomplish this with the **-N** flag for mm-commands. For example:

- To mount a file system on *nodeA*:

```
mmmount <filesystem> -N nodeA
```

- To mount a file system on all nodes in *nodeclass1*:

```
mmmount <filesystem> -N nodeclass1
```

The IBM Storage Scale native REST API feature introduces the **target_nodes** field to achieve similar behavior. If the **target_nodes** field is not set, the node that process the request handles it locally.

The scalectl CLI provides the **-N or --target-nodes <string>** option for specifying target nodes on which to mount the filesystem.

Examples

The following examples show how to mount a file system on specific nodes or node classes by using the REST API:

- To mount the file system on nodes *scale-11* and *scale-12*, execute the following request:

```
curl -X 'POST' 'https://localhost:8000/scalemgmt/v3/filesystems/fs0:mount' -H 'accept: application/json' -H 'Content-Type: application/json' -d '{
  "target_nodes": ["scale-11", "scale-12"]
}' -k
```

- To mount the filesystem on all node classes, execute the following request:

```
curl -X 'POST' \ . 'https://localhost:8000/scalemgmt/v3/filesystems/fs0:mount' \ -H
'accept: application/json' \ -H 'Content-Type: application/json' \ -d '{
  "target_nodes": ["all"]
}'
```

The following examples show how to mount the file system on nodes and node classes by using the **scalectl** command:

- To mount the file system on nodes *scale-11* and *scale-12*, run the following command:

```
scalectl filesystem mount fs0 -N scale-11,scale-12
```

A sample output is as follows:

```
Successfully mounted filesystem fs0
```

- To mount the filesystem on all node classes, run the following command:

```
scalectl filesystem mount fs0 -N all
```

A sample output is as follows:

```
Successfully mounted filesystem fs0
```

Pagination in IBM Storage Scale native REST API

The IBM Storage Scale native REST API feature includes pagination support to optimize performance and handle large collections. Listing resource collections, such as filesets, can be an expensive operation when the collection contains many items. Pagination helps improve efficiency by reducing lookup times and controlling the size of responses as collections grow.

The following characteristics apply to IBM Storage Scale native REST API endpoints that implement pagination:

- The IBM Storage Scale native REST API endpoints that implement pagination accept the **page_size** and **page_token**. Both of these parameters are optional for requests.
 - **page_size**: Specifies the number of items that need to be returned for the request. If not specified, a default value is used, dependent on the collection being listed.
 - **page_token**: Specifies a pointer to the starting item in the collection for the request. If not specified, the default is an empty page token, which means that the output begins with the first item in the collection and returns up to **page_size** items or the end of the collection if the size of the collection is smaller than **page_size**.

Note: The page token is base64-encoded. The contents of the token might change at any time, so decoded contents must not be relied upon.

- The IBM Storage Scale native REST API endpoints that implement pagination return a **next_page_token** field when collections contain items beyond the page returned. If the **next_page_token** field is empty, the request has reached the end of the collection.

Examples

- The following example displays the **LIST** request with **page_size =1** query parameter for two total items in the collection:

```
curl -X 'GET' 'https://stgbladeh99:8000/scalemgmt/v3/filesystems/fs0/filesets?
page_size=1' -H 'accept: application/json' -k
{
  "filesets": [
    {
      "name": "root",
      "id": 0,
      "status": "FILESETSTATUS_LINKED",
      "path": "/fs0"
    }
  ],
  "next_page_token": "MQ=="
}
```

- The following example displays the **LIST** request with **page_size =** and **page_token=MQ==** query parameter for two total items in collection:

```
curl -X 'GET' 'https://example.com:8000/scalemgmt/v3/filesystems/fs0/filesets?
page_size=1&page_token=MQ%3D%3D' -H 'accept: application/json' -k
{
  "filesets": [
    {
      "name": "fset1",
      "id": 1,
      "status": "FILESETSTATUS_UNLINKED",
      "path": "--"
    }
  ]
}
```



```
]
}
```

- The following examples show the **LIST** request for no query parameters for two total items in the collection:

```
curl -X 'GET' \
  'https://example.com:8000/scalemgmt/v3/filesystems/fs0/filesets' \
  -H 'accept: application/json' -k
{
  "filesets": [
    {
      "name": "root",
      "id": 0,
      "status": "FILESETSTATUS_LINKED",
      "path": "/fs0"
    },
    {
      "name": "fset1",
      "id": 1,
      "status": "FILESETSTATUS_UNLINKED",
      "path": "--"
    }
  ]
}
```

scalectl command

All commands that implement pagination support following flags to control how the command paginates through the collection:

--no-pagination

Returns only one page of results. Does not follow **next_page_token**.

--page-size int

Specifies the number of items per API request.

--max-items int

Specifies the maximum number of items to list.

--page-token string

Specifies the token to retrieve the next set of items. If the number of items (**--max-items**) in the output is fewer than the total number of items in the collection, the response includes a **next_page_token** value, which can be used in a subsequent request.

Examples

- To list all items in the collection, issue the following command:

```
scalectl fileset list fs0
```

A sample output is as follows:

Fileset name	Id	Status	Path
root	0	link	/fs0
fset1	1	unlink	--

- To list at most one item in the collection, issue the following command:

```
scalectl fileset list fs0 --max-items 1
```

A sample output is as follows:

Fileset name	Id	Status	Path
root	0	link	/fs0

- To disable the pagination, issue the following command:

```
scalectl fileset list fs0 --no-pagination --page-size 1
```

A sample output is as follows:

```
Fileset name | Id | Status | Path
=====
root         | 0  | link   | /fs0
NEXT PAGE TOKEN: MQ==
```

- To follow the page token, issue the following command:

```
scalectl fileset list fs0 --no-pagination --page-size 1 --page-token MQ==
```

A sample output is as follows:

```
Fileset name | Id | Status | Path
=====
fset1        | 1  | unlink | --
NEXT PAGE TOKEN: <END OF LIST>
```

Views and fields options in IBM Storage Scale native REST API

Describes the views and fields options when listing the resources in IBM Storage Scale native REST API.

When listing individual resources through the **GET** method or listing a collection of resources with the **LIST** method, you might be interested only in a subset of the fields rather than the full output. IBM Storage Scale native REST API provides views and fields to customize responses and optimize performance.

View

Views return a predefined subset of the fields in the resource. Endpoints and commands that support views accept a **view** option as input. In the API request, specify the view as a string through the URL query parameters. In the **scalectl** command, specify the view by using the **--view** flag. The **view** option is an enumerated type that determines which fields are displayed. For example, the **fileset list** or **get** endpoints support views.

Examples

- The following example displays the attributes view with the **scalectl** command:

```
scalectl fileset list fs0 --view attributes
```

A sample output is as follows:

```
Fileset name | Id | Root Inode | Status | Path | Creation time | Comment | Inode Space | Max Inodes | Allocated inodes | Change flag | Perm Inherit Flag | IAM Mode
=====
root         | 0  | 3          | link   | /fs0 | 2023-10-18T15:42:53Z | root fileset | 0           | 286720     | 286720          | chmodAndSetAcl | inheritAclOnly | off
fset1        | 1  | 106496     | unlink | --   | 2023-10-23T21:32:56Z |              | 0           |             | 0               | chmodAndSetAcl | inheritAclOnly | off
```

- The following example displays the basic, which is the default view with the API request:

```
curl -X 'GET' \
'https://example.com:8000/scalemgmt/v3/filesystems/fs0/filesets?view=BASIC' \
-H 'accept: application/json'
```

A sample output is as follows:

```
{
  "filesets": [
```

```

    {
      "name": "root",
      "id": 0,
      "status": "FILESETSTATUS_LINKED",
      "path": "/fs0"
    },
    {
      "name": "fset1",
      "id": 1,
      "status": "FILESETSTATUS_UNLINKED",
      "path": "--"
    }
  ]
}

```

Fields

The **--fields** option allows more control over the specific fields that are returned in a response. If the **--fields** option is nonempty, the response includes only the fields that are specified by this option. For example, the file system endpoints support a **fields** option.

The **--fields** option supports both JSON field names and pretty-printed field names.

Note:

- IBM Storage Scale native REST API uses Protocol Buffers to define resources. Protocol Buffers do not guarantee field order in serialized output, and the field order might vary across repeated invocations. For more information about the field ordering, see [Field Order](#) in the Protocol Buffers documentation.
- The JSON response generated from the Protocol Buffers schema might not exactly match the output of the corresponding mm-command with the **-Y** option. Differences might include field names, structure, or field ordering.

Examples

- The following example shows how to use **--fields** option with the pretty-printed field names in the **scalectl** command:

```
scalectl filesystem get fs0 --fields "Filesystem name",uid,"Default Mount Point"
```

A sample output is as follows:

attribute	value
Filesystem name	fs0
Uid	4ED10B09:652FFCE9
Default Mount Point	/fs0

- The following example shows how to use **--fields** option with the JSON field names in the **scalectl** command:

```
scalectl filesystem get fs0 --fields name,uid,default_mount_point
```

A sample output is as follows:

attribute	value
Filesystem name	fs0
Uid	4ED10B09:652FFCE9
Default Mount Point	/fs0

- The following example shows how to use **--fields** option with both pretty-printed and JSON field names in the **scalectl** command:

```
scalectl filesystem list --fields "Inode Size",logfile_size,"Auto Inode Limit",local_version
```

A sample output is as follows:

attribute	value
Filesystem name	fs1
Inode Size	4096
Logfile Size	33554432
Auto Inode Limit	no
Local Version	37.00 (5.2.3.0)

- The following example displays getting only the file system name, uid, and default mount point with the REST API:

```
curl -X 'GET' 'https://example.com:8000/scalemgmt/v3/filesystems/fs0?fields=name&fields=uid&fields=default_mount_point' -H 'accept: application/json'
```

A sample output is as follows:

```
{
  "name": "fs0",
  "uid": "4ED10B09:652FFCE9",
  "default_mount_point": "/fs0"
}
```

Base64-encoded fields in IBM Storage Scale native REST API

When you use the IBM Storage Scale native REST API, some fields return **Base64** encoded strings.

The REST API returns data in JSON format, which supports only UTF-8 strings. To represent binary data or non-UTF-8 characters, the REST API uses repeated byte fields instead of string fields. These byte values are base64-encoded.

When a field is a repeated byte fields:

- For Rest API, it returns a base64-encoded string.
- For **scalectl** with **--json** flag, it returns a base64-encoded string.
- For **scalectl** without **--json** flag, by default, it decodes and displays human-readable text by default by using pretty printing.

For REST API and **scalectl --json** output, you can decode base64-encoded strings to view human-readable text with the inline decoding **jq** parameter. For example, **scalectl policy get fs1 --json** command returns a base64-encoded strings as follows:

```
"policy_contents":
"LyogREVGVQVVMVCAqLwovKiBTdG9yZSBkYXRhIGluIHRoZSBmaXJzdCBkYXRhIHBvb2wgb3Igc3lzdGVtIHBvb2wgi8K"
```

To decode the string inline, issue the following command:

```
scalectl policy get fs1 --json | jq -r '.policy_contents | @base64d'
```

A sample decoded output is as follows:

```
/* DEFAULT */
/* Store data in the first data pool or system pool */
```

Alternatively, use the **base64** utility on Linux to decode the base64-encoded strings as follows:

```
echo
"LyogREVGVQVVMVCAqLwovKiBTdG9yZSBkYXRhIGluIHRoZSBmaXJzdCBkYXRhIHBvb2wgb3Igc3lzdGVtIHBvb2wgi8K"
| /usr/bin/base64 -d
```

A sample decoded output is as follows:

```
/* DEFAULT */
/* Store data in the first data pool or system pool */
```

Update masks in IBM Storage Scale native REST API

Describes the update masks option used in the IBM Storage Scale native REST API.

Update masks specify the fields to update in an update request. When updating a resource, you might want only to update a subset of the fields. API update requests accept the resource to be updated as part of the body. By using update masks, you can instruct the server to ignore certain fields even if they are modified.

Examples

The following examples show the usage of update mask where the request is sent to update only the fileset name:

1. To get the details of the fileset *fset1*, run the following request:

```
curl -X 'GET' \ 'https://example.com:8000/scalemgmt/v3/filesystems/fs0/filesets/fset1?view=ATTRIBUTES' \ -H 'accept: application/json'
```

A sample output is as follows:

```
{
  "name": "fset1",
  "id": 1,
  "root_inode": "106496",
  "status": "FILESETSTATUS_UNLINKED",
  "path": "--",
  "create_time": "2023-10-23T21:32:56Z",
  "comment": "12345",
  "inode_space": 0,
  "alloc_inodes": "0",
  "perm_change_flag": "PERMCHANGEFLAG_CHMODANDSETACL",
  "perm_inherit_flag": "PERMINHERITFLAG_INHERIT_ACL_ONLY",
  "iam_mode": "IAMMODE_OFF",
  "afm_associated": false,
  "inode_space_designation": "root"
}
```

2. To update the fileset name from *fset1* to *fset2*, run the following request:

```
curl -X 'PATCH' 'https://example.com:8000/scalemgmt/v3/filesystems/fs0/filesets/fset1'?fields=name -H 'accept: application/json' -H 'Content-Type: application/json' -d '{
  "comment": "",
  "name": "fset2"
}'
```

A sample output is as follows:

```
{
  "name": "fset2",
  "fs_name": "fs0",
  "id": 1,
  "root_inode": "106496",
  "status": "FILESETSTATUS_UNLINKED",
  "path": "--",
  "create_time": "2023-10-23T21:32:56Z",
  "comment": "sample comment",
  "inode_space": 0,
  "is_inode_space_owner": false,
  "alloc_inodes": "0",
  "perm_change_flag": "PERMCHANGEFLAG_CHMODANDSETACL",
  "prevent_snapshot_restore": false,
  "perm_inherit_flag": "PERMINHERITFLAG_INHERIT_ACL_ONLY",
  "fal_status": "FALSTATUS_DEFAULT",
  "iam_mode": "IAMMODE_OFF",
  "inode_space_designation": "root"
}
```

List of IBM Storage Scale native REST API endpoints

The administration daemon supports a REST API that allows applications to communicate with it directly, without the use of **scalectl** command-line interface. The following table lists the currently supported REST API operations for each type of resource and resource element. The first and second columns list the resource type and the operation. The third column lists the resource and resource element terms that appear in the HTTPS request. The braces `{}` are a notation for indicating a resource element in text and do not appear in the actual HTTPS request. The fourth column lists the HTTP method that is called to do the operation.

Note: You can access the documentation corresponding to each API endpoint from the OpenAPI specification at `https://x.x.x.x:46443/openapi/`, where x.x.x.x is the public IP address of the API server node. Ensure that the trailing slash (/) is included in the URL.

Table 17. Operations supported for resources and resource elements in API endpoints			
Resource type	Operation	Resource/{element}	Method
API-health	Lists the health status of the IBM Storage Scale native REST API services on nodes.	/scalemgmt/v3/apihealth	GET
	Lists the health status of the IBM Storage Scale native REST API services on the specified node.	/scalemgmt/v3/apihealth/{node_name}	GET
Authorization	Checks whether the current user is authorized to perform an action on a resource in the domain.	/scalemgmt/v3/authorization/cani	POST
	Checks whether the specified user is authorized to perform an action on a resource in the domain.	/scalemgmt/v3/authorization/cani:impersonate	GET
	Lists information about all domains.	/scalemgmt/v3/authorization/domains	GET
	Creates a role-based access control (RBAC) domain.	/scalemgmt/v3/authorization/domains	POST
	Lists information about the RBAC domain for the specified domain name.	/scalemgmt/v3/authorization/domains/{name}	GET
	Deletes the domain for the specified name.	/scalemgmt/v3/authorization/domains/{name}	DELETE
	Updates the domain for the specified name.	/scalemgmt/v3/authorization/domains/{name}	PATCH
	Updates the customer-defined policy rule set.	/scalemgmt/v3/authorization/module	PATCH
	Gets details of the CLI commands that are issued on the cluster.	/scalemgmt/v3/authorization/module	GET
Cluster	Lists information about the cluster.	/scalemgmt/v3/clusters	GET
	Creates the IBM Storage Scale cluster.	/scalemgmt/v3/clusters	POST
	Migrates a classic IBM Storage Scale cluster to the use of IBM Storage Scale native REST API.	/scalemgmt/v3/clusters:migrate	POST
Configuration	Lists information about the current configuration data for the administration daemon that is stored on a disk.	/scalemgmt/v3/config/admin	GET
	Lists information about the configuration data of the administration daemon for the specified attribute.	/scalemgmt/v3/config/admin/{name}	GET
	Updates the configuration value of an administration daemon.	/scalemgmt/v3/config/admin:batchUpdate	PATCH
	Lists the current configuration data for the I/O daemon that is stored on a disk.	/scalemgmt/v3/config/cluster	GET
	Lists information about the I/O daemon configuration value for the specified attribute.	/scalemgmt/v3/config/cluster/{name}	GET
	Updates the configuration value of an IO daemon.	/scalemgmt/v3/config/cluster:batchUpdate:	PATCH
Diagnostics	Lists mmfsd version of the node.	/scalemgmt/v3/nodes/{node}/diagnostics/version	GET

Table 17. Operations supported for resources and resource elements in API endpoints (continued)

Resource type	Operation	Resource/{element}	Method
File system	Lists information about all IBM Storage Scale file systems in an IBM Storage Scale cluster.	/scalemgmt/v3/filesystems	GET
	Creates an IBM Storage Scale file system.	/scalemgmt/v3/filesystems	POST
	Lists information about the specified IBM Storage Scale file system.	/scalemgmt/v3/filesystems/{filesystem}	GET
	Lists the mount state of an IBM Storage Scale file system.	/scalemgmt/v3/filesystems/{filesystem}:mount	GET
	Mounts an IBM Storage Scale file system.	/scalemgmt/v3/filesystems/{filesystem}:mount	POST
	Unmounts an IBM Storage Scale file system.	/scalemgmt/v3/filesystems/{filesystem}:unmount	POST
	Deletes an IBM Storage Scale file system.	/scalemgmt/v3/filesystems/{name}	DELETE
	Updates the attributes of an IBM Storage Scale file system.	/scalemgmt/v3/filesystems/{name}	PATCH
	Mounts all IBM Storage Scale file systems.	/scalemgmt/v3/filesystems:mount	POST
	Unmounts all IBM Storage Scale file systems.	/scalemgmt/v3/filesystems:unmount	POST
File system disk	Lists information about all disks in an IBM Storage Scale file system	/scalemgmt/v3/filesystems/{filesystem}/disks	GET
	Adds disks in an IBM Storage Scale file system.	/scalemgmt/v3/filesystems/{filesystem}/disks	POST
	Lists information about the specified disk in an IBM Storage Scale file system.	/scalemgmt/v3/filesystems/{filesystem}/disks/{disk_name}	GET
	Deletes a disk in an IBM Storage Scale file system.	/scalemgmt/v3/filesystems/{filesystem}/disks/{disk_name}	DELETE
	Add a batch of disks from an IBM Storage Scale file system.	/scalemgmt/v3/filesystems/{filesystem}/disks:batchAdd	POST
	Deletes a batch of disks in an IBM Storage Scale file system.	/scalemgmt/v3/filesystems/{filesystem}/disks:batchDelete	POST
	Lists information about the disk quorum.	/scalemgmt/v3/filesystems/{filesystem}/disksquorum	GET
Fileset	Lists filesets in an IBM Storage Scale file system.	/scalemgmt/v3/filesystems/{filesystem}/filesets	GET
	Creates an IBM Storage Scale fileset.	/scalemgmt/v3/filesystems/{filesystem}/filesets	POST
	Gets details of the fileset for the specified name.	/scalemgmt/v3/filesystems/{filesystem}/filesets/{fileset_name}	GET
	Deletes the fileset for the specified name.	/scalemgmt/v3/filesystems/{filesystem}/filesets/{fileset_name}	DELETE
	Updates attributes of the fileset for the specified name.	/scalemgmt/v3/filesystems/{filesystem}/filesets/{fileset_name}	PATCH
	Retrieves the fileset inode usage, data usage, or both.	/scalemgmt/v3/filesystems/{filesystem}/filesets/{fileset_name}/usage	GET
	Links an IBM Storage Scale fileset.	/scalemgmt/v3/filesystems/{filesystem}/filesets/{fileset_name}:link	POST
	Unlinks an IBM Storage Scale fileset.	/scalemgmt/v3/filesystems/{filesystem}/filesets/{fileset_name}:unlink	POST
Manager	Updates the cluster manager node.	/scalemgmt/v3/clusters/manager	POST
	Update the manager node of an existing file system.	/scalemgmt/v3/filesystems/{name}/manager	POST
Node	Adds a node to the IBM Storage Scale cluster.	/scalemgmt/v3/nodes	POST
	Starts the mmfsd service on a node.	/scalemgmt/v3/nodes:start	POST
	Get the status of the node.	/scalemgmt/v3/nodes/status	GET
	Stops mmfsd on a cluster node.	/scalemgmt/v3/nodes:stop	POST
	Lists the configuration details of the node.	/scalemgmt/v3/nodes/config	GET

Table 17. Operations supported for resources and resource elements in API endpoints (continued)

Resource type	Operation	Resource/{element}	Method
Node class	Lists all node classes that are registered to the cluster.	/scalemgmt/v3/nodeclasses	GET
	Creates user-defined node classes.	/scalemgmt/v3/nodeclasses	POST
	Lists the members of the specified node class.	/scalemgmt/v3/nodeclasses/{name}	GET
	Deletes an existing node class.	/scalemgmt/v3/nodeclasses/{name}	DELETE
	Updates an existing node class.	/scalemgmt/v3/nodeclasses/{name}	PATCH
NSD	Lists information about all network-shared disks (NSDs) that are part of the system.	/scalemgmt/v3/nsds	GET
	Creates a network-shared disk in an IBM Storage Scale cluster.	/scalemgmt/v3/nsds	POST
	Lists information about the specified NSD.	/scalemgmt/v3/nsds/{nsd_name}	GET
	Deletes the specified NSD.	/scalemgmt/v3/nsds/{nsd_name}	DELETE
	Updates an existing NSD.	/scalemgmt/v3/{nsd_name}	PATCH
	Creates one or more NSDs that are used by IBM Storage Scale.	/scalemgmt/v3/nsds:batchCreate	POST
	Deletes one or more network-shared disks (NSD) that are used by IBM Storage Scale.	/scalemgmt/v3/nsds:batchDelete	POST
Operations	Lists information about all long-running operations (LROs).	/scalemgmt/v3/operations	GET
	Lists details of the LRO for the specified ID.	/scalemgmt/v3/operations/{id}	GET
	Delete the LRO for the specified ID.	/scalemgmt/v3/operations/{id}	DELETE
	Gets the output details of the LRO.	/scalemgmt/v3/operations/{id}/output	GET
	Cancels the running LRO.	/scalemgmt/v3/operations/{id}:cancel	POST
AFM	Lists AFM state information for all filesets in a file system.	/scalemgmt/v3/filesystems/{filesystem}/filesets/afm/state	GET
	Lists AFM state information for a specific fileset.	/scalemgmt/v3/filesystems/{filesystem}/filesets/{fileset}/afm/state	GET
	Checks for dirty data in an AFM fileset.	/scalemgmt/v3/filesystems/{filesystem}/filesets/{fileset}/afm/checkdirty	GET
	Checks for uncached data in an AFM fileset.	/scalemgmt/v3/filesystems/{filesystem}/filesets/{fileset}/afm/checkuncached	GET
	Resets the local AFM fileset to a clean state.	/scalemgmt/v3/filesystems/{filesystem}/filesets/{fileset}/afm:resetlocal	PATCH
	Sets the AFM fileset to local mode.	/scalemgmt/v3/filesystems/{filesystem}/filesets/{fileset}/afm:setlocal	PATCH
	Flushes the AFM queue for a fileset.	/scalemgmt/v3/filesystems/{filesystem}/filesets/{fileset}/afm/flushqueue	POST
	Resumes requeued AFM operations for a fileset.	/scalemgmt/v3/filesystems/{filesystem}/filesets/{fileset}/afm/resumerequeued	POST
	Resynchronizes an AFM fileset with its source.	/scalemgmt/v3/filesystems/{filesystem}/filesets/{fileset}/afm/resync	POST
	Starts AFM operations for a fileset.	/scalemgmt/v3/filesystems/{filesystem}/filesets/{fileset}/afm:start	POST
	Stops AFM operations for a fileset.	/scalemgmt/v3/filesystems/{filesystem}/filesets/{fileset}/afm:stop	POST

Table 17. Operations supported for resources and resource elements in API endpoints (continued)

Resource type	Operation	Resource/{element}	Method
AFM COS	Deletes cloud object storage keys for an AFM COS bucket.	/scalemgmt/v3/buckets/{bucket_name}/afmcos/delcoskeys	DELETE
	Gets cloud object storage keys for an AFM COS bucket.	/scalemgmt/v3/buckets/{bucket_name}/afmcos/getcoskeys	GET
	Deletes data from cloud object storage for an AFM COS fileset.	/scalemgmt/v3/filesystems/{filesystem}/filesets/{fileset}/afmcos/delete	POST
	Downloads data from cloud object storage to an AFM COS fileset.	/scalemgmt/v3/filesystems/{filesystem}/filesets/{fileset}/afmcos/download	POST
	Evicts data from an AFM COS fileset.	/scalemgmt/v3/filesystems/{filesystem}/filesets/{fileset}/afmcos/evict	POST
	Reconciles an AFM COS fileset with cloud object storage.	/scalemgmt/v3/filesystems/{filesystem}/filesets/{fileset}/afmcos/reconcile	POST
	Sets cloud object storage keys for an AFM COS bucket.	/scalemgmt/v3/buckets/{bucket_name}/afmcos/setcoskeys	PUT
	Configures AFM COS settings for a fileset.	/scalemgmt/v3/filesystems/{filesystem}/filesets/{fileset}/afmcos/configure	PUT
	Uploads data from an AFM COS fileset to cloud object storage.	/scalemgmt/v3/filesystems/{filesystem}/filesets/{fileset}/afmcos/upload	PUT
Quota	Lists quota information for a fileset.	/scalemgmt/v3/filesystems/{filesystem}/filesets/{fileset}/quotas	GET
	Sets quota limits for a fileset.	/scalemgmt/v3/filesystems/{filesystem}/filesets/{fileset}/quotas	PUT
	Checks quota usage for a fileset.	/scalemgmt/v3/filesystems/{filesystem}/filesets/{fileset}/quotas	POST
	Lists quota information for a file system.	/scalemgmt/v3/filesystems/{filesystem}/quotas	GET
	Sets quota limits for a file system.	/scalemgmt/v3/filesystems/{filesystem}/quotas	PUT
	Checks quota usage for a file system.	/scalemgmt/v3/filesystems/{filesystem}/quotas	POST
Policy	Lists information about the file system policy.	/scalemgmt/v3/filesystems/{filesystem}/policy	GET
	Updates the file system policy.	/scalemgmt/v3/filesystems/{filesystem}/policy	PATCH
Remote cluster	Lists information about remote clusters.	/scalemgmt/v3/clusters/remote	GET
	Adds an owning cluster to the set of remote clusters known to the accessing cluster.	/scalemgmt/v3/clusters/remote	POST
	Authorizes an accessing cluster to access resources on the owning cluster.	/scalemgmt/v3/clusters/remote/authorized	POST
	Deletes an accessing cluster authorization from the owning cluster.	/scalemgmt/v3/clusters/remote/authorized/{name}	DELETE
	Retrieves details of a remote cluster.	/scalemgmt/v3/clusters/remote/{name}	GET
	Deletes an owning cluster definition from the accessing cluster.	/scalemgmt/v3/clusters/remote/{name}	DELETE
	Updates information associated with a remote cluster in this cluster.	/scalemgmt/v3/clusters/remote/{name}	PATCH
	Refreshes the information of an owning cluster on the accessing cluster.	/scalemgmt/v3/clusters/remote/{name}/refresh	POST
Remote file system	Adds a remote file system.	/scalemgmt/v3/filesystems/remote	POST
	Deletes a remote file system.	/scalemgmt/v3/filesystems/remote/{filesystem}:	DELETE
	Updates the information associated with a remote file system.	/scalemgmt/v3/filesystems/remote/{filesystem}	PATCH

Table 17. Operations supported for resources and resource elements in API endpoints (continued)

Resource type	Operation	Resource/{element}	Method
Snapshot	Lists global snapshots in an IBM Storage Scale file system.	/scalemgmt/v3/filesystems/{filesystem}/snapshots	GET
	Creates a global snapshot in an IBM Storage Scale file system.	/scalemgmt/v3/filesystems/{filesystem}/snapshots	POST
	Gets information for an IBM Storage Scale global snapshot.	/scalemgmt/v3/filesystems/{filesystem}/snapshots/{snapshot_name}	GET
	Delete a global snapshot in an IBM Storage Scale file system.	/scalemgmt/v3/filesystems/{filesystem}/snapshots/{snapshot_name}	DELETE
	Displays the current snapshot directory settings.	/scalemgmt/v3/filesystems/{filesystem}/snapshots:snapdir.	GET
	Deletes one or more global snapshots.	/scalemgmt/v3/filesystems/{filesystem}/snapshots:batchDelete	POST
	Creates one or more fileset snapshots.	/scalemgmt/v3/filesystems/{filesystem}/filesets/snapshots:batchCreate	POST
	Deletes one or more fileset snapshots.	/scalemgmt/v3/filesystems/{filesystem}/filesets/snapshots:batchDelete	POST
	Lists fileset snapshots in an IBM Storage Scale fileset.	/scalemgmt/v3/filesystems/{filesystem}/filesets/{fileset}/snapshots	GET
	Creates an IBM Storage Scale fileset snapshot.	/scalemgmt/v3/filesystems/{filesystem}/filesets/{fileset}/snapshots	POST
	Gets details about the fileset snapshot in an IBM Storage Scale fileset for the specified snapshot name.	/scalemgmt/v3/filesystems/{filesystem}/filesets/{fileset}/snapshots/{snapshot_name}	GET
	Deletes the fileset snapshot in an IBM Storage Scale fileset for the specified snapshot name.	/scalemgmt/v3/filesystems/{filesystem}/filesets/{fileset}/snapshots/{snapshot_name}	DELETE
Storage pool	Lists details of storage pools of an IBM Storage Scale file system.	/scalemgmt/v3/filesystems/{filesystem}/storagepools	GET
	Lists details of the specified storage pool.	/scalemgmt/v3/filesystems/{filesystem}/storagepools/{pool_name}	GET
	Updates an existing storage pool.	/scalemgmt/v3/filesystems/{filesystem}/storagepools/{pool_name}	PATCH
Troubleshooting	Reads persistent reserve registration key values from a device.	/scalemgmt/v3/troubleshooting/persistentReserve/{device}/keys	GET
	Clears the persistent reserve registration key values from a device.	/scalemgmt/v3/troubleshooting/persistentReserve/{device}/keys	DELETE
	Deletes the NSD volume ID from a device.	/scalemgmt/v3/troubleshooting/nsds/clearID	DELETE
Version	Lists the version of the scaleadm service.	/scalemgmt/v3/version	GET
For more information about the endpoints, see <i>IBM Storage Scale native REST API endpoints</i> in the <i>IBM Storage Scale: Command and Programming Reference Guide</i> .			

Limitations of IBM Storage Scale native REST API

Describes the limitations of IBM Storage Scale native REST API.

IBM Storage Scale native REST API has the following limitations:

- The native REST API is an evolving feature. To ensure access to the latest updates and capabilities, plan to regularly apply the latest PTFs and upgrade to the latest major releases of IBM Storage Scale.
- This feature is not supported on ARM architecture, Windows and AIX operating system nodes within a cluster. The entire IBM Storage Scale cluster must consist of supported operating system and architecture nodes. Do not intermix unsupported nodes with native REST API clusters. For more information about supported operating systems and architectures, see *Planning for IBM Storage Scale native REST API* section in the *IBM Storage Scale: Concepts, Planning, and Installation Guide*.
- The administration daemon-based remote execution and transfer capabilities are used for control plane operations and do not eliminate the requirement for root access during installation, upgrade, or support activities.

- When the administration daemon starts or restarts, there might be a brief delay before it can serve requests. As a result, the first CLI or REST API request might be rejected if the daemon startup routines are still completing. This behavior is similar to the delay that occurs before the I/O daemon can serve all requests.

The administration daemon restart can introduce this delay in the following situations:

- Immediately after installation of the native REST API rpm (the administration daemon starts at the end of the installation process).
- After nodes are added to the cluster (the remote node administration daemon internal services restart on the added nodes).
- After you manual start or restart the administration daemon by using the **systemd** or a similar command
- After node deletion (the administration daemon restarts on the node that is deleted).

A common error status that is returned is **UNAVAILABLE**, but other errors might also occur. If a request fails after the administration daemon restarts under these conditions, wait briefly (a few seconds) and retry the command. Repeat this process as needed.

- IBM Storage Scale native REST API does not support the **mmdsh** function. Remote operations across nodes must be performed by using tools such as SSH or SCP, or by using the administration daemon-based remote execution and transfer capabilities where applicable.
- If an IBM Storage Scale cluster has the native REST API enabled, any remote file system that is mounted must be at file system version 5.2.3.0 or later.
- If an IBM Storage Scale cluster is mounting a remote file system, the file system manager on the owning cluster must conform to the native REST API architecture and operating system requirements. Specifically, the file system manager must not be an AIX node, Windows node, or ARM node.
- All **scaleadmd** services start only on quorum nodes. The nonquorum nodes can start only a subset of **scaleadmd** services. Most requests are processed only by quorum nodes.
- The **scalectl cluster create** command does not currently support the **config_file** and **profile file** options that are available with the **mmcrcluster** command. To set the required options, use the **mmchcluster** or **mmchconfig** commands after the native REST API cluster is created. Alternatively, create the cluster with **mmcrcluster** using the configuration or profile files, and then migrate the cluster to the native REST API feature.
- Behavior of the **mmaddnode** command depends on the cluster level.
 - **For clusters at or above release 6.0.1.0:** The **mmaddnode** command is supported after the cluster is configured for the native REST API (created or migrated). When you add a node, the system automatically migrates the node during the operation to maintain coexistence between mm-commands and the native REST API. Before you add a node, ensure that the native REST API prerequisites are met on the new node, including configuring node identity for node-to-node access.
 - **For clusters below release 6.0.1.0:** When a native REST API cluster is created or an existing IBM Storage Scale cluster is migrated to the native REST API framework, the **mmaddnode** command is not supported unless the following conditions are met:
 - The IBM Storage Scale cluster software version is at least 6.0.0.0. To verify the version, run the **mmdiag --version** command after the I/O daemon starts.
 - Remote shell (**scaleadmremoteexecute**) and copy services (**scaleadmremotetransfer**) are enabled. Run the **mmlscluster** or **scalectl cluster list** command to verify the services.
- When you migrate a classic IBM Storage Scale cluster to the native REST API by using the **scalectl cluster migrate** command or REST API, you must restart the I/O daemons if they are running to complete the migration.
- The **scalectl node add** command displays the **FPO** option as part of the **--license-type** parameter, but this option is not supported. This option is deprecated and is not supported in native REST API.

- The **scalectl config cluster** command and the **/scalemgmt/v3/config/cluster** REST APIs do not support all attributes that are available with the IBM Storage Scale **mmchconfig** command. Unsupported attributes include, but are not limited to: **release**, **minrelease**, **profile**, **exitscript.*elect**, and **exitscript.*.members**.
- When you use the native REST API with IPv6, the native remote transfer protocol (**scaleadmremotetransfer**) is temporarily not supported. After you create a cluster by using the native REST API with IPv6, change the protocol from **scaleadmremotetransfer** by using the **mmchcluster** command.
- You must continue to use the existing mm-commands for the functions that are not yet implemented in the native REST API. The following features are not implemented through the scalectl CLI or native REST API endpoints in IBM Storage Scale 5.2.3.0:
 - Backup of a file system
 - Call home
 - Clustered watch folder
 - Encryption of data at rest
 - File audit logging
 - HDFS
 - IBM Storage Scale Container Native deployment
 - IBM Storage Scale CSI Driver
 - Installation toolkit
 - Integrated protocols (CES)
 - Performance monitoring
 - Quotas
 - Transparent cloud tiering
- If the following errors occur immediately after upgrading the native REST API, restart the administration daemons on the nodes that are contacted during the endpoint operation:
 - During **mmdelnode**, an error is returned that includes the message:

```
rpc error: code = Unavailable desc = connection error: desc = "transport: Error while dialing: dial unix..."
```

- During **scalectl node add** or a node add REST API request to add a node that is designated as a quorum node, an error is returned that includes the message:

```
failed to run 'mmccr chnodes' to update the cluster quorum nodes
```

Introduction to transparent cloud tiering (discontinued)

transparent cloud tiering is a feature of IBM Storage Scale that provides a native cloud storage tier. cloud services have the following two components:

- Transparent cloud tiering
- cloud data sharing

transparent cloud tiering allows data center administrators to free up IBM Storage Scale storage capacity, by moving out cooler data to the cloud storage, reducing capital and operational expenditures. Tiering can also be used to archive an extra copy of your data by using pre-migration, a function that copies data rather than moving it. The transparent cloud tiering feature leverages the existing ILM policy query language semantics available in IBM Storage Scale, and administrators can define policies to tier data to cloud storage. On an IBM Storage Scale cluster with multiple storage tiers configured, this external cloud storage can be used as the cooler storage tier to store infrequently accessed data from a cool storage pool. *For performance reasons, it is recommended not to move any active or hot data to this*

external storage pool, as it drives excessive data traffic on the transparent cloud tiering which in turn can cause delays, leading to problems like application timeouts. Copying hot data to the cloud by using the pre-migration function is acceptable when the file is not likely to be updated or deleted soon.

cloud data sharing is an IBM Storage Scale cloud services that allows a way to set up sharing between IBM Storage Scale and various types of object storage, including IBM Cloud Object Storage. Furthermore, for export of data to object storage the service can be invoked from an ILM policy to allow for periodic sharing of data based on when the policy is run. For import, you can specify a list of files that you can use to move object storage (by using your IBM Cloud data sharing service) into the IBM Storage Scale cluster.

This is useful when data needs to be distributed across multiple domains or when large amounts of data need to be transferred.

You can associate your cloud services file systems and file sets with a cloud services node group. Up to four node groups are supported in a cluster to allow for scaling (of the solution). For each file system or file set, you can associate up to two cloud storage tiers. For example, you could define an IBM Cloud Object Storage service locally to act as your cool tier. You could use IBM's Public Cloud Object Storage (on IBM Cloud) for your cold tier. The following diagram is an example of a cloud services configuration that is on a cluster:

Figure 29 on page 241 illustrates these features.




Figure 29. transparent cloud tiering and cloud data sharing features

Note:

- For performance reasons, the median file size to be migrated to the cloud tier must be greater than 1 MB. Migration is supported for file size less than one MB, but performance is slower because of the overhead that is associated with small files.
- For transparent cloud tiering, data on the Cloud Object Storage is opaque and cannot be accessed directly by applications. All I/O operations must happen through your IBM Storage Scale system.

- transparent cloud tiering works with IBM Storage Scale on multi-site stretch clusters to allow continued access to cloud storage data even if failure of an entire site.
- For applications that require tiering of data that scales beyond 100 million files per file system, you can create extra containers on cloud storage to hold your files. For applications that require more I/O bandwidth than a single maximum node group (four nodes) is capable of, you can create extra node groups.

Note: You must create a new container when the existing container has approximately 100,000,000 files.

Unsupported use cases

Transparent cloud tiering does not support the following use cases:

- Using transparent cloud tiering to migrate or recall hot (active) data.
- Using transparent cloud tiering as a backup mechanism.
- Using transparent cloud tiering to restore data in disaster recovery scenarios.

Note: IBM Storage Protect cloud-container storage pools can be used to back up data to cloud storage providers.

Note: To enable cloud services nodes, you must first enable the transparent cloud tiering feature. This feature provides a new level of storage tiering capability to IBM Storage Scale customers. Contact your IBM Client Technical Specialist (or send an email to <mailto:scale@us.ibm.com>) to review your use case of the transparent cloud tiering feature and to obtain the instructions to enable the feature in your environment.

How transparent cloud tiering works

This topic describes how you can use transparent cloud tiering (which tiers files to object storage) feature in the IBM Storage Scale cluster.

transparent cloud tiering is an IBM Storage Scale feature that allows the IBM Storage Scale cluster to tier files to object storage. The transparent cloud tiering service stores a file as two objects out in the cloud in two separate containers. These containers are also called vaults or buckets depending on which cloud storage vendor you use. One object holds the file's data and the other object holds the file's metadata. The metadata contains items such as ACLs, extended attributes, and other information that allows transparent cloud tiering to do a full restore of the file from the cloud tier.

When a file is migrated to the cloud, the metadata of the file is retained in the IBM Storage Scale file system. The only reason for copying the metadata to the cloud is to manage scenarios where a file is destroyed on the cluster for some reason (such as accidental deletion) and the file needs to be restored from the cloud.

For the usual case, only the file's data is recalled from the cloud. You can tier the files by migrating the data to the cloud. This means that the data is no longer on the IBM Storage Scale cluster and is stored only on the cloud. Such files are considered to be in the non-resident state. Alternatively, you can pre-migrate the data to the object storage so that the data is retained in the IBM Storage Scale cluster, but is also copied out to object storage so that the data is on both tiers. Such files are considered to be in the co-resident state.

transparent cloud tiering stores information about files that are migrated to the cloud tier in a database called *cloud directory*. A separate database or cloud directory is kept for every transparent cloud tiering container pair (the pair of containers that hold the file data and metadata). This database contains a list and versions of all files that are migrated to the cloud. The metadata contains items such as ACLs, extended attributes, and other information that allows transparent cloud tiering to do a full restore of the file from the cloud tier in the future.

Because some applications access the front end of the file frequently, there is an option by specifying "thumbnail-size" to choose how much of the front end of a file is to be kept in the file system when a file is migrated out to the object storage. Some applications such as Windows Explorer and Linux GNOME only access a very small amount of data on the front end as a part of directory services and other applications

such as media streamers might want to cache hundreds of megabytes. By specifying the size it is possible to efficiently accommodate both such applications.

transparent cloud tiering has maintenance activities that remove data from the cloud for files that are deleted or reverted. Backup and reconciliation services are performed to deal with disaster recovery and other unusual error clean-up cases. You can change the times and frequencies of these activities.

Additionally, if a file is migrated and only the metadata is changed later, a subsequent migration copies only the metadata to the cloud. This metadata references the original data object.

Note: transparent cloud tiering data is migrated to the cloud in a way that is not designed for direct use in the cloud. If data must be consumed in the cloud, you must consider the cloud data sharing service.

Data migration and configuration activities happen through the defined cloud services nodes. You can allow transparent recalls to be handled directly (and thus more efficiently and with better performance) by client nodes by setting up this option when installing the client.

How cloud data sharing works

This topic describes how the cloud data sharing feature works in the IBM Storage Scale cluster.

cloud data sharing allows you to import data from object storage into your IBM Storage Scale file system using the **import** command. You can also export data to the cloud for use in the cloud by using the **export** command. Optionally, there is a manifest that can be built that allows those sharing to very quickly ascertain what files are being shared.

Cloud data sharing is different from transparent cloud tiering because it is meant as a means of sharing and distributing data, whereas with transparent cloud tiering the data that is migrated to object storage is concealed so it can be consumed only by recalling the data back into the IBM Storage Scale system.

Note: There is no implicit consistency link that is established by the export or import of a file. The import or export is just a transfer of a file. Once the transfer is complete, if the original file is changed at one place it has no bearing on its copies. And if a copy of a file is changed or deleted it has no bearing on the original.

How exporting file system data to cloud storage works




Figure 30. Exporting file system data to a cloud storage tier

You can use the cloud data sharing service to export data to the cloud. Data can be exported manually, but typically it is done leveraging the ILM Policy Manager that comes with IBM Storage Scale. When data transfer requests are generated, they are distributed evenly between all nodes in the cloud services node group providing the cloud data sharing service.

Also, optionally as a part of the export request, a list of what was sent is also put in a manifest file that can be used to track what was exported. Export is very flexible because for each command the target container can be specified. Any container accessible by the given cloud account can be used. For each **export** command, you can also specify the target manifest file. The manifest file is not automatically stored in the cloud storage but rather remains stored within the cloud data sharing service. The manifest is exported on demand to the desired target location.

How importing object storage data into the IBM Storage Scale file system works




Figure 31. Importing object storage data into the file system

You can import data directly by providing a list of objects to be imported in the command itself or by providing a pointer to a manifest that contains the list of objects to be imported. The files are imported and placed in a target directory in the file system. Since the files are not yet present in the file system, there is no way to use the policy engine to initially import the files. You can also import the stub of the file and then use the policy engine to look at the file stub information that you use to import the desired data subset. You can delete the stubs that you are not interested in.

Note: The import of the stub only works if the data on the object storage that is being imported was originally created using the IBM Storage Scale export service so that the stub has the appropriate format. (Native object storage stubs are not mapped or normalized as of yet).

How to be aware of and exploit what data is shared

A manifest might be kept which lists information on objects that are in cloud storage available for sharing. When the cloud data sharing service exports files to the object storage, it can be directed to add entries on what it exported to a manifest. Other applications that generate cloud data can also generate a manifest. The location of where the manifest is stored and when it is exported to be looked at is the decision of the application – it is not automatic. A typical time to export the manifest to object storage

is in the same cron job as the policy invocation immediately following the policy execution completion so that it represents all the files that are exported by that policy run.

The `mmcloudmanifest` manifest utility program is available that can read a manifest and provide a comma-separated value (CSV) output stream of all the files in the manifest, or a subset of those files as provided by some simple filtering. This manifest utility is written in Python and can run separately from the cloud data sharing service most anywhere python can run. The utility is also built into cloud data sharing and can be used to get input for import operations.

Currently, there is no built-in way to do asynchronous notification or to have a database that tracks all updates at a scale larger than is reasonable with manifest files, but it is possible to add notification in the script where the export policy is invoked. Also, manifests can be readily consumed by a database (for example the Elastic Search, Logstash, Kibana ELK Stack) since the data structure is very simple. On larger deployments, such an approach with asynchronous notification that is built into the scripts (and uses a database) is worth consideration.

How a native cloud storage application can prepare a set of objects for sharing

There is a simple way for native cloud storage applications or other generators of cloud object data to generate a manifest. The manifest utility can build the manifest from a list of objects that it is passed.

How Write Once Read Many (WORM) storage works

This topic describes how you can use Write Once Read Many (WORM) solutions to take advantage of functions that are offered by IBM Storage Scale, transparent cloud tiering, and IBM Cloud Object Storage.

WORM (Write Once-Read Many) storage solutions leverage IBM Storage Scale immutability, transparent cloud tiering, and IBM Cloud Object Storage locked vaults. Essentially, you can create the immutable file sets on the IBM Storage Scale file system, and you can set files as immutable either through IBM Storage Scale commands or through POSIX interface. You can also set a retention period on the immutable files.

A locked vault on the IBM Cloud Object Storage cannot be deleted by the IBM Cloud Object Storage administrator, and its Access Control Lists (ACLs) cannot be changed. Additionally, you cannot rename it or enable the proxy settings. RSA private key and private certificate are used to create and access the locked vaults. transparent cloud tiering is configured with the RSA private key and private certificate to create locked vaults. Once configured, you can use configured private keys and certificates to use REST APIs against the Accesser nodes.

You can perform a migration and recall operation on immutable files. Since the immutable files cannot be deleted, the data on the IBM Cloud Object Storage locked vaults is only deleted according to locked vault deletion policy, which provides a WORM solution for compliance archival.

The following diagram provides an overview of the feature:




Figure 32. WORM storage overview

Supported cloud providers

This topic describes the cloud providers that the IBM Storage Scale supports.

The cloud services use the following cloud providers:

- IBM Cloud Object Storage (version 3.7.3.2, 3.8.2.19, 3.9, and 3.13.4.40)
- IBM Cloud Object Storage on IBM SoftLayer® and IBM Bluemix.
- OpenStack Swift with Swift version 2.13 and below
- Amazon Web Services S3
- Swift3 on OpenStack Swift version 2.13, 1.1 and above

Note: OpenStack Swift is supported only with Keystone version 2 only. You need to install Apache libcloud 1.3.0 for mmcloudmanifest to function. Run **pip install apache-libcloud==1.3.0** to install libcloud version 1.3.0.

- Microsoft Azure object storage service

Note: Cloud services move data to a cool tier as a block 'blob' only. Moving data to other tiers and other blob types is not supported.

Interoperability of transparent cloud tiering with other IBM Storage Scale features

This topic describes the interoperability of transparent cloud tiering with other IBM Storage Scale features.

IBM Spectrum Archive (LTFS) and IBM Storage Protect

Running IBM Spectrum Archive and transparent cloud tiering on the same file system is not supported.

However, both IBM Storage Protect and transparent cloud tiering can co-exist on the same systems if they are configured with different file systems.

It is advised not to enable the Data Management API (DMAPI) on the file system that is managed by the transparent cloud tiering.

AFM

Running transparent cloud tiering service on AFM Gateway nodes is not supported.

Data from AFM or AFM DR filesets cannot be accessed by transparent cloud tiering.

Using transparent cloud tiering on AFM home is not supported.

Multi-clusters

Multi-cluster support is limited to remotely mounted clients.

File Placement Optimizer (FPO)

Transparent cloud tiering is supported on FPO cluster. Recalled files might not have optimal placement.

IBM Storage Scale Object

Transparent cloud tiering can be configured on Object file sets.

Snapshots

Transparent cloud tiering should not be used to migrate/recall snapshots. Space that is contained in snapshots is not freed if a file is migrated to cloud object storage. Data tiered to the cloud is not designed for proper behavior relative to a snapshot. It is not recommended to take a snapshot on a file set or filesystem that contains data tiered by transparent cloud tiering service.

Sparse files

Transparent cloud tiering can be used to migrate/recall sparse files. However, sparseness is not maintained and on recall, full blocks are allocated.

Encryption

Transparent cloud tiering can be used with Scale file system level encryption feature (available in the Advanced and Data Management Editions). However, all the data that is migrated to the cloud object storage is migrated with the key that is configured for transparent cloud tiering. Essentially, when data is read from file system, the files get decrypted, and encrypted again at user space by transparent cloud tiering and pushed into the cloud storage.

Compression

Transparent cloud tiering can be used along with Scale file system level compression capability. Essentially, when data is read from the file system, the files are uncompressed, transparent cloud tiering push uncompressed, but by default encrypted, files onto the cloud storage.

CES (protocol services)

Transparent cloud tiering can co-exist with active or inactive NFS, SMB, or Object services on the CES nodes.

IBM Storage Protect

For the file systems that are managed by IBM Storage Protect system, ensure that hot data is backed up to IBM Storage Protect by using the **mmbackup** command, and as the data gets cooler, migrate them to the cloud storage tier. This ensures that the **mmbackup** command has already backed up the cooler files that are migrated to the cloud.

Elastic Storage Server (ESS)

Transparent cloud tiering cannot be deployed directly on ESS nodes. However, it can be deployed on X86 protocol nodes that can be attached to ESS. On a Power Linux cluster, a few X86 protocol nodes can be added and transparent cloud tiering can be configured on those nodes. Remotely mounted client support is also helpful when running with ESS configurations since multi-cluster configurations with ESS are very common.

Mixed-node cluster configuration

Transparent cloud tiering service runs on x86, Power LE Linux nodes, and IBM Z nodes. Transparent cloud tiering does not run on Windows or Power BE Linux nodes. No mixed-node cluster support with Windows. Both x86 Linux and Power Linux nodes can initiate migrations/recalls of data, and these nodes can initiate a transparent recall on file access.

transparent cloud tiering is supported for use only in IBM Storage Scale clusters and on any associated remotely mounted clients, with Intel x86 Linux and Power LE Linux nodes. Use of Windows, PowerPC® Big Endian, Linux on Z or AIX nodes that access file systems, where transparent cloud tiering is used (either within the cluster or via a remotely mounted client), is untested and unsupported. Inclusion of such nodes is allowed only if there is no possibility that they will ever access transparent cloud tiering; for example, such nodes are tested and supported for PowerPC Big Endian nodes in Elastic Storage Server (ESS) BE & LE servers, where no user applications are supported on Elastic Storage Server (ESS).

SELinux

Supported

SOBAR

SOBAR backup and restore support is available.

IPV6 Support

Not supported

IBM Storage Scale Stretch Clusters

This service can be used in conjunction with stretch clusters. cloud services node classes can be set up all on one site or, most usually, can be split across sites to allow for continued access to cloud data even through failure of an entire site.

installation toolkit

Not supported.

Linux on Power8 Little Endian

Supported.

Note: For information on known limitations, see the *Known Limitations for transparent cloud tiering* topic in the *IBM Storage Scale: Administration Guide*.

Interoperability of cloud data sharing with other IBM Storage Scale features

This topic describes the interoperability of cloud data sharing with other IBM Storage Scale features.

IBM Spectrum Archive (LTFS) and IBM Spectrum Protect

- Running IBM Spectrum Archive and cloud data sharing on the same file system is not supported.
- However, both IBM Spectrum Protect and cloud data sharing can co-exist on the same systems if they are configured with different file systems.
- It is advised not to enable the Data Management API (DMAPI) on the file system that is managed by the cloud data sharing.

AFM

- Data from AFM or AFM DR filesets might be accessed by cloud data sharing.
- Using cloud data sharing on AFM home is supported.

Multi-clusters

Running cloud data sharing on multiple clusters to the same or different cloud accounts is supported.

File Placement Optimizer (FPO)

cloud data sharing is supported on FPO cluster. Transferred files might not have an optimal placement.

IBM Storage Scale Object

- cloud data sharing works only with Unified File Object based IBM Storage Scale object storage at this time.
- Running cloud data sharing in the same file system with another data set is supported.

Snapshots

cloud data sharing can be used with the snapshot function to export snapshot files to the cloud since it is simply doing the file replication.

Sparse files

cloud data sharing can be used to import or export sparse files. However, sparseness is not maintained and on re-import, full blocks are allocated.

Encryption

cloud data sharing can be used with the IBM Storage Scale file system level encryption feature (available in the Advanced Edition or IBM Storage Scale Erasure Code Edition). However, all the data that is exported to the cloud object storage is unencrypted or if cloud data sharing encryption is enabled it is exported with the encryption key configured for cloud services. Essentially, when data is read from file system, the files get decrypted, and may be encrypted again at user space by cloud data sharing and pushed into the cloud storage.

Compression

cloud data sharing can be used along with IBM Storage Scale file system level compression capability. Essentially, when data is read from the file system, the files will get uncompressed, cloud data sharing will push uncompressed file onto the cloud storage.

CES (protocol services)

cloud data sharing can co-exist with active or inactive NFS, SMB, or Object services on the CES nodes.

IBM Storage Scale stretched cluster

IBM Storage Scale stretched cluster and cloud data sharing on the same file system is supported.

Spectrum Protect (TSM)

IBM Storage Scale file systems backed up using Spectrum Protect may leverage cloud data sharing.

Elastic Storage Server (ESS)

cloud data sharing cannot be deployed directly on ESS nodes. However, it can be deployed on x86 Protocol (CES) nodes that can be attached to ESS that effectively allows sharing of ESS data.

Mixed-node cluster configuration

- cloud data sharing service runs only on x86 and Power Linux nodes. cloud data sharing is not supported to be deployed on Power Linux or Windows. However, it can co-exist in a Linux cluster that contains both x86 and Power nodes, as long as the Power nodes are not running the cloud services.
- No mixed-node cluster support with Windows. Only x86 Linux nodes can initiate migrations/recalls of data and only x86 Linux nodes can initiate a transparent recall on file access.

SELinux

Supported

SOBAR

Running SOBAR and cloud data sharing on the same file system is supported.

IPV6 Support

Not supported

Linux on Power8 Little Endian

Supported

installation toolkit

Not supported.

Note: For information on known limitations, see the *Known Limitations for cloud services* topic in the *IBM Storage Scale: Administration Guide*.

File audit logging

File audit logging captures file operations on a file system and logs them to a retention enabled fileset.

Each file operation is generated as a local event on the node that serves the file operation. These events are produced to the audit fileset. These events are called lightweight events. Lightweight events occur at the file system level and capture all accesses to a monitored file system from protocol exports to even root access that occurs directly on nodes. For more information, see [“Producers in file audit logging” on page 250](#).

The most common file operations such as open, close, close with write, rename, unlink, create, remove directory, extended attribute change, ACL change, and GPFS attribute change are the events that are captured. Events are created in a highly parseable JSON formatted string as they are written to the designated fileset. For each file system enabled for file audit logging, a fileset is designated for the audit logs. This fileset keeps the logs currently being written to in append-only mode and, as it rotates to a new log file, compresses the old log file and makes it immutable for the retention period.

Configurable options for file audit logging filesets include their name, whether or not the audit fileset should be IAM mode compliant, and the retention period in days. An entire file system can be audited, a subset of filesets within the file system can be audited, or the file system can be audited while skipping events in selected filesets. File audit logging is integrated into the system health infrastructure, so alerts are generated for the producer and state changes of the producer.

Note: IBM Storage Scale file audit logging is the preferred method for logging file system activities.

For more information about file audit logging, see *Monitoring file audit logging* in *IBM Storage Scale: Administration Guide*.

Producers in file audit logging

Producers create lightweight events and publish them to the audit fileset.

Lightweight events are generated when file operations are served locally on a node in the data path. These lightweight events are generated as node local events that are then published to the audit fileset. They reside in the GPFS daemon and are activated and configured by IBM Storage Scale policy code.

The file audit logging fileset

The audited events are logged in an independent file audit logging fileset.

The user can also supply a previously created fileset if it is already in IAM mode noncompliant (or IAM mode compliant if the **--compliant** flag is given). Using IAM mode noncompliant or compliant is a requirement so that audit logs are not altered accidentally. This requirement also allows a retention period to be set on the files within the audit fileset. By default, the file audit logging fileset is IAM mode noncompliant, so the root user can change the expiration date on files and then delete them if needed to free up space. If file audit logging is enabled with the optional **--compliant** flag, then not even the root user can change the expiration date on files within the audit fileset. In that case, no files can be deleted before the expiration date. This previously created fileset must also be linked directly from the root mount point: `<FS_Mount_Point>/<Fileset_Name>`. The user specifies this fileset during enablement by using the **--log-fileset** option on the **mmaudit** command.

File audit logging records for a file system are arranged in the following manner:

```
<FS_Mount_Point>/<Fileset_Name>/Topic/Year/Month/Day
```

- **FS_Mount_Point:** The default mount point of the device being audited is designated to hold the file audit logging destination fileset.
- **Fileset_Name:** The fileset name indicates the file audit logging destination fileset. If no other fileset name is given, then the default is `.audit_log`.
- **Topic:** The topic that is associated with the file system that is being audited. It consists of the following structure:

```
<Device_Minor Number>_<Cluster_ID>_<Generation_Number>_audit
```

- **Year/Month/Day:** Details when the individual file that is being written was first created.

Note: The **mmbackupconfig** command does not save the file audit logging enabled state of a file system. However, the IAM noncompliant or IAM-compliant mode of the file audit logging fileset is saved. Therefore, when you restore the configuration information with the **mmrestoreconfig** command, file audit logging must be reenabled on the file system with the **mmaudit** command. You need to specify **--log-fileset** and any other enable flags if you are not using the default values. For more information, see *IAM modes and their effects on file operations on immutable files* in the *IBM Storage Scale: Administration Guide*.

File audit logging records

Files that hold the JSON formatted event records.

Files in the file audit logging fileset are named in the following format:

```
auditLogFile_<IO_Node_Name >_YYYY-MM-DD_HH:MM:SS
```

When a file audit logging file is created, it is put in append only mode and an expiration date is set for it. The expiration date corresponds to the number of days after the current date given in the retention option of the **mmaudit** command. The default retention period is 365 days, which means that a given file audit logging file is set to expire one year from when it is created.

When auditing events are generated by file system operations, they are appended to the files in the file audit logging fileset. After approximately 5,000,000 entries have been made in a file or the file size exceeds 400MB, whichever comes first, it is made immutable, GPFS native compression is enabled on it, and a new file with the current date and time is created for the next batch of audit logs. Therefore, new files are created at a rate that corresponds to the file system activity of the device that is being audited.

Important: Since files in the file audit logging fileset are created in the directory based on their creation date, there might be audit records for several days within the same file if there is not much file system activity.

Furthermore, since the events are written to these files in batches, the file audit logging event information might be placed in a previous day's or next day's directory. In order to improve usability, there are soft links that are created in the audit fileset topic subdirectory that link each current audit record file. These links are named in the following format: `auditLogFile.latest_<IO_Node_Name>`.

For more information, see [“JSON attributes in file audit logging” on page 252](#).

File audit logging events

Use this information to learn more about which I/O operations result in the 11 events for file audit logging.

Table 18. File audit logging events		
Event name	Description	Examples
ACCESS_DENIED	A user was denied access to operate on a file.	<code>open()</code> with <code>O_WRONLY</code> where user has no write permission.
ACLCHANGE*	A file's or directory's ACL permissions were modified.	mmputacl, chown, chgrp, chmod
CLOSEWRITE*	A file was opened for writing and then closed.	<code>cp</code> , <code>echo</code> , <code>write()</code> , <code>close()</code>
CLOSE	A file was closed.	<code>close()</code> , <code>touch</code> , policy MIGRATE rule.
CREATE*	A file or directory was created.	<code>open(create flag)</code> , <code>vi</code> , <code>ln</code> , <code>dd</code> , <code>mkdir</code>
GPFSATTRCHANGE*	A file's or directory's IBM Storage Scale attributes were changed.	mmchattr -i -e --indefinite-retention
OPEN	A file or directory was opened for reading, writing, or creation.	<code>open()</code> , <code>mmlsattr</code> , <code>cat</code> , <code>cksum</code> , <code>ls</code> (only for directories), policy LIST rule
RENAME*	A file or directory was renamed.	<code>rename()</code> , <code>mv</code>
RMDIR*	A directory was removed.	<code>rmdir()</code> , <code>rm</code> , <code>rmdir</code>
UNLINK*	A file or directory was unlinked from its parent directory. When the linkcount = 0, the file is deleted.	<code>unlink()</code> , <code>rm</code> <code>hardlink/softlink</code>
XATTRCHANGE*	A file's or directory's extended attributes were changed.	mmchattr --set-attr --delete-attr
* These events are not applicable to a file system mounted as read-only.		

Note: If a file is opened for write and only the CLOSE event is enabled, only a CLOSE event is reported. If both CLOSEWRITE and CLOSE are enabled, both events are recorded as appropriate.

For more information, see *JSON reporting issues in file audit logging* in *IBM Storage Scale: Problem Determination Guide*.

JSON attributes in file audit logging

Use this information to learn more about the JSON attributes that are associated with the 11 events in file audit logging.

LWE_JSON

The version of the record.

path

The path name of the file that is involved in the event.

oldPath

The previous path name of the file during the **RENAME** event. For all other events, it is not displayed.

clusterName

The name of the cluster where the event took place.

nodeName

The name of the node where the event took place.

nfsClientIp

The IP address of the remote client that is involved in the event.

fsName

The name of the file system that is involved in the event.

event

This is one of the following events: **OPEN**, **CREATE**, **CLOSE**, **CLOSEWRITE**, **RENAME**, **UNLINK**, **XATTRCHANGE**, **ACLCHANGE**, **RMDIR**, **GPFSATTRCHANGE**, or **ACCESS_DENIED**.

inode

The inode number of the file that is involved in the event.

linkCount

The Unix link count of the file that is involved in the event.

openFlags

The open flags that are specified during the event. For example:

```
fcntl.h ( O_RDONLY,O_WRONLY,O_RDWR, O_CREAT, ...)
```

For example:

```
"openFlags": "32962" = 0x80C2 = 0100302 translates to ( O_RDWR | O_CREAT | O_EXCL | O_LARGEFILE)
```

poolName

The pool name where the file resides.

fileSize

The current size of the file in bytes.

ownerUserId

The owner ID of the file that is involved in the event.

ownerGroupId

The group ID of the file that is involved in the event.

atime

The time in UTC format of the last access of the file that is involved in the event.

ctime

The time in UTC format of the last status change of the file that is involved in the event.

mtime

The time in UTC format of the last modification to the file that is involved in the event.

eventTime

The time in UTC format of the event.

clientUserId

The user ID of the process that is involved in the event.

clientGroupId

The group ID of the process that is involved in the event.

accessMode

The access type for which the operation was denied for the **ACCESS_DENIED** event.

processId

The process ID that is involved in the event.

bytesRead

The bytes read from a file.

bytesWritten

The bytes written to a file.

minReadOffset

The starting position of bytes read from a file.

maxReadOffset

The ending position of bytes read from a file.

minWriteOffset

The starting position of bytes written to a file.

maxWriteOffset

The ending position of bytes written to a file.

permissions

The permissions on the file that is involved in the event.

acls

The access control lists that are involved in the event.

xattrs

The extended attributes that are involved in the event.

subEvent

The type of IBM Storage Scale attribute change. Only applies to the immutability and appendOnly flags.

The following table describes the JSON attributes that are provided for the 10 events in file audit logging:

Table 19. JSON attributes in file audit logging											
Attribute	OPEN	CREATE	CLOSE	CLOSE WRITE	RENAM E			UNLINK	RMDIR		ACCESS_DENIED
LWE_JSON	X	X	X	X	X	X	X	X	X	X	X
path	X	X	X	X	X	X	X	X	X	X	X
oldPath					X						
clusterName	X	X	X	X	X	X	X	X	X	X	X
nodeName	X	X	X	X	X	X	X	X	X	X	X
nfsClientIp	X ¹	X ¹	X ¹	X ¹	X ¹	X ^{1,2}	X ¹	X ¹	X ¹		X ¹
fsName	X	X	X	X	X	X	X	X	X	X	X
event	X	X	X	X	X	X	X	X	X	X	X
inode	X	X	X	X	X	X	X	X	X	X	X
linkCount	X	X	X	X	X	X	X	X	X	X	
openFlags	X	0	X	X	0	0	0	0	0	0	X
poolName	X	X	X	X	X	X	X	X	X	X	X
fileSize	X	X	X	X	X	X	X	X	X	X	X
ownerUserId	X	X	X	X	X	X	X	X	X	X	X
ownerGroupId	X	X	X	X	X	X	X	X	X	X	X
atime	X	X	X	X	X	X	X	X	X	X	X
ctime	X	X	X	X	X	X	X	X	X	X	X
mtime	X	X	X	X	X	X	X	X	X	X	X
eventTime	X	X	X	X	X	X	X	X	X	X	X

Table 19. JSON attributes in file audit logging (continued)											
Attribute	OPEN	CREATE	CLOSE	CLOSE WRITE	RENAM E			UNLINK	RMDIR		ACCESS_DENIED
clientId	X	X	X	X	X	X	X	X	X	X	X
clientGroupId	X	X	X	X	X	X	X	X	X	X	X
processId	X	X	X	X	X	X	X	X	X	0	X
bytesRead	0	Null	x ⁵	x ⁵	Null	Null	0	Null	Null	Null	Null
bytesWritten	0	Null	x ⁵	x ⁵	Null	Null	0	Null	Null	Null	Null
minReadOffset	MAX INT	Null	x ⁵	x ⁵	Null	Null	MAX INT	Null	Null	Null	Null
maxReadOffset	0	Null	x ⁵	x ⁵	Null	Null	0	Null	Null	Null	Null
minWriteOffset	MAX INT	Null	x ⁵	x ⁵	Null	Null	MAX INT	Null	Null	Null	Null
maxWriteOffset	0	Null	x ⁵	x ⁵	Null	Null	0	Null	Null	Null	Null
permissions	X	X	X	X	X	X	X	X	X	X	X
acls	Null	Null	Null	Null	Null	Null	X	Null	Null	Null	Null
xattrs	Null	Null	Null	Null	Null	x ³	Null	Null	Null	Null	Null
subEvent	NONE	NONE	NONE	NONE	NONE	NONE	NONE	NONE	NONE	APPENDONLY IMMUTABILITY	NONE
accessMode	Null	Null	Null	Null	Null	Null	Null	Null	Null	Null	X

Note: In the above table, 0, Null, or MAX INT represents that the attribute is not applicable for that particular event.

For more information about some of the issues that might occur with the events and when they might occur, see in the *IBM Storage Scale: Problem Determination Guide*.

Note:

- The **nfsClientId** attribute is provided for NFS clients that use Ganesha. The value is NULL for kernel NFS versions and SMB.
- The **nfsClientId** attribute is populated for an **XATTRCHANGE** event when SELinux is enabled and a **CREATE** event is generated via NFS.
- The **xattrs** attribute only shows the **xattr** that was changed.
- The best effort is made to provide both **path** and **nfsClientId** attributes for files accessed via NFS, but it is not guaranteed.
- Results might be inaccurate with mmap IO, mmap IO via SMB, or IO via NFS.

Remotely mounted file systems in file audit logging

In IBM Storage Scale, events are recorded for file systems that are remotely mounted when file audit logging is enabled on the **owningCluster**.

Events are only recorded on nodes that meet the file audit logging prerequisites. For more information, see [“Requirements, limitations, and support for file audit logging” on page 647](#). File audit logging commands do not work on file systems that are not owned by the local cluster. File audit logging can be enabled for both local and remotely mounted file systems at the same time with audit records that are stored by the **owningCluster** of the file system. To determine whether a file system is under audit from an **accessingCluster**, submit the **mmfs FS --file-audit-log** command.

Clustered watch folder

You can use the clustered watch folder to monitor file systems, filesets, and inode spaces for file accesses.

With the clustered watch folder feature, it is possible to watch file operations across clusters by using a centralized tool that has scalability and resiliency built-in. An entire file system, fileset, or inode space can be watched. Clustered watch folder captures file system activities at the IBM Storage Scale file system level, generates event notifications for those activities, and streams the notifications to topics in an external sink that you can manage.

Producers in clustered watch folder

Producers generate file access event notifications.

Producers are located in the IBM Storage Scale daemon. For a file system with an enabled clustered watch, a producer exists in the daemon on each node that mounts the file system. The producer generates events for a clustered watch if the producer is on the IBM Storage Scale cluster that owns the file system and the clusters with the file system remotely mounted. In addition, both clusters must meet the minimum IBM Storage Scale level versions. The producer code requires one package to be installed. For more information, see [“Manually installing clustered watch folder”](#) on page 652.

Interaction between the clustered watch folder and the external Kafka sink

Clustered watch folder supports sending watch events to an external Kafka sink.

Before the events are sent to the external Kafka sink, they are compressed by using the LZ4 algorithm. Make sure that the lz4 library is installed on the system from which the events are meant to be consumed. Each consumer library is different and may bundle LZ4 directly or have a preferred method of installation; therefore, to get the appropriate LZ4 installation instructions, consult the consumer library documentation.

To send watch events to an external Kafka sink, a minimum of two attributes must be present when you enable a clustered watch:

- A list of accessible broker addresses with ports to access the external Kafka queue.
- The topic name on the external Kafka queue where the clustered watch publishes events.

In addition to the two required attributes, authentication or authorization can also be specified. If authentication or authorization is not given when you enable a clustered watch, it is assumed that it is not needed. The following types of authentication or authorization are supported:

- **NONE**: This is the default. It can also be specified by excluding any type of authentication configuration.
- **PLAINTEXT**: Use Kafka plain text authentication. You must provide a PRODUCER_USERNAME and PRODUCER_PASSWORD with the authentication information for the producer to write to the external Kafka sink.
- **SASL**: Use SASL-based authentication between the IBM Storage Scale cluster that hosts the clustered watch and the external Kafka sink. You must provide a PRODUCER_USERNAME and PRODUCER_PASSWORD with the authentication information for the producer to write to the external Kafka sink; and you must provide the specific mechanism to use:
 - SCRAM256
 - SCRAM512
- **SASL_SSL**: Use SASL_SSL-based authentication between the IBM Storage Scale cluster that hosts the clustered watch and the external Kafka sink. You must provide a PRODUCER_USERNAME and PRODUCER_PASSWORD with the authentication information for the producer to write to the external Kafka sink; and you must provide the specific mechanism to use:
 - SCRAM256
 - SCRAM512

For the producer to write to the external Kafka sink, you must also provide three extra parameters (the fourth one is optional) when you specify this type of authentication:

- **CA_CERT_LOCATION:** Full path (including the actual file) to the location of the ca-cert file. This field is required.
- **CLIENT_PEM_CERT_LOCATION:** Full path (including the actual file) to the location of the client certificate (.pem format) file. This field is required.
- **CLIENT_KEY_FILE_LOCATION:** Full path to the location of the client key (client.key) file. This field is required.
- **CLIENT_KEY_FILE_PASSWORD:** Key file password. This field is optional.
- **CERT:** Use certificate-based authentication and encryption of data in flight between the IBM Storage Scale cluster hosting the clustered watch and the external Kafka sink. For the producer to write to the external Kafka sink, you must also provide three extra parameters (the fourth one is optional) when you specify this type of authentication:
 - **CA_CERT_LOCATION:** Full path (including the actual file) to the location of the ca-cert file. This field is required.
 - **CLIENT_PEM_CERT_LOCATION:** Full path (including the actual file) to the location of the client certificate (.pem format) file. This field is required.
 - **CLIENT_KEY_FILE_LOCATION:** Full path to the location of the client key (client.key) file. This field is required.
 - **CLIENT_KEY_FILE_PASSWORD:** Key file password. This field is optional.

For more information about specifying these parameters, see the **mmwatch** command in the *IBM Storage Scale: Command and Programming Reference Guide*.

Note: For the producer to write to the external Kafka queue, the firewall ports must be open between the source IBM Storage Scale cluster and the external Kafka queue.

Note: All of these parameters are used in the **--sink-auth-config** flag of the **mmwatch** command. This parameter is optional. When it is used, you must pass a configuration file with specific parameters. For more information, see the following examples. The first example is of a clustered watch folder setup with SCRAM512 to the external Kafka sink.

```
SINK_AUTH_TYPE:SASL
SINK_AUTH_MECHANISM:SCRAM512
PRODUCER_USERNAME:<will be found in external kafka config>
PRODUCER_PASSWORD:<will be found in external kafka config>
```

The second example is of a CERT-based authentication setup between the IBM Storage Scale cluster and the external Kafka sink that the clustered watch folder uses.

```
SINK_AUTH_TYPE:CERT
CA_CERT_LOCATION:<path to certs>
CLIENT_PEM_CERT_LOCATION:<path to pem cert>
CLIENT_KEY_FILE_LOCATION:<path to key>
CLIENT_KEY_FILE_PASSWORD:<password from certificate setup>
```

Clustered watch folder events

Descriptions of the events that are supported in clustered watch folder.

The following events are supported:

Table 20. File access events that are supported by clustered watch folder	
File access event	Notes
IN_ACCESS ²	A file was accessed (read or ran).

Table 20. File access events that are supported by clustered watch folder (continued)

File access event	Notes
IN_ATTRIB ²	Metadata was changed (for example, chmod , chown , setxattr etc.).
IN_CLOSE_NOWRITE	A file or folder that was not opened for writing was closed.
IN_CLOSE_WRITE ²	A file that was opened for writing was closed.
IN_CREATE ²	A file or folder was created in a watched folder.
IN_DELETE ²	A file or folder was deleted from a watched folder.
IN_DELETE_SELF ²	A watched file or folder was deleted.
IN_IGNORED ¹	Event that is triggered when a file system is unmounted. Always follows the IN_UNMOUNT event.
IN_ISDIR ¹	A directory is listed.
IN_MODIFY ²	A file was modified (for example, write or truncate).
IN_MOVE_SELF ²	A folder that was being watched was moved.
IN_MOVED_FROM ²	A file within the watched folder was renamed.
IN_MOVED_TO ²	A file was moved or renamed to this folder.
IN_OPEN	A file or folder was opened.
IN_UNMOUNT ¹	A file system that is being watched is unmounted.
¹ These events are always watched.	
² These events are not applicable to a file system mounted as read-only.	

JSON attributes in clustered watch folder

Use this information to learn more about the JSON attributes that are associated with the 12 events that can be specified when enabling a clustered watch and the other events that might be received due to other file system activity.

WF_JSON

The version of the record.

wd

The watch descriptor.

cookie

A unique integer that connects related events. It allows the resulting pair of **IN_MOVED_FROM** and **IN_MOVED_TO** events to be connected.

mask

A hex value representing the event.

event

The event type. One of the following events:

- IN_ACCESS
- IN_ATTRIB
- IN_CLOSE_NOWRITE
- IN_CLOSE_WRITE

- IN_CREATE
- IN_DELETE
- IN_DELETE_SELF
- IN_MODIFY
- IN_MOVED_FROM
- IN_MOVED_TO
- IN_MOVE_SELF
- IN_OPEN

path

The path name of the file that is involved in the event.

clusterName

The name of the cluster where the event took place.

nodeName

The name of the node where the event took place.

nfsClientIp

The IP address of the client involved in the event.

fsName

The name of the file system that is involved in the event.

inode

The inode number of the file that is involved in the event.

filesetID

The fileset ID of the file.

linkCount

The link count of the file.

openFlags

The open flags that are specified during the event (**O_RDONLY**, **O_WRONLY**, **O_RDWR**, **O_CREAT**, etc.) as defined in `fcntl.h`.

poolName

The pool name where the file resides.

fileSize

The current size of the file in bytes.

ownerUserId

The owner ID of the file that is involved in the event.

ownerGroupId

The group ID of the file that is involved in the event.

atime

The time in UTC format of the last access of the file that is involved in the event.

ctime

The time in UTC format of the last inode or metadata change of the file that is involved in the event.

mtime

The time in UTC format of the last time that the file was modified.

eventTime

The time in UTC format of the event.

clientUserId

The user ID of the process that is involved in the event.

clientGroupId

The group ID of the process that is involved in the event.

processId

The process ID that is involved in the event.

bytesRead¹

The bytes read from a file.

bytesWritten¹

The bytes written to a file.

minReadOffset¹

The starting position of bytes read from a file.

maxReadOffset¹

The ending position of bytes read from a file.

minWriteOffset¹

The starting position of bytes written to a file.

maxWriteOffset¹

The ending position of bytes written to a file.

permissions

The permissions on the file that is involved in the event.

acls

The access control lists that are involved in the event (only in case of an acl change event).

xattrs

The extended attributes that are involved in the event (only in case of an Xattr change event).

subEvent

None.

You can start a watch with any of the following events:

- IN_OPEN - 0x20
- IN_ACCESS - 0x01
- IN_MODIFY - 0x02
- IN_ATTRIB - 0x04
- IN_CLOSE_WRITE - 0x8
- IN_CLOSE_NOWRITE - 0x10
- IN_CREATE - 0x100
- IN_DELETE - 0x200
- IN_MOVED_FROM - 0x40
- IN_MOVED_TO - 0x80
- IN_DELETE_SELF - 0x400
- IN_MOVE_SELF - 0x800

You might see these events during some of the previously mentioned watch types:

- IN_IGNORED - 0x8000
- IN_UNMOUNT - 0x2000
- IN_OPEN IN_ISDIR - 0x40000020
- IN_ATTRIB IN_ISDIR - 0x40000004
- IN_CLOSE IN_ISDIR - 0x40000010
- IN_CREATE IN_ISDIR - 0x40000100
- IN_DELETE IN_ISDIR - 0x40000200
- IN_MOVE_SELF IN_ISDIR - 0x40000800
- IN_MOVED_FROM - 0x40000040

- IN_MOVED_TO - 0x40000080
- IN_DELETE_SELF - 0x40000400

Note: ¹ Results might be inaccurate with mmap IO, mmap IO via SMB, or IO via NFS.

Overview of call home

IBM Storage Scale users can use the call home feature to share the basic configuration data of a cluster with IBM support and development team. The shared details ensure that the cluster is configured to function according to its maximum potential.

With the call home feature enabled, the IBM Storage Scale users can upload the daily or weekly cluster configuration or debugging data from the customer site to a specific IBM server. The IBM server then sends the data to the IBM backend, Enhanced Customer Data Repository (IBM ECuRep). For more information, see [ECuRep](#).

The uploaded data is further analyzed by the IBM support team to determine whether there is a common misconfiguration, best practice violation, or any other issues. This enables the IBM support team to provide solutions to any problems encountered on the customer site or leverage the data to provide more benefits to the customer. For more information on understanding the call home benefits and supported use cases, see *Benefits of enabling call home in IBM Storage Scale: Problem Determination Guide*.

The call home component in IBM Storage Scale must be configured before it can be used for manual or automated data uploads to the IBM server. The call home configuration process consists of the following steps:

1. Configuring the call home settings.
2. Creating call home groups.

For more information, see *Configuring call home* in the *IBM Storage Scale: Administration Guide*.

Call home architecture

The following figure shows the basic call home group structure:




Figure 33. Call home architecture

Call Home Group

The Call Home Group is a group of call home nodes that are configured by using the **mmcallhome group** command. A call home group consists of at least one child node, which also acts as its call home node. A call home group can have more than one child node but has only one call home node. Multiple call home groups can be configured within an IBM Storage Scale cluster. You can automate the call home group creation by using the **mmcallhome group auto** command.

Call home groups help to distribute the data-gather and data-upload workload to prevent bottlenecks. You can create groups of any size between one and the number of nodes in your cluster. The larger the group, the higher the workload on the call home node. Therefore, you are recommended to limit the group size to 1024 nodes. Larger groups are possible, but it might lead to performance issues on the call home node.

Call Home Node

The Call Home Node is the node that performs the data upload. If a scheduled data gathering is enabled, then the call home node initiates the data collection within its call home group and uploads the data package to the IBM support center. A gather-send task runs on the call home node, which collects the data from the call home child nodes and uploads that data to a specific IBM server. The IBM server then sends the data to the IBM backend, ECuRep. For more information, see [ECuRep](#). The gather-send configuration file includes information about the data that is collected from the child nodes.

Note: The call home node is also a child node of the group. If the call home node becomes unavailable, then the whole call home group is unable to perform any data uploads until the call home node is online again.

Important: The call home node needs to have access to the external network through 443 port number. For more information on network-related requirements, see the *Installing call home* section in the *IBM Storage Scale: Concepts, Planning, and Installation Guide*.

Call Home Child Node

The Call Home Child Node is a member of a call home group. The call home node can collect data from all the call home child nodes in a call home group.

Note: If a call home child node, which is not a call home node, becomes unavailable, then the rest of the call home group continues to work. However, the scheduled uploads do not include the details of the unavailable child node.

mmcallhome commands impact

The call home data upload and share are managed by using the **mmcallhome** command.

The **mmcallhome** command provides options to configure, enable, run, schedule, and monitor call home related tasks in the IBM Storage Scale cluster. Information from each node within a call home group is collected and securely uploaded to the IBM ECuRep server.

The uploaded packages that contain the daily or weekly scheduled uploads and non-ticket-related sent files are saved on the corresponding call home node for an inspection by the customer.

You can use the **mmcallhome status list --verbose** command to find the exact location of the uploaded packages. A maximum of 1 GB of space is used by the packages of each of these categories:

- Sent packages
- Failed uploads

A maximum of 48 such packages can be kept.

The uploaded data is stored for at least two weeks on IBM ECuRep, and can be identified by using your customer information. If you need to access this data, contact IBM support. For more information, see [ECuRep](#). The PMR or Salesforce case-related data is saved and the PMR or Salesforce case is still open.

The following **mmcallhome** command options respond in the same way when run on nodes, which belong or do not belong to a call home group:

- **mmcallhome group**
- **mmcallhome capability**
- **mmcallhome info**
- **mmcallhome proxy**

Note: For compatibility reasons, in a mixed cluster configuration, the **mmcallhome capability**, **mmcallhome info**, and **mmcallhome proxy** commands are available only to the global settings if the corresponding nodes are not part of a call home group. If the call home node of this group also has IBM Storage Scale 4.2.3 PTF 6 or older nodes, then they also manage a separate settings configuration for their group.

All other mmcallhome command options

All other **mmcallhome** commands can be run only from a node, which is a member of a call home group.

For more information, see the *mmcallhome* command in the *IBM Storage Scale: Command and Programming Reference Guide*.

Types of call home data upload

The call home feature allows data to be uploaded either manually or automatically to the IBM server.

The manual and automatic data upload options can upload the data to the IBM ECuRep server. The data which is uploaded to the IBM ECuRep server is not analyzed automatically and gets deleted after a specified period. The default period is set to two weeks, but it can be changed. For more information about the usage of the uploaded data, contact IBM support.

Note: You must consider the space requirements of your system before the call home data is uploaded. For more information, see *Space requirements for call home data upload* in the *IBM Storage Scale: Concepts, Planning, and Installation Guide*.

Manual data upload

The call home feature provides the option to manually upload a specific file or multiple data packages to the IBM server. To upload data manually, issue the **mmcallhome run** command in one of the following ways:

- To manually initiate the **daily data** upload, issue the following command:

```
mmcallhome run GatherSend --task DAILY
```

- To manually upload a **specific data file** to the IBM server, issue one of the following commands:

```
mmcallhome run SendFile --file myfile
```

The **mmcallhome run SendFile --file <file name>** command uploads a specific file to a common directory, which is available for support to all customers.

Note: The IBM support team must be informed that a specific file is to be considered and all IBM support representatives can read the specific file.

```
mmcallhome run SendFile --file myFile --pmr TS123456789
```

The **mmcallhome run SendFile --file <file name> -pmr <number>** command uploads a specific file to a specific location, which is linked to the specified PMR or Salesforce team.

Note: The IBM support team automatically notices the changes in a PMR or Salesforce case if they are working on the same ticket. Only a specific set of IBM support representatives who are allowed to work on the specific PMR or Salesforce case can read the contents of the specific file.

Automatic data upload

The call home feature provides three different ways to automatically upload a specific file or data packages to the IBM server.

- **Event-based data upload**

This type of data upload collects the debugging information after an occurrence of a specific failure is detected. For more information, see [“Event-based uploads” on page 264](#).

- **Scheduled data upload**

This type of data upload is enabled by using the **mmcallhome schedule** command to schedule a weekly or daily upload of the predefined data. For more information, see [“Scheduled data upload” on page 270](#).

- **Heartbeats feature**

This type of data upload is done only once a day by each call home master node. The call home master nodes send the basic cluster information to IBM ECuRep server. For more information, see [“Heartbeats feature for call home” on page 270](#).

Event-based uploads

If the call home feature is enabled and one of the specific RAS events occur, which degrades the current state of an **mmhealth** component, then the corresponding debugging data is collected and automatically uploaded to IBM ECuRep server for a detailed problem analysis. This event-based data upload feature is called **FTDC2CallHome**.

The **FTDC2CallHome** feature provides the following benefits:

- Allows the IBM support representatives to receive and analyze the relevant debugging data in a faster and easier manner, which reduces the duration of detected outages.
- Directs the efforts of IBM support representatives to the areas that are facing the maximum outages. This increases the stability of IBM Storage Scale in the areas that face outage issues more often.
- Introduces features for protection against data flooding. In this way, this feature ensures minimal CPU usage and preserves the network bandwidth, which might be needed to upload the debugging data if any issue occurs.

You can run the following command to see the uploaded data files: `mmcallhome status list --task sendfile --verbose`

Note: If you disable the event-based uploads, then the unified call home feature for Elastic Storage Server (ESS) systems is also disabled to ensure that service tickets are not automatically created for any hardware failures.

The following RAS events trigger data collection:

Table 21. RAS events that trigger data collection	
Event name	Snap
ads_down	auth.snap.py
ads_failed	auth.snap.py
bootdrive_mirror_degraded	esa.snap.py
bootdrive_mirror_failed	esa.snap.py
bootdrive_missing	esa.snap.py

<i>Table 21. RAS events that trigger data collection (continued)</i>	
Event name	Snap
bootdrive_power_mode_degraded	esa.snap.py
bootdrive_smart_failed	esa.snap.py
can_fan_failed	esa.snap.py
can_fan_missing	esa.snap.py
can_power_supply_failed	esa.snap.py
can_power_supply_missing	esa.snap.py
can_temp_bus_failed	esa.snap.py
can_temp_sensor_failed	esa.snap.py
canister_thermal_ok	commonsnap.py
canister_thermal_shutdown	esa.snap.py
ccr_auth_keys_fail	commonsnap.py
ccr_comm_dir_fail	commonsnap.py
ccr_paxos_12_fail	commonsnap.py
ces_many_rx_errors	commonsnap.py
ces_many_tx_errors	commonsnap.py
ces_network_ips_not_assignable	ces.snap.py
coin_battery_low	esa.snap.py
coin_battery_missing	esa.snap.py
cpu_inspection_failed	esa.snap.py
cpu_speed_wrong	esa.snap.py
cpu_unit_speed_wrong	esa.snap.py

Table 21. RAS events that trigger data collection (continued)

Event name	Snap
ctdb_down	smb.snap.py
ctdb_state_down	smb.snap.py
ctdb_version_mismatch	smb.snap.py
dcm_failed	esa.snap.py
dimm_inspection_failed	esa.snap.py
dimm_module_size_wrong	esa.snap.py
dimm_module_speed_wrong	esa.snap.py
dimm_size_wrong	esa.snap.py
dimm_speed_wrong	esa.snap.py
disk_call_home	esa.snap.py
disk_call_home2	esa.snap.py
esm_failed	esa.snap.py
expander_failed	esa.snap.py
fan_absent	esa.snap.py
fan_failed	esa.snap.py
fan_fault_indicated	esa.snap.py
fan_speed_high	esa.snap.py
fan_speed_low	esa.snap.py
fserrallocblock	commonsnap.py
fserrbadaclref	commonsnap.py
fserrbadidirblock	commonsnap.py

Table 21. RAS events that trigger data collection (continued)	
Event name	Snap
fserrbaddiskaddrindex	commonsnap.py
fserrbaddiskaddrsector	commonsnap.py
fserrbaddittoaddr	commonsnap.py
fserrbadinodeorgen	commonsnap.py
fserrbadinodestatus	commonsnap.py
fserrbadptrreplications	commonsnap.py
fserrbadreplicationcounts	commonsnap.py
fserrbadxattrblock	commonsnap.py
fserrcheckheaderfailed	commonsnap.py
fserrclonetree	commonsnap.py
fserrdealloccblock	commonsnap.py
fserrdotdotnotfound	commonsnap.py
fserrgennummismatch	commonsnap.py
fserrinconsistentfilesetrootdir	commonsnap.py
fserrinconsistentfilesetsnapshot	commonsnap.py
fserrinconsistentinode	commonsnap.py
fserrindirectblock	commonsnap.py
fserrindirectionlevel	commonsnap.py
fserrinodecorrupted	commonsnap.py
fserrinodenummismatch	commonsnap.py
fserrinvalid	commonsnap.py

Table 21. RAS events that trigger data collection (continued)

Event name	Snap
fserrinvalidfilesetmetadatarecord	commonsnap.py
fserrinvalidsnapshotstates	commonsnap.py
fserrsnapinodemodified	commonsnap.py
fserrvalidate	commonsnap.py
gnr_da_out_of_space	esa.snap.py
gnr_da_space_critical	fcm.snap.py
kafka_failed	msgqueue.snap.py
ks_failed	object.snap.py
logdrv_is_missing	esa.snap.py
many_rx_errors	commonsnap.py
many_tx_errors	commonsnap.py
mmfsd_abort_warn	commonsnap.py
nfs_not_active	nfsdebug.snap.py
nfs_not_active_pod	nfsdebug.snap.py
nfs_unresponsive	nfsdebug.snap.py
nfsd_down	nfs.snap.py
nfsd_restart	nfs.snap.py
nic_temperature_exceeded	commonsnap.py
node_call_home	esa.snap.py
node_call_home2	esa.snap.py
nvmeof_raw_disk_smart_failed	esa.snap.py

Table 21. RAS events that trigger data collection (continued)

Event name	Snap
nvmeof_raw_disk_smart_unknown	esa.snap.py
pair_canister_failed	esa.snap.py
postgresql_failed	object.snap.py
power_high_current	esa.snap.py
power_high_voltage	esa.snap.py
power_supply_config_mismatch	esa.snap.py
power_supply_failed	esa.snap.py
proxy-server_failed	object.snap.py
raid_root_disk_bad	esa.snap.py
raid_sas_adapter_bad	esa.snap.py
ring_checksum_failed	object.snap.py
sas_adapter_down_trained	esa.snap.py
sas_adapter_failed	esa.snap.py
sas_adapter_missing	esa.snap.py
sas_cable_failed	esa.snap.py
scaleadmd_fail	msgqueue.snap.py
smbd_down	smb.snap.py
temp_bus_failed	esa.snap.py
temp_sensor_failed	esa.snap.py
test_call_home	esa.snap.py
voltage_bus_failed	esa.snap.py

Table 21. RAS events that trigger data collection (continued)	
Event name	Snap
voltage_sensor_failed	esa.snap.py
zookeeper_failed	msgqueue.snap.py

Heartbeats feature for call home

If the call home feature is installed and enabled on the IBM Storage Scale cluster, then the CALLHOME service runs on the call home master nodes. As soon as the CALLHOME service is enabled, it starts collecting heartbeat data once a day.

The heartbeats feature uploads a snap of the basic cluster information to the IBM ECURep server. The call home schedule settings do not affect the heartbeats feature.

The basic snap script, `basicsnap.py` collects only the following basic cluster information:

- `mmdia -Y`
- `mmsysmonc d cfgshow`
- `df -k -t gpfs`
- `lsblk`
- `lscpu`
- `/etc/os-release`
- `/proc/cpuinfo`
- `/proc/meminfo`
- `/proc/uptime`
- `mmlslicense -Y`
- `mmlsconfig -Y`

Note: The CALLHOME component is a part of the **mmhealth** command monitoring list. For more information, see the *mmhealth command* in the *IBM Storage Scale: Command and Programming Reference Guide*, and the *Call home events* section, and *Monitoring the health of a node* section in the *IBM Storage Scale: Problem Determination Guide*.

Scheduled data upload

You can run the **mmcallhome schedule** command to scheduled data uploads.

The call home feature has a built-in data collection mechanism that collects predefined data on a daily or weekly basis. You can find the schedules that are defined for this data collection in `/usr/lpp/mmfs/lib/mmsysmon/callhome/callhomeSchedules.json`.

The following data is collected and uploaded regularly by the call home feature on a daily or weekly basis according to the defined schedule:

Table 22. Data collected and uploaded by call home					
COMMAND	Machine type	Node	OS	Product	Schedule
<code>/usr/bin/lscpu</code>	all	all	all	all	all

Table 22. Data collected and uploaded by call home (continued)

COMMAND	Machine type	Node	OS	Product	Schedule
/usr/bin/cat /proc/interrupts	all	all	all	all	all
/usr/bin/lsblk	all	all	all	all	all
/sbin/lsdasd -a -u	s390x	all	all	all	all
/sbin/lszfcp -H -P -D	s390x	all	all	all	all
/sbin/lsdasd -b -l	s390x	all	all	all	all
/usr/bin/lsscsi -li	all	all	all	all	all
multipath -ll	all	all	all	all	all
/usr/lpp/mmfs/bin/mmdiag --config -Y	all	all	all	all	all
/usr/lpp/mmfs/bin/mmdiag --version -Y	all	all	all	all	all
/usr/lpp/mmfs/bin/mmlslicense -L -Y	all	CALLHOME_SERVERS	all	all	all
/usr/lpp/mmfs/bin/mmlsnodeclasses --all -Y	all	CALLHOME_SERVERS	all	all	all
/usr/lpp/mmfs/bin/mmlscluster -Y	all	CALLHOME_SERVERS	all	all	all
/usr/lpp/mmfs/bin/mmlscluster --ces -Y	all	CALLHOME_SERVERS	all	all	all

Table 22. Data collected and uploaded by call home (continued)

COMMAND	Machine type	Node	OS	Product	Schedule
<code>/usr/lpp/ mmfs/bin/ mmces address list -Y</code>	all	CALLHOME_SE RVERS	all	all	all
<code>/usr/lpp/ mmfs/bin/ mmlsnode -a</code>	all	CALLHOME_SE RVERS	all	all	all
<code>/usr/lpp/ mmfs/bin/ mmlsmgr -Y</code>	all	CALLHOME_SE RVERS	all	all	all
<code>/usr/lpp/ mmfs/bin/ mmlsnsd -L -Y</code>	all	CALLHOME_SE RVERS	all	all	all
<code>/usr/lpp/ mmfs/bin/ mmlscallback -Y</code>	all	CALLHOME_SE RVERS	all	all	all
<code>/usr/lpp/ mmfs/bin/ tsctl nqStatus -Y</code>	all	all	all	all	all
<code>/usr/lpp/ mmfs/bin/ mmremotecclus ter show all -Y</code>	all	CALLHOME_SE RVERS	all	all	all
<code>/usr/lpp/ mmfs/bin/ mmremotefs show all -Y</code>	all	CALLHOME_SE RVERS	all	all	all
<code>/usr/lpp/ mmfs/bin/ mmauth show -Y</code>	all	CALLHOME_SE RVERS	all	all	all
<code>/usr/lpp/ mmfs/bin/ tsstatus</code>	all	CALLHOME_SE RVERS	all	all	all
<code>/usr/lpp/ mmfs/bin/ mmsdrquery sdrq_fs_info all</code>	all	CALLHOME_SE RVERS	all	all	all

Table 22. Data collected and uploaded by call home (continued)

COMMAND	Machine type	Node	OS	Product	Schedule
<code>/usr/lpp/ mmfs/bin/ mmlsfs <fs> -aABdDEfiIkK LnoPQStTuVwY z --create- time -- encryption --fastea -- filesetdf -- inode-limit -- is4KAaligned --log- replicas -- mount- priority -- perfileset --rapid- repair -- write-cache- threshold -- striped- logs -- fileset- count --afm --snc --uid --snapid</code>	all	CALLHOME_SE RVERS	all	all	all
<code>/usr/lpp/ mmfs/bin/ tsslspolicy <fs> -L -Y --ptn</code>	all	CALLHOME_SE RVERS	all	all	all
<code>/usr/lpp/ mmfs/bin/ mmlsfileset <fs> -L -Y</code>	all	CALLHOME_SE RVERS	all	all	all
<code>/usr/lpp/ mmfs/bin/ mmlssnapshot <fs> -Y</code>	all	CALLHOME_SE RVERS	all	all	all
<code>/usr/lpp/ mmfs/bin/ mmdf <fs> -Y</code>	all	CALLHOME_SE RVERS	all	all	all
<code>echo -e "ver io_s fs_io_s vio_s rhist_s" /usr/lpp/ mmfs/bin/ mmpmon 2> /dev/null</code>	all	CALLHOME_SE RVERS	all	all	all

Table 22. Data collected and uploaded by call home (continued)

COMMAND	Machine type	Node	OS	Product	Schedule
/usr/lpp/ mmfs/bin/ tsgskkm print -- enddate- plus -- cert /var/ mmfs/ssl/ id_rsa_commi tted.cert	all	all	all	all	all
<GET_FILES_R EADABLE_BY_A LL> /var/ mmfs/ccr/ committed/*	all	all	all	all	all
<GET_FILES_R EADABLE_BY_A LL> /var/ mmfs/ hadoop/etc/ hadoop/*- site.xml	all	cesNodes	all	all	all
<GET_FILES_R EADABLE_BY_A LL> /var/ mmfs/ hadoop/etc/ hadoop/*- env.sh	all	cesNodes	all	all	all
/usr/bin/rpm -qi gpfs.smb gpfs.gss.pmc ollector gpfs.gss.pms ensors pmswift gpfs.pm- ganesha gpfs.nfs- ganesha gpfs.nfs- ganesha-doc gpfs.nfs- ganesha gpfs nfs- ganesha nfs- ganesha gpfs nfs- ganesha- utils	all	all	RHEL,SLES	all	all

Table 22. Data collected and uploaded by call home (continued)

COMMAND	Machine type	Node	OS	Product	Schedule
<pre> /usr/bin/rpm -q gpfs.license .adv gpfs.gui gpfs.kafka gpfs.scst gpfs.librdka fka gpfs.gpl gpfs.protoco ls-support gpfs.tct.cli ent gpfs.msg.en_ US gpfs.base gpfs.java gpfs.docs gpfs.adv gpfs.compres sion gpfs.gskit gpfs.hdfs- protocol gpfs.tct.ser ver gpfs.crypto gpfs.scaleap i </pre>	all	all	RHEL,SLES	all	all
<pre> /usr/bin/rpm -q gpfs.ess.too ls gpfs.ess.fir mware gpfs.ess.uti lity.tools </pre>	all	all	RHEL,SLES	ess	all
<pre> /usr/bin/rpm -q esagent.pLin ux </pre>	all	CALLHOME_SE RVERS	RHEL,SLES	ess	all
<pre> /opt/ibm/esa /bin/ isactivated </pre>	all	CALLHOME_SE RVERS	RHEL,SLES	ess	all
<pre> /usr/bin/cat /opt/ibm/es a/data/ IBM_ESAconfi g.properties </pre>	all	CALLHOME_SE RVERS	RHEL,SLES	ess	all

Table 22. Data collected and uploaded by call home (continued)

COMMAND	Machine type	Node	OS	Product	Schedule
<pre> /usr/bin/ dpkg-query --list gpfs.nfs- ganesha gpfs.nfs- ganesha-doc gpfs.nfs- ganesha- gpfs nfs- ganesha nfs- ganesha-doc nfs-ganesha- fsal python- nfs-ganesha gpfs.smb gpfs.gss.pms ensors pmswift gpfs.pm- ganesha gpfs.gss.pmc ollector </pre>	all	all	DEBIAN	all	all
<pre> /usr/bin/ dpkg-query --list gpfs.compres sion gpfs.gui gpfs.adv gpfs.msg.en- us gpfs.protoco ls-support gpfs.java gpfs.license .adv gpfs.crypto gpfs.docs gpfs.gskit gpfs.gpl gpfs.base </pre>	all	all	DEBIAN	all	all
<pre> /usr/lpp/ mmfs/bin/ mmlsvdisk -Y </pre>	all	CALLHOME_SE RVERS	all	ess	all
<pre> /usr/lpp/ mmfs/bin/ mmvdisk pdisk list --smart- data -- recovery- group all -L -Y </pre>	all	CALLHOME_SE RVERS	all	ess	all
<pre> /usr/lpp/ mmfs/bin/ mmvdisk nvmeof list -Y </pre>	all	CALLHOME_SE RVERS	all	ess	all

Table 22. Data collected and uploaded by call home (continued)

COMMAND	Machine type	Node	OS	Product	Schedule
/usr/lpp/mmfs/bin/mmvdisk recoverygroup list --declustered-array -Y	all	CALLHOME_SERVERS	all	ess	all
/usr/lpp/mmfs/bin/mmlscomp -Y	all	CALLHOME_SERVERS	all	ess	all
/usr/lpp/mmfs/bin/mmlsenclosure all -Y -N GUI_RG_SERVERS	all	CALLHOME_SERVERS	all	ess	all
/usr/lpp/mmfs/bin/mmlspdisk all -Y	all	CALLHOME_SERVERS	all	ess	all
/usr/lpp/mmfs/bin/tlsencslot -adY	all	CALLHOME_SERVERS	all	ess	all
/usr/lpp/mmfs/bin/mmlsrecoverygroup -Y	all	CALLHOME_SERVERS	all	ess	all
/usr/lpp/mmfs/bin/mmlsrecoverygroupevents <rg> --days 2	all	CALLHOME_SERVERS	all	ess	daily
/usr/lpp/mmfs/bin/mmlscluster --cnfs -Y	all	CALLHOME_SERVERS	all	all	all
/usr/lpp/mmfs/bin/mmlspool <fs> all -Y	all	CALLHOME_SERVERS	all	all	all
/usr/lpp/mmfs/bin/mmccr vlist	all	CALLHOME_SERVERS	all	all	all

Table 22. Data collected and uploaded by call home (continued)

COMMAND	Machine type	Node	OS	Product	Schedule
<pre>/usr/bin/cat /var/ mmfs/gen/ mmfs.cfg.sho w 2> /dev/ null { /usr/lpp/ mmfs/bin/ mmlsconfig > /dev/null && cat /var/ mmfs/gen/ mmfs.cfg.sho w; }</pre>	all	CALLHOME_SE RVERS	all	all	all
<pre>/usr/lpp/ mmfs/bin/ mmsysmonc d cfgshow</pre>	all	all	all	all	all
<pre>/usr/lpp/ mmfs/bin/ mmsysmonc d cmdstats</pre>	all	all	all	all	all
<pre>/usr/lpp/ mmfs/bin/ mmsysmonc callhome ensureAndGet CCRConfig -- json</pre>	all	CALLHOME_SE RVERS	all	all	all
<pre>/usr/lpp/ mmfs/bin/ mmsysmonc callhome getClusterOv erview -- wait 300</pre>	all	CALLHOME_SE RVERS	all	all	all
<pre>/usr/lpp/ mmfs/bin/ mmsysmonc noderoles list -Y</pre>	all	all	all	all	all
<pre>/usr/bin/cat /var/ mmfs/gen/ scaleadmd_co nfig.toml 2> /dev/null</pre>	all	all	all	all	all
<pre>/usr/bin/cat /var/ mmfs/etc/ scaleadmd_co nfig.toml 2> /dev/null</pre>	all	all	all	all	all

Table 22. Data collected and uploaded by call home (continued)

COMMAND	Machine type	Node	OS	Product	Schedule
<code>/usr/lpp/mmfs/bin/mmlsrecovery groupevents <rg> --days 8</code>	all	CALLHOME_SERVERS	all	ess	weekly
<code>/usr/bin/numactl -- hardware</code>	all	all	all	all	all
<code>/usr/bin/numastat -nm</code>	all	all	all	all	all
<code>/usr/sbin/ip route</code>	all	all	all	all	all
<code>/usr/bin/cat / proc/net/dev</code>	all	all	all	all	all
<code>/usr/sbin/ip a</code>	all	all	all	all	all
<code>lsqeth -p</code>	s390x	all	all	all	all
<code>/usr/lpp/mmfs/bin/ network.call home.py <lcdir></code>	all	all	all	all	all
<code>/usr/sbin/sysctl -a</code>	all	all	all	all	all
<code>/usr/bin/uptime</code>	all	all	all	all	all
<code>/usr/bin/df -kl</code>	all	all	all	all	all
<code>lspci -vv</code>	all	all	all	all	all
<code>/usr/sbin/lsmmod</code>	all	all	all	all	all

Table 22. Data collected and uploaded by call home (continued)

COMMAND	Machine type	Node	OS	Product	Schedule
<code>/usr/bin/cat /proc/device-tree/system-id</code>	power	all	all	all	all
<code>ppc64_cpu --smt;ppc64_cpu --cores-present;ppc64_cpu --cores-on</code>	power	all	all	all	all
<code>/usr/bin/cat /proc/cpuinfo</code>	all	all	all	all	all
<code>/usr/bin/cat /proc/meminfo</code>	all	all	all	all	all
<code>/sbin/lscss --vpm --avail</code>	s390x	all	all	all	all
<code>/usr/sbin/dmidecode</code>	x86	all	all	all	all
<code>/usr/lpp/mmfs/bin/mmces service list --verbose -Y</code>	all	CALLHOME_SERVERS	all	all	all
<code>/usr/lpp/mmfs/bin/mmuserauth service list -Y</code>	all	CALLHOME_SERVERS	all	all	all
<code>/usr/lpp/mmfs/bin/ece.callhome.py</code>	all	all	all	ess	all
<code>/usr/bin/ofed_info -s</code>	all	all	all	all	all
<code>/usr/sbin/ibstat</code>	all	all	all	all	all
<code>/usr/sbin/ibdev2netdev</code>	all	all	all	all	all

Table 22. Data collected and uploaded by call home (continued)

COMMAND	Machine type	Node	OS	Product	Schedule
<code>/usr/bin/cat /etc/os-release</code>	all	all	all	all	all
<code>/usr/bin/uname -a</code>	all	all	all	all	all
<code>/usr/bin/ls -l /sys/class/sas_host /usr/bin/sort -k1.5 -n while read a ; do echo "# \$a" ; /usr/bin/cat /sys/class/sas_host/\${a}/device/scsi_host/\${a}/version* ; done</code>	all	all	all	all	all
<code>/usr/lpp/mmfs/bin/mmdia --memory -Y</code>	all	all	all	all	all
<code>/usr/lpp/mmfs/bin/mmdia --health -Y</code>	all	all	all	all	all
<code>/usr/lpp/mmfs/bin/mmdia --lroc -Y</code>	all	all	all	all	all
<code>/usr/lpp/mmfs/bin/mmdia --commands -Y</code>	all	all	all	all	all
<code>/usr/lpp/mmfs/bin/mmdia --waiters -Y</code>	all	all	all	all	all
<code>/usr/lpp/mmfs/bin/mmdia --deadlock -Y</code>	all	all	all	all	all

Table 22. Data collected and uploaded by call home (continued)

COMMAND	Machine type	Node	OS	Product	Schedule
<code>/usr/lpp/ mmfs/bin/ mmdiag -- rpc 24h -Y</code>	all	all	all	all	all
<code>/usr/lpp/ mmfs/bin/ mmdiag -- network -Y</code>	all	all	all	all	all
<code>/usr/lpp/ mmfs/bin/ mmdiag -- nsd -Y</code>	all	all	all	all	all
<code>/usr/lpp/ mmfs/bin/ mmdiag -- afm -Y</code>	all	all	all	all	all
<code>/usr/lpp/ mmfs/bin/ mmdiag -- stats -Y</code>	all	all	all	all	all
<code>/usr/lpp/ mmfs/bin/ mmdiag -- tokenmgr -Y</code>	all	all	all	all	all
<code>/usr/lpp/ mmfs/bin/ mmhealth cluster show -- verbose -Y</code>	all	CALLHOME_SE RVERS	all	all	all
<code>/usr/lpp/ mmfs/bin/ mmperfmon config show</code>	all	CALLHOME_SE RVERS	all	all	all
<code>/usr/lpp/ mmfs/bin/ mmperfmon query -- list=keys /usr/bin/w c</code>	all	CALLHOME_SE RVERS	all	all	all
<code>/usr/lpp/ mmfs/bin/ mmperfmon query -- list=keystat</code>	all	CALLHOME_SE RVERS	all	all	all

Table 22. Data collected and uploaded by call home (continued)

COMMAND	Machine type	Node	OS	Product	Schedule
<code>/usr/lpp/mmfs/bin/perfmon.callhome.py <lcdir></code>	all	all	all	all	daily
<code>/containermon/callhome/perfmon.container.callhome.py <lcdir></code>	all	all	all	all	all
<code>/usr/lpp/mmfs/lib/mmsysmon/container/openshift.collect.py Pod DaemonSet Deployment StatefulSet</code>	all	CALLHOME_SERVERS	all	all	all
<code>/usr/lpp/mmfs/samples/vdisk/gnrhealthcheck</code>	all	CALLHOME_SERVERS	all	ess	weekly
<code>/usr/lpp/mmfs/bin/mmdiag --eventproducer -Y</code>	all	all	all	all	all
<code>/usr/lpp/mmfs/bin/mmwatch all functional --list-clustered-watch-config</code>	all	CALLHOME_SERVERS	all	all	all
<code>/usr/lpp/mmfs/bin/mmaudit all functional --list-all-file systems-config</code>	all	CALLHOME_SERVERS	all	all	all
<code>/usr/lpp/mmfs/bin/gui.callhome.py <lcdir></code>	all	all	all	all	all

Table 22. Data collected and uploaded by call home (continued)

COMMAND	Machine type	Node	OS	Product	Schedule
/ containermon /callhome/ gui.callhome .py <lcdir>	all	all	all	all	all
nstat -asj	all	all	all	all	all
/usr/bin/ tail -n 10000 /var/a dm/ras/ mmfs.log.la test	all	all	all	all	all
/usr/bin/ journalctl -k -p 5 -n 10000 -r -- since="yeste rday"	all	all	all	all	all
/usr/bin/ systemctl show --no- pager mmsysmon	all	all	all	all	all
/usr/bin/ systemctl show --no- pager pmcollector	all	all	all	all	all
/usr/bin/ systemctl show --no- pager pmsensors	all	all	all	all	all
/usr/lpp/ mmfs/bin/ mmhealth node show -v -Y	all	all	all	all	all
/usr/lpp/ mmfs/bin/ mmhealth thresholds list -Y	all	CALLHOME_SE RVERS	all	all	all

Table 22. Data collected and uploaded by call home (continued)

COMMAND	Machine type	Node	OS	Product	Schedule
<code>/usr/lpp/mmfs/bin/mmhealth config interval -Y</code>	all	CALLHOME_SERVERS	all	all	all
<code>/usr/lpp/mmfs/bin/mmhealth event list hidden -Y</code>	all	CALLHOME_SERVERS	all	all	all
<code>/usr/lpp/mmfs/bin/mmhealth node eventlog --day -Y</code>	all	all	all	all	daily
<code>/usr/bin/systemctl show --no-pager mmsysmon</code>	all	all	all	all	all
<code>/usr/bin/systemctl show --no-pager pmcollector</code>	all	all	all	all	all
<code>/usr/bin/systemctl show --no-pager pmsensors</code>	all	all	all	all	all
<code>/usr/bin/tail -n 10000 /var/adm/ras/mmaudit.log</code>	all	all	all	all	all
<code>/usr/bin/tail -n 10000 /var/adm/ras/mmwf.log</code>	all	all	all	all	all
<code>/usr/lpp/mmfs/bin/mmhealth node eventlog --week -Y</code>	all	all	all	all	weekly

Table 22. Data collected and uploaded by call home (continued)

COMMAND	Machine type	Node	OS	Product	Schedule
<code>/usr/lpp/mmfs/bin/tsplatformstat -a</code>	all	all	all	ess3000,ess3200,ess3500	all
<code>/opt/ibm/ess/tools/bin/ess3kplt --local -Y</code>	all	all	all	ess3000	weekly
<code>echo /install/?ss/manifest*; /usr/bin/cat /install/?ss/manifest*</code>	all	all	all	ess	all
<code>/usr/lpp/mmfs/bin/mmsmb export list --all -Y</code>	all	CALLHOME_SERVERS	all	all	all
<code>/usr/lpp/mmfs/bin/mmsmb config list -Y</code>	all	CALLHOME_SERVERS	all	all	all
<code>/usr/lpp/mmfs/bin/mmqs device status <fs> -Y</code>	all	CALLHOME_SERVERS	all	all	all
<code>/usr/lpp/mmfs/bin/mmqs report list <fs> -Y</code>	all	CALLHOME_SERVERS	all	all	all
<code>/usr/lpp/mmfs/bin/mmnfs config list -Y</code>	all	CALLHOME_SERVERS	all	all	all
<code>/usr/lpp/mmfs/bin/mmnfs export list -Y</code>	all	CALLHOME_SERVERS	all	all	all

Table 22. Data collected and uploaded by call home (continued)

COMMAND	Machine type	Node	OS	Product	Schedule
<code>/usr/lpp/ mmfs/lib/ mmsysmon/ noobaa_api.py noobaa_call_ home > <lcdir>/ noobaa_data. txt</code>	all	all	all	all	all
<code>/usr/bin/cat /etc/ ganesha/ config/ gpfs.ganesha .main.conf</code>	all	all	all	all	all
<code>/ containermon / get_ganesha_ stats.py</code>	all	all	all	all	all
<code>/usr/lpp/ mmfs/bin/ mmsysmonc d haldump -- drive_slot all</code>	all	all	all	ess hal	all
<code>/opt/ibm/ess /hal/bin/ sastopology -j</code>	all	all	all	ess3500	all
<code>/opt/ibm/ess /hal/bin/ vpd_dump</code>	all	all	all	ess hal	all
<code>/usr/lpp/ mmfs/bin/ mmsysmonc d haldump</code>	all	all	all	ess hal	all
<code>/usr/bin/cat /etc/ noobaa.conf. d/ config_dir_r edirect</code>	all	all	all	all	all
<code>/usr/bin/cat <noobaa_conf _path>/ system.json</code>	all	all	all	all	all

Table 22. Data collected and uploaded by call home (continued)

COMMAND	Machine type	Node	OS	Product	Schedule
<code>/usr/bin/cat <noobaa_conf _path>/ config.json</code>	all	all	all	all	all
<code>/usr/bin/ls -l <noobaa_conf _path>/ accounts</code>	all	all	all	all	all
<code>/usr/lpp/ mmfs/bin/ mms3 account list</code>	all	all	all	all	all
<code>/usr/bin/cat /etc/ rsyslog.d/ noobaa_syslo g.conf</code>	all	all	all	all	all
<code>/usr/bin/cat /etc/ rsyslog.d/ noobaa_syslo g.conf</code>	all	all	all	all	all
<code>/usr/bin/cat /etc/ logrotate.d/ noobaa- logrotate</code>	all	all	all	all	all
<code>/usr/bin/ systemctl show --no- pager rsyslog</code>	all	all	all	all	all
<code>/usr/bin/ systemctl show --no- pager noobaa</code>	all	all	all	all	all
<code>/usr/bin/ systemctl show --no- pager syslog-ng</code>	all	all	all	all	all

Table 22. Data collected and uploaded by call home (continued)

COMMAND	Machine type	Node	OS	Product	Schedule
<code>/usr/bin/journalctl -u noobaa.service</code>	all	all	all	all	all
<code>/usr/local/bin/noobaa-cli diagnose health</code>	all	all	all	all	all
<code>/usr/bin/rpm -qi gpfs.mms3 noobaa-core</code>	all	all	RHEL,SLES	all	all
<code>/usr/lpp/mmfs/bin/mmcallservice ticket list -Y</code>	all	CALLHOME_SERVERS	all	ess	all
<code>/usr/bin/tail -n 10000 /var/adm/ras/mmsysmonitor.log</code>	all	all	all	all	all

Inspecting call home data uploads

The uploaded call home data is generally used by the IBM service and support teams for diagnosing issues and best practice violations, which can be reported at a customer site. However, a customer can opt to self-inspect a sample call home data upload to manually review the data without sending out any information to IBM.

Using review mode to inspect call home data uploads

Review mode allows the call home feature to run tasks such as daily, weekly, and SendFile without automatically uploading the collected data to IBM. The collected jobs are marked as under review so that you can inspect the generated data collection files and explicitly upload them later.

Ensure that the review mode is enabled by running the following command:

```
mmcallservice capability enable --feature REVIEW
```

Follow these steps to run call home tasks in review mode and review the collected data before uploading it:

1. Run call home tasks in review mode so that the collected data is not uploaded automatically.

```
mmcallservice run GatherSend --task daily
```

A sample output is as follows:

```
JobID 20260221145014.189 marked as under-review
Run mmcallhome run uploadReviewed --jobid 20260221145014.189 to upload the job
```

2. Run a SendFile task.

```
mmcallhome run SendFile --file /tmp/mmfs/
basic.snap.cluster_and_local.scale-51.20260219_130024.tar.gz
```

A sample output is as follows:

```
Running sendFile... (In case of network errors, it may take over 20 minutes for retries.)
Successfully marked the file upload for a review with the jobID 20260306090859.145
Run mmcallhome run uploadReviewed --jobid 20260306090859.145 to upload the file
```

3. To list all pending jobs and identify the data collection files to review, run the following command:

```
mmcallhome status list --verbose
```

A sample output is as follows:

Group RC or Step Name	Task	Start Time Package File	Updated Time	Review Time	Status
Filename		Original			
-----		-----			
-----		-----			
autoGroup_1	daily	20260418020014.702	2026-04-18 02:00:53		under_review
Marked as under-review		/tmp/mmfs/callhome/rsENHomeQ/ 16443032091191.6_0_1_0.123456.US.ibmtest.autoGroup_1.gat_daily.g_daily.scale.2026041802001470 2.cl0.DC			
autoGroup_1	sendfile	20260420131922.669	2026-04-20 13:19:22		under_review
Marked as under-review		/tmp/mmfs/callhome/rsENHomeQ/ 16443032091191.6_0_1_0.123456.US.ibmtest.autoGroup_1.HEARTBEAT.s_file.scale.20260420131922669 .cl0.DC			
		/tmp/mmfs/basic.snap.cluster_and_local.scale-51.20260420_131922.tar.gz			

Jobs with an under_review status are pending review, and their file paths are listed in the Package File Name column.

4. Review the data collection files.

The data collection files listed in the Package File Name column are located on the call home server node of the corresponding call home group.

For more information about extracting and analyzing these files, see [“Extracting and inspecting the data collection file” on page 292](#).

5. Upload the reviewed job:

- To upload a single job with the job ID:

```
mmcallhome run uploadReviewed --jobid <jobid>
```

For example:

```
mmcallhome run uploadReviewed --jobid 20260306090742.705
```

A sample output is as follows:

```
weekly job 20260306090742.705 uploaded successfully.
```

- To upload all jobs that are currently marked as under review:

```
mmcallhome run uploadReviewed
```

Sample outputs are as follows:

```
weekly job 20260221145106.931 uploaded successfully.  
sendfile job 20260308091042.401 uploaded successfully.
```

6. To view all executed call home tasks:

```
mmcallhome status list
```

A sample output is as follows:

Group	Task	Start Time	Review Time	Status
autoGroup_2	weekly	20260306090742.705	2026-03-06 09:09:17	success
autoGroup_2	sendfile	20260306091233.739		under_review

You can use the `--verbose` option for more detailed information. When this option is specified, only the pathnames of the files to review are displayed.

Configuring incorrect proxy settings on purpose

You can configure incorrect proxy settings to prevent automatic uploads. This approach helps simulate failed uploads while still collecting the data for inspection.

1. Set up the call home feature with incorrect proxy settings.

For example:

```
mmcallhome proxy change --proxy-location wrong_proxy --proxy-port 1111  
Call home proxy-location has been set to wrong_proxy  
Call home proxy-port has been set to 1111  
  
mmcallhome proxy enable  
Call home proxy-auth-enabled has been set to false  
Call home proxy-enabled has been set to true
```

2. Start the uploaded data collection, such as daily uploaded data. The daily data upload overrides the standard 24-hour wait time and delivers the results immediately.

For example:

```
mmcallhome run GatherSend --task daily  
One time run completed with failure: sending failed
```

As expected, the data uploads to the server fails as you provided incorrect proxy settings.

3. To find the path names of the data collection files on the call home server node, run the following command:

```
mmcallhome status list --verbose
```

For more information about extracting and analyzing these files, see [“Extracting and inspecting the data collection file”](#) on page 292.

4. Run the **mmcallhome proxy** command to enable the call home feature and do one of the following actions:

- If you are using a proxy, then provide correct proxy settings for call home usage by using the following command:

```
mmcallhome proxy change --proxy-location <CORRECT_PROXY> --proxy-port <CORRECT_PORT>
```

- If you are using a proxy with authentication, then issue the following command:

```
mmcallhome proxy enable --with-proxy-auth  
Call home proxy-auth-enabled has been set to true  
Call home proxy-enabled has been set to true
```

- If you are not using a proxy, then disable the proxy usage for call home by using the following command:

```
mmcallhome proxy disable
Call home proxy-auth-enabled has been set to false
Call home proxy-enabled has been set to false
```

Extracting and inspecting the data collection file

After the call home data collection file is generated, you can locate the file on the call home server node, extract its contents, and inspect the collected diagnostic data before it is uploaded.

1. Locate the data collection file that contains the data collected by the call home feature.

For example:

```
mmcallhome status list --task daily --verbose --numbers 1
=== Executed call home tasks ===

Group      Task      Start Time          Updated Time      Status  RC or Step
Package File
Name
-----
Original Filename
-----
autoGroup_1 daily  20220126153252.022  20220126153310   failed  Uploading operation
failed /tmp/mmfs/callhome/rsENFailedQ/
14559237835643.5_1_3_0.123456.US.ibmtest.autoGroup_1.gat_daily.g_daily.scale.2022012615325202
2.cl0.DC
```

2. Extract the data collection file, which is a *.tar.gz file.

For example:

```
mkdir wrong_proxy
cd wrong_proxy/
tar xvf /tmp/mmfs/callhome/rsENFailedQ/
14559237835643.5_1_3_0.123456.US.ibmtest.autoGroup_1.gat_daily.g_daily.scale.2022012615325202
2.cl0.DC
CH_20220126153252.022_daily/
CH_20220126153252.022_daily/HEADER
CH_20220126153252.022_daily/LC_g5130-14.tar.gz
CH_20220126153252.022_daily/LC_g5130-13.tar.gz
CH_20220126153252.022_daily/LC_g5130-12.tar.gz
CH_20220126153252.022_daily/MC_g5130-11.tar.gz
tar xvf CH_20220126153252.022_daily/MC_g5130-11.tar.gz
MC_g5130-11/
MC_g5130-11/cat_etcosrelease.txt
...
```

3. Inspect every detail in the data collection file.

Benefits of enabling call home

You should configure and enable the call home feature as it improves the user experience when IBM Storage Scale is used.

The call home feature when enabled provides the following benefits:

1. **Improves customer service response time**

- *Scheduled data uploads:*

If the call home feature is enabled, then the IBM support representatives can start with the analysis of the issue that is reported without any delay. The IBM support team can use the cluster data from the call home scheduled uploads to do an immediate analysis of the issue. For more information, see *Scheduled data upload* in the *IBM Storage Scale: Concepts, Planning, and Installation Guide*.

However, if the call home feature is not enabled, then the IBM support team first collects debugging data from command outputs like **gpfs.snap** command. The whole process takes time to request, collect, and deliver the data, which might cause a delay in resolving the issue on the site.

- *Selected failure events trigger automatic call home uploads:*

If the call home feature is enabled, then the event-based uploads feature automatically collects the relevant component-specific information and uploads it to the IBM ECuRep server. The uploaded data enables the IBM support team to immediately find the root cause of the issue that is reported on the customer site. For more information, see *Event-based uploads* in the *IBM Storage Scale: Concepts, Planning, and Installation Guide*.

2. Proactively detects issues

- Detects violations against best practices and common misconfiguration.
- Finds out which specific customer site is affected by the reported issue by using tools like High Impact Programming Error Authorized Program Analysis Report (HIPER APARs).
- Constantly extends the list of automatically detected issues, best practice violations, and common misconfiguration over time.

Note: The `mmhealth` command output and IBM Storage Scale GUI display the list of events that are resulted from proactively detected issues. For more information, see *Proactive system health alerts* in the *IBM Storage Scale: Problem Determination Guide*.

3. Improves customer support experience in areas of consumability and ease of usage

- *Uploads gpfs snaps directly from the IBM Storage Scale system:*

If the call home feature is enabled, then a customer can easily find out the customer number. The requested debugging data is transferred from the customer site cluster or node to the IBM support team by using the `mmcallhome` command.

- *Views system configuration changes that take places between call home data uploads:*

If the call home feature is enabled, then you can run the `mmcallhome status diff` command to easily detect the system configuration changes that took place between different time spans.

4. Generates reports based on customer insights for support and development

- Improves test and development for the features that are used by your cluster.
- Provides more data to analyze the issues that are reported on the customer site. For example, the root cause of an issue can be diagnosed by using the historic call home data.

5. Creates service tickets automatically for reported failures

In Elastic Storage Server (ESS), if the call home feature is enabled, then the system automatically creates Salesforce service tickets for any hardware failures that are reported on customer sites.

Note: A user can configure the call home settings on IBM Storage Scale or on Elastic Storage Server (ESS). This feature enhancement is applicable only to Elastic Storage Server (ESS) and automatically opens a Salesforce service ticket to report a hardware issue.

Data privacy with call home

A licensee can configure the IBM Storage Scale to automatically send specific cluster information to the IBM support team by using the Call Home feature.

When a licensee enables the Call Home feature, the licensee provides the support contact information, such as names, phone numbers, or email addresses, which are needed during the Call Home feature configuration and activation process.

By default, when a licensee decides to activate and use the Call Home feature, the licensee agrees to allow IBM and its subsidiaries to store and use the licensee's support contact information. This information is processed and used with IBM business relationship and might be shared with third-parties under the direction of IBM. For example, IBM support center representatives or assignees of IBM and its subsidiaries.

The support contact information can be used for processing business orders, business promotions, problem determination, or market research.

When the call home feature is configured and activated, the IBM Storage Scale collects all monitoring information that is related to system utilization, performance, capacity planning, and service

maintenance. The service information includes system failure logs, part numbers, machine serial number, software version, maintenance levels, installed patches, and configuration values.

When the Call Home feature is enabled, the licensee allows IBM to use and share the data, which is gathered from system monitoring functions, within IBM and with IBM business partners and third-parties, such as IBM support center sub-contractors or assignees. The shared data is used only under the direction of IBM for the following defined purposes:

- Determining a problem.
- Assisting licensee with performance and capacity planning.
- Assisting IBM to enhance IBM products and services.
- Notifying licensee about the licensee's system status and available solutions.

Note: The licensee information excludes the collection and transmission of licensee's financial, statistical, and personal data, and licensee's business plans.

When the Call Home feature is enabled, a licensee agrees to the fact that the licensee's support contact information can be transferred to countries that might not be a member of the European Union. A licensee can disable the Call Home feature at any time.

Call home monitors for PTF updates

IBM Storage Scale provides temporary fixes to its code in between planned releases. The temporary fixes to the codes are called Program Temporary Fix (PTF).

Types of PTF monitors

The call home function has the following two monitors that provide PTF information to a user:

Callhome Update Request

The **Callhome Update Request** monitor is scheduled to run only once a day and exclusively on the first call home master node of a cluster. The monitor requests the PTF update information from the eSupport server.

The monitor output list contains information about the PTF versions that are available for the IBM Storage Scale version that is installed on the master node. The PTF information is stored in the Cluster Configuration Repository (CCR) to make it available to all other cluster nodes.

The following events show up during the run time of the **Callhome Update Request** monitor:

callhome_ptfupdates_ok

Informs the user that the PTF update check is processed without any errors.

callhome_ptfupdates_disabled

Informs the user that the call home function is disabled, or the monitor is disabled through configuration.

callhome_ptfupdates_ccr_failed

Informs the user that an error occurred, and the update information cannot be written to the CCR.

callhome_ptfupdates_failed

Informs the user about a connectivity error when a PTF update information is requested from the eSupport server.

callhome_ptfupdates_noop

Informs the user that an event is raised on a call home master node that is not the first node in the list.

Note: When the first call home master node disappears, the next call home master node in the list automatically takes over the master role.

PTF Updates Check

The **PTF Updates Check** monitor is scheduled to run only once a day on all call home child nodes, including the master nodes of the cluster. The monitor reads the update information from the CCR and compares it against the IBM Storage Scale version that is installed on the local node. When an appropriate PTF update is found in the information during this check, a system health tip is raised to inform the system administrator about the software update that can be installed.

Note: The software update packages can be downloaded from the Fix Central. However, the software update package download is not done automatically.

The following events show up during the runtime of the **PTF Updates Check** monitor:

scale_up_to_date

Informs the user that the installed packages are up to date. No new PTF packages are available.

scale_updatecheck_disabled

Informs the user that either the call home function or the PTF monitor is disabled by setting the value of **scale_updatecheck_disabled** as true.

scale_ptf_update_available

Informs the user that the PTF updates are available for the current node.

Important:

The call home PTF update notification function can be disabled by changing the configuration value of the **mmhealthPtFUpdatesMonitorEnabled** to no. The **mmhealthPtFUpdatesMonitorEnabled** value can be set by using the following **mmchconfig** command.

```
mmchconfig mmhealthPtFUpdatesMonitorEnabled=no
```

Limitations of PTF monitors

The PTF monitors display the following limitations:

- The two PTF monitors consider only the PTF updates for the major version that is installed on the first call home master node. Other versions that might be installed on the other cluster nodes are not considered.
- The PTF monitors consider only the system architecture of the first call home master node.
- The PTF software update process does not run automatically. You must download the applicable PTF updates from the Fix Central before running them.

IBM Storage Scale in an OpenStack cloud deployment

OpenStack is a cloud operating system that controls large pools of compute, storage, and networking resources throughout a datacenter. Using IBM Storage Scale in an OpenStack cloud environment offers many benefits, including the enterprise features of IBM Storage Scale and consolidated storage options.


An introduction to OpenStack

OpenStack is an open source software platform that is widely used as the base to build cloud infrastructure as a service solution. OpenStack is typically deployed on commodity hardware and is used to virtualize various parts of the infrastructure (compute, storage, and network) to ease the sharing of the infrastructure across applications, use cases and workloads. IBM Storage Scale also supports the OpenStack Cinder component with Remote IBM Storage Scale Access Deployment mode on Red Hat OpenStack Platform 16.1.3 onwards. For more information see the topic *Integrating IBM Storage Scale Cinder driver with Red Hat OpenStack Platform 16.1* in the *IBM Storage Scale: Administration Guide*.

Using IBM Storage Scale in an OpenStack cloud deployment

Deploying OpenStack over IBM Storage Scale offers benefits that are provided by many enterprise features in IBM Storage Scale as well as the ability to consolidate storage for various OpenStack components and applications running on top of the OpenStack infrastructure under a single storage

management plane. One key benefit of IBM Storage Scale is that it provides uniform access to data under a single namespace with integrated analytics.



OpenStack components supported by IBM Storage Scale

- Cinder: Provides virtualized block storage for virtual machines. The IBM Storage Scale Cinder driver, also known as the GPFS driver, is written to take full advantage of the IBM Storage Scale enterprise features.
- Glance: Provides the capability to manage virtual machine images. When Glance is configured to use the same IBM Storage Scale fileset that stores Cinder volumes, bootable images can be created almost instantly by using the copy-on-write file clone capability.
- Swift: Provides Swift object storage to any user or application that requires access to data through a RESTful API. The Swift object storage configuration is optimized for the IBM Storage Scale environment, providing high availability and simplified management. Swift object storage also supports native the Swift APIs and Amazon S3 APIs for accessing data. Finally, the Swift object storage also supports access to the same data through either Swift object interface or file interface (POSIX, NFS, SMB) without creating a copy.
- Keystone: Although not a storage component, internal keystone with in-built HA is provided by IBM Storage Scale as part of the Swift object protocol. In deployments that already have keystone support, the Swift object protocol can be configured to use the external keystone server rather than the internal one.

IBM Storage Scale supports OpenStack Cinder component from Ussuri release in OpenStack Community.

The following table lists the available features that IBM Storage Scale supports for deploying the OpenStack cloud storage:

<i>Table 23. Features that IBM Storage Scale supports for deploying the OpenStack cloud storage</i>	
Feature	Support
Volume and Image Management (Cinder and Glance)	
Volume Creation and Deletion	Yes
Volume Creation from Snapshot	Yes
Image Creation from Volume	Yes
Volume Creation from another Volume	Yes
Live migration of instances	Yes
Encryption of volumes	Yes
Consistency Group Snapshot Creation and Deletion	Yes
Volume Snapshot Creation and Deletion	Yes
Extend (non-attached) volumes	Yes
Consistency Group Creation and Deletion	Yes
Volume Creation from Image	Yes
Instantaneous Boot Volume Create From Glance Repo	Yes
Backup of volumes	Yes
Quality of service using multi-tier storage (with Flash support)	Yes
Tiering of volumes	Yes
Volume Attach and Detach to a VM instance	Yes
Instantaneous Boot Volume Create From Glance Repo	Yes
Identity Management (Keystone)	
Integrated High Availability across the IBM Storage Scale protocol nodes	Yes
AD and LDAP Support	Yes
Easy configuration, management, and monitoring	Yes
External keystone support	Yes
Object store features (Swift)	
Unified file and object support	Yes
High performance and High Availability	Yes
Object encryption	Yes
WAN caching with Active File Management (AFM)	Yes
Easy install, configuration, and management	Yes

Table 23. Features that IBM Storage Scale supports for deploying the OpenStack cloud storage (continued)

Feature	Support
Swift and S3 API Support	Yes
Support for OpenStack Swift object store features	Yes
In-place analytics with Hadoop compatibility	Yes
Object compression	Yes
Multi-region support	Yes
Policy based information life cycle management	Yes
Integrated monitoring	Yes
Large object support (5 TB)	Yes

For more information on OpenStack, see the [OpenStack Redpaper](#).

IBM Storage Scale product editions

IBM Storage Scale offers different editions that are based on functional levels.

Note: All IBM Storage Scale software nodes in a cluster must have the same edition installed. Capacity licensing and socket licensing cannot be mixed in the same cluster, with the exception of ESS. ESS can be mixed with IBM Storage Scale software in the same cluster when one of the following requirements is met:

- All nodes in the cluster run Data Access Edition or Standard Edition.
- All nodes in the cluster run Data Management Edition or Advanced Edition.

For information about capacity licensing, see “Capacity-based licensing” on page 304. For information about the features available in respective IBM Storage Scale editions, see [Table 24 on page 299](#).

IBM Storage Scale Data Access Edition

Available on AIX, Linux, and Windows. This edition provides base IBM Storage Scale functions. On AIX and Linux, the available features include Information Lifecycle Management (ILM), Active File Management (AFM), and Clustered NFS (CNFS). CNFS is not available on Linux on Z. On Windows, the available features include limited Information Lifecycle Management (ILM).

On Red Hat Enterprise Linux 7.x and 8.x, SLES 15, Ubuntu 20.04, and Ubuntu 22.04, the available features include the ability to enable and use the additional protocol functionality integration (NFS and SMB). On Red Hat Enterprise Linux 8.x, the object protocol functionality is also supported.

IBM Storage Scale Data Management Edition

Available on AIX and Linux. This edition provides all the features of the Data Access Edition and certain additional features. For more information, see [Table 24 on page 299](#).

IBM Storage Scale Developer Edition

Available on Red Hat Enterprise Linux on x86_64. This edition provides all the features of the Data Management Edition and it is limited to 12 TB per cluster. You can check your licensed usage on IBM Storage Scale Developer Edition by using the **mmlslicense --licensed-usage** command. For more information, see *mmlslicense command* in *IBM Storage Scale: Command and Programming Reference Guide*.

Note: This edition is made available to the customers free of cost to try IBM Storage Scale features in test setups. Its use in production is not allowed. There is no support from IBM for IBM Storage Scale Developer Edition.

You can sign up for IBM Storage Scale Developer Edition [here](#).

IBM Storage Scale Erasure Code Edition

This edition provides identical functionality to IBM Storage Scale Data Management Edition plus the support for storage rich servers. IBM Storage Scale Erasure Code Edition provides network-dispersed erasure coding, distributing data, and metadata across the internal disks of a cluster of servers. This allows IBM Storage Scale to use internal disks as reliable storage with low overhead and high performance.

IBM Storage Scale Backup

IBM Storage Scale Backup delivers deeper synergies with IBM Storage Protect to offer a licensing bundle for IBM Storage Scale clients looking to protect their data.

Legacy editions:

IBM Storage Scale Standard Edition

This edition provides comparable functionality to IBM Storage Scale Data Access Edition under socket-based licensing. Standard Edition does not provide erasure coding and it is not supported with ESS.

IBM Storage Scale Advanced Edition

This edition provides comparable functionality to IBM Storage Scale Data Management Edition under socket-based licensing. Advanced Edition does not provide erasure coding and it is not supported with ESS.

Note:

- IBM Storage Scale Standard Edition licenses can only be purchased by existing licensees of Standard Edition. IBM Storage Scale Advanced Edition licenses can only be purchased by existing licensees of Advanced Edition. All others must purchase Data Access Edition, Data Management Edition, or Erasure Code Edition.
- IBM Storage Scale Standard Edition and IBM Storage Scale Advanced Edition can be licensed by IBM Storage Scale Client license, IBM Storage Scale Server license, and IBM Storage Scale FPO license. For more information, see [“IBM Storage Scale license designation”](#) on page 300.

Table 24 on page 299 lists the availability of key features in the IBM Storage Scale editions.

Table 24. Features in IBM Storage Scale editions			
Feature	Data Access	Data Management ¹	Erasure Code Edition
Multi-protocol scalable file service with simultaneous access to a common set of data	✓	✓	✓
Facilitate data access with a global namespace, massively scalable file system, quotas and snapshots, data integrity and availability, and filesets	✓	✓	✓
Simplify management with GUI	✓	✓	✓
Improved efficiency with QoS and compression	✓	✓	✓
Create optimized tiered storage pools based on performance, locality, or cost	✓	✓	✓

Table 24. Features in IBM Storage Scale editions (continued)

Feature	Data Access	Data Management ¹	Erasure Code Edition
Simplify data management with Information Lifecycle Management (ILM) tools that include policy based data placement and migration	✓	✓	✓
Enable worldwide data access using AFM asynchronous replication	✓	✓	✓
Asynchronous multi-site Disaster Recovery		✓	✓
Multi-site replication with AFM to cloud object storage		✓	✓
Protect data with native software encryption and secure erase, NIST compliant and FIPS certified		✓	✓
Tenant management		✓	✓
File audit logging		✓	✓
Watch folder		✓	✓
Erasure coding	ESS only	ESS only	✓
IBM Storage Fusion Data Cataloging service		✓	✓

Note: ¹IBM Storage Scale Developer Edition provides the same features as Data Management Edition and it is limited to 12 TB per cluster.

Consult the [IBM Storage Scale FAQ in IBM Documentation](#) for the latest features that are included in each edition.

IBM Storage Scale license designation

According to the IBM Storage Scale Licensing Agreement, each server that is licensed with IBM Storage Scale Standard Edition or IBM Storage Scale Advanced Edition in the cluster must be designated as possessing an IBM Storage Scale Client license, an IBM Storage Scale File Placement Optimizer (FPO) license, or an IBM Storage Scale Server license.

There are three basic storage cluster models, which are storage area network (SAN), network shared disk (NSD), and shared nothing. IBM Storage Scale supports all three of them. For SAN and NSD clusters, various Server licenses are required in addition to Client licenses. For shared nothing clusters of storage rich servers, FPO licenses are required on each of those servers.

The type of license that is associated with any one server depends on the functional roles that the server is designated to perform.

IBM Storage Scale Client license

The IBM Storage Scale Client license permits exchange of data between servers that locally mount the same IBM Storage Scale file system. No other export of the data is permitted.

IBM Storage Scale Server license

The IBM Storage Scale Server license permits the licensed server to perform management functions, such as cluster configuration manager, quorum node, manager node, AFM gateway node, and Network Shared Disk (NSD) server. In addition, the IBM Storage Scale Server license permits the licensed server to share IBM Storage Scale data directly through any application, service protocol or method such as Network File System (NFS), Server Message Block (SMB), File Transfer Protocol (FTP), Hypertext Transfer Protocol (HTTP), Object Protocol (OpenStack Swift, Amazon S3 API). Therefore, protocol nodes also require an IBM Storage Scale Server license.

IBM Storage Scale FPO license

The IBM Storage Scale FPO license permits the licensed server to perform the NSD server functions for sharing IBM Storage Scale data with other servers that have an IBM Storage Scale FPO or an IBM Storage Scale Server license.

The full text of the Licensing Agreement can be found at the [IBM Terms website \(https://www.ibm.com/support/customer/csol/terms/\)](https://www.ibm.com/support/customer/csol/terms/).

Use the guidance in [Figure 34 on page 302](#) to decide which IBM Storage Scale license to buy for your requirements.




Figure 34. Guidance on which license to buy

These licenses are all valid for use in IBM Storage Scale Standard Edition and IBM Storage Scale Advanced Edition. For more information, see [“IBM Storage Scale product editions”](#) on page 298.

With IBM Storage Scale Data Access Edition, IBM Storage Scale Data Management Edition, and IBM Storage Scale Erasure Code Edition, you must use capacity-based licensing. For more information, see [“Capacity-based licensing”](#) on page 304.

- You can view the number and type of licenses that are currently in effect for the cluster by using the `mmfslslicense` command.
- If needed, you can change the type of license that is assigned to a node by using the `mmchlicense` command.

For more information, see **`mmfslslicense` command** and **`mmchlicense` command** in *IBM Storage Scale: Command and Programming Reference Guide*.

The following are IBM Storage Scale licensing considerations including considerations for using IBM Storage Scale with other offerings.

AFM-based Async Disaster Recovery (AFM DR) with multicluster

When using AFM-based Async Disaster Recovery (AFM DR) in a multi-cluster environment, both the home and the cache cluster require the IBM Storage Scale Advanced Edition, IBM Storage Scale Data Management Edition, or IBM Storage Scale Developer Edition.

AFM to cloud object storage with multi-site replication

When using AFM to cloud object storage multi-site replication, the cache cluster requires IBM Storage Scale Advanced Edition, IBM Storage Scale Data Management Edition, or IBM Storage Scale Erasure Code Edition.

Encryption

The encryption function that is available in IBM Storage Scale Advanced Edition, IBM Storage Scale Data Management Edition, or IBM Storage Scale Developer Edition requires a separate IBM Security Key Lifecycle Manager (ISKLM) license.

Encryption with multi-cluster

When using the IBM Storage Scale encryption function in a multicluster environment, each server in an IBM Storage Scale cluster requiring access to another cluster's encrypted file system must be licensed with IBM Storage Scale Advanced Edition (Client, Server, or FPO as appropriate) or IBM Storage Scale Data Management Edition.

Hadoop access

The IBM Storage Scale Hadoop connector can be used with all license types (Client, Server, FPO) and all Editions (Standard, Advanced, and Data Management). There is no additional license requirement for Hadoop access. A Hadoop node using the connector needs no IBM Storage Scale license. The Hadoop node can connect to a node in the IBM Storage Scale cluster, which is licensed as Server because it is exporting data, and access the file system directly by using the Hadoop connector.

ILMT service

The IBM License Metric Tool (ILMT) ensures efficiency in inventory deduction and better estimation of license consumption. To use the ILMT service, you must schedule the command **mmlslicense --ilmt-data** to run once a week, by using mechanisms like cron. For more information, see *mmlslicense* in the *IBM Storage Scale: Command and Programming Reference Guide*

IBM Storage Scale with IBM Storage Protect

When using IBM Storage Scale with IBM Storage Protect for backup and restore, each IBM Storage Scale node performing data movement requires an IBM Storage Protect license.

IBM Storage Scale with IBM Storage Protect for Space Management

When using IBM Storage Scale with IBM Storage Protect for Space Management for migration and recall, each IBM Storage Scale node performing data movement requires an IBM Storage Protect for Space Management license.

IBM Storage Scale with IBM Storage Archive Enterprise Edition (EE)

In addition to an IBM Storage Archive Enterprise Edition (EE) license, a server in the IBM Storage Scale cluster which is being used for IBM Storage Archive Enterprise Edition (EE) requires IBM Storage Scale to be installed with the correct license (ordered separately). The IBM Storage Archive Enterprise Edition (EE) server requires an IBM Storage Scale Server license, if the IBM Storage Archive Enterprise Edition (EE) server is also being used for any of the functions requiring an IBM Storage Scale Server license, such as NSD server, quorum, management, etc. Otherwise, the IBM Storage Archive Enterprise Edition (EE) server can be licensed with an IBM Storage Scale client license. The IBM Storage Archive Enterprise Edition (EE) server's function of moving data from the file system to a tape drive is not considered as a type of data export that requires an IBM Storage Scale server license.

Virtualization licensing

In an IBM Storage Scale environment containing VMs or LPARs, licenses are required for the sockets available to the IBM Storage Scale VMs or LPARs on a physical server. The number of licenses required

is the number of cores available to IBM Storage Scale. For more information, see *Virtualization questions* in *IBM Storage Scale FAQ* in IBM Documentation.

Capacity-based licensing

You can use capacity-based licensing to license IBM Storage Scale based on the storage capacity being managed in an IBM Storage Scale cluster.

The storage capacity to be licensed is the capacity in Terabytes (TiB) from all NSDs in the IBM Storage Scale cluster. The capacity to be licensed is not affected by using functions such as replication or compression or by doing tasks such as creating or deleting files, file systems, or snapshots.

For example, if you have 100 TiB of capacity from all NSDs in the IBM Storage Scale cluster and if you have set the replication factor to 2, for 50 TiB of application data, 100 TiB of disk capacity is consumed. In this case, the capacity to be licensed is 100 TiB, not 50 TiB.

The capacity to be licensed changes only when an NSD is added or deleted from the IBM Storage Scale cluster.

You can view the cluster capacity for licensing by using one of the following methods:

- Select the **About** option from the menu that is available in the upper right corner of the IBM Storage Scale management GUI.
- Issue the **mmlslicense --capacity** command.

Capacity-based licensing can be used only with IBM Storage Scale Data Access Edition, IBM Storage Scale Data Management Edition, and IBM Storage Scale Erasure Code Edition. The Data Access Edition provides identical functionality as the Standard Edition. The Data Management Edition provides identical functionality as the Advanced Edition. The Erasure Code Edition provides identical functionality as the Data Management Edition plus support for storage rich servers.

Capacity-based licensing can also be used to license an Elastic Storage Server (ESS) system. For ESS, the capacity to be licensed is the capacity after applying IBM Storage Scale RAID. The exact capacity to be licensed depends on the RAID code that is selected, 8+2p or 8+3p.

Dynamic pagepool

The dynamic pagepool feature in IBM Storage Scale adjusts the size of the pagepool memory dynamically. This feature is available only on IBM Storage Scale that runs on Linux nodes. The pagepool is managed between the minimum size and the maximum size.

When IBM Storage Scale detects shortage of the pagepool memory, it attempts to increase the pagepool size, but not beyond the configured maximum size. When the Linux kernel detects the memory pressure, it requests IBM Storage Scale to shrink the size of the pagepool. IBM Storage Scale attempts to shrink the size of the pagepool in response, but the pagepool is not shrunk less than the configured minimum size.

Dynamic pagepool limitations

- Nodes with ECE devices do not use the dynamic pagepool and revert to a static pagepool.
- The dynamic pagepool with RDMA uses large RDMA memory registrations (one for each 256 KiB of pagepool memory). Some older versions of RDMA driver stacks showed problems with large RDMA memory registrations. It is recommended to run a current version of the RDMA stack. If large memory registrations causes any problems, contact your RDMA vendor.
- Omni-path RDMA is not supported if the dynamic pagepool is enabled on a node.
- It is recommended to set the **pagepool** parameter to **default** on the nodes where the dynamic pagepool is enabled. If the **pagepool** parameter is set, it serves as an extra limit to the maximum size of the dynamic pagepool. The effective maximum-allowed size of the dynamic pagepool is a smaller value of **pagepool** and the calculated value from **pagepoolMaxPhysMemPct**.

- In systems with a heavy file system activity, where applications also consume memory, memory can be insufficient for a node. Occasionally, the increased rate of memory demand might exceed the rate of reduction in the use of pagepool (as the pagepool size is shrunk), and memory in the system might be exhausted. If this happens, reduce the value of the **pagepoolMaxPhysMemPct** parameter.
- The dynamic pagepool is not supported on systems with less than 6 GiB of memory. The dynamic pagepool ensures that at least 4 GiB of memory is available outside the pagepool. However, some of this memory is used by other parts of GPFS. When the dynamic pagepool is enabled, the minimum size is adjusted to at least 1 GiB.
- A check of the Linux kernel, which is used in older Linux distributions (RHEL7 and RHEL8), does not allow single memory allocations if they are larger than available memory on a node. Because Linux does not know that GPFS shrinks its pagepool. This limitation is not applicable for Linux distributions that are based on kernel 5.2 or higher.

Workarounds to solve this issue are as follows:

- Free swap space on a node is also considered as available memory, so swap space can be a workaround.
- You can set the **vm.overcommit_memory sysctl** parameter value to 1 to disable the Linux kernel check.

Other considerations

- Memory that is used in the dynamic pagepool is shown as the used memory. This used memory might be misleading because the dynamic pagepool relinquishes memory whenever it is requested from the Linux kernel.
- This consideration is for job scheduling software that queries available memory on the cluster nodes and schedule jobs depending on the available memory. If the dynamic pagepool is close to its maximum size, large jobs might not be scheduled on the node.
- If the job scheduling software allows modifying the memory check, the software can be used to consider only the minimum size of the pagepool for scheduling decisions. Otherwise, it might be necessary to adjust the dynamic pagepool size to meet the scheduling needs.

Shared memory communications direct on Linux on Z

The shared memory communication (SMC) protocol is an addition to TCP/IP and can be used transparently for shared memory communications.

The SMC can be used for connections through Shared Memory Communications Direct (SMC-D) with ISM devices. Shared Memory Communications through RDMA (SMC-R) with RoCE devices are not supported.

SMC-D is supported in IBM Storage Scale on Linux on Z. SMC-R is not supported.

The benefits of SMC-D are as follows:

- Low latency
- Less CPU usage compared to native TCP/IP

SMC-D can be used for the workloads with large data, high throughput, and many parallel network connections.

For more information about SMC-D, see [SMC protocol support](#).

Plugins for NVIDIA Nsight Systems

An NVIDIA Nsight Systems plugin that collects and visualizes IBM Storage Scale System filesystem performance metrics during application profiling sessions.

This plugin enables real-time collection of performance counters and integrates them into NVIDIA Nsight Systems profiling reports. It allows developers and system administrators to correlate application I/O behavior with filesystem performance, helping identify bottlenecks and optimize storage access patterns.

Key features

- Real-time Metrics: Collects performance data at 1ms intervals
- Seamless Integration: Metrics appear directly in Nsight Systems timeline
- Comprehensive I/O Metrics: Throughput, latency, and operation sizes
- Zero Code Changes: No application modification required

Performance Considerations

Sampling Rate: 1ms intervals (1000 samples/second)

Supported Platforms

- Architectures: x86_64, aarch64
- Operating Systems: All Linux distributions supported by IBM Storage Scale 6.0.1.0
- GPFS Versions: IBM Storage Scale 6.0.1.0 and later
- NVIDIA Nsight Systems: 2025.5.1 and later

Data tiering overview

Data tiering enables IBM Storage Scale to place blocks from a single file across different storage tiers based on file offset ranges.

With this feature, files can have their data blocks distributed between high-performance storage and standard storage pools based on configurable file offset ranges. This approach helps optimize performance while reducing the amount of expensive high-speed storage that is required.

Data tiering provides per-file control over block placement across storage tiers. A storage pool that is designated as a fast tier pool is required for this feature. The fast tier pool is specified by using the `fastTierPool=yes` attribute when the pool is created. Without a fast tier pool, data tiering cannot be used. The fast tier pool can be either a `dataOnly` pool or a `dataAndMetadata` pool.

Important:

- Only one fast tier pool is supported in a file system.
- File systems that use both the `fastTierPool` and performance pool are not supported.

Files that use data tiering can have their data blocks distributed between the following storage pools:

• Fast Tier Pool

Stores the initial blocks of a file, beginning at offset 0. This pool is typically configured with high-performance storage devices such as NVMe drives or SSDs.

• Default Pool

Stores the remaining blocks that are beyond the configured threshold. This pool is the regular data pool and is typically configured with standard-capacity storage devices.

The number of blocks that are allocated to the fast tier pool is determined for each file by the `gpfs.DataTiering` extended attribute. This attribute stores the fast tier range threshold in bytes. Blocks from offset 0 up to the offset corresponding to the fast tier range threshold, rounded up to the file system block size, are allocated to the fast tier pool. Blocks beyond this range are allocated to the default pool. Because the threshold is set on a per-file basis, different files can use different data tiering ranges based on workload requirements, policy rules, or manual configuration.

IBM Storage Scale calculates the number of `nFastTierBlocks` by rounding the `dataTieringRange` value up to a multiple of the file system block size and using the following formula:

```
nFastTierBlocks = (dataTieringRange + blockSize - 1) / blockSize
```

Block allocation behavior

IBM Storage Scale evaluates each block allocation request and selects the appropriate storage pool.

1. Blocks within the configured fast tier range are allocated in the fast tier pool.
2. Blocks beyond the configured threshold are allocated in the default pool.
3. When a file grows beyond the threshold, additional blocks are allocated in the default pool.
4. Indirect blocks continue to be allocated in the metadata pool.
5. Files that are less than the fast tier range threshold in size have all their blocks allocated in the fast tier pool.

Applications continue to access files as single logical files. IBM Storage Scale manages block placement across storage pools transparently, without requiring application changes.



Warning: Assigning a fast tier range to existing files that were previously backed up by using **mmbackup** can cause the files to be backed up again. To prevent unnecessary backups, ensure that the **UPDATECTIME** and **SKIPACLUPDATECHECK** options are set. For more information, see Options in the IBM Storage Protect configuration file **dsm.opt**.

Benefits of data tiering

Data tiering can provide the following benefits:

- Improves performance by placing frequently accessed data on high-performance storage.
- Reduces storage costs by limiting the use of expensive fast storage devices.
- Enables flexible per-file tiering through commands or policy rules.
- Supports dynamic block migration when tiering settings are changed.
- Helps optimize storage placement for workloads with mixed performance and capacity requirements.

Chapter 2. Planning for IBM Storage Scale

Planning for GPFS

Although you can modify your GPFS configuration after it is set, a little consideration before installation and initial setup will reward you with a more efficient and immediately useful file system.

About this task

During configuration, GPFS requires you to specify several operational parameters that reflect your hardware resources and operating environment. During file system creation, you can specify parameters that are based on the expected size of the files or you can let the default values take effect.

The installation toolkit is also available to assist with GPFS installation on Linux nodes. For more information, see [“Overview of the installation toolkit” on page 509](#).

Hardware requirements

You can validate that your hardware meets GPFS requirements by taking the steps that are outlined in this topic.

1. Consult the [IBM Storage Scale FAQ in IBM Documentation](#) for the latest list of:

- Supported server hardware
- Tested disk configurations
- Maximum cluster size
- Additional hardware requirements for protocols

2. Provide enough disks to contain the file system. Disks can be:

- SAN-attached to each node in the cluster
- Attached to one or more NSD servers
- A mixture of directly attached disks and disks that are attached to NSD servers

For more information, see the *Network Shared Disk (NSD) creation considerations* topic in *IBM Storage Scale: Administration Guide*.

3. When doing network-based NSD I/O, GPFS passes a large amount of data between its daemons. For NSD server-to-client traffic, it is suggested that you configure a dedicated high-speed network solely for GPFS communications when the following are true:

- There are NSD disks configured with servers providing remote disk capability.
- Multiple GPFS clusters are sharing data using NSD network I/O.

For more information, see the *Disk considerations* topic in *IBM Storage Scale: Administration Guide*.

GPFS communications require static IP addresses for each GPFS node. IP address takeover operations that transfer the address to another computer are not allowed for the GPFS network. Other IP addresses within the same computer that are not used by GPFS can participate in IP takeover. To provide availability or additional performance, GPFS can use virtual IP addresses created by aggregating several network adapters using techniques such as EtherChannel or channel bonding.

Cluster Export Services (CES) have dedicated network addresses to support the SMB, NFS, HDFS, and object protocols. These network addresses are appropriately assigned to CES nodes, and they facilitate CES node failover and failback operations. File and object clients use the public IPs to access data on GPFS™ file systems.

For additional information, see *CES network configuration* in *IBM Storage Scale: Administration Guide*.

Software requirements

IBM Storage Scale planning includes understanding the latest software requirements.

Before you start the installation of IBM Storage Scale, make sure to meet the following software requirements.

Note: This topic lists the additional software requirements to use specific functions after successful installation. It does not list the dependencies for installing IBM Storage Scale packages.

- GPFS is supported on AIX, Linux, and Windows.
- On Linux systems, the Linux ksh package is required (on Ubuntu, this package is provided only in the Universe repository).
- Kernel development files and compiler utilities are required to build the GPFS portability layer on Linux nodes. The required RPMs or packages for each supported Linux distribution are as follows:

- **Red Hat Enterprise Linux RPMs**

- kernel-devel
- cpp
- gcc
- gcc-c++
- binutils
- elfutils-libelf-devel (required for Red Hat Enterprise Linux 8.x)
- which command (required for mmbuildgpl, included in the which RPM)

- **SLES Linux RPMs**

- kernel-default-devel
- cpp
- gcc
- gcc-c++
- binutils
- libelf-devel
- which command (required for mmbuildgpl, included in the which RPM)

- **Ubuntu Linux packages**

- linux-generic
- cpp
- gcc
- g++
- binutils
- libelf-dev
- which command (required for mmbuildgpl, included in the debianutils package)

Additional software requirements for specific IBM Storage Scale functions

After the successful installation of IBM Storage Scale, the following functions need the additional software requirements that are detailed in the next lists.

Required package for active file management (AFM)

nfs-utils

Required packages for Cluster Export Services (CES) and CNFS

- ethtool

- nfs-utils (on Ubuntu: nfs-common)
- rpcbind
- psmisc
- iputils (on Ubuntu: iputils-arping)
- ifupdown (only on Ubuntu; there is no additional package requirement for RHEL/SLES)
- python36-dbus (only on RHEL 7.9)

Remember:

- The python36-dbus package is not part of the official RHEL repository, it is part of the Extra Packages for Enterprise Linux (EPEL). You must manually download the python36-dbus package from the [EPEL repository](#) and install it.
- The python36-dbus package is not part of the IBM Storage Scale repositories because it is not part of RHEL itself.
- The python36-dbus RPM is not included in the IBM Storage Scale Systems packages.

Note: On a Linux node that is not a protocol node running the NFS service, it is recommended to disable the port mapper service (also called rpcbind).

Required package for SLES 15 CES nodes

gdb

Requirements for file authentication

Before configuring file authentication, ensure that prerequisites are met to avoid installation failure. For more information, see the topic *Configuring authentication and ID mapping for file access* in *IBM Storage Scale: Administration Guide*.

You can use Active Directory (AD) authentication. For more information, see the topic *Integrating with AD server* in *IBM Storage Scale: Administration Guide*.

For more information, see the *Integrating with LDAP server* topic in *IBM Storage Scale: Administration Guide*.

Required packages as per OS and Authentication schemes

```
[ad:RHEL]="bind-utils"
[ad:SLES]="bind-utils libarchive13"
[ad:UBUNTU]="dnsutils"

[ldap:RHEL]="openldap-clients sssd sssd-tools"
[ldap:SLES]="openldap2-client sssd sssd-tools sssd-ldap"
[ldap:UBUNTU]="ldap-utils sssd sssd-tools"

[nis:RHEL]="sssd ybind sssd-proxy"
[nis:SLES]="sssd ybind sssd-proxy"
[nis:UBUNTU]="sssd nis sssd-proxy"

[kerberos:RHEL]="krb5-workstation"
[kerberos:SLES]="krb5-client sssd-krb5"
[kerberos:UBUNTU]="krb5-user"
```

Required packages for the performance monitoring tool

Performance monitoring pmcollector package requirements for sqlite3 shared libraries are distro-specific. The name of the library package differs between distros.

- RHEL 8 and 9: sqlite-libs
- SLES 15: libsqlite3-0
- UBUNTU 22.04: libsqlite3-0

Important: Ensure that these packages are installed on the system before the installation of the performance monitoring tool.

Required package for the IBM Storage Scale monitoring service on AIX and Linux

- Python 3

Note:

- If you are using a RHEL operating system, downloading, installing, or setting Python ≥ 3.12 as the default path causes issues. The latest version of Python may cause failures in the installation toolkit operations because RHEL operating systems do not usually work with the latest version of Python. For more information, see [“Preparing to use the installation toolkit” on page 528](#) and [“Limitations of the installation toolkit” on page 518](#).
- On Linux, Python 3 is usually installed automatically.
- Python 3 is shipped with AIX 7.3.

On AIX 7.2, you must manually install Python 3.10 or later. For more information, see [AIX Toolbox for Linux Applications](#). To manually install, one of the ways is to configure Yum with the AIX toolbox. For more information, see [Configuring YUM and creating local repositories on IBM AIX](#).

- On Windows, manually install Python 3 under the Cygwin environment as described in the following steps. Windows native (non-Cygwin) distributions of Python 3 are not supported.
 1. From <http://www.cygwin.com>, download and run the Cygwin 64-bit setup program setup-x86_64.exe.
 2. In the "Select Packages" window, click **View > Category**.
 3. Click **All > Python > Python3** and select the latest level.
 4. Follow the instructions to complete the installation.

IBM Storage Scale 6.0.x supports Python 3.10 or later. It is recommended that Python 3.10 is installed through the OS package manager (For example, **yum install python3.12**). If you install Python 3.10 by other means, then unexpected results might occur, such as failure to install `gpfs.base` for prerequisite checks, and workarounds might be required.

Requirements for the management GUI

The following packages must be installed on the node on which GUI needs to be installed:

- PostgreSQL server including contrib package. The PostgreSQL server must be specific to the Linux distribution.
- IBM Storage Scale Java™ Runtime Environment (JRE) (`gpfs.java` RPM).
- IBM Storage Scale performance collector (`gpfs.pmcollector` RPM).
- OpenSSL for HTTPS certificate management.
- `gpfs.gskit` to use CCR.

Required packages for clustered watch folder

Every node that is mounting the watched file system must have the following packages installed:

- For RHEL, the `librdkafka` package requires the `openssl-devel` and `cyrus-sasl-devel` packages.
- For Ubuntu, the `librdkafka` package requires the `libssl-dev` and `libsasl2-dev` packages.
- `librdkafka` (`gpfs.librdkafka` rpm/package).

For more information, see [“Requirements, limitations, and support for clustered watch folder” on page 651](#).

Required package to use the `mmchconfig numaMemoryInterleave` parameter

`numactl`

Required package to use the `mmhealth` command on Ubuntu Linux

`sqlite3`

Required packages for `mmprotocoltrace`

Some of the advanced tracing components of the `mmprotocoltrace` command require specific packages to be installed on all nodes that need to participate in tracing related operations.

- To enable network tracing with `mmprotocoltrace`, you need to install the `tcpdump` package.
- To enable the `syscalls-tracing` for SMB, you need to install the `strace` package.

Prerequisites for the installation toolkit

For information on prerequisites for using the installation toolkit, see [“Preparing to use the installation toolkit” on page 528](#).

For additional prerequisites on Linux, see [“Installation prerequisites” on page 470](#).

Consult the [IBM Storage Scale FAQ in IBM Documentation](#) for the latest list of:

- AIX environments
- Linux distributions
- Linux kernel versions
- Windows environments

Recoverability considerations

Good file system planning requires several decisions about recoverability. After you make these decisions, GPFS parameters enable you to create a highly available file system with rapid recovery from failures.

- At the disk level, consider preparing disks for use with your file system by specifying failure groups that are associated with each disk. With this configuration, information is not vulnerable to a single point of failure. See [“Network Shared Disk \(NSD\) creation considerations” on page 329](#).
- At the file system level, consider replication through the metadata and data replication parameters. See [“File system replication parameters” on page 355](#).

Also, GPFS provides several layers of protection against failures of various types.

Node failure

If a node fails, GPFS prevents the continuation of I/O from the failing node and replays the file system metadata log for the failed node.

GPFS prevents the continuation of I/O from a failing node through a GPFS-specific fencing mechanism called *disk leasing*. When a node has access to file systems, it obtains disk leases that allow it to submit I/O. However, when a node fails, that node cannot obtain or renew a disk lease. When GPFS selects another node to perform recovery for the failing node, it first waits until the disk lease for the failing node expires. This allows for the completion of previously submitted I/O and provides for a consistent file system metadata log. Waiting for the disk lease to expire also avoids data corruption in the subsequent recovery step.

To reduce the amount of time it takes for disk leases to expire, you can use Persistent Reserve (SCSI-3 protocol). If Persistent Reserve (configuration parameter: `usePersistentReserve`) is enabled, GPFS prevents the continuation of I/O from a failing node by fencing the failed node that uses a feature of the disk subsystem called Persistent Reserve. Persistent Reserve allows the failing node to recover faster because GPFS does not need to wait for the disk lease on the failing node to expire. For additional information, refer to [“Reduced recovery time by using Persistent Reserve” on page 318](#). For further information about recovery from node failure, see *Installation and configuration issues in IBM Storage Scale: Problem Determination Guide*.

There is a temporary impact to I/O during file system recovery from node failure. Recovery involves rebuilding metadata structures that might be under modification at the time of the failure. If the failing node is acting as the file system manager when it fails, the delay is longer and proportional to the level of activity on the file system at the time of failure. In this case, the failover file system management task happens automatically to a surviving node.

Managing node failures also involves sizing the solution adequately so that remaining nodes in the cluster can support a node down situation such as a planned system maintenance or an unplanned node failure in terms of bandwidth and throughput. For protocols, this includes supporting SMB, NFS, or S3 connections that have to fail over to another CES node in the cluster if a node fails.

Quorum

GPFS uses a cluster mechanism called quorum to maintain data consistency in the event of a node failure.

Quorum operates on the principle of majority rule. This means that a majority of the nodes in the cluster must be successfully communicating before any node can mount and access a file system. This keeps any nodes that are cut off from the cluster (for example, by a network failure) from writing data to the file system.

During node failure situations, quorum needs to be maintained in order for the cluster to remain online. If quorum is not maintained due to node failure, GPFS unmounts local file systems on the remaining nodes and attempts to reestablish quorum, at which point file system recovery occurs. For this reason it is important that the set of quorum nodes be carefully considered (refer to [“Selecting quorum nodes”](#) on page 316 for additional information).

GPFS quorum must be maintained within the cluster for GPFS to remain active. If the quorum semantics are broken, GPFS performs recovery in an attempt to achieve quorum again. GPFS can use one of two methods for determining quorum:

- Node quorum
- Node quorum with tiebreaker disks

Node quorum

Node quorum is the default quorum algorithm for GPFS.

With node quorum:

- Quorum is defined as one plus half of the *explicitly defined* quorum nodes in the GPFS cluster.
- There are no default quorum nodes; you must specify which nodes have this role.

For example, in [Figure 35 on page 314](#), there are three quorum nodes. In this configuration, GPFS remains active as long as there are two quorum nodes available.




Figure 35. GPFS configuration using node quorum

Node quorum with tiebreaker disks

When running on small GPFS clusters, you might want to have the cluster remain online with only one surviving node.

To achieve this, you need to add a tiebreaker disk to the quorum configuration. Node quorum with tiebreaker disks allows you to run with as little as one quorum node available as long as you have access to a majority of the quorum disks (refer to [Figure 36 on page 315](#)). Enabling node quorum with tiebreaker

disks starts by designating one or more nodes as quorum nodes. Then one to three disks are defined as tiebreaker disks using the **tiebreakerDisks** parameter on the `mmchconfig` command. You can designate any disk to be a tiebreaker.

When utilizing node quorum with tiebreaker disks, there are specific rules for cluster nodes and for tiebreaker disks.

Cluster node rules:

1. There is a maximum of eight quorum nodes.
2. All quorum nodes need to have access to all of the tiebreaker disks.
3. You may have an unlimited number of non-quorum nodes.
4. If a network connection fails, which causes the loss of quorum, and quorum is maintained by tiebreaker disks, the following rationale is used to re-establish quorum. If a group has the cluster manager, it is the "survivor". The cluster manager can give up its role if it communicates with fewer than the minimum number of quorum nodes as defined by the **minQuorumNodes** configuration parameter. In this case, other groups with the minimum number of quorum nodes (if they exist) can choose a new cluster manager.

Changing quorum semantics:

When using the cluster configuration repository (CCR) to store configuration files, the total number of quorum nodes is limited to eight, regardless of quorum semantics, but the use of tiebreaker disks can be enabled or disabled at any time by issuing an `mmchconfig tiebreakerDisks` command. The change will take effect immediately, and it is not necessary to shut down GPFS when making this change.

Tiebreaker disk rules:

- You can have one, two, or three tiebreaker disks. However, you should use an odd number of tiebreaker disks.
- Among the quorum node groups that appear after an interconnect failure, only those having access to a majority of tiebreaker disks can be candidates to be the survivor group.
- Tiebreaker disks must be connected to all quorum nodes.
- In a CCR-based cluster, when adding tiebreaker disks:
 - GPFS should be up and running, if tiebreaker disks are part of the file system.
 - GPFS can be either running or shut down, if tiebreaker disks are not a part of the file system.

In [Figure 36 on page 315](#) GPFS remains active with the minimum of a single available quorum node and two available tiebreaker disks.




Figure 36. GPFS configuration using node quorum with tiebreaker disks

When a quorum node detects loss of network connectivity, but before GPFS runs the algorithm that decides if the node will remain in the cluster, the `tiebreakerCheck` event is triggered. This event is

generated only in configurations that use quorum nodes with tiebreaker disks. It is also triggered on the cluster manager periodically by a challenge-response thread to verify that the node can still continue as cluster manager.

Selecting quorum nodes

To configure a system with efficient quorum nodes, follow these rules.

About this task

- Select nodes that are likely to remain active
 - If a node is likely to be rebooted or require maintenance, do not select that node as a quorum node.
- Select nodes that have different failure points such as:
 - Nodes that are located in different racks
 - Nodes that are connected to different power panels
- You should select nodes that GPFS administrative and serving functions rely on such as:
 - Network Shared Disk (NSD) servers
- Select an odd number of nodes as quorum nodes
 - Clusters with no tiebreaker disks can be configured with up to nine quorum nodes. Clusters with tiebreaker disks are limited to a maximum of eight quorum nodes.
- Having multiple quorum nodes might increase the time required for startup and failure recovery.
 - Having more than seven quorum nodes does not guarantee higher availability.
 - All quorum nodes must have access to all of the tiebreaker disks.

The `/var/mmfs` directory on quorum nodes cannot be on `tmpfs` or `ramfs` file system.

Note: When adding a quorum node by using either `mmcrcluster` or `mmchnode` or when you enable CCR from server-based cluster configuration method, then the system checks whether the `/var/mmfs` directory on the designated quorum node is not on `tmpfs` or `ramfs`. If `/var/mmfs` is from `tmpfs` or `ramfs`, the command output prints an error message.

Network Shared Disk server and disk failure

The three most common reasons why data becomes unavailable are disk failure, disk server failure with no redundancy, and failure of a path to the disk.

In the event of a disk failure in which GPFS can no longer read or write to the disk, GPFS will discontinue use of the disk until it returns to an available state. You can guard against loss of data availability from disk failure by:

- Utilizing hardware data protection as provided by a Redundant Array of Independent Disks (RAID) device (see [Figure 37 on page 317](#))




Figure 37. An example of a highly available SAN configuration for a GPFS file system

- Utilizing the GPFS data and metadata replication features (see [“Increased data availability”](#) on page 2) along with the designation of failure groups (see [“Network Shared Disk \(NSD\) creation considerations”](#) on page 329 and [Figure 38 on page 317](#))




Figure 38. Configuration using GPFS replication for improved availability

It is suggested that you consider RAID as the first level of redundancy for your data and add GPFS replication if you desire additional protection.

In the event of an NSD server failure in which a GPFS client can no longer contact the node that provides remote access to a disk, GPFS discontinues the use of the disk. You can guard against loss of an NSD server availability by using common disk connectivity on multiple NSD server nodes and specifying multiple Network Shared Disk servers for each common disk.

Note: In the event that a path to a disk fails, GPFS reports a disk failure and marks the disk down. To bring the disk back online, first follow the directions supplied by your storage vendor to determine and repair the failure.

Guarding against loss of data availability due to path failure

You can guard against loss of data availability from failure of a path to a disk by doing the following:

- Creating multiple NSD servers for each disk

As GPFS determines the available connections to disks in the file system, it is recommended that you always define more than one NSD server for each disk. GPFS allows you to define up to eight NSD servers for each NSD. In a SAN configuration where NSD servers have also been defined, if the physical connection is broken, GPFS dynamically switches to the next available NSD server (as defined on the server list) and continues to provide data. When GPFS discovers that the path has been repaired, it moves back to local disk access. This is the default behavior, which can be changed by designating file

system mount options. For example, if you never want a node to use the NSD server path to a disk, even if the local path fails, you can set the `-o useNSDserver` mount option to `never`. You can set the mount option using the `mmchfs`, `mmmout`, `mmremotefs`, and `mount` commands.

Important: In Linux on Z, it is mandatory to have multiple paths to one SCSI disk (LUN) to avoid single path of failure. The coalescing of the paths to one disk is done by the kernel (via the device-mapper component). As soon as the paths are coalesced, a new logical, multipathed device is created, which is used for any further (administration) task. (The single paths can no longer be used.)

The multipath device interface name depends on the distribution and is configurable:

SUSE

`/dev/mapper/Unique_WW_Identifier`

For example: `/dev/mapper/36005076303ffc56200000000000010cc`

Red Hat

`/dev/mapper/mpath*`

To obtain information about a multipathed device, use the `multipath` tool as shown in the following example:

```
# multipath -ll
```

The system displays output similar to this:

```
36005076303ffc56200000000000010cc dm-0 IBM,2107900
[size=5.0G][features=1 queue_if_no_path][hwhandler=0]
\_ round-robin 0 [prio=2][active]
  \_ 1:0:0:0 sdb 8:16 [active][ready]
     \_ 0:0:0:0 sda 8:0 [active][ready]
```

See the question, "What considerations are there when setting up DM-MP multipath service" in the [IBM Storage Scale FAQ in IBM Documentation](#).

- Using an I/O driver that provides multiple paths to the disks for failover purposes

Failover is a path-management algorithm that improves the reliability and availability of a device because the system automatically detects when one I/O path fails and reroutes I/O through an alternate path.

Reduced recovery time by using Persistent Reserve

Persistent Reserve (PR) provides a mechanism for reducing recovery times from node failures.

To enable PR and to obtain recovery performance improvements, your cluster requires a specific environment:

- All disks must be PR-capable. For a list of devices supported by PR, see [What devices does IBM Storage Scale support with SCSI-3 persistent reservations?](#) in [IBM Storage Scale FAQ in IBM Documentation](#).
- On AIX, all disks must be hdisks. Starting with 3.5.0.16, it is also possible to use a logical volume as a descOnly disk without disabling the use of Persistent Reserve. For more information, see the [IBM Storage Scale FAQ in IBM Documentation](#).

On Linux, typically all disks must be generic (`/dev/sd*`) or DM-MP (`/dev/dm-*`) disks.

However, for Linux on Z, multipath device names are required for SCSI disks, and the names depend on the distribution and are configurable. For more information, see [“Guarding against loss of data availability due to path failure” on page 317](#).

- If the disks have NSD servers that are defined, all NSD server nodes must be running the same operating system (AIX or Linux).
- If the disks are SAN-attached to all nodes, all nodes in the cluster must be running the same operating system (AIX or Linux).

For quicker recovery times when you use PR, set the **failureDetectionTime** configuration parameter on the **mmchconfig** command. For example, for quick recovery a recommended value would be 10:
`mmchconfig failureDetectionTime=10`

Explicitly enable PR by specifying the **usePersistentReserve** parameter on the **mmchconfig** command. If you set `usePersistentReserve=yes`, GPFS attempts to set up PR on all of the PR capable disks. All subsequent NSDs are created with PR enabled if they are PR capable. However, PR is only supported in the home cluster. Therefore, access to PR-enabled disks from another cluster must be through an NSD server that is in the home cluster and not directly to the disk (for example, through a SAN).

GPFS cluster creation considerations

A GPFS cluster is created that uses the `mmcrcluster` command. On supported Linux distributions, you can use the installation toolkit to create a GPFS cluster. For more information, see [“Overview of the installation toolkit”](#) on page 509.

Table 25 on page 319 details the cluster creation options, how to change the options, and the default values for each option.

Table 25. GPFS cluster creation options		
Options	Command to change the option	Default value
“Creating an IBM Storage Scale cluster” on page 321	<code>mmaddnode</code> <code>mmdelnode</code>	None
Node designation: manager or client. See “Creating an IBM Storage Scale cluster” on page 321.	<code>mmchnode</code>	client
Node designation: quorum or nonquorum. See “Creating an IBM Storage Scale cluster” on page 321.	<code>mmchnode</code>	nonquorum
“Remote shell command” on page 323	<code>mmchcluster</code>	/usr/bin/ssh
“Remote file copy command” on page 324	<code>mmchcluster</code>	/usr/bin/scp
“Cluster name” on page 325	<code>mmchcluster</code>	The node name of the primary GPFS cluster configuration server.
Cluster configuration information. See “IBM Storage Scale cluster configuration information” on page 322.	<code>mmchcluster</code>	The Cluster Configuration Repository (CCR) is enabled.
GPFS administration adapter port name. See “GPFS node adapter interface names” on page 320	<code>mmchnode</code>	Same as the GPFS communications adapter port name.
GPFS communications adapter port name. See “GPFS node adapter interface names” on page 320.	<code>mmchnode</code>	None
“User ID domain for the cluster” on page 326.	<code>mmchconfig</code>	The name of the GPFS cluster.
“Starting GPFS automatically” on page 327	<code>mmchconfig</code>	no

Table 25. GPFS cluster creation options (continued)		
Options	Command to change the option	Default value
“Cluster configuration file” on page 327	Not applicable	None

GPFS node adapter interface names

An adapter interface name refers to the hostname or IP address that GPFS uses to communicate with a node. Specifically, the hostname or IP address identifies the communications adapter over which the GPFS daemons or GPFS administration commands communicate.

The administrator can specify two node adapter interface names for each node in the cluster:

GPFS node name

Specifies the name of the node adapter interface to be used by the GPFS daemons for internode communication.

GPFS admin node name

Specifies the name of the node adapter interface to be used by GPFS administration commands when communicating between nodes. If not specified, the GPFS administration commands use the same node adapter interface used by the GPFS daemons.

These names can be specified by means of the node descriptors passed to the `mmaddnode` or `mmcrcluster` command and can later be changed with the `mmchnode` command.

If multiple adapters are available on a node, this information can be communicated to GPFS by means of the `subnets` parameter on the `mmchconfig` command.

Related concepts

[Creating an IBM Storage Scale cluster](#)

When you create an IBM Storage Scale cluster, you can either create a small cluster of one or more nodes and later add nodes to it or you can create a cluster with all its nodes in one step.

[IBM Storage Scale cluster configuration information](#)

You can use the Cluster Configuration Repository (CCR) to maintain cluster configuration information.

[Remote shell command](#)

GPFS commands need to be able to communicate across all nodes in the cluster. To achieve this, the GPFS commands use the remote shell command that you specify on the `mmcrcluster` command or the `mmchcluster` command.

[Remote file copy command](#)

The GPFS commands must maintain several configuration files across all nodes in the cluster. To achieve this, the GPFS commands use the remote file copy command that you specify on the `mmcrcluster` command or the `mmchcluster` command.

[Cluster name](#)

Provide a name for the cluster by issuing the `-C` option on the `mmcrcluster` command.

[User ID domain for the cluster](#)

This option is the user ID domain for a cluster when accessing a file system remotely.

[Cluster configuration file](#)

GPFS provides default configuration values, so a cluster configuration file is not required to create a cluster.

Related tasks

[Starting GPFS automatically](#)

You can specify whether to start the GPFS daemon automatically on a node when it is started.

Creating an IBM Storage Scale cluster

When you create an IBM Storage Scale cluster, you can either create a small cluster of one or more nodes and later add nodes to it or you can create a cluster with all its nodes in one step.

Two methods are available for creating an IBM Storage Scale cluster:

- Run the `mmcrcluster` command to create a cluster that contains one or more nodes, and later run the `mmaddnode` command as needed to add nodes to the cluster. This method is more flexible and is appropriate when you want to build a cluster step by step.
- Run the `mmcrcluster` command to create a cluster and at the same time to add a set of nodes to the cluster. This method is quicker when you already know which nodes you want to add to the cluster.

Whichever method you choose, you can later add nodes to the cluster with the `mmaddnode` command or delete nodes with the `mmdelnode` command. For more information, see the topics *mmcrcluster command* and *mmaddnode command* in the *IBM Storage Scale: Command and Programming Reference Guide*.

With both commands, you specify node descriptors to identify the nodes that you are adding to the cluster. You can list node descriptors either on the command line or in a separate node descriptor file. A node descriptor has the following format:

```
NodeName:NodeDesignations:AdminNodeName
```

NodeName is a required parameter. *NodeDesignations* and *AdminNodeName* are optional parameters.

NodeName

The host name or IP address of the node for GPFS daemon-to-daemon communication.

The host name or IP address that is used for a node must refer to the communication adapter over which the GPFS daemons communicate. Alias names are not allowed. You can specify an IP address at NSD creation, but it is converted to a host name that must match the GPFS node name. You can specify a node by using any of these forms:

- Short hostname (for example, h135n01)
- Long hostname (for example, h135n01.frf.ibm.com)
- IP address (for example, 7.111.12.102)

Note: Host names must always include at least one alpha character, and they should not start with a hyphen (-).

Whichever form you specify, the other two forms must be defined correctly in DNS or the hosts file.

NodeDesignations

An optional, "-" separated list of node roles.

- `manager | client` – Indicates whether a node is part of the node pool from which file system managers and token managers can be selected. The default is `client`, which means do not include the node in the pool of manager nodes. For detailed information on the manager node functions, see [“File system manager” on page 10](#).

In general, it is a good idea to define more than one node as a manager node. How many nodes you designate as manager depends on the workload and the number of GPFS server licenses you have. If you are running large parallel jobs, you may need more manager nodes than in a four-node cluster supporting a Web application. As a guide, in a large system there must be a different file system manager node for each GPFS file system.

- `quorum | nonquorum` – This designation specifies whether or not the node should be included in the pool of nodes from which quorum is derived. The default is `nonquorum`. You need to designate at least one node as a quorum node. It is recommended that you designate at least the primary and secondary cluster configuration servers and NSD servers as quorum nodes.

How many quorum nodes you designate depends upon whether you use node quorum or node quorum with tiebreaker disks. See [“Quorum” on page 314](#).

AdminNodeName

Specifies an optional field that consists of a node name to be used by the administration commands to communicate between nodes.

If *AdminNodeName* is not specified, the *NodeName* value is used.

Follow these rules when you add nodes to an IBM Storage Scale cluster:

- While a node may mount file systems from multiple clusters, the node itself may only reside in a single cluster. Nodes are added to a cluster using the `mmcrcluster` or `mmaddnode` command.
- The nodes must be available when they are added to a cluster. If any of the nodes listed are not available when the command is issued, a message listing those nodes is displayed. You must correct the problem on each node and then issue the `mmaddnode` command to add those nodes.
- Designate at least one but not more than seven nodes as quorum nodes. When not using tiebreaker disks, you can designate more quorum nodes, but it is recommended to use fewer than eight if possible. When using server-based configuration repository, it is recommended that you designate the cluster configuration servers as quorum nodes. How many quorum nodes altogether you will have depends on whether you intend to use the node quorum with tiebreaker algorithm or the regular node based quorum algorithm. For more details, see [“Quorum” on page 314](#).

Related concepts

GPFS node adapter interface names

An adapter interface name refers to the hostname or IP address that GPFS uses to communicate with a node. Specifically, the hostname or IP address identifies the communications adapter over which the GPFS daemons or GPFS administration commands communicate.

IBM Storage Scale cluster configuration information

You can use the Cluster Configuration Repository (CCR) to maintain cluster configuration information.

Remote shell command

GPFS commands need to be able to communicate across all nodes in the cluster. To achieve this, the GPFS commands use the remote shell command that you specify on the `mmcrcluster` command or the `mmchcluster` command.

Remote file copy command

The GPFS commands must maintain several configuration files across all nodes in the cluster. To achieve this, the GPFS commands use the remote file copy command that you specify on the `mmcrcluster` command or the `mmchcluster` command.

Cluster name

Provide a name for the cluster by issuing the `-C` option on the `mmcrcluster` command.

User ID domain for the cluster

This option is the user ID domain for a cluster when accessing a file system remotely.

Cluster configuration file

GPFS provides default configuration values, so a cluster configuration file is not required to create a cluster.

Related tasks

Starting GPFS automatically

You can specify whether to start the GPFS daemon automatically on a node when it is started.

IBM Storage Scale cluster configuration information

You can use the Cluster Configuration Repository (CCR) to maintain cluster configuration information.

It is a good practice to use the Cluster Configuration Repository (CCR) to maintain the cluster configuration information. By default, the `mmcrcluster` command enables the CCR. For more information, see [“Cluster configuration repository” on page 27](#).



Attention: The primary and secondary configuration server functionality is deprecated. For more information, see the notes in the topics *mmcrcluster command* and *mmchcluster command* in the *IBM Storage Scale: Command and Programming Reference Guide*.

Related concepts

GPFS node adapter interface names

An adapter interface name refers to the hostname or IP address that GPFS uses to communicate with a node. Specifically, the hostname or IP address identifies the communications adapter over which the GPFS daemons or GPFS administration commands communicate.

Creating an IBM Storage Scale cluster

When you create an IBM Storage Scale cluster, you can either create a small cluster of one or more nodes and later add nodes to it or you can create a cluster with all its nodes in one step.

Remote shell command

GPFS commands need to be able to communicate across all nodes in the cluster. To achieve this, the GPFS commands use the remote shell command that you specify on the `mmcrcluster` command or the `mmchcluster` command.

Remote file copy command

The GPFS commands must maintain several configuration files across all nodes in the cluster. To achieve this, the GPFS commands use the remote file copy command that you specify on the `mmcrcluster` command or the `mmchcluster` command.

Cluster name

Provide a name for the cluster by issuing the `-C` option on the `mmcrcluster` command.

User ID domain for the cluster

This option is the user ID domain for a cluster when accessing a file system remotely.

Cluster configuration file

GPFS provides default configuration values, so a cluster configuration file is not required to create a cluster.

Related tasks

Starting GPFS automatically

You can specify whether to start the GPFS daemon automatically on a node when it is started.

Remote shell command

GPFS commands need to be able to communicate across all nodes in the cluster. To achieve this, the GPFS commands use the remote shell command that you specify on the `mmcrcluster` command or the `mmchcluster` command.

The default remote shell command is `ssh`. You can designate the use of a different remote shell command by specifying its fully qualified path name on the `mmcrcluster` command or the `mmchcluster` command. The remote shell command must adhere to the same syntax as `ssh`, but it can implement an alternate authentication mechanism.

Clusters that include both UNIX and Windows nodes must use `ssh` for the remote shell command. For more information, see “[Installing and configuring OpenSSH on Windows nodes](#)” on page 594.

Clusters that only include Windows nodes may use the `mmwinrsh` utility that comes with GPFS. The fully qualified path name is `/usr/lpp/mmfs/bin/mmwinrsh`. For more information about configuring Windows GPFS clusters, see **`mmwinseructl command`** in *IBM Storage Scale: Command and Programming Reference Guide*.

By default, you can issue GPFS administration commands from any node in the cluster. Optionally, you can choose a subset of the nodes that are capable of running administrative commands. In either case, the nodes that you plan to use for administering GPFS must be able to run remote shell commands on any other node in the cluster as user `root` without the use of a password and without producing any extraneous messages. You can also use `sudo` wrappers for this purpose.

For additional information, see *Requirements for administering a GPFS file system* in *IBM Storage Scale: Administration Guide*.

Related concepts

[GPFS node adapter interface names](#)

An adapter interface name refers to the hostname or IP address that GPFS uses to communicate with a node. Specifically, the hostname or IP address identifies the communications adapter over which the GPFS daemons or GPFS administration commands communicate.

[Creating an IBM Storage Scale cluster](#)

When you create an IBM Storage Scale cluster, you can either create a small cluster of one or more nodes and later add nodes to it or you can create a cluster with all its nodes in one step.

[IBM Storage Scale cluster configuration information](#)

You can use the Cluster Configuration Repository (CCR) to maintain cluster configuration information.

[Remote file copy command](#)

The GPFS commands must maintain several configuration files across all nodes in the cluster. To achieve this, the GPFS commands use the remote file copy command that you specify on the `mmcrcluster` command or the `mmchcluster` command.

[Cluster name](#)

Provide a name for the cluster by issuing the `-C` option on the `mmcrcluster` command.

[User ID domain for the cluster](#)

This option is the user ID domain for a cluster when accessing a file system remotely.

[Cluster configuration file](#)

GPFS provides default configuration values, so a cluster configuration file is not required to create a cluster.

Related tasks

[Starting GPFS automatically](#)

You can specify whether to start the GPFS daemon automatically on a node when it is started.

Remote file copy command

The GPFS commands must maintain several configuration files across all nodes in the cluster. To achieve this, the GPFS commands use the remote file copy command that you specify on the `mmcrcluster` command or the `mmchcluster` command.

The default remote file copy program is **scp**. You can designate the use of a different remote file copy command by specifying its fully qualified path name on the **mmcrcluster** command or the **mmchcluster** command. The remote file copy command must adhere to the same syntax as **scp**, but it can implement an alternate authentication mechanism.

Clusters that include both UNIX and Windows nodes must use **scp** for the remote copy command. For more information, see [“Installing and configuring OpenSSH on Windows nodes”](#) on page 594.

The file copy command **rsh** cannot be used in a cluster that contains Windows Server nodes. However, clusters that include only Windows nodes might use the **mmwinrsh** utility that comes with GPFS. The fully qualified path name is `/usr/lpp/mmfs/bin/mmwinrsh`. For more information about configuring Windows GPFS clusters, see *mmwinservctl* command in *IBM Storage Scale: Command and Programming Reference Guide*.

The nodes that you plan to use for administering GPFS must be able to copy files by using the remote file copy command to and from any other node in the cluster without the use of a password and without producing any extraneous messages.

For additional information, see *Requirements for administering a GPFS file system* in *IBM Storage Scale: Administration Guide*.

Related concepts

[GPFS node adapter interface names](#)

An adapter interface name refers to the hostname or IP address that GPFS uses to communicate with a node. Specifically, the hostname or IP address identifies the communications adapter over which the GPFS daemons or GPFS administration commands communicate.

Creating an IBM Storage Scale cluster

When you create an IBM Storage Scale cluster, you can either create a small cluster of one or more nodes and later add nodes to it or you can create a cluster with all its nodes in one step.

IBM Storage Scale cluster configuration information

You can use the Cluster Configuration Repository (CCR) to maintain cluster configuration information.

Remote shell command

GPFS commands need to be able to communicate across all nodes in the cluster. To achieve this, the GPFS commands use the remote shell command that you specify on the `mmcrcluster` command or the `mmchcluster` command.

Cluster name

Provide a name for the cluster by issuing the `-C` option on the `mmcrcluster` command.

User ID domain for the cluster

This option is the user ID domain for a cluster when accessing a file system remotely.

Cluster configuration file

GPFS provides default configuration values, so a cluster configuration file is not required to create a cluster.

Related tasks

Starting GPFS automatically

You can specify whether to start the GPFS daemon automatically on a node when it is started.

Cluster name

Provide a name for the cluster by issuing the `-C` option on the `mmcrcluster` command.

If the user-provided name contains dots, it is assumed to be a fully qualified domain name. Otherwise, to make the cluster name unique in a multiple cluster environment, GPFS appends the domain name of the first node in the list primary cluster configuration server. If the `-C` option is not specified, the cluster name defaults to the hostname of the primary cluster configuration server. The name of the cluster may be changed at a later time by issuing the `-C` option on the `mmchcluster` command.

The cluster name is applicable when GPFS file systems are mounted by nodes belonging to other GPFS clusters. See the `mmauth` and the `mmremoteccluster` commands.

Related concepts

GPFS node adapter interface names

An adapter interface name refers to the hostname or IP address that GPFS uses to communicate with a node. Specifically, the hostname or IP address identifies the communications adapter over which the GPFS daemons or GPFS administration commands communicate.

Creating an IBM Storage Scale cluster

When you create an IBM Storage Scale cluster, you can either create a small cluster of one or more nodes and later add nodes to it or you can create a cluster with all its nodes in one step.

IBM Storage Scale cluster configuration information

You can use the Cluster Configuration Repository (CCR) to maintain cluster configuration information.

Remote shell command

GPFS commands need to be able to communicate across all nodes in the cluster. To achieve this, the GPFS commands use the remote shell command that you specify on the `mmcrcluster` command or the `mmchcluster` command.

Remote file copy command

The GPFS commands must maintain several configuration files across all nodes in the cluster. To achieve this, the GPFS commands use the remote file copy command that you specify on the `mmcrcluster` command or the `mmchcluster` command.

User ID domain for the cluster

This option is the user ID domain for a cluster when accessing a file system remotely.

Cluster configuration file

GPFS provides default configuration values, so a cluster configuration file is not required to create a cluster.

Related tasks

Starting GPFS automatically

You can specify whether to start the GPFS daemon automatically on a node when it is started.

User ID domain for the cluster

This option is the user ID domain for a cluster when accessing a file system remotely.

This option is further explained in *IBM Storage Scale: Administration Guide* and the IBM white paper *UID Mapping for GPFS in a Multi-cluster Environment* (https://www.ibm.com/docs/en/storage-scale?topic=STXKQY/uid_gpfs.pdf).

Related concepts

GPFS node adapter interface names

An adapter interface name refers to the hostname or IP address that GPFS uses to communicate with a node. Specifically, the hostname or IP address identifies the communications adapter over which the GPFS daemons or GPFS administration commands communicate.

Creating an IBM Storage Scale cluster

When you create an IBM Storage Scale cluster, you can either create a small cluster of one or more nodes and later add nodes to it or you can create a cluster with all its nodes in one step.

IBM Storage Scale cluster configuration information

You can use the Cluster Configuration Repository (CCR) to maintain cluster configuration information.

Remote shell command

GPFS commands need to be able to communicate across all nodes in the cluster. To achieve this, the GPFS commands use the remote shell command that you specify on the `mmcrcluster` command or the `mmchcluster` command.

Remote file copy command

The GPFS commands must maintain several configuration files across all nodes in the cluster. To achieve this, the GPFS commands use the remote file copy command that you specify on the `mmcrcluster` command or the `mmchcluster` command.

Cluster name

Provide a name for the cluster by issuing the `-C` option on the `mmcrcluster` command.

Cluster configuration file

GPFS provides default configuration values, so a cluster configuration file is not required to create a cluster.

Related tasks

Starting GPFS automatically

You can specify whether to start the GPFS daemon automatically on a node when it is started.

Starting GPFS automatically

You can specify whether to start the GPFS daemon automatically on a node when it is started.

About this task

Whether or not GPFS automatically starts is determined using the `autoload` parameter of the `mmchconfig` command. The default is *not* to automatically start GPFS on all nodes. You may change this by specifying `autoload=yes` using the `mmchconfig` command. This eliminates the need to start GPFS by issuing the `mmstartup` command when a node is booted.

The `autoload` parameter can be set the same or differently for each node in the cluster. For example, it may be useful to set `autoload=no` on a node that is undergoing maintenance since operating system upgrades and other software can often require multiple reboots to be completed.

Related concepts

GPFS node adapter interface names

An adapter interface name refers to the hostname or IP address that GPFS uses to communicate with a node. Specifically, the hostname or IP address identifies the communications adapter over which the GPFS daemons or GPFS administration commands communicate.

Creating an IBM Storage Scale cluster

When you create an IBM Storage Scale cluster, you can either create a small cluster of one or more nodes and later add nodes to it or you can create a cluster with all its nodes in one step.

IBM Storage Scale cluster configuration information

You can use the Cluster Configuration Repository (CCR) to maintain cluster configuration information.

Remote shell command

GPFS commands need to be able to communicate across all nodes in the cluster. To achieve this, the GPFS commands use the remote shell command that you specify on the `mmcrcluster` command or the `mmchcluster` command.

Remote file copy command

The GPFS commands must maintain several configuration files across all nodes in the cluster. To achieve this, the GPFS commands use the remote file copy command that you specify on the `mmcrcluster` command or the `mmchcluster` command.

Cluster name

Provide a name for the cluster by issuing the `-C` option on the `mmcrcluster` command.

User ID domain for the cluster

This option is the user ID domain for a cluster when accessing a file system remotely.

Cluster configuration file

GPFS provides default configuration values, so a cluster configuration file is not required to create a cluster.

Cluster configuration file

GPFS provides default configuration values, so a cluster configuration file is not required to create a cluster.

You can provide an optional cluster configuration file at the time of cluster creation. This optional file can be useful if you already know the correct parameter values based on previous testing or if you are restoring a cluster and have a backup copy of configuration values that apply to most systems. However, this option is not used at cluster creation time, and configuration parameters are modified after the cluster is created (by using the `mmchconfig` command).

For more information about the cluster configuration file, see the description of the **`mmcrcluster -c ConfigFile`** option in *IBM Storage Scale: Command and Programming Reference Guide*.

Related concepts

GPFS node adapter interface names

An adapter interface name refers to the hostname or IP address that GPFS uses to communicate with a node. Specifically, the hostname or IP address identifies the communications adapter over which the GPFS daemons or GPFS administration commands communicate.

Creating an IBM Storage Scale cluster

When you create an IBM Storage Scale cluster, you can either create a small cluster of one or more nodes and later add nodes to it or you can create a cluster with all its nodes in one step.

IBM Storage Scale cluster configuration information

You can use the Cluster Configuration Repository (CCR) to maintain cluster configuration information.

Remote shell command

GPFS commands need to be able to communicate across all nodes in the cluster. To achieve this, the GPFS commands use the remote shell command that you specify on the `mmcrcluster` command or the `mmhccluster` command.

Remote file copy command

The GPFS commands must maintain several configuration files across all nodes in the cluster. To achieve this, the GPFS commands use the remote file copy command that you specify on the `mmcrcluster` command or the `mmhccluster` command.

Cluster name

Provide a name for the cluster by issuing the `-C` option on the `mmcrcluster` command.

User ID domain for the cluster

This option is the user ID domain for a cluster when accessing a file system remotely.

Related tasks

Starting GPFS automatically

You can specify whether to start the GPFS daemon automatically on a node when it is started.

Disk considerations

Designing a proper storage infrastructure for your IBM Storage Scale file systems is key to achieving performance and reliability goals. When deciding what disk configuration to use, consider three key areas: infrastructure, performance, and disk access method.

Infrastructure

- Make sure that you have sufficient disks to meet the expected I/O load. In IBM Storage Scale terminology, a disk may be a physical disk or a RAID device.
- Make sure that you have sufficient connectivity (adapters and buses) between disks and network shared disk servers.
- Determine whether you are within IBM Storage Scale limits. Starting with GPFS 3.1, the structural limit on the maximum number of disks in a file system is 2048, which is enforced by IBM Storage Scale. (However, the number of disks in your system may be constrained by products other than IBM Storage Scale.)
- For a list of storage devices tested with IBM Storage Scale, see the [IBM Storage Scale FAQ in IBM Documentation](#).
- For Linux on Z, see the "Storage" topics "DASD device driver" and "SCSI-over-Fibre Channel device driver" in [Device Drivers, Features, and Commands](#), in the [Linux on IBM Systems documentation](#).

Disk access method

- Decide how your disks will be connected. Supported types of disk connectivity include the following configurations:
 1. All disks SAN-attached to all nodes in all clusters that access the file system

In this configuration, every node sees the same disk simultaneously and has a corresponding disk device entry.

2. Each disk connected to multiple NSD server nodes (up to eight servers), as specified on the server list

In this configuration, a single node with connectivity to a disk performs data shipping to all other nodes. This node is the first NSD server that is specified on the NSD server list. You can define additional NSD servers on the server list. Having multiple NSD servers guards against the loss of a single NSD server. When using multiple NSD servers, all NSD servers must have connectivity to the same disks. In this configuration, all nodes that are not NSD servers will receive their data over the local area network from the first NSD server on the server list. If the first NSD server fails, the next available NSD server on the list controls data distribution.

3. A combination of SAN-attached and an NSD server configuration.

Configuration consideration:

- If the node has a physical attachment to the disk and that connection fails, the node switches to using a specified NSD server to perform I/O. For this reason, it is recommended that you define NSDs with multiple servers, even if all nodes have physical attachments to the disk.
 - Configuring IBM Storage Scale disks without an NSD server stops the serving of data when the direct path to the disk is lost. This may be a preferable option for nodes that require a higher speed data connection to be provided through a SAN, as opposed to a lower speed network NSD server connection. Parallel jobs that use MPI often have this characteristic.
 - The `-o useNSDserver` file system mount option on the `mmmount`, `mount`, `mmchfs`, and `mmremote fs` commands can be used to specify the disk discovery, and limit or eliminate switching from local access to NSD server access, or the other way around.
- Decide whether you will use storage pools to manage your disks.

Storage pools allow you to manage your file system's storage in groups. You can partition your storage based on such factors as performance, locality, and reliability. Files are assigned to a storage pool based on defined policies.

Policies provide for the following:

- Placing files in a specific storage pool when the files are created
- Migrating files from one storage pool to another
- File deletion based on file characteristics

For more information, see *Storage pools* in *IBM Storage Scale: Administration Guide*.

Network Shared Disk (NSD) creation considerations

You must prepare each physical disk you intend to use with GPFS by first defining it as a Network Shared Disk (NSD) using the `mmcrnsd` command.

On Windows, GPFS will only create NSDs from empty disk drives. `mmcrnsd` accepts Windows *Basic* disk or *Unknown/Not Initialized* disks. It always re-initializes these disks so that they become *Basic GPT Disks* with a single *GPFS partition*. NSD data is stored in GPFS partitions. This allows other operating system components to recognize that the disks are used. `mmde1nsd` deletes the partition tables created by `mmcrnsd`.

A new NSD format was introduced with GPFS 4.1. The new format is referred to as NSD v2, and the old format is referred to as NSD v1. The NSD v1 format is compatible with GPFS releases prior to 4.1. The latest GPFS release recognizes both NSD v1 and NSD v2 formatted disks.

The NSD v2 format provides the following benefits:

- On Linux, includes a partition table so that the disk is easily recognized as a GPFS device
- Adjusts data alignment to support disks with a 4 KB physical block size
- Adds backup copies of some key GPFS data structures
- Expands some reserved areas to allow for future growth

Administrators do not need to select one format or the other when managing NSDs. GPFS will always create and use the correct format based on the `minReleaseLevel` for the cluster and the file system version. When `minReleaseLevel` (as reported by `mmfsconfig`) is less than 4.1.0.0, `mmcrnsd` will only create NSD v1 formatted disks. When `minReleaseLevel` is at least 4.1.0.0, `mmcrnsd` will only create NSD v2 formatted disks. In this second case, however, the NSD format may change dynamically when the NSD is added to a file system so that the NSD is compatible with the file system version.

On Linux, NSD v2 formatted disks include a GUID Partition Table (GPT) with a single partition. The GPT allows other operating system utilities to recognize when a disk is owned by GPFS, which helps prevent inadvertent data corruption. After running `mmcrnsd`, Linux utilities like `parted` can show the partition table. When an NSD v2 formatted disk is added to a 3.5 or older file system, its format is changed to NSD v1 and the partition table is converted to an MBR (MS-DOS compatible) type.

Note: Leftover persistent reserve (PR) keys can cause problems such as reservation conflicts in multipath, which can in turn cause I/O failure. In such cases, it is necessary to clean up leftover PR keys on a fresh installation. For a detailed procedure, see *Clearing a leftover Persistent Reserve reservation in IBM Storage Scale: Problem Determination Guide*.

The `mmcrnsd` command expects a stanza file as input. For details, see the following topics in *IBM Storage Scale: Command and Programming Reference Guide* and *IBM Storage Scale: Administration Guide*:

- *Stanza files*
- **`mmchdisk`** command
- **`mmchnsd`** command
- **`mmcrfs`** command
- **`mmcrnsd`** command

Related concepts

[IBM Storage Scale with data reduction storage devices](#)

Data reduction is a feature in storage arrays that optimizes storage usage with techniques such as compression, deduplication, and thin provisioning to provide storage space to applications that exceeds the amount of physical space in the device.

[NSD server considerations](#)

If you plan to use NSD servers to remotely serve disk data to other nodes, as opposed to having disks SAN-attached to all nodes, you should consider the total computing and I/O load on these nodes.

[File system descriptor quorum](#)

A GPFS structure called the *file system descriptor* is initially written to every disk in the file system and is replicated on a subset of the disks as changes to the file system occur, such as the adding or deleting of disks.

[Preparing direct access storage devices \(DASD\) for NSDs](#)

[Planning for GPUDirect Storage](#)

IBM Storage Scale support for GPUDirect Storage (GDS) enables a direct path between GPU memory and storage. You need to ensure that certain conditions are met before you start installing the feature.

IBM Storage Scale with data reduction storage devices

Data reduction is a feature in storage arrays that optimizes storage usage with techniques such as compression, deduplication, and thin provisioning to provide storage space to applications that exceeds the amount of physical space in the device.

Related concepts

[Network Shared Disk \(NSD\) creation considerations](#)

You must prepare each physical disk you intend to use with GPFS by first defining it as a Network Shared Disk (NSD) using the `mmcrnsd` command.

[NSD server considerations](#)

If you plan to use NSD servers to remotely serve disk data to other nodes, as opposed to having disks SAN-attached to all nodes, you should consider the total computing and I/O load on these nodes.

File system descriptor quorum

A GPFS structure called the *file system descriptor* is initially written to every disk in the file system and is replicated on a subset of the disks as changes to the file system occur, such as the adding or deleting of disks.

Preparing direct access storage devices (DASD) for NSDs

Planning for GPUDirect Storage

IBM Storage Scale support for GPUDirect Storage (GDS) enables a direct path between GPU memory and storage. You need to ensure that certain conditions are met before you start installing the feature.

IBM Storage Scale with thin-provisioned devices

IBM Storage Scale provides features to reduce the risks of thin provisioning and to optimize storage use in thin-provisioned storage systems.

With thin provisioning, storage space can be optimized by allocating space on demand and reclaiming when the application that is using a given volume declares that it is no longer using a given block of data. The storage system allocates storage to applications, including file systems, only when the storage is needed, from a pool of free storage. Thin provisioning addresses the problem of under-utilization of available storage system capacity. System administrators do not have to monitor and manage how much storage is required by each application. Instead, thin provisioning lets administrators provision large "thin" LUNs to a host based on estimated and anticipated usage without reserving physical space in advance. When the application writes data, the storage system allocates physical storage from the free pool on the storage system to the thin-provisioned LUNs.

Thin provisioning of storage for file systems can provide an excellent alternative in increasingly cost-conscious data centers by allowing storage hardware purchases to be delayed and capacity to be purchased just in time.

A significant concern with data reduction storage is the possibility that the device might run out of physical space on a thin-provisioned volume. For a file system running on such a device, the expectation is that the volume has been fully allocated, but if the storage has been over-provisioned, it is possible that no physical space exists to accommodate a write request. The result is that the device returns an I/O error, an event that can cause the file system to unmount and become inaccessible. In the case of an IBM Storage Scale file system, remounting the file system requires running through log recovery, a process that itself requires writing into the file system, and that might again fail if the device has no physical disk space available.

To minimize the possibility of the device running out of physical space, the device needs to be vigilantly monitored; otherwise, the benefits of thin provisioning might be offset by the possibility that the device runs out of space and brings the file system offline.



Attention: Data reduction in IBM Storage Scale, including thin-provisioned volumes, is supported only after a technical compatibility review by IBM. Ask your sales representative to contact IBM Storage Scale development about the RPQ or SCORE process.

Enhancement in IBM Storage Scale

IBM Storage Scale provides several features to mitigate the risks of supporting thin provisioning and to optimize storage use on thin-provisioned storage systems:

- Space reclamation: This operation ensures that disk blocks that are not used by the file system are returned to the storage device.
- Emergency disk space management: Space is reserved in the file system, in the form of metadata files with disk blocks filled with incompressible data, which can be freed back to the storage on demand.

Table 26. Thin provisioning storage solutions in IBM Storage Scale

Feature	Description	Benefits
Automatic recovery space. For more information, see the description of the three types of recovery space in the subtopic “Deployment recommendation” on page 332.	Automatic recovery space is reserved by the file system, which enables IBM Storage Scale to recover the storage system if it runs low on physical space.	This feature reduces system administrator workload to estimate how much additional physical space is needed to prevent a file system from becoming inaccessible when a storage system runs low on physical space.
The recovery log flush operation is thin-space aware.	The recovery log flush operation ensures that enough reserved space is always available to run log recovery.	This feature ensures that a file system can be recovered after a storage system runs out of physical space.
Thin space reclamation is supported.	The mmreclaimspace command reclaims unused space in a file system and returns it to the storage system.	This feature improves storage space utilization and savings.
Thin space emergency recovery.	Free the previously reserved space to allow the storage system to acquire enough space for file system recovery.	This feature recovers file systems after a failure caused by an out-of-physical-space condition in the storage system.

Deployment recommendation

Three types of emergency recovery space are used:

- *Required recovery space:* You must reserve this space before you create an IBM Storage Scale file system. Create several volumes in the storage system and fill them with uncompressible data. In case the storage system runs out of physical space, the required recovery space enables IBM Storage Scale to recover the volumes in the storage system back to a read/write state. For more information, see the subtopic [“During system configuration”](#) on page 332 later in this help topic.
- *Optional recovery space:* It is a good practice to also create several additional recovery volumes and fill them with uncompressible data. If the storage system runs nearly out of physical space, the optional recovery space enables IBM Storage Scale to prevent the storage system from running completely out of space. For more information, see the subtopic [“During system configuration”](#) on page 332 later in this help topic.
- *Automatic recovery space:* IBM Storage Scale reserves this space for a file system the first time that the file system is mounted. This type of space enables IBM Storage Scale to recover the file system if the storage system runs out of space. However, this type of space can be released to IBM Storage Scale only while the device is operational and in read/write mode. It cannot be released if the file system has become either read-only or offline because the storage system has run completely out of space. For more information, see [“Emergency recovery”](#) on page 334.

The following subtopics describe the steps that you should take for thin provisioning support during system configuration, during ongoing operations (production), and during emergency recovery after the storage system has run nearly or completely out of physical space.

During system configuration

During system configuration you must reserve a *required recovery space* and you can also create an *optional recovery space*. For the meaning of these terms, see the previous subtopic.

To create the required recovery space, create several emergency recovery volumes in the storage system and fill them with uncompressible data, such as fully random data that cannot be subject to deduplication. The size of this space varies depending on the type and brand of the storage system. For more information about the amount of reserved space that is needed, contact the vendor of the storage system. If you are using an IBM FlashSystem A9000, A9000R or any IBM Storage Virtualize systems that use IBM FlashCore Modules (FCMs), and the system runs out of space, contact IBM Storage Virtualize support team for assistance.

If the storage system runs out of physical space, IBM Storage Scale uses this required recovery space to bring the storage system back online. For more information, see [“Emergency recovery” on page 334](#).

After you create the required recovery space, you must create LUNs in the storage system, export them, and specify each LUN in an `nsd` stanza in the IBM Storage Scale stanza file. See the following example:

```
%nsd:
  nsd=gpfs1nsd
  usage=metadataOnly
  pool=system
  thinDiskType={scsi | nvme-thin}
```

The line `thinDiskType={scsi | nvme-thin}` is required and indicates that the device is a thin-provisioned disk.

Note: The device must support the SCSI WRITE SAME command.

For more information, see the topics *mmcrnsd command* and *mmcrfs command* in the *IBM Storage Scale: Command and Programming Reference Guide*.

To create the optional recovery space, create several volumes in the storage system and fill them with uncompressible data. A suggested size for the emergency recovery space is 20 GiB times the sum of all the local and remote nodes that can mount each of the file systems that is on the storage system. In other words:

$$20 \text{ GiB} \times ((\text{the number of local and remote nodes that can mount FS1}) + (\text{the number of local nodes that can mount FS2}) + \dots + (\text{the number of local and remote nodes that can mount FSN}))$$

If the storage system runs nearly out of physical space, IBM Storage Scale can use this optional recovery space to prevent the storage system from running completely out of physical space. For more information, see the subtopic [“Emergency recovery” on page 334](#) later in this help topic.

During production

During production you must carefully monitor the amount of physical space that is available in the storage system. The vendor of the storage system can probably provide utility programs for monitoring available physical space. If you determine that the storage system is running nearly out of physical space, follow the next steps. (If the storage system runs completely out of physical space, follow the instructions in the [“Emergency recovery” on page 334](#) subtopic later in this help topic.)

1. Quiesce all applications that are writing data into the storage system.
2. Delete unused volumes from the storage system.
3. Delete unused files and directories in the IBM Storage Scale file systems.
4. If you created optional recovery space during system configuration, free it. With this freed space IBM Storage Scale can recover the storage system before it runs completely out of space.
5. Issue the **mmreclaimspace** command to reclaim physical space in the storage system that is not in use but is not marked as available:

```
mmreclaimspace Device --reclaim-threshold 0
```

For more information, see the topic *mmreclaimspace command* in the *IBM Storage Scale: Command and Programming Reference Guide*.

Note: Because the space-reclaiming operation consumes a great deal of CPU, memory, network, and I/O system resources, it is a good idea to run the **mmreclaimspace** command at a time when the system is not heavily loaded.

6. Expand the free physical space in the storage system by adding new disks or deleting existing volumes.

If the storage system now has enough available physical space, you can resume all applications. If you freed the optional recovery space in Step 4, it is a good practice to re-create it as soon as possible.

Emergency recovery

If the storage system runs completely out of physical space, the volumes that use the storage system are likely to go offline or become read-only. To recover the system storage and to recover the IBM Storage Scale file systems, follow the procedure that is outlined next. This procedure has three main parts:

Part 1: Bring the file system to a state in which data can be read although data is not written. Follow these steps:

1. Stop all applications that are writing into the storage system.
2. Disable the operating system automount feature by issuing the appropriate system command:

- On Linux, issue the following command:

```
systemctl disable autofs
```

- In AIX, issue the following command:

```
stopsrc -s automountd
```

3. Issue the following command to unmount all IBM Storage Scale file systems:

```
/usr/lpp/mmfs/bin/mmumount all -a
```

4. Release the required recovery space that you reserved during system configuration. All volumes should be operational and in read/write mode.
5. Issue the following command to correct the states of the NSD disks. This action allows the mount operation in Step 7 to succeed across the cluster:

```
/usr/lpp/mmfs/bin/mmnsddiscover -a -N all
```

6. If you created optional recovery space during system configuration, free it.
7. Remount the file system in read-only mode. If the mount succeeds, skip the next step and go on to Part 2. IBM Storage Scale reserves the automatic recovery space as the file system is mounted for the first time.
8. If the mount operation in the preceding step fails, follow these steps to mount the file system:

- a. Issue the following command to mount the file system in restricted mode:

```
/usr/lpp/mmfs/bin/mmmount Device -o rs
```

- b. Issue the following command to release the automatic recovery space that IBM Storage Scale reserved when it created the file system:

```
/usr/lpp/mmfs/bin/mmreclaimspace Device --emergency-reclaim
```

Note: The **--emergency-reclaim** option is valid only if the thin-provisioned storage is read-writable and the file system is mounted in restrict mode. For more information, see the topic *mmreclaimspace* command in the *IBM Storage Scale: Command and Programming Reference Guide*.

- c. Remount the file system in read-only mode as in Step 7. The file system should be mounted successfully.

Part 2: Expand the free physical space of the storage system. You can add disks, delete data, or move data to other file systems.

Part 3: Bring the file system to a state in which data can be both read and written. Follow these steps:

1. Reserve required recovery space in the physical storage system as you did during system configuration.
2. If you freed optional recovery space in Step 6 of Part 1, re-create it as you did during system configuration.
3. Issue the following command to mount the file system in read/write mode:

```
/usr/lpp/mmfs/bin/mmmount Device
```

If you ran the **mmreclaimspace** command in Step 8(b) of Part 1, IBM Storage Scale reserves automatic recovery space during this mount operation.

4. Enable the operating system automount feature by issuing the appropriate system command:

- On Linux, issue the following command:

```
systemctl enable autofs
```

- In AIX, issue the following command:

```
startsrc -s automountd
```

If Step 3 is completed without any error, the file system is writable again. If an error occurred, the free physical space that you expanded in Part 2 is not enough. Repeat the emergency recovery process, beginning with Part 1, Step 1.

Graceful emergency recovery

Starting with release 5.2.0, IBM Storage Scale offers this additional emergency recovery feature for thin-provisioned storage systems, LUNs, and volumes created from FCM drives.

The graceful emergency recovery procedure provides a mounting mode for space reclaim (mount option **-o sr**), which allows system administrators and users to delete files to expand the free physical space for the system.

To recover the storage system and the IBM Storage Scale file system by using the graceful emergency recovery, follow the procedure outlined next. This procedure has three main parts.

Part 1: Bring the file system to a state in which it can be mounted in restricted mode: then, in preparation for the next step (where files can be accessed for deletion and removal), release the automatic recovery space. Follow these steps:

1. Stop all applications that are writing into the file system.
2. Issue the following command and make sure that the file system is unmounted (if it is not already unmounted by the **SGPanic()** procedure).

```
# /usr/lpp/mmfs/bin/mmumount Device -a
```

3. Mount the file system in restricted mode. Then, release the automatic recovery space that gets reserved when the file system is mounted in regular I/O mode.

```
# /usr/lpp/mmfs/bin/mmmount Device -o sr
```

To release the automatic recovery space:

```
# /usr/lpp/mmfs/bin/mmreclaimspace Device --emergency-reclaim
```

4. Unmount the file system.

```
# /usr/lpp/mmfs/bin/mmumount Device
```

Part 2: Mount the file system with the space reclaim mode, **-o sr**. Then, still in this mode, users must delete files and expand the free physical space of the storage system. Follow these steps:

1. Mount the file system in space reclaim mode.

```
# /usr/lpp/mmfs/bin/mmmount Device -o sr
```

2. Expand the free physical space of the file system by deleting (or moving out) the files in the file system.

Note: If a snapshot is created, make sure to also delete all of them (for more information, see the *mmdeletesnapshot* command in the *IBM Storage Scale: Command and Programming Reference Guide*). Otherwise, the free physical space is not recognized properly due to the snapshot effects.

3. Reclaim free spaces so that all of them are recognized by the storage system.

Note: The **mmreclaimspace** command would fail while the background space reclaim is running.

Note: If background space reclaim is enabled on the file systems, make sure to wait until the operation is completed. Check `mmfs.log.latest`, where you must see a message that displays an output similar to the following one:

```
# grep background /var/adm/ras/mmfs.log.latest
2024-02-02_13:06:37.543-0500: [I] Initializing background space reclamation for
'gpfs'...
2024-02-02_13:10:55.015-0500: [I] Exit background space reclamation for 'gpfs'...
```

4. After enough physical space is free, unmount the file system.

Note: The size of the free physical space can be subjective. It is important to make *enough* free space (or add more disks to the storage system) for the file system to prevent encountering the OOS (Out-Of-Space) condition soon.

```
# /usr/lpp/mmfs/bin/mmumount Device
```

Part 3: Bring the file system to a state in which data can be both read and written. Follow these steps:

1. Issue the following command to mount the file system in regular read/write mode.

```
# /usr/lpp/mmfs/bin/mmmount Device -a
```

2. If preferred, enable the operating system automount feature by issuing the appropriate system command.

```
# systemctl enable autofs
```

If the graceful emergency recovery procedure is completed without any error, the file system is writable again. If an error occurs or if the free physical space that you expanded in [Part 2](#) is not enough, repeat the graceful recovery procedure, beginning with [Part 1, Step 1](#).

IBM Storage Scale with TRIM-supporting SSDs and NVMe SSDs

IBM Storage Scale provides features that address the performance issues of flash memory.

File systems that create and remove files and directories often reuse storage blocks by overwriting the same storage blocks with new data contents. However, an NVMe or a general solid-state drive (SSD) device cannot overwrite a page of flash storage without first erasing the entire block of storage in which the page is located. This behavior creates a performance issue for I/O writes to previously used blocks of data when compared with I/O writes to unused or erased blocks. Also, write amplification occurs when the drive must save a copy of existing pages that are unaffected by the I/O write operation, then erase the entire block, and then restore the unaffected pages to the block after the erase.

To improve performance, the file system can issue a TRIM operation to an SSD notifying it about which blocks of data are no longer in use and can therefore be erased. The SSD erases the unused blocks before the blocks are required for the next reuse, which improves the performance of future I/O writes to the device. The TRIM operation also reduces fragmentation, because unused blocks are erased.

For more information about write amplification, garbage collection, and performance issues, see the vendor specifications for the flash memory devices that you are using.

Deployment recommendation

In the IBM Storage Scale stanza file, specify the `thinDiskType={scsi-recl | nvme}` line in each NSD stanza that describes a TRIM-supported device:

```
%nsd:
  nsd=gpfs1nsd
  usage=metadataOnly
  pool=system
  thinDiskType={scsi-recl | nvme}
```

This line indicates to IBM Storage Scale that the disk is trim-capable.

Enhancement in IBM Storage Scale

IBM Storage Scale provides the **mmreclaimspace** command, which calls the TRIM command to notify the specified SSD device as to which blocks are no longer in use and can be erased. For more information, see the topic *mmreclaimspace* in the *IBM Storage Scale: Command and Programming Reference Guide*.

Note: Because the space-reclaiming operation consumes a great deal of CPU, memory, network, and I/O system resources, it is a good idea to run the **mmreclaimspace** command at a time when the system is not heavily loaded.

Support matrix

The following conditions must be met for thin provisioning support to be activated:

- The file system must be upgraded to file system format version 5.0.4 or later.
- The file system must be upgraded to IBM Storage Scale 5.2.3.2 or later to recognize the additional `nvme-thin` and `scsi-recl` for the **thinDiskType** attribute.
- The stanza file must include the following line to add thin-provisioned disks or TRIM-supported SSDs into the file system. See the example NSD stanza in the preceding subtopic, *Deployment recommendation*:

```
thinDiskType={scsi-recl | scsi | nvme | nvme-thin}
```

- Both thin-provisioned disks and TRIM-supported SSDs must be connected to nodes that are running the Linux operating system.

Table 27. Supported operating systems and device connections		
Operating system	A node can act as an NSD server or as a node with a disk directly attached	A node can act as a client
Linux (RHEL 7.3 and later SLES 11 and later, Ubuntu 16 and later)	Yes	Yes
Linux on Z	No	Yes
AIX	No	Yes
Windows	No	Yes

Table 28. Supported devices

Supported devices

The following storage devices are supported:

- IBM A9000/R
- IBM v7K
- Intel SSDs

The following devices are believed to operate correctly with IBM Storage Scale but support has not been fully verified:

- IBM FS9100
- General SSDs

IBM Storage Scale with IBM FlashCore Modules (FCMs) drives

Starting from version 5.1.6, IBM Storage Scale supports compression-enabled FCM 3.0 drives (or higher).

FCM drives are capable of data compression with no performance penalty. The data compression ratio can vary depending on the generation of the FCM and the data or file types that are being written to the drive. Due to the indeterminable nature of data compression, it is possible to run out of physical space for the file system created from the FCM.

IBM Storage Scale FCMs are a family of high-performance flash drives in a standard form factor of 2.5 inches, 15 millimeters. The FCM design uses the NVMe protocol, a PCIe Gen 3 U.2 interface, and high-speed NAND memory to provide high throughput and input/output operations per second with consistent and predictable latency.

FCM drives support of IBM Storage Scale is built on IBM Storage Scale RAID technology so that users can only deploy an IBM Storage Scale file system on FCM drives by using an IBM Storage Scale System. This support introduces space monitor capability, which is meant to monitor the physical space utilization of the declustered array (DA) that is created from FCM drives.

Note: The space monitor is available and enabled for FCM drives only on an IBM Storage Scale System.

The FCM features hardware-based, real-time inline data compression and self-encryption that provides consistent compression and encryption across the full range of workloads and with no performance penalties. By using FCM, usable capacity can be significantly increased while maintaining consistent application performance. FCM can substantially reduce costs for storage acquisition, rack space, power, and cooling; FCM also contributes to extend the life of existing storage assets.

When connected, the IBM Storage Scale System recognizes the presence of FCM drives and starts space monitoring. This capability monitors the space utilization of the DA that was created from the virtual disk. When a file system is created out of those virtual disks, IBM Storage Scale treats the FCM drives the same way as it treats the NSDs with thin-provisioning enabled, and the emergency space gets reserved for the file system.

Note: For the space monitor capability, refer to IBM Storage Scale System 3500 and IBM Storage Scale System 6000 documentation.

The following example shows how to specify a stanza for the NSD-based of an FCM drive:

```
%nsd:
nsd=gpfs1nsd
usage-metadataOnly
pool=system
thinDiskType=nvme-thin
```

For an IBM Storage Scale System with FCM drives, a stanza file can be used. However, creating a file system with the stanza file is not a typical practice; instead, the deployment is done through automated scripts and the file system is created then. The script has `thinDiskType` automatically configured so that the administrator does not need to manually set up `thinDiskType`. For more information, see [IBM](#)

[Storage Scale System Deployment Guide](#) (and instructions for the new deployment of an Elastic Storage Server).

Configuration

For a file system that consists of FCM drives, a new set of control variables are introduced to control the behavior of the file system with FCM drives in the context of space monitoring and space reclaim.

For the DA, physical space states are defined for the space monitor. These states are incremental and defined as: Good (first physical space state), Info, Warning, Critical, and OOS (out-of-space). As the space of each DA is filled with data, a state progresses higher up until it reaches the next state.

```
2024-05-10_15:36:36.830-0400: [I] Command: successful mount essfs1
2024-05-10_15:36:39.338-0400: [I] DiskSpaceState transitioned from Invalid to Good for essfs1
pool 0, bg reclaim config (mode/threshold/interval/rate) changed to: 1/10/60/1
```

Note: For more details about the physical space states, see IBM Storage Scale System 3500 and IBM Storage Scale System 6000 documentation.

As the compression ratio can vary depending on the file types that are being written to the FCM drives, it is important to manage and monitor the physical space utilization and the state.

The space monitor, when IBM Storage Scale is initializing, gets automatically enabled and started if the presence of FCM drives is detected on an IBM Storage Scale System.

Deployment recommendation

Data reduction storage types like FCM and thin-provisioned storage systems share the concern about the possibility that the device might run out of physical space on a drive (or a LUN or a volume that is created from the drive).

In addition to existing features that mitigate the out-of-space handling or recovery, IBM Storage Scale provides a feature called graceful emergency recovery for FCM.

For more information, see [“Graceful emergency recovery” on page 335](#).

NSD server considerations

If you plan to use NSD servers to remotely serve disk data to other nodes, as opposed to having disks SAN-attached to all nodes, you should consider the total computing and I/O load on these nodes.

- Will your Network Shared Disk servers be dedicated servers or will you also be using them to run applications? If you will have non-dedicated servers, consider running less time-critical applications on these nodes. If you run time-critical applications on a Network Shared Disk server, servicing disk requests from other nodes might conflict with the demands of these applications.
- The special functions of the file system manager consume extra processing time. If possible, avoid using a Network Shared Disk server as the file system manager. The Network Shared Disk server consumes both memory and processor cycles that could impact the operation of the file system manager. See [“File system manager” on page 10](#).
- The actual processing capability required for Network Shared Disk service is a function of the application I/O access patterns, the type of node, the type of disk, and the disk connection. You can later run `iostat` on the server to determine how much of a load your access pattern will place on a Network Shared Disk server.
- Providing sufficient disks and adapters on the system to yield the required I/O bandwidth. Dedicated Network Shared Disk servers should have sufficient disks and adapters to drive the I/O load you expect them to handle.
- Knowing approximately how much storage capacity you will need for your data.

You should consider what you want as the default behavior for switching between local access and NSD server access in the event of a failure. To set this configuration, use the `-o useNSDserver` file system mount option of the `mmount`, `mount`, `mmchfs`, and `mmremotefs` commands to:

- Specify the disk discovery behavior
- Limit or eliminate switching from either:
 - Local access to NSD server access
 - NSD server access to local access

You should consider specifying how long to wait for an NSD server to come online before allowing a file system mount to fail because the server is not available. The `mmchconfig` command has these options:

nsdServerWaitTimeForMount

When a node is trying to mount a file system whose disks depend on NSD servers, this option specifies the number of seconds to wait for those servers to come up. If a server recovery is taking place, the wait time you are specifying with this option starts after recovery completes.

Note: The decision to wait for servers is controlled by the `nsdServerWaitTimeWindowOnMount` option.

nsdServerWaitTimeWindowOnMount

Specifies a window of time (in seconds) during which a mount can wait for NSD servers as described for the `nsdServerWaitTimeForMount` option. The window begins when quorum is established (at cluster startup or subsequently), or at the last known failure times of the NSD servers required to perform the mount.

Notes:

1. When a node rejoins a cluster, it resets all the failure times it knew about within that cluster.
2. Because a node that rejoins a cluster resets its failure times within that cluster, the NSD server failure times are also reset.
3. When a node attempts to mount a file system, GPFS checks the cluster formation criteria first. If that check falls outside the window, it will then check for NSD server fail times being in the window.

Related concepts

Network Shared Disk (NSD) creation considerations

You must prepare each physical disk you intend to use with GPFS by first defining it as a Network Shared Disk (NSD) using the `mmcrnsd` command.

IBM Storage Scale with data reduction storage devices

Data reduction is a feature in storage arrays that optimizes storage usage with techniques such as compression, deduplication, and thin provisioning to provide storage space to applications that exceeds the amount of physical space in the device.

File system descriptor quorum

A GPFS structure called the *file system descriptor* is initially written to every disk in the file system and is replicated on a subset of the disks as changes to the file system occur, such as the adding or deleting of disks.

Preparing direct access storage devices (DASD) for NSDs

Planning for GPUDirect Storage

IBM Storage Scale support for GPUDirect Storage (GDS) enables a direct path between GPU memory and storage. You need to ensure that certain conditions are met before you start installing the feature.

File system descriptor quorum

A GPFS structure called the *file system descriptor* is initially written to every disk in the file system and is replicated on a subset of the disks as changes to the file system occur, such as the adding or deleting of disks.

Based on the number of failure groups and disks, GPFS creates one to five replicas of the descriptor:

- If there are at least five different failure groups, five replicas are created.
- If there are at least three different disks, three replicas are created.

- If there are only one or two disks, a replica is created on each disk.

Once it decides how many replicas to create, GPFS picks disks to hold the replicas, so that all replicas are in different failure groups, if possible, to reduce the risk of multiple failures. In picking replica locations, the current state of the disks is taken into account. Stopped or suspended disks are avoided. Similarly, when a failed disk is brought back online, GPFS might rebalance the file system descriptors in order to assure reliability across the failure groups. The disks used to hold the file system descriptor replicas can be seen by running the `mmfslsdisk fsname -L` command and looking for the string desc in the remarks column.

GPFS requires that a majority of the replicas on the subset of disks remain available to sustain file system operations:

- If there are at least five different replicas, GPFS can tolerate a loss of two of the five replicas.
- If there are at least three replicas, GPFS can tolerate a loss of one of the three replicas.
- If there are fewer than three replicas, a loss of one replica might make the descriptor inaccessible.

The loss of all disks in a disk failure group might cause a majority of file systems descriptors to become unavailable and inhibit further file system operations. For example, if your file system is backed up by three or more disks that are assigned to two separate disk failure groups, one of the failure groups will be assigned two of the file system descriptor replicas, while the other failure group will be assigned only one replica. If all of the disks in the disk failure group that contains the two replicas were to become unavailable, the file system would also become unavailable. To avoid this particular scenario, you might want to introduce a third disk failure group consisting of a single disk that is designated as a descOnly disk. This disk would exist solely to contain a replica of the file system descriptor (that is, it would not contain any file system metadata or data). This disk should be at least 128MB in size.

For more information, see [“Network Shared Disk \(NSD\) creation considerations”](#) on page 329 and [Establishing disaster recovery for your GPFS cluster](#) in *IBM Storage Scale: Administration Guide*.

Related concepts

[Network Shared Disk \(NSD\) creation considerations](#)

You must prepare each physical disk you intend to use with GPFS by first defining it as a Network Shared Disk (NSD) using the `mmcrnsd` command.

[IBM Storage Scale with data reduction storage devices](#)

Data reduction is a feature in storage arrays that optimizes storage usage with techniques such as compression, deduplication, and thin provisioning to provide storage space to applications that exceeds the amount of physical space in the device.

[NSD server considerations](#)

If you plan to use NSD servers to remotely serve disk data to other nodes, as opposed to having disks SAN-attached to all nodes, you should consider the total computing and I/O load on these nodes.

[Preparing direct access storage devices \(DASD\) for NSDs](#)

[Planning for GPUDirect Storage](#)

IBM Storage Scale support for GPUDirect Storage (GDS) enables a direct path between GPU memory and storage. You need to ensure that certain conditions are met before you start installing the feature.

Preparing direct access storage devices (DASD) for NSDs

When preparing direct access storage devices (DASD) for NSDs, see the "Disk hardware tested with GPFS for Linux on Z" table in the [IBM Storage Scale FAQ](#) in IBM Documentation.

Related concepts

[Network Shared Disk \(NSD\) creation considerations](#)

You must prepare each physical disk you intend to use with GPFS by first defining it as a Network Shared Disk (NSD) using the `mmcrnsd` command.

[IBM Storage Scale with data reduction storage devices](#)

Data reduction is a feature in storage arrays that optimizes storage usage with techniques such as compression, deduplication, and thin provisioning to provide storage space to applications that exceeds the amount of physical space in the device.

NSD server considerations

If you plan to use NSD servers to remotely serve disk data to other nodes, as opposed to having disks SAN-attached to all nodes, you should consider the total computing and I/O load on these nodes.

File system descriptor quorum

A GPFS structure called the *file system descriptor* is initially written to every disk in the file system and is replicated on a subset of the disks as changes to the file system occur, such as the adding or deleting of disks.

Planning for GPUDirect Storage

IBM Storage Scale support for GPUDirect Storage (GDS) enables a direct path between GPU memory and storage. You need to ensure that certain conditions are met before you start installing the feature.

Preparing your environment for use of extended count key data (ECKD) devices

If your GPFS cluster includes Linux on Z instances, do not use virtual reserve/release.

Instead, follow the process that is described in [Sharing DASD without Using Virtual Reserve/Release \(https://www.ibm.com/docs/en/zvm/7.1?topic=sharing-dasd-without-using-virtual-reserverelease\)](https://www.ibm.com/docs/en/zvm/7.1?topic=sharing-dasd-without-using-virtual-reserverelease). Data integrity is handled by GPFS itself.

Preparing an ECKD device for GPFS

To prepare an ECKD device for GPFS, complete these steps on a single node:

1. Make sure that the ECKD device is online. To set it online, issue the following command:

```
chccwdev -e device_bus_id
```

Where *device_bus_id* identifies the device to be configured. *device_bus_id* is a device number with a leading 0.n, where *n* is the subchannel set ID. For example:

```
chccwdev -e 0.0.3352
```

2. Low-level format the ECKD using one of the following commands.

Note: GPFS supports ECKD disks in either compatible disk layout (CDL) format or Linux disk layout (LDL) format. The DASD must be formatted with a block size of 4096.

- To specify CDL format, issue the following command:

```
dasdfmt -d cdl device
```

There is no need to specify a block size value, as the default value is 4096.

- To specify LDL format, issue the following command:

```
dasdfmt -d ldl device
```

There is no need to specify a block size value, as the default value is 4096.

In both of these commands, *device* is the node of the device. For example:

```
dasdfmt -d cdl /dev/dasda
```

3. This step is for *CDL disks only*. It is an *optional* step because partitioning is optional for CDL disks.

If you want to partition the ECKD and create a single partition that spans the entire device, use the following command:

```
fdasd -a device
```

Note: This step is not required for LDL disks because the `dasdfmt -d 1d1` command issued in the previous step automatically creates a single Linux partition on the disk.

4. If the ECKD device is shared with other cluster nodes, then it is a good idea to set the ECKD device offline and then back online on each node where the ECKD is shared. This procedure helps to ensure that the updated partition information is spread among all shared nodes.

For more information about all of these commands, see the following documentation.

- "Commands for Linux on Z" at [Device Drivers, Features, and Commands](#), in the [Linux on IBM Systems documentation](#).
- "Getting started with Elastic Storage for Linux on Z based on GPFS technology" white paper, available on the Welcome Page for [IBM Storage Scale in IBM Documentation](#).

Repeat these steps for each ECKD to be used with GPFS.

After preparing the environment, set the ECKD devices online on the other nodes.

Note: After partitioning and formatting an ECKD device, the partition information is not distributed automatically to other cluster nodes. If the ECKD device is shared with different cluster nodes when they are online. It is recommended to set the ECKD device offline and then back online. This procedure updates the partition information for the specific ECKD device on the cluster node.

Always make sure that the ECKD devices are online before you start GPFS. To automatically set ECKD devices online at system start, see the documentation for your Linux distribution.

Planning for GPUDirect Storage

IBM Storage Scale support for GPUDirect Storage (GDS) enables a direct path between GPU memory and storage. You need to ensure that certain conditions are met before you start installing the feature.

For more information on GDS, see [“GPUDirect Storage support for IBM Storage Scale” on page 28](#).

The following list provides the prerequisites for installing and using GDS:

- **IBM Storage Scale version:** Full GDS support including accelerated writes requires IBM Storage Scale 5.1.6 on the GPU client and on the NSD server side. The earlier versions starting with IBM Storage Scale 5.1.2 (IB) or 5.1.3 (RoCE) support accelerated reads and writes in compatibility mode but not accelerated writes. For more details on supported versions, see [IBM Storage Scale FAQ in IBM Documentation](#). Moreover, if you need to configure GDS for a remote file system, IBM Storage Scale 5.1.2 or later is required on all NSD servers of the remote cluster.
- **Data path:** GDS requires the NSD path for storage access. The storage must not be locally attached to the GDS clients.
- **Storage:** ESS and non-ESS storage servers are supported. The storage servers can be in the same or different cluster as the GDS clients.
- GDS requires Mellanox RDMA over InfiniBand or RoCE between GDS clients and storage servers.
- **Adapters:** Mellanox ConnectX-5 or newer adapters are required to use GDS with RoCE.

Supported hardware

The following list provides the hardware requirements:

- **GDS clients:** x86 with a GPU model that supports GDS. For more details, see [NVIDIA GDS documentation](#).
- **Network:** EDR or HDR InfiniBand, Ethernet (RoCE).
- **InfiniBand adapter:** CX5 or CX6. (Mellanox CX4 is OK for IB only if the CX4 firmware is 12.27.4000 or higher.)

Supported software for GDS clients

GDS client must be an IBM Storage Scale NSD client. For information about the supported versions of the required software for GDS clients, see [Components required for GDS in IBM Storage Scale FAQ in IBM Documentation](#).

Supported software for GDS servers

The following list provides the software requirements:

- A storage server for GDS must either be an IBM Storage Scale NSD server or an IBM ESS or ECE.
- Any operating system that supports IBM Storage Scale can be used.
- MOFED: Use the Mellanox OFED stack for the OS.
- For details about driver versions, see [Components required for GDS in IBM Storage Scale FAQ in IBM Documentation](#).

Related concepts

[Network Shared Disk \(NSD\) creation considerations](#)

You must prepare each physical disk you intend to use with GPFS by first defining it as a Network Shared Disk (NSD) using the `mmcrnsd` command.

[IBM Storage Scale with data reduction storage devices](#)

Data reduction is a feature in storage arrays that optimizes storage usage with techniques such as compression, deduplication, and thin provisioning to provide storage space to applications that exceeds the amount of physical space in the device.

[NSD server considerations](#)

If you plan to use NSD servers to remotely serve disk data to other nodes, as opposed to having disks SAN-attached to all nodes, you should consider the total computing and I/O load on these nodes.

[File system descriptor quorum](#)

A GPFS structure called the *file system descriptor* is initially written to every disk in the file system and is replicated on a subset of the disks as changes to the file system occur, such as the adding or deleting of disks.

[Preparing direct access storage devices \(DASD\) for NSDs](#)

[“GPUDirect Storage support for IBM Storage Scale” on page 28](#)

IBM Storage Scale's support for NVIDIA's GPUDirect Storage (GDS) enables a direct path between GPU memory and storage. This solution addresses the need for higher throughput and lower latencies. File system storage is directly connected to the GPU buffers to reduce latency and load on the CPU. For IBM Storage Scale, this means that data can be read or written directly from / to an NSD server's pagepool and it is sent to the GPU buffer of the IBM Storage Scale clients by using RDMA. IBM Storage Scale with GDS requires an InfiniBand or RoCE fabric. In IBM Storage Scale, the `mmdiag` command is enhanced to print diagnostic information for GPUDirect Storage.

Related tasks

[“Installing GPUDirect Storage for IBM Storage Scale” on page 639](#)

IBM Storage Scale support for GPUDirect Storage (GDS) enables a direct path between GPU memory and storage.

[“Upgrading GPUDirect Storage” on page 692](#)

The recommended version for use of GDS with IBM Storage Scale is 5.1.6 as it supports accelerated writes.

File system creation considerations

File system creation involves anticipating usage within the file system and considering your hardware configurations. Before creating a file system, consider how much data will be stored and how great is the demand for the files in the system.

Each of these factors can help you to determine how much disk resource to devote to the file system, which block size to choose, where to store data and metadata, and how many replicas to maintain. For the latest supported file system size, see the [IBM Storage Scale FAQ](#) in *IBM Documentation*.

Your GPFS file system is created by issuing the `mmcrfs` command. [Table 29 on page 345](#) details the file system creation options that are specified on the `mmcrfs` command, which options can be changed later with the `mmchfs` command, and what the default values are.

To move an existing file system into a new GPFS cluster, see *Exporting file system definitions between clusters* in *IBM Storage Scale: Administration Guide*.

Table 29. File system creation options			
Options	mmcrfs	mmchfs	Default value
<code>--auto-inode-limit</code> automatically increases the maximum number of inodes per inode space in the file system. See the topic <i>mmcrfs command</i> in the <i>IBM Storage Scale: Command and Programming Reference Guide</i> .	X	X	<code>--noauto-inode-limit</code>
Device name of the file system. See “Device name of the file system” on page 349 .	X	X	none
DiskDesc for each disk in your file system. Note: The use of disk descriptors is discouraged. See “Disks for your file system” on page 349 .	X	Issue the <code>mmadddisk</code> or <code>mmdeledisk</code> command to add or delete disks from the file system.	none
<code>-F StanzaFile</code> specifies a file that contains a list of NSD stanzas and pool stanzas. For more information, see <i>Stanza files</i> in <i>IBM Storage Scale: Administration Guide</i> .	X	Issue the <code>mmadddisk</code> or <code>mmdeledisk</code> command to add or delete disks as indicated in the stanza file.	none
<code>-A {yes no automount}</code> to determine when to mount the file system. See “Deciding how the file system is mounted” on page 350 .	X	X	yes

Table 29. File system creation options (continued)

Options	mmcrfs	mmchfs	Default value
-B <i>BlockSize</i> By default, all the data blocks and metadata blocks in a file system are set to the same block size with the same subblock size. See “Block size” on page 350 and <i>mmcrfs</i> command in <i>IBM Storage Scale: Command and Programming Reference Guide</i>.	X	The block size, subblock size, and number of subblocks per block of a file system are set when the file system is created and cannot be changed later.	4 MiB
-D {posix nfs4} semantics for a deny-write open lock See “NFSv4 deny-write open lock” on page 349.	X	X	nfs4
-E {yes no} to report exact mtime values. See “mtime values” on page 354.	X	X	yes
--flush-on-close specifies the automatic flushing of disk buffers when closing files that were opened for write operations on the device. See the topic <i>mmcrfs</i> command in the <i>IBM Storage Scale: Command and Programming Reference Guide</i>.	X	X	--noflush-on-close
-i <i>InodeSize</i> to set the size of inodes: 512, 1024, or 4096 bytes.	X	This value cannot be changed.	4096
-j {cluster scatter} to determine the block allocation map type. See “Block allocation map” on page 354.	X	NA	See “Block allocation map” on page 354 .
-k {posix nfs4 all} to determine the authorization types supported by the file system. See “File system authorization” on page 354.	X	X	all
-K {no whenpossible always} to enforce strict replication. See “Type of replication” on page 355.	X	X	whenpossible

Table 29. File system creation options (continued)

Options	mmcrfs	mmchfs	Default value
-L <i>LogFileSize</i> to specify the size of the internal log files. See “GPFS recovery logs” on page 15.	X	X	See the topic <i>mmcrfs</i> command in the <i>IBM Storage Scale: Command and Programming Reference Guide</i> .
-m <i>DefaultMetadataReplicas</i> See “File system replication parameters” on page 355.	X	X	1
-M <i>MaxMetadataReplicas</i> See “File system replication parameters” on page 355.	X	This value cannot be changed.	2
-n <i>NumNodes</i> that will mount the file system. See “Number of nodes mounting the file system” on page 357.	X	X	32
--nfs4-owner-write-acl specifies whether object owners are given implicit NFSv4 WRITE_ACL permission. See the topic <i>mmcrfs</i> command in the <i>IBM Storage Scale: Command and Programming Reference Guide</i> .	X	X	yes
-o <i>MountOptions</i> to be passed to the mount command. See “Assign mount command options” on page 357.	NA	X	none
-Q {yes no} to activate quota. See “Enabling quotas” on page 357.	X	X	no
-r <i>DefaultDataReplicas</i> See “File system replication parameters” on page 355.	X	X	1
-R <i>MaxDataReplicas</i> See “File system replication parameters” on page 355.	X	This value cannot be changed.	2
-S {yes no relatime} to control how the atime value is updated. See “atime values” on page 353.	X	X	The default value is relatime if minReleaseLevel is 5.0.0 or later when the file system is created. Otherwise the default value is no.

Table 29. File system creation options (continued)

Options	mmcrfs	mmchfs	Default value
-t <i>DriveLetter</i> See “Windows drive letter” on page 357.	X	X	none
-T <i>Mountpoint</i> See “Mountpoint directory” on page 357.	X	X	/gpfs/DeviceName
-V {full compat} to change the file system format to the latest level. See “Changing the file system format to the latest level” on page 359	NA	X	none
-v {yes no} to verify disk usage. See “Verifying disk usage” on page 359.	X	NA	yes
-W <i>NewDeviceName</i> to assign a new device name to the file system.	NA	X	none
-z {yes no} to enable DMAPI See “Enabling DMAPI” on page 359.	X	X	no
--filesetdf to specify (when quotas are enforced for a fileset) whether the df command will report numbers based on the quotas for the fileset and not for the total file system.	X	X	--nofilesetdf
--inode-limit <i>MaxNumInodes</i> [: <i>NumInodesToPreallocate</i>] to determine the maximum number of files in the file system. See “Specifying the maximum number of files that can be created” on page 360.	X	X	file system size/1 MiB
--log-replicas <i>LogReplicas</i> to specify the number of recovery log replicas.	X	X	none
--mount-priority <i>Priority</i> to control the order in which the individual file systems are mounted at daemon startup or when one of the all keywords is specified on the mmmount command.	X	X	0

Table 29. File system creation options (continued)

Options	mmcrfs	mmchfs	Default value
--perfilesset-quota to set the scope of user and group quota limit checks to the individual filesset level.	X	X	--noperfilesset-quota
--rapid-repair to keep track of incomplete replication on an individual file block basis (as opposed to the entire file).	NA	X	none
--version <i>VersionString</i> to enable only the file system features that are compatible with the specified release. See “Enabling file system features” on page 359.	X	NA	The default value is defined by the current committed function level for the cluster minReleaseLevel, for example 4.2.2.0 or other values, depending on the cluster.
Notes: 1. X – indicates that the option is available on the command. 2. NA (not applicable) – indicates that the option is not available on the command.			

Device name of the file system

File system names must be unique within a GPFS cluster. However, two different clusters can have two distinct file systems with the same name.

The device name of the file system does not need to be fully qualified. `fs0` is as acceptable as `/dev/fs0`. The name cannot be the same as an existing entry in `/dev`. The file system name must be no longer than 128 characters.

Note: If your cluster includes Windows nodes, the file system name must be no longer than 32 characters.

NFSv4 deny-write open lock

You can specify whether a deny-write open lock blocks writes, which is expected and required by NFSv4, Samba, and Windows.

For more information, see *Managing GPFS access control lists* in *IBM Storage Scale: Administration Guide*.

nfs4

Must be specified for file systems that support NFSv4 and file systems that are mounted on Windows. NFS4 is the default.

posix

Specified for file systems that support NFSv3 or file systems that are not NFS exported.

Allows NFS writes even in the presence of a deny-write open lock.

Disks for your file system

Disks must be defined as network shared disks (NSDs) before they can be added to a GPFS file system.

NSDs are created using the `mmcrnsd` command. Use the `mmllnsd -F` command to display a list of available NSDs. Use the `mmadddisk` to add the NSD after you create it.

See “Disk considerations” on page 328.

Deciding how the file system is mounted

About this task

Specify when the file system is to be mounted:

yes

When the GPFS daemon starts. This is the default.

no

Manual mount.

automount

When the file system is first accessed.

This can be changed at a later time by using the `-A` option on the `mmchfs` command.

Considerations:

1. GPFS mount traffic may be lessened by using the automount feature to mount the file system when it is first accessed instead of at GPFS startup. Automatic mounts only produce additional control traffic at the point that the file system is first used by an application or user. Mounts at GPFS startup on the other hand produce additional control traffic at every GPFS startup. Thus startup of hundreds of nodes at once may be better served by using automatic mounts.
2. Automatic mounts will fail if the node does not have the operating system's automount support enabled for the file system.
3. When exporting file systems for NFS mounts, it may be useful to mount the file system when GPFS starts.

Block size

Choose the file system block size based on the projected workload of the file system and the type of storage that it uses.

[“General” on page 350](#)

[“Test actual performance with different block sizes” on page 351](#)

[“Factors that can affect performance” on page 352](#)

General

In a file system, a *block* is the largest contiguous amount of disk space that can be allocated to a file and also the largest amount of data that can be transferred in a single I/O operation. The block size determines the maximum size of a read request or write request that a file system sends to the I/O device driver. Blocks are composed of an integral number of *subblocks*, which are the smallest unit of contiguous disk space that can be allocated to a file. Files larger than one block are stored in some number of full blocks plus any subblocks that might be required after the last block to hold the remaining data. Files smaller than one block size are stored in one or more subblocks.

For information about setting block size and subblock size for a file system, see the descriptions of the `-B BlockSize` parameter and the `--metadata-block-size` parameter in help topic *mmcrfs command* in the *IBM Storage Scale: Command and Programming Reference Guide*. Here are some general facts from those descriptions:

- The block size, subblock size, and number of subblocks per block of a file system are set when the file system is created and cannot be changed later.
- All the data blocks in a file system have the same block size and the same subblock size. Data blocks and subblocks in the system storage pool and those in user storage pools have the same sizes. An example of a valid block size and subblock size is a 4 MiB block with an 8 KiB subblock.
- All the metadata blocks in a file system have the same block size and the same subblock size. The metadata blocks and subblocks are set to the same sizes as data blocks and subblocks, unless the `--metadata-block-size` parameter is specified.

Note: The `--metadata-block-size` parameter that is used to specify a different metadata block size than the data block size is being deprecated. This option is no longer required to use for performance improvements for file systems with file system format 5.0.0 or later and it will be removed in a future release.

- If the system storage pool contains only metadataOnly NSDs, the metadata block can be set to a different size than the data block size with the `--metadata-block-size` parameter.

Note: This setting can result in a change in the data subblock size and in the number of subblocks in a data block, if the block size (`-B` parameter) is different from the `--metadata-block-size`. For an example, see Scenario 3 in a later bullet in this list.

- The data blocks and metadata blocks must have the same number of subblocks, even when the data block size and the metadata block size are different. See Scenario 3 in the next bullet.
- The number of subblocks per block is derived from the smallest block size of any storage pool in the file system, including the system metadata pool. Consider the following example scenarios:

Note: For a table of the valid block sizes and subblock sizes, see *mmcrfs command* in the *IBM Storage Scale: Command and Programming Reference Guide*.

- Scenario 1: The file system is composed of a single system storage pool with all the NSD usage configured as dataAndMetadata. The file system block size is set with the `-B` parameter to 16 MiB. As a result, the block size for both metadata and data blocks is 16 MiB. The metadata and data subblock size is 16 KiB.
 - Scenario 2: The file system is composed of multiple storage pools with system storage pool NSD usage configured as metadataOnly and user storage pool NSD usage configured as dataOnly. The file system block size is set (`-B` parameter) to 16 MiB. The `--metadata-block-size` is also set to 16 MiB. As a result, the metadata and data block size is 16 MiB. The metadata and data subblock size is 16 KiB.
 - Scenario 3: The file system is composed of multiple storage pools with the system storage pool NSD usage configured as metadataOnly and the user storage pool NSD usage configured as dataOnly. The file system block size is set (`-B` parameter) to 16 MiB, which has a subblock size of 16 KiB, but the `--metadata-block-size` is set to 1 MiB, which has a subblock size of 8 KiB. The number of subblocks across the pools of a file system needs to be the same and this is calculated based on the storage pool with smallest block size. In this case, the system pool has the smallest block size (1 MiB). The number of subblocks per block in the system storage pool is 128 (1 MiB block size / 8 KiB subblock size = 128 subblocks per block). The other storage pools inherit the 128-subblocks-per-block setting and their subblock size is recalculated based on 128 subblocks per block. In this case the subblock size of the user storage pool is recalculated as 128 KiB (16 MiB / 128 subblocks per block = 128 KiB subblock size)
- The block size cannot exceed the value of the cluster attribute `maxblocksize`, which can be set by the `mmchconfig` command.

Select a file system block size based on the workload and the type of storage. For a list of supported block sizes with their subblock sizes, see the description of the `-B BlockSize` parameter in help topic *mmcrfs command* in the *IBM Storage Scale: Command and Programming Reference Guide*.



Attention: In IBM Storage Scale, the default block size of 4 MiB with an 8 KiB subblock size provides good sequential performance, makes efficient use of disk space, and provides good or similar performance for small files compared to the previous default block size of 256 KiB with 32 subblocks per block. It works well for the widest variety of workloads.

Test actual performance with different block sizes

The ideal file system block size can be determined by running performance tests with different file system block sizes using actual workloads or representative benchmarks that match the file sizes that you expect to use in production.

Factors that can affect performance

For more performance information, see the IBM Storage Scale white papers in the [Techdocs Library](http://www.ibm.com/support/techdocs/atsmastr.nsf/Web/WhitePapers) (www.ibm.com/support/techdocs/atsmastr.nsf/Web/WhitePapers).

RAID stripe size

The RAID stripe size is the size of the sequential block of data that a disk array writes to or reads from each storage volume (the block device corresponding to an NSD). For better performance, it is a good idea to set the file system block size to the same value as either the RAID stripe size or a multiple of the RAID stripe size. If the block size is not equal to or a multiple of the RAID stripe size, then the file system performance can be severely degraded, especially for write requests, because of the increase in read-modify-write operations that occur in the underlying hardware RAID controllers.

Note: The block size for IBM Storage Scale RAID that is implemented with vdisk is specifically designed for optimal behavior. For IBM Storage Scale RAID, the block size must be equal to the vdisk track size. For more information, see the online documentation available for IBM Storage Scale RAID Documentation.

File system size

For file systems larger than 100 TiB, it is a good idea to set the block size to at least 256 KiB. The default block size is 4 MiB in IBM Storage Scale. Generally larger block sizes provide better performance.

Large block size and page pool

For block sizes larger than the default size of 4 MiB, it is a good idea to increase the page pool size in proportion to the block size. The reason is that the efficiency of internal optimizations that rely on caching file data in the GPFS page pool depends more on the number of blocks that are cached than on the amount data that is cached. A larger block size results in fewer cached blocks.

Variation in file size

For a file system that contain files of many different sizes, the file system delivers better overall performance from selecting a larger block size, 4 MiB or greater, rather than a smaller one. It is true that with a larger block size some space is wasted when a small file is written into a large subblock, because the unused space in the subblock cannot be written to with data from another file unless the block is freed.

However, the amount of waste in the general case is likely to be insignificant overall, because the smaller files occupy a smaller percentage of the storage space in the file system compared to the space occupied by the larger files (files on the order of GiBs).

Application I/O patterns

The effect of block size on file system performance greatly depends on the application I/O pattern:

- A larger block size is often beneficial for large sequential read and write workloads.
- A smaller block size can offer better performance for applications that do small random writes to sparse files or small random writes to large files that are subject to frequent snapshots.

Metadata performance

The choice of block size affects the performance of certain metadata operations, in particular, block allocation performance. The IBM Storage Scale block allocation map is stored in blocks, similar to regular files. When the block size is small:

- More blocks are required to store the same amount of data, which results in more work to allocate those blocks
- One block of allocation map data contains less information

Metadata-only system pool

The `--metadata-block-size` option on the `mmcrfs` command allows a different block size to be specified for the system storage pool, provided its usage is set to `metadataOnly`. Valid values are the same as the ones that are listed for the `-B` parameter.

Note: Setting the metadata block size to a different value than the data block size can have the effect of changing the data subblock size and the number of subblocks per data block. For more information see [Scenario 3](#) earlier in this help topic.

atime values

`atime` is a standard file attribute that represents the time when the file was last accessed.

The file attribute **atime** is updated locally in memory, but the value is not visible to other nodes until after the file is closed. To get an accurate value of **atime**, an application must call subroutine **gpfs_stat** or **gpfs_fstat**.

You can control how **atime** is updated with the **-S** option of the `mmcrfs` and `mmchfs` commands:

-S no

The **atime** attribute is updated whenever the file is read. This setting is the default if the minimum release level (**minReleaseLevel**) of the cluster is earlier than 5.0.0 when the file system is created.

-S yes

The **atime** attribute is not updated. The subroutines **gpfs_stat** and **gpfs_fstat** return the time that the file system was last mounted with **-S no**.

Note: This setting can cause a problem if you have a file management policy that depends on the **ACCESS_TIME** attribute. For more information, see the topic *Exceptions to Open Group technical standards* in the *IBM Storage Scale: Administration Guide*.

-S relatime

The **atime** attribute is updated whenever the file is read, but only if one of the following conditions is true:

- The current file access time (**atime**) is earlier than the file modification time (**mtime**).
- The current file access time (**atime**) is greater than the **atimeDeferredSeconds** attribute. For more information, see the topic *mmchconfig command* in the *IBM Storage Scale: Command and Programming Reference Guide*.

This setting is the default if the minimum release level (**minReleaseLevel**) of the cluster is 5.0.0 or later when the file system is created.

You can temporarily override the value of the **-S** setting by specifying a mount option specific to IBM Storage Scale when you mount the file system. The mount option persists until the file system is unmounted. The following table shows how the mount options correspond to the **-S** settings:

Table 30. Correspondence between mount options and the -S option in <code>mmcrfs</code> and <code>mmchfs</code>		
-S option of <code>mmcrfs</code> and <code>mmchfs</code> commands.	Equivalent mount option; persists until file system is unmounted.	Effect – see topic text for details.
no	norelatime	Allow atime to be updated.
yes	noatime	Do not allow atime to be updated.
relatime	relatime	Allow atime to be updated if a condition is met.

For more information, see the topic *Mount options specific to IBM Storage Scale* in the *IBM Storage Scale: Administration Guide*.

mtime values

mtime is a standard file attribute that represents the time when the file was last modified.

The -E parameter controls when the mtime is updated. The default is -E yes, which results in standard interfaces including the stat() and fstat() calls reporting exact mtime values. Specifying -E no results in the stat() and fstat() calls reporting the mtime value available at the completion of the last sync period. This may result in the calls not always reporting the exact mtime. Setting -E no can affect backup operations, AFM, and AFM DR functions that rely on the last modified time or the operation of policies using the MODIFICATION_TIME file attribute.

For more information, see *Exceptions to the Open Group technical standards* in *IBM Storage Scale: Administration Guide*.

Block allocation map

GPFS has two different methods of allocating space in a file system. The -j parameter specifies the block allocation map type to use when creating a file system. The block allocation map type cannot be changed once the file system is created.

When allocating blocks for a given file, GPFS first uses a round-robin algorithm to spread the data across all of the disks in the file system. After a disk is selected, the location of the data block on the disk is determined by the block allocation map type.

The two types of allocation methods are **cluster** and **scatter**:

cluster

GPFS attempts to allocate blocks in clusters. Blocks that belong to a given file are kept next to each other within each cluster.

This allocation method provides better disk performance for some disk subsystems in relatively small installations. The benefits of clustered block allocation diminish when the number of nodes in the cluster or the number of disks in a file system increases, or when the file system free space becomes fragmented. The cluster allocation method is the default for GPFS clusters with eight or fewer nodes and for file systems with eight or fewer disks.

scatter

GPFS chooses the location of the blocks randomly.

This allocation method provides more consistent file system performance by averaging out performance variations due to block location (for many disk subsystems, the location of the data relative to the disk edge has a substantial effect on performance). This allocation method is appropriate in most cases and is the default for GPFS clusters with more than eight nodes or file systems with more than eight disks.

This parameter for a given file system is specified at file system creation by using the -j option on the mmcrfs command, or allowing it to default. This value *cannot* be changed after the file system has been created.

File system authorization

The type of authorization for the file system is specified on the -k option on the mmcrfs command or changed at a later time by using the -k option on the mmchfs command.

posix

Traditional GPFS access control lists (ACLs) only (NFSv4 and Windows ACLs are not allowed).

nfs4

Support for NFSv4 and Windows ACLs only. Users are not allowed to assign traditional ACLs to any file system objects.

all

Allows for the coexistence of POSIX, NFSv4, and Windows ACLs within a file system. This is the default.

If the file system will be used for the CES stack, `nfs4` is required.

Type of replication

You can control the type of replication that IBM Storage Scale uses.

Run the `mmchfs` command with the `-K` option set to the preferred type of replication. The `-K` option can have the following values:

always

Indicates that strict replication is enforced.

whenpossible

Strict replication is enforced if the disk configuration allows it. If the number of failure groups is insufficient, strict replication is not enforced. This is the default value.

no

Strict replication is not enforced. GPFS tries to create the needed number of replicas, but returns an `errno` of `EOK` if it can allocate at least one replica.

For more information on changing the attributes for replication, see the `mmchfs` command.

Internal log file

You can specify the internal log file size. Refer to [“GPFS recovery logs” on page 15](#) for additional information.

File system replication parameters

The metadata (inodes, directories, and indirect blocks) and data replication parameters are set at the file system level and apply to all files. They are initially set for the file system when issuing the `mmcrfs` command.

They can be changed for an existing file system using the `mmchfs` command. When the replication parameters are changed, files created after the change are affected. To apply the new replication values to existing files in a file system, issue the `mmrestripefs` command.

Metadata and data replication are specified independently. Each has a default replication factor of 1 (no replication) and a maximum replication factor with a default of 2. Although replication of metadata is less costly in terms of disk space than replication of file data, excessive replication of metadata also affects GPFS efficiency because all metadata replicas must be written. In general, more replication uses more space.

Default metadata replicas

The default number of copies of metadata for all files in the file system may be specified at file system creation by using the `-m` option on the `mmcrfs` command or changed at a later time by using the `-m` option on the `mmchfs` command.

This value must be equal to or less than *MaxMetadataReplicas*, and cannot exceed the number of failure groups with disks that can store metadata. The allowable values are 1, 2, or 3, with a default of 1 (no replication).

Related concepts

Maximum metadata replicas

The maximum number of copies of metadata for all files in the file system can be specified at file system creation by using the `-M` option on the `mmcrfs` command.

Default data replicas

The default replication factor for data blocks may be specified at file system creation by using the `-r` option on the `mmcrfs` command or changed at a later time by using the `-r` option on the `mmchfs` command.

Maximum data replicas

The maximum number of copies of data blocks for a file can be specified at file system creation by using the -R option on the `mmcrfs` command. The default is 2.

Maximum metadata replicas

The maximum number of copies of metadata for all files in the file system can be specified at file system creation by using the -M option on the `mmcrfs` command.

The default is 2. The allowable values are 1, 2, or 3, but it cannot be less than the value of *DefaultMetadataReplicas*. This value cannot be changed after the file system is created.

Related concepts

Default metadata replicas

The default number of copies of metadata for all files in the file system may be specified at file system creation by using the -m option on the `mmcrfs` command or changed at a later time by using the -m option on the `mmchfs` command.

Default data replicas

The default replication factor for data blocks may be specified at file system creation by using the -r option on the `mmcrfs` command or changed at a later time by using the -r option on the `mmchfs` command.

Maximum data replicas

The maximum number of copies of data blocks for a file can be specified at file system creation by using the -R option on the `mmcrfs` command. The default is 2.

Default data replicas

The default replication factor for data blocks may be specified at file system creation by using the -r option on the `mmcrfs` command or changed at a later time by using the -r option on the `mmchfs` command.

This value must be equal to or less than *MaxDataReplicas*, and the value cannot exceed the number of failure groups with disks that can store data. The allowable values are 1, 2, and 3, with a default of 1 (no replication).

If you want to change the data replication factor for the entire file system, the data disk in each storage pool must have a number of failure groups equal to or greater than the replication factor. For example, you will get a failure with error messages if you try to change the replication factor for a file system to 2 but the storage pool has only one failure group.

Related concepts

Default metadata replicas

The default number of copies of metadata for all files in the file system may be specified at file system creation by using the -m option on the `mmcrfs` command or changed at a later time by using the -m option on the `mmchfs` command.

Maximum metadata replicas

The maximum number of copies of metadata for all files in the file system can be specified at file system creation by using the -M option on the `mmcrfs` command.

Maximum data replicas

The maximum number of copies of data blocks for a file can be specified at file system creation by using the -R option on the `mmcrfs` command. The default is 2.

Maximum data replicas

The maximum number of copies of data blocks for a file can be specified at file system creation by using the -R option on the `mmcrfs` command. The default is 2.

The allowable values are 1, 2, and 3, but cannot be less than the value of *DefaultDataReplicas*. This value cannot be changed after the file system is created.

Related concepts

Default metadata replicas

The default number of copies of metadata for all files in the file system may be specified at file system creation by using the `-m` option on the `mmcrfs` command or changed at a later time by using the `-m` option on the `mmchfs` command.

Maximum metadata replicas

The maximum number of copies of metadata for all files in the file system can be specified at file system creation by using the `-M` option on the `mmcrfs` command.

Default data replicas

The default replication factor for data blocks may be specified at file system creation by using the `-r` option on the `mmcrfs` command or changed at a later time by using the `-r` option on the `mmchfs` command.

Number of nodes mounting the file system

The estimated number of nodes that will mount the file system may be specified at file system creation by using the `-n` option on the `mmcrfs` command or allowed to default to 32.

When creating a GPFS file system, over-estimate the number of nodes that will mount the file system. This input is used in the creation of GPFS data structures that are essential for achieving the maximum degree of parallelism in file system operations (see “GPFS architecture” on page 9). Although a larger estimate consumes a bit more memory, insufficient allocation of these data structures can limit the ability to process certain parallel requests efficiently, such as the allotment of disk space to a file. If you cannot predict the number of nodes, allow the default value to be applied. Specify a larger number if you expect to add nodes, but avoid wildly overestimating as this can affect buffer operations.

You can change the number of nodes using the `-n` option on the `mmchfs` command. Changing this value affects storage pools created after the value was set; so, for example, if you need to increase this value on a storage pool, you could change the value, create a new storage pool, and migrate the data from one pool to the other.

Windows drive letter

In a Windows environment, you must associate a drive letter with a file system before it can be mounted. The drive letter can be specified and changed with the `-t` option of the `mmcrfs` and `mmchfs` commands. GPFS does not assign a default drive letter when one is not specified.

The number of available drive letters restricts the number of file systems that can be mounted on Windows.

Note: Certain applications give special meaning to drive letters A :, B :, and C :, which could cause problems if they are assigned to a GPFS file system.

Mountpoint directory

Every GPFS file system has a default mount point associated with it. This mount point can be specified and changed with the `-T` option of the `mmcrfs` and `mmchfs` commands.

If you do not specify a mount point when you create the file system, GPFS will set the default mount point to `/gpfs/DeviceName`.

Assign mount command options

Options may be passed to the file system mount command using the `-o` option on the `mmchfs` command.

In particular, you can choose the option to perform quota activation automatically when a file system is mounted.

Enabling quotas

The IBM Storage Scale quota system can help you control file system usage.

Quotas can be defined for individual users, groups of users, or filesets. Quotas can be set on the total number of files and on the total amount of data space that is used.

Quota accounting depends on a consistent mapping between user names and their numeric identifiers. This means that a single user who is accessing a quota-enabled file system from different nodes must be mapped to the same numeric user identifier from each node. Within a local cluster, the mapping is done by ensuring that `/etc/passwd` and `/etc/group` are identical across the cluster. When accessing quota-enabled file systems from other clusters, ensure that individual users who are accessing the quota-enabled file system have equivalent entries in `/etc/passwd` and `/etc/group`. Or, use the user identity mapping facility that is outlined in the [UID Mapping for GPFS in a Multi-cluster Environment IBM white paper](https://www.ibm.com/docs/en/storage-scale?topic=STXKQY/uid_gpfs.pdf) (https://www.ibm.com/docs/en/storage-scale?topic=STXKQY/uid_gpfs.pdf).

If data replication is configured, be sure to include the size of the replicated data in your calculation of the proper size for a quota limit. For more information, see the topic *Listing quotas* in the *IBM Storage Scale: Administration Guide*.

To have IBM Storage Scale automatically enable quotas when a file system is mounted, choose one of the following options:

- When you create the file system, issue the **mmcrfs** command with the **-Q** option included.
- After the file system is created, issue the **mmchfs** command with the **-Q** option included.

After the file system is mounted, you can set quota values by issuing the **mmedquota** command and activate quotas by issuing the **mmquotaon** command. By default, quota values are not automatically activated when they are set.

Quota levels are defined at three limits that you can set with the **mmedquota** and **mmdefedquota** commands:

Soft limit

Defines levels of disk space and files below which the user, group, or fileset can safely operate.

Specified in units of KiB (k or K), MiB (m or M), or GiB (g or G). If no suffix is provided, the number is assumed to be in bytes.

Hard limit

Defines the maximum amount of disk space and number of files that the user, group, or fileset can accumulate.

Specified in units of KiB (k or K), MiB (m or M), or GiB (g or G). If no suffix is provided, the number is assumed to be in bytes.

Grace period

Allows the user, group, or fileset to exceed the soft limit for a specified period of time. The default period is one week. If usage is not reduced to a level below the soft limit during that time, the quota system interprets the soft limit as the hard limit and no further allocation is allowed. The user, group, or fileset can reset this condition by reducing usage enough to fall below the soft limit; or the administrator can increase the quota levels with the **mmedquota** or **mmdefedquota**.

The use of SMB and NFS protocols affects the operation of quotas. For more information, see *Implications of quotas for different protocols* in the *IBM Storage Scale: Administration Guide*.

Default quotas

Applying default quotas provides all new users, groups of users, or filesets with established minimum quota limits. If default quota values are not enabled, new users, new groups, or new filesets have a quota value of zero, which establishes no limit to the amount of space that can be used.

Default quota limits can be set or changed only if the **-Q yes** option is in effect for the file system. To set default quotas at the fileset level, the **--perfileset-quota** option must also be in effect. The **-Q yes** and **--perfileset-quota** options are specified when creating a file system with the **mmcrfs** command or changing file system attributes with the **mmchfs** command. Use the **mmllsfs** command to display the current settings of these quota options. Default quotas may then be enabled by issuing the **mmdefquotaon** command. Default values are established by issuing the **mmdefedquota** command.

Enabling DMAPI

Whether or not the file system can be monitored and managed by the GPFS Data Management API (DMAPI) may be specified at file system creation by using the `-z` option on the `mmcrfs` command or changed at a later time by using the `-z` option on the `mmchfs` command.

About this task

The default is *not* to enable DMAPI for the file system.

For more information about DMAPI for GPFS, see *IBM Storage Scale: Command and Programming Reference Guide*.

Verifying disk usage

The `-v` option controls whether the `mmcrfs` command checks whether the specified disks can safely be added to the file system.

The default (`-v yes`) is to perform the check and fail the command if any of the disks appear to belong to some other GPFS file system. You should override the default behavior and specify `-v no` only if the `mmcrfs` command rejects the disks and you are certain that *all* of the disks indeed do not belong to an active GPFS file system. An example for an appropriate use of `-v no` is the case where an `mmcrfs` command is interrupted for some reason and you are reissuing the command. Another example would be if you are reusing disks from an old GPFS file system that was not formally destroyed with the `mmdel` command.

Important: Using `mmcrfs -v no` on a disk that already belongs to a file system will corrupt that file system.

Changing the file system format to the latest level

You can change the file system format to the latest format supported by the currently-installed level of GPFS by issuing the `mmchfs` command with the `-V full` option or the `-V compat` option.

The `full` option enables all new functionality that requires different on-disk data structures. This may cause the file system to become permanently incompatible with earlier releases of GPFS. The `compat` option enables only changes that are backward compatible with the previous GPFS release. If all GPFS nodes that are accessing a file system (both local and remote) are running the latest level, then it is safe to use `full` option. Certain features may require you to run the `mmigratefs` command to enable them.

For more information, see *File system format changes between versions of GPFS* in *IBM Storage Scale: Administration Guide*.

Enabling file system features

By default, new file systems are created with all currently available features enabled.

Since this may prevent clusters that are running earlier GPFS releases from accessing the file system, you can enable only the file system features that are compatible with the specified release by issuing the `mmcrfs` command with the `--version Version` option. For more information, see *mmcrfs command* in *IBM Storage Scale: Command and Programming Reference Guide*.

Specifying whether the df command will report numbers based on quotas for the fileset

You can specify (when quotas are enforced for a fileset) whether the `df` command will report numbers based on the quotas for the fileset and not for the total file system. If `ignoreReplicationOnStatfs` is enabled, `df` command on fileset output ignores data replication factor.

To do so, use the `--filesetdf | --nofilesetdf` option on either the `mmchfs` command or the `mmcrfs` command. This option affects the `df` command behavior only on Linux nodes.

If quota is disabled and `filesetdf` is enabled in IBM Storage Scale 5.1.1 or later with file system version 5.1.1 or later, then the **df** command reports inode space capacity and inode usage at the independent fileset-level. However, the **df** command reports the block space at the file system-level because the block space is shared with the whole file system.

Note: If quota is enabled, then no behavior change for **df** command with **filesetdf** enabled, regardless of the cluster and file system versions.

For more information, see *mmchfs command* and *mmcrfs command* in *IBM Storage Scale: Command and Programming Reference Guide*.

Specifying the maximum number of files that can be created

The maximum number of files that can be created can be specified by using the `--inode-limit` option on the `mmcrfs` command and the `mmchfs` command.

Allowable values, which range from the current number of created inodes (determined by issuing the `mmdf` command with the `-F` option) through the maximum number of files that are supported, are constrained by the formula:

```
maximum number of files = (total file system space) / (inode size + subblock size)
```

You can determine the inode size (`-i`) and subblock size (`-f`) of a file system by running the `mmfsfs` command. The maximum number of files in a file system may be specified at file system creation by using the `--inode-limit` option on the `mmcrfs` command, or it may be increased at a later time by using `--inode-limit` on the `mmchfs` command. This value defaults to the size of the file system at creation divided by 1 MB and cannot exceed the architectural limit. When a file system is created, 4084 inodes are used by default; these inodes are used by GPFS for internal system files.

For more information, see *mmcrfs command* and *mmchfs command* in *IBM Storage Scale: Command and Programming Reference Guide*.

The `--inode-limit` option applies only to the root fileset. Preallocated inodes created using the **mmcrfs** command are allocated only to the root fileset and these inodes cannot be deleted or moved to another independent fileset. When there are multiple inode spaces, use the `--inode-limit` option of the `mmchfileset` command to alter the inode limits of independent filesets. The `mmchfileset` command can also be used to modify the root inode space. The `--inode-limit` option of the `mmfsfs` command shows the sum of all inode spaces.

Inodes are allocated when they are used. When a file is deleted, the inode is reused, but inodes are never deallocated. When setting the maximum number of inodes in a file system, there is the option to preallocate inodes. However, in most cases there is no need to preallocate inodes because, by default, inodes are allocated in sets as needed. If you do decide to preallocate inodes, be careful not to preallocate more inodes than will be used; otherwise, the allocated inodes will unnecessarily consume metadata space that cannot be reclaimed, and it might result in excessive memory use associated with the objects allocated with the inode segments.

These options limit the maximum number of files that may actively exist within a file system. However, the maximum number of files in the file system may be restricted further by GPFS so the control structures associated with each file do not consume all of the file system space.

Further considerations when managing inodes:

1. For file systems that are supporting parallel file creates, as the total number of free inodes drops below 5% of the total number of inodes, there is the potential for slowdown in file system access. Take this into consideration when creating or changing your file system. Use the `mmdf` command to display the number of free inodes.
2. Excessively increasing the value for the maximum number of files may cause the allocation of too much disk space for control structures.

Controlling the order in which file systems are mounted

You can control the order in which the individual file systems are mounted at daemon startup or when using the `mmmount` command with one of the all keywords specified for the file system.

For more information, see *mmlsfs command* and *mmmount command* in *IBM Storage Scale: Command and Programming Reference Guide*.

To do so, use the `--mount-priority` *Priority* option on the `mmcrfs`, the `mmchfs`, or the `mmremotefs` command.

The shared root file system for protocols must be mounted before other file systems that will be used to export protocol data.

A sample file system creation

To create a file system called `gpfs2` with the following properties:

- The disks for the file system that are listed in the file `/tmp/gpfs2dsk`
- Automatically mount the file system when the GPFS daemon starts (`-A yes`)
- Use a block size of 256 KB (`-B 256K`)
- Expect to mount it on 32 nodes (`-n 32`)
- Set both the default metadata replication and the maximum metadata replication to two (`-m 2 -M 2`)
- Set the default data replication to one and the maximum data replication to two (`-r 1 -R 2`)
- Use a default mount point of `/gpfs2` (`-T /gpfs2`)

Enter:

```
mmcrfs /dev/gpfs2 -F /tmp/gpfs2dsk -A yes -B 256K -n 32 -m 2 -M 2 -r 1 -R 2 -T /gpfs2
```

The system displays information similar to:

```
The following disks of gpfs2 will be formatted on node k194p03.tes.nnn.com:
  hd25n09: size 17796014 KB
  hd24n09: size 17796014 KB
  hd23n09: size 17796014 KB
Formatting file system ...
Disks up to size 59 GB can be added to storage pool system.
Creating Inode File
  56 % complete on Mon Mar  3 15:10:08 2014
 100 % complete on Mon Mar  3 15:10:11 2014
Creating Allocation Maps
Clearing Inode Allocation Map
Clearing Block Allocation Map
  44 % complete on Mon Mar  3 15:11:32 2014
  90 % complete on Mon Mar  3 15:11:37 2014
 100 % complete on Mon Mar  3 15:11:38 2014
Completed creation of file system /dev/gpfs2.
mmcrfs: Propagating the cluster configuration data to all
        affected nodes. This is an asynchronous process.
```

To confirm the file system configuration, issue the command:

```
mmlsfs gpfs2
```

The system displays information similar to:

flag	value	description
-f	262144	Minimum fragment (subblock) size in bytes
-i	512	Inode size in bytes
-I	32768	Indirect block size in bytes
-m	2	Default number of metadata replicas
-M	2	Maximum number of metadata replicas
-r	1	Default number of data replicas
-R	2	Maximum number of data replicas
-j	scatter	Block allocation type
-D	nfs4	File locking semantics in effect

-k	all	ACL semantics in effect
-n	32	Estimated number of nodes that will mount file
system		
-B	262144	Block size
-Q	none	Quotas accounting enabled
	none	Quotas enforced
	none	Default quotas enabled
--perfilesset-quota	yes	Per-fileset quota enforcement
--filesetdf	yes	Fileset df enabled?
-V	15.01 (4.2.0.0)	File system version
--create-time	Wed Jan 18 17:22:25 2017	File system creation time
-z	no	Is DMAPI enabled?
-L	262144	Logfile size
-E	yes	Exact mtime mount option
-S	yes	Suppress atime mount option
-K	whenpossible	Strict replica allocation option
--fastea	yes	Fast external attributes enabled?
--encryption	no	Encryption enabled?
--inode-limit	2015232	Maximum number of inodes
--log-replicas	0	Number of log replicas (max 2)
--is4KAligned	yes	is4KAligned?
--rapid-repair	yes	rapidRepair enabled?
-P	system	Disk storage pools in file system
-d	gpfs1001nsd;gpfs1002nsd	Disks in file system
-A	yes	Automatic mount option
-o	none	Additional mount options
-T	/gpfs2	Default mount point
--mount-priority	0	Mount priority

For more information, see *mmcrfs* command and *mmlsfs* command in *IBM Storage Scale: Command and Programming Reference Guide*.

Backup considerations for using IBM Storage Protect

If you are planning to use IBM Storage Protect to back up IBM Storage Scale file systems, some considerations apply.

Considerations for provisioning IBM Storage Protect servers to handle backup of IBM Storage Scale file systems

When backing up IBM Storage Scale file systems, the IBM Storage Protect server must be configured with adequate resources to handle a larger load of client data, sessions, and mount points than typically encountered with single-user file systems such as workstations.

For cluster file system backup, a system architect must plan for multiple nodes (storage hosts) sending data in multiple streams, at high data rates to the server. Therefore, the IBM Storage Protect server must be equipped with the following:

- High bandwidth network connections to the IBM Storage Scale cluster
- The sufficient disk space organized into an IBM Storage Protect storage pool to be allocated to backup data from the cluster file system
- Sufficient number of mount points (or infinite for a disk class pool) accorded to the proxy client used for the cluster file system
- Sufficient number of IBM Storage Protect sessions permitted to the proxy client used for the cluster file system
- A sufficient number of physical tape drives to offload the backup data once it arrives into the disk pool

Related concepts

[IBM Storage Protect data storage model](#)

When using IBM Storage Protect for backing up IBM Storage Scale file systems, the IBM Storage Protect data storage architecture and its implications need to be considered.

[How to identify backup and migration candidates](#)

When using IBM Storage Protect for backing up IBM Storage Scale file systems, there are several choices for the method used to identify backup and migration candidates.

[Comparison of snapshot based backups and backups from live system](#)

Backing up large file systems can take many hours or even days. When using the IBM Storage Scale command **mmbackup**, time is needed for the following steps.

File system and fileset backups with IBM Storage Protect

mmbackup supports backup from independent filesets in addition to backup of the whole file system. Independent filesets have similar capabilities as a file system in terms of quota management, dependent filesets, snapshots, and having their own inode space.

Considerations for using fileset backup with IBM Storage Protect

IBM Storage Protect stores backup data that is indexed by file or directory object path names in its database. However, file system objects are identified for IBM Storage Protect by their object ID extended attributes. In both cases, IBM Storage Protect is not aware of the fileset entity. Therefore, the fileset configuration of an IBM Storage Scale cluster is neither backed up nor it can be restored by IBM Storage Protect alone, unless separate action is taken.

Considerations for backing up file systems that are managed with IBM Storage Protect for Space Management

When IBM Storage Protect for Space Management is used to migrate data to secondary storage pools (for example, tape volumes), this migrated data might need to be recalled to disk if the file becomes a backup candidate due to certain changes a user may apply to it.

IBM Storage Protect data storage model

When using IBM Storage Protect for backing up IBM Storage Scale file systems, the IBM Storage Protect data storage architecture and its implications need to be considered.

In IBM Storage Scale, the following kinds of data are stored on disk:

- The file data that is the content stored in a file.
- The file metadata that includes all attributes that are related to the file. For example:
 - Create, access, and modify times
 - Size of the file, size that is occupied in the file system, and number of blocks used
 - Inode information, owner user ID, owning group ID, and mode bits
 - POSIX rights or access control lists (ACLs)
 - Flags to indicate whether the file is immutable or mutable, read only, or append only
 - Extended attributes (EAs)

In IBM Storage Protect, the same file data and metadata is stored but the method of storing this data differs. The file content is stored in an IBM Storage Protect storage pool such as on disk or on tape while some of the metadata is stored in the IBM Storage Protect database. The primary reason for storing metadata in the IBM Storage Protect database is to provide fast access to information useful for backup requests.

However, not all metadata is stored in the IBM Storage Protect database. Access control lists (ACLs) and extended attributes (EAs) are stored with the file content in the storage pool (media depends on storage pool type). This has the following implications:

- When the ACL or EA of a file changes then the next backup job backs up the whole file again. This occurs because the file content, ACL, and EA are stored together in the IBM Storage Protect data pool, for example on tape and they need to be updated as one entity.

Note: ACLs are inherited in the IBM Storage Scale file system. Therefore, an ACL on a top-level directory object can be inherited to all the descendant objects. ACL changes to a top-level directory object are therefore propagated down through the object tree hierarchy, rippling the change through all objects that inherited the original ACL. The number of files to be backed up increases even though nothing else in these files has changed. A surprisingly large backup workload can be induced by a seemingly small change to an ACL to a top level directory object.

When the IBM Storage Protect for Space Management capability is enabled, an ACL change such as this occurs to objects that are currently migrated to offline storage as well. These files will then need to be

recalled during the next backup cycle to enable the updated ACL to be stored with the file data once again in their IBM Storage Protect backup storage pool.

- Renaming of a file also leads to a backup of the whole file because the IBM Storage Protect database is indexed by file object path name.

You can use the following approaches to mitigate the size of a backup workload when widely inherited ACLs are likely to be changed frequently.

- Avoid renaming directories that are close to the file system root.
- Avoid ACL and EA changes in migrated files as much as possible.
- Consider using the `skipacl` or `skipaclupdatecheck` options of the IBM Storage Protect client.

Important: Be certain to note the implications of using these options by referring to *Clients options reference* in the *Backup-archive Client options and commands* section of IBM Storage Protect documentation.

Note: Using the `skipacl` option also omits EAs from the backup data store in the IBM Storage Protect backup pool. Using this option can be considered when static ACL structures are used that can be reestablished through another tool or operation external to the IBM Storage Protect restore operation. If you are using this approach, ensure that the ACL is restored either manually or automatically, by inheritance, to avoid an unauthorized user getting access to a file or a directory after it is restored.

Related concepts

[Considerations for provisioning IBM Storage Protect servers to handle backup of IBM Storage Scale file systems](#)

When backing up IBM Storage Scale file systems, the IBM Storage Protect server must be configured with adequate resources to handle a larger load of client data, sessions, and mount points than typically encountered with single-user file systems such as workstations.

[How to identify backup and migration candidates](#)

When using IBM Storage Protect for backing up IBM Storage Scale file systems, there are several choices for the method used to identify backup and migration candidates.

[Comparison of snapshot based backups and backups from live system](#)

Backing up large file systems can take many hours or even days. When using the IBM Storage Scale command **mmbackup**, time is needed for the following steps.

[File system and fileset backups with IBM Storage Protect](#)

mmbackup supports backup from independent filesets in addition to backup of the whole file system. Independent filesets have similar capabilities as a file system in terms of quota management, dependent filesets, snapshots, and having their own inode space.

[Considerations for using fileset backup with IBM Storage Protect](#)

IBM Storage Protect stores backup data that is indexed by file or directory object path names in its database. However, file system objects are identified for IBM Storage Protect by their object ID extended attributes. In both cases, IBM Storage Protect is not aware of the fileset entity. Therefore, the fileset configuration of an IBM Storage Scale cluster is neither backed up nor it can be restored by IBM Storage Protect alone, unless separate action is taken.

[Considerations for backing up file systems that are managed with IBM Storage Protect for Space Management](#)

When IBM Storage Protect for Space Management is used to migrate data to secondary storage pools (for example, tape volumes), this migrated data might need to be recalled to disk if the file becomes a backup candidate due to certain changes a user may apply to it.

How to identify backup and migration candidates

When using IBM Storage Protect for backing up IBM Storage Scale file systems, there are several choices for the method used to identify backup and migration candidates.

Limitations when using IBM Storage Protect Backup-Archive client to identify backup candidates

Using IBM Storage Protect Backup-Archive client to traverse IBM Storage Scale file systems to identify backup candidates does not scale well. For this reason, using the IBM Storage Scale **mmapplypolicy** engine is preferable because it is much faster to scan the file system for identifying backup candidates than traversing the file system.

Therefore, for processing backups on larger file systems, use the IBM Storage Scale command **mmbackup** instead of using the IBM Storage Protect Backup-Archive client commands such as **dsmc expire** or **dsmc selective** or **dsmc incremental** directly. Using the **mmbackup** command also provides the following benefits:

- Backup activities are run in parallel by using multiple IBM Storage Scale cluster nodes to send backup data in parallel to the IBM Storage Protect server.
- **mmbackup** creates a local shadow of the IBM Storage Protect database in the file system and uses it along with the policy engine to identify candidate files for backup. The IBM Storage Protect server does not need to be queried for this information saving time when calculating the backup candidate list.
 - **mmbackup** and its use of the policy engine can select candidates faster than the **dsmc** progressive incremental operation that is bounded by walk of the file system using the POSIX directory and file status reading functions.
 - Using **dsmc selective** with lists generated by **mmbackup** is also faster than using **dsmc incremental** even with similar lists generated by **mmbackup**.

Note: It is recommended that scheduled backups of an IBM Storage Scale file system use **mmbackup** because **mmbackup** does not actively query the IBM Storage Protect server to calculate backup candidates. However, events such as file space deletion or file deletion executed on IBM Storage Protect server are not recognized until the user triggers a synchronization between the **mmbackup** shadow database and the IBM Storage Protect database.

The following table contains a detailed comparison of **mmbackup** and IBM Storage Protect Backup-Archive client backup commands:

<i>Table 31. Comparison of mmbackup and IBM Storage Protect Backup-Archive client backup commands</i>		
	IBM Storage Scale policy-driven backup (mmbackup)	IBM Storage Protect progressive incremental backup (dsmc incremental)
Detects changes in files and sends a new copy of the file to the server.	Yes	Yes
Detects changes in metadata and updates the file metadata on the server or sends a new copy of the file to the server (for ACL/EA changes).	Yes	Yes

Table 31. Comparison of **mmbackup** and IBM Storage Protect Backup-Archive client backup commands (continued)

	IBM Storage Scale policy-driven backup (mmbackup)	IBM Storage Protect progressive incremental backup (dsmc incremental)
Detects directory move, copy, or rename functions, and sends a new copy of the file to the server.	Yes	Yes
Detects local file deletion and expires the file on the server.	Yes	Yes
Detects IBM Storage Protect file space deletion or node/policy changes, and sends a new copy of the file to the server.	No*	Yes
Detects file deletion from the IBM Storage Protect server and sends a new copy of the file to the server.	No*	Yes
Detects additions of new exclude rules and expires the file on the server.	Yes	Yes
Detects policy changes made to include rules and rebinds the file to the new storage pool.	No**	Yes
Detects copy mode and copy frequency configuration changes.	No*	Yes
Detects migration state changes (IBM Storage Protect for Space Management) and updates the server object.	Yes	Yes
Detects that a file wasn't processed successfully during a backup operation and attempts again at the next backup.	Yes	Yes
Supports IBM Storage Protect Virtual Mount Points to divide a file system into smaller segments to reduce database size and contention.	No	Yes
* The mmbackup command queries the IBM Storage Protect server only once at the time of the first backup. Changes that are performed on the IBM Storage Protect server by using the IBM Storage Protect administrative client cannot be detected by mmbackup processing. You must rebuild the mmbackup shadow database if the IBM Storage Protect server file space changes. ** IBM Storage Protect includes rules with associated management class bindings that cannot be detected by mmbackup processing. Therefore, mmbackup processing does not rebind a file if a management class changes include rules.		

If you use IBM Storage Protect Backup-Archive client backup commands on file systems that are otherwise handled by using **mmbackup**, the shadow database maintained by **mmbackup** loses its

synchronization with the IBM Storage Protect inventory. In such cases, you need to resynchronize with the IBM Storage Protect server which will inform **mmbackup** of the recent backup activities conducted with the **dsmc** command. Resynchronization might be a very time-consuming activity for large file systems with a high number of backed up items. To avoid these scenarios, use the **mmbackup** command only.

If you have used **dsmc selective** or **dsmc incremental** since starting to use **mmbackup** and need to manually trigger a synchronization between the **mmbackup** maintained shadow database and the IBM Storage Protect server:

- Use the **mmbackup --rebuild** if you need to do a synchronization only.
- Use the **mmbackup -q** if you need to do a synchronization followed by a backup of the corresponding file system.

Using the IBM Storage Protect for Space Management client to identify migration candidates

Using IBM Storage Protect for Space Management clients for traversing IBM Storage Scale file system to identify migration candidates does not scale well. The IBM Storage Protect **automigration** daemons consume space in the file system and also consume CPU resources. They do not have access to the internal structures of the file system in the way that the IBM Storage Scale **mmapplypolicy** command does, and so they cannot scale. Use the following steps instead:

1. Set the following environment variable before the installation of the IBM Storage Protect for Space Management client to prevent the **automigration** daemons from starting during the installation:

```
export HSMINSTALLMODE=SCOUTFREE
```

It is recommended to place this setting into the profile file of the root user.
2. Add the following option to your IBM Storage Protect client configuration file, `dsm.opt`, to prevent the **automigration** daemons from starting after every system reboot:

```
HSMDISABLEAUTOMIGDAEMONS YES
```
3. Add the following option to your IBM Storage Protect client configuration file, `dsm.opt`, to ensure that the object ID is added to the inode, so that the file list based reconciliation (two-way-orphan-check) can be used:

```
HSMEXTObjidattr YES
```
4. While the **automigration** daemons are disabled, changes such as removal of migrated files are not automatically propagated to the IBM Storage Protect server. For housekeeping purposes, you must run the IBM Storage Protect reconciliation either manually or in a scheduled manner. For more information, see *Reconciling by using a GPFS policy* in IBM Storage Protect for Space Management documentation.

Related concepts

[Considerations for provisioning IBM Storage Protect servers to handle backup of IBM Storage Scale file systems](#)

When backing up IBM Storage Scale file systems, the IBM Storage Protect server must be configured with adequate resources to handle a larger load of client data, sessions, and mount points than typically encountered with single-user file systems such as workstations.

[IBM Storage Protect data storage model](#)

When using IBM Storage Protect for backing up IBM Storage Scale file systems, the IBM Storage Protect data storage architecture and its implications need to be considered.

[Comparison of snapshot based backups and backups from live system](#)

Backing up large file systems can take many hours or even days. When using the IBM Storage Scale command **mmbackup**, time is needed for the following steps.

[File system and fileset backups with IBM Storage Protect](#)

mmbackup supports backup from independent filesets in addition to backup of the whole file system. Independent filesets have similar capabilities as a file system in terms of quota management, dependent filesets, snapshots, and having their own inode space.

Considerations for using fileset backup with IBM Storage Protect

IBM Storage Protect stores backup data that is indexed by file or directory object path names in its database. However, file system objects are identified for IBM Storage Protect by their object ID extended attributes. In both cases, IBM Storage Protect is not aware of the fileset entity. Therefore, the fileset configuration of an IBM Storage Scale cluster is neither backed up nor it can be restored by IBM Storage Protect alone, unless separate action is taken.

Considerations for backing up file systems that are managed with IBM Storage Protect for Space Management

When IBM Storage Protect for Space Management is used to migrate data to secondary storage pools (for example, tape volumes), this migrated data might need to be recalled to disk if the file becomes a backup candidate due to certain changes a user may apply to it.

Comparison of snapshot based backups and backups from live system

Backing up large file systems can take many hours or even days. When using the IBM Storage Scale command **mmbackup**, time is needed for the following steps.

- Scanning the system to identify the objects that need to be backed up or expired.
- Expiring all objects removed from the system.
- Backing up all new or changed objects.

If the backup is run on the live file system while it is active, objects selected for the backup job can be backed up at different points in time. This can lead to issues when temporary or transient files that were present during the scan time are removed by the time the backup command tries to send them to the IBM Storage Protect server. The attempt to back up a file that is removed fails and the need to back up this object is still recorded in the shadow database.

Instead of backing up from a live file system, an alternative is to use snapshot based backups. Using the snapshot adds the additional actions of creating or reusing a snapshot and removing it when the overall backup process completes. However, this approach provides several advantages because a snapshot is a point in time view of the file system that is read only and it can be used for backing up the file system for as long as is necessary to complete the backup job. The advantages of following this approach are as follows:

- Transient or temporary files are backed up, provided they existed at the time the snapshot was taken.
- Protection against failures to back up due to server-side faults such as the IBM Storage Protect server running out of space. For example, if the database or storage pool becomes full or if the IBM Storage Protect server crashes, etc. In this case, a retry of the backup is possible for the point in time when the snapshot has been taken with no loss of function or backup.
- Retention of a backup within the system. Snapshots can be kept for a period of time providing an online backup copy of all files. This can protect against accidental deletions or modifications, and can be used to retrieve an earlier version of a file, etc.
- A means to fulfill a data protection policy even if the backup activity to IBM Storage Protect exceeds the nominal time window allotted. The snapshot can be kept for several days until backups are complete, and multiple snapshots can be kept until backup completes for each of them.

IBM Storage Scale provides the capability to create snapshots for a complete file system, known as global snapshots, and for independent filesets, known as fileset snapshots.

While snapshot based backups provide several advantages, the following considerations apply when using the snapshot capability:

- Snapshots might consume space; usually a snapshot used for backup is removed shortly after the backup operation finishes. Long lived snapshots retain their copy of the data blocks, taken from the file system's pool of free blocks. As they age, the snapshots consume more data blocks owing to the changes made in the read or write view of the file system.

- Snapshot deletion can take time depending on the number of changes that need to be handled while removing the snapshot. In general, the older a snapshot is, the more work it will require to delete it.
- Special consideration for use of IBM Storage Protect for Space Management:
 - When a migrated file stub that is already part of a snapshot is removed, a recall is initiated to keep the snapshot consistent. This is required because removal of the stub invalidates the offline references to the stored data. The recall is to fill blocks on disk and assign them to the snapshot view. Once the stub is removed from the file system and a reconcile process removes this file from the Space Management pool on the IBM Storage Protect server, there are no longer any references to the file data except the snapshot copy.

Related concepts

Considerations for provisioning IBM Storage Protect servers to handle backup of IBM Storage Scale file systems

When backing up IBM Storage Scale file systems, the IBM Storage Protect server must be configured with adequate resources to handle a larger load of client data, sessions, and mount points than typically encountered with single-user file systems such as workstations.

IBM Storage Protect data storage model

When using IBM Storage Protect for backing up IBM Storage Scale file systems, the IBM Storage Protect data storage architecture and its implications need to be considered.

How to identify backup and migration candidates

When using IBM Storage Protect for backing up IBM Storage Scale file systems, there are several choices for the method used to identify backup and migration candidates.

File system and fileset backups with IBM Storage Protect

mmbackup supports backup from independent filesets in addition to backup of the whole file system. Independent filesets have similar capabilities as a file system in terms of quota management, dependent filesets, snapshots, and having their own inode space.

Considerations for using fileset backup with IBM Storage Protect

IBM Storage Protect stores backup data that is indexed by file or directory object path names in its database. However, file system objects are identified for IBM Storage Protect by their object ID extended attributes. In both cases, IBM Storage Protect is not aware of the fileset entity. Therefore, the fileset configuration of an IBM Storage Scale cluster is neither backed up nor it can be restored by IBM Storage Protect alone, unless separate action is taken.

Considerations for backing up file systems that are managed with IBM Storage Protect for Space Management

When IBM Storage Protect for Space Management is used to migrate data to secondary storage pools (for example, tape volumes), this migrated data might need to be recalled to disk if the file becomes a backup candidate due to certain changes a user may apply to it.

File system and fileset backups with IBM Storage Protect

mmbackup supports backup from independent filesets in addition to backup of the whole file system. Independent filesets have similar capabilities as a file system in terms of quota management, dependent filesets, snapshots, and having their own inode space.

Fileset **mmbackup** provides a finer granularity for backup purposes and permits the administrator to:

- Have different backup schedules for different filesets. For example, providing the flexibility that filesets containing more important data to be backed up more often than other filesets.
- Have a fileset dedicated for temporary or transient files or files that do not need to be backed up.
- Use file system backups in conjunction with fileset backup to implement a dual data protection scheme such that file system backup goes to one IBM Storage Protect server while fileset backup goes to a different IBM Storage Protect server.

Backups of file systems and independent filesets are controlled by using the **mmbackup** option `--scope`.

In the examples used in this topic, the following environment is assumed:

```
tsm server name          => tsm1,tsm2,...
file system device name  => gpfs0 or /dev/gpfs0
file system mountpoint   => /gpfs0

fileset names/junction path
-----
depfset1,depfset2,...    /gpfs0/depfset1,/gpfs0/depfset2      for dependent filesets
indepfset1,indepfset2,... /gpfs0/indepfset1,/ibm/gpfs0/indepfset2 for independent filesets
```

File system backup

File system backup protects the whole file system. The default value for option `--scope` is `filesystem` and thus it can be omitted.

To back up the `gpfs0` file system in its entirety to the IBM Storage Protect server named `tsm1`, use the following command:

```
mmbackup gpfs0 --tsm-servers tsm1
mmbackup gpfs0 --tsm-servers tsm1 --scope filesystem
mmbackup /dev/gpfs0 --tsm-servers tsm1
mmbackup /dev/gpfs0 --tsm-servers tsm1 --scope filesystem
mmbackup /gpfs0 --tsm-servers tsm1
mmbackup /gpfs0 --tsm-servers tsm1 --scope filesystem
```

This example shows the file system backup of `gpfs0` using the short or long device name (first 4 lines) or the mount point as a directory input (last 2 lines).

Independent fileset backup

Independent fileset backup protects all files that belong to the inode space of a single independent fileset. The backup of an independent fileset might include other dependent filesets and folders, but not nested independent filesets because each independent fileset has its own inode space.

The following information describes the `mmbackup` capability to support independent fileset. The examples with IBM Storage Protect show the potential/likely use of this capability.

To back up the independent fileset `indepfset1` of the file system `gpfs0` to IBM Storage Protect server named `tsm1`, use the following command:

```
mmbackup /gpfs0/indepfset1 --tsm-servers tsm1 --scope inodespace
```

Fileset backup is an additional option to existing IBM Storage Scale backup or data protection options. Therefore, starting with a file system in which the whole file system has already been backed up on a regular basis, administrators can also start protecting data with fileset backup on some or all filesets. Administrators may choose to utilize fileset backup at any time, with no limit. It is not required to back up the whole file system in either one mode or another unless the chosen data protection approach requires it. Also, administrators can choose different IBM Storage Protect servers for some or all fileset backups if it is desired to have these on separate servers. For example, a valid configuration may have one IBM Storage Protect server that only gets whole file system backups, and a different one for fileset backups.

Note: When mixing file system and fileset backup on the same IBM Storage Protect server, consider that a target file that is handled by both backup approaches gets backed up twice. Each backup activity consumes one storage version of the file. Thus, the same file version is stored twice on the IBM Storage Protect server.

If migrating to use only fileset backup to protect the whole file systems, ensure that all independent filesets are backed up and keep the backup of all filesets current. Remember to include a backup of the root fileset as well by using a command such as:

```
mmbackup /gpfs/gpfs0 --scope inodespace --tsm-servers tsm1
```

To verify the completeness of the data protection using fileset backup only, use the following steps:

1. Identify all independent filesets available in the file system by using the following command:

```
mmlsfileset device -L
```

Identify the independent filesets by the InodeSpace column. The value identifies the corresponding inode space or independent fileset while the first occurrence refers to the corresponding fileset name and fileset path.

Note: Ensure that a backup is maintained of the fileset called `root` that is created when the file system is created and that can be seen as the first fileset of the file system.

2. For every identified independent fileset, verify that a shadow database exists as follows:
 - a. Check for the file `.mmbackupShadow.<digit>.<TSMserverName>.fileset` in the fileset junction directory.
 - b. If the file exists, determine the time of last backup by looking at the header line of this file. The backup date is stored as last value of this line.

For example:

```
head -n1 /gpfs0/indepset1/.mmbackupShadow.*.fileset
%%mshadow0%%:00_BACKUP_FILES_41:1400:/gpfs0:mmbackup:1:Mon Apr 20 13:48:45 2015
```

Where *Mon Apr 20 13:48:45 2015* in this example is the time of last backup taken for this fileset.

3. If any independent filesets are missing their corresponding `.mmbackupShadow.*` files, or if they exist but are older than the data protection limit for their backup time, then start or schedule a backup of these filesets.

Note: This action needs to be done for every file system for which the fileset backup approach is chosen.

Note: Backup of a nested independent fileset is not supported. To work around this limitation, unlink the nested fileset and link it at another location in the root fileset prior to beginning to use fileset backup pervasively.

Related concepts

[Considerations for provisioning IBM Storage Protect servers to handle backup of IBM Storage Scale file systems](#)

When backing up IBM Storage Scale file systems, the IBM Storage Protect server must be configured with adequate resources to handle a larger load of client data, sessions, and mount points than typically encountered with single-user file systems such as workstations.

[IBM Storage Protect data storage model](#)

When using IBM Storage Protect for backing up IBM Storage Scale file systems, the IBM Storage Protect data storage architecture and its implications need to be considered.

[How to identify backup and migration candidates](#)

When using IBM Storage Protect for backing up IBM Storage Scale file systems, there are several choices for the method used to identify backup and migration candidates.

[Comparison of snapshot based backups and backups from live system](#)

Backing up large file systems can take many hours or even days. When using the IBM Storage Scale command **mmbackup**, time is needed for the following steps.

[Considerations for using fileset backup with IBM Storage Protect](#)

IBM Storage Protect stores backup data that is indexed by file or directory object path names in its database. However, file system objects are identified for IBM Storage Protect by their object ID extended attributes. In both cases, IBM Storage Protect is not aware of the fileset entity. Therefore, the fileset configuration of an IBM Storage Scale cluster is neither backed up nor it can be restored by IBM Storage Protect alone, unless separate action is taken.

[Considerations for backing up file systems that are managed with IBM Storage Protect for Space Management](#)

When IBM Storage Protect for Space Management is used to migrate data to secondary storage pools (for example, tape volumes), this migrated data might need to be recalled to disk if the file becomes a backup candidate due to certain changes a user may apply to it.

Considerations for using fileset backup with IBM Storage Protect

IBM Storage Protect stores backup data that is indexed by file or directory object path names in its database. However, file system objects are identified for IBM Storage Protect by their object ID extended attributes. In both cases, IBM Storage Protect is not aware of the fileset entity. Therefore, the fileset configuration of an IBM Storage Scale cluster is neither backed up nor it can be restored by IBM Storage Protect alone, unless separate action is taken.

IBM Storage Scale does provide commands such as **mmbackupconfig** and **mmrestoreconfig** to backup and restore file system configuration data.

Even though backups of independent filesets might be stored on a plurality of IBM Storage Protect servers, IBM Storage Protect for Space Management can only utilize one server per managed file system if **mmbackup** is used on that file system and sends the backup data to that same server. IBM Storage Protect for Space Management multiserver feature is not integrated with the **mmbackup** command.

If space management is enabled for an IBM Storage Scale file system, and a plurality of IBM Storage Protect servers are used for backup, the following limitations apply to the use of IBM Storage Protect for Space Management:

- The setting of `migrequiresbkup` cannot be utilized since the server for space management might not have access to the database for the server doing backup on a particular object.
- The *inline copy* feature that allows the backup of a migrated object through copying data internally from the IBM Storage Protect for Space Management pool to the backup pool inside the IBM Storage Protect server cannot be used since migrated data might reside on a different server than the one performing the backup.

The following special considerations apply when using the fileset backup with IBM Storage Protect:

- IBM Storage Protect Backup-Archive client does not restore fileset configuration information. Therefore, a full file system backup is not sufficient for disaster recovery when the whole file system structure needs to be recovered.
- Nested independent filesets cannot be backed up when using the fileset backups. Resolve the nesting before starting to protect the data with **mmbackup** on each fileset.
- The following operations are discouraged: fileset unlink, fileset junction path change via unlink and link operations, fileset deletion. These operations will incur a heavy load on the IBM Storage Protect server at the next backup activity due to the significant alteration of the path name space that result from these operations. If a fileset must be moved in the file system, be prepared for the extra time that might be required to run the next backup.
- If only utilizing fileset backup, be sure to include every fileset, including the root fileset to ensure complete data protection.
- When both file system and fileset backups use the same IBM Storage Protect server, new and changed files may be unintentionally backed up twice. This consumes two of the file versions stored in the IBM Storage Protect server with the same file data.

Loss of data protection may occur if a fileset is unlinked or deleted for an extended period of time and the objects formerly contained in that fileset are allowed to expire from the backup sets.

- Do not move, rename, or relink a fileset, unless a new full backup of the fileset is expected to be made immediately after the restructuring occurs. IBM Storage Protect server only recognizes objects in backup only as long as they remain in their original path name from the root of the file tree (mount point).

The following special considerations apply when using fileset backup and IBM Storage Protect for Space Management:

- When using fileset backup to different IBM Storage Protect servers, the server setting **migrequiresbkup** must be set to NO to be able to migrate files.
- Automatic server selection for backup by **mmbackup** is not supported and therefore the IBM Storage Protect *multiple space management server* feature is not integrated with **mmbackup** data protection of IBM Storage Scale. Using fileset backup to different IBM Storage Protect servers in conjunction with the *multiple space management server* feature requires special handling and configuration because backup and space management rules must match so that files are handled by the same IBM Storage Protect server.

Related concepts

Considerations for provisioning IBM Storage Protect servers to handle backup of IBM Storage Scale file systems

When backing up IBM Storage Scale file systems, the IBM Storage Protect server must be configured with adequate resources to handle a larger load of client data, sessions, and mount points than typically encountered with single-user file systems such as workstations.

IBM Storage Protect data storage model

When using IBM Storage Protect for backing up IBM Storage Scale file systems, the IBM Storage Protect data storage architecture and its implications need to be considered.

How to identify backup and migration candidates

When using IBM Storage Protect for backing up IBM Storage Scale file systems, there are several choices for the method used to identify backup and migration candidates.

Comparison of snapshot based backups and backups from live system

Backing up large file systems can take many hours or even days. When using the IBM Storage Scale command **mmbackup**, time is needed for the following steps.

File system and fileset backups with IBM Storage Protect

mmbackup supports backup from independent filesets in addition to backup of the whole file system. Independent filesets have similar capabilities as a file system in terms of quota management, dependent filesets, snapshots, and having their own inode space.

Considerations for backing up file systems that are managed with IBM Storage Protect for Space Management

When IBM Storage Protect for Space Management is used to migrate data to secondary storage pools (for example, tape volumes), this migrated data might need to be recalled to disk if the file becomes a backup candidate due to certain changes a user may apply to it.

Considerations for backing up file systems that are managed with IBM Storage Protect for Space Management

When IBM Storage Protect for Space Management is used to migrate data to secondary storage pools (for example, tape volumes), this migrated data might need to be recalled to disk if the file becomes a backup candidate due to certain changes a user may apply to it.

Recall will automatically occur synchronously if the user attempts to change the file data. However, no synchronous recall occurs in the following scenarios:

- If a user changes the owner, group, mode, access control list, or extended attributes of the file.
- If the user renames the file or any of the parent directories of the file.

Any of these changes, however, does make the file a candidate for the next **mmbackup** invocation.

When **mmbackup** supplies the path to a migrated file to IBM Storage Protect for backup, synchronous recall will occur. This can lead to a "Recall Storm" in which tape library and IBM Storage Protect server resources become overburdened handling many simultaneous recall requests. Disk space is also consumed immediately by recalled data. These are undesirable outcomes and can interfere with the normal processing of the customer's workload.

Since **mmbackup** is able to discern that the data was migrated for such files, it will defer the backup, and instead append a record referring to the file's path name into a special file in the root of the file system, or fileset in the case of fileset-level backup. This list can then be used by the system administrator

to schedule recall of the deferred files data to permit the backup to occur on the next invocation of the **mmbackup** command. By deferring the recall to the system administrator, **mmbackup** prevents unexpected, excessive tape library activity which might occur if many such migrated files were nominated for backup due to user actions such as renaming a high level directory or changing owner, group, or modes on many files.

The system administrator must schedule the recall for the migrated objects omitted by **mmbackup** according to the availability of tape and disk resources in the cluster. For information about optimized tape recall, see *Optimized tape recall processing* in IBM Storage Protect for Space Management documentation.

Related concepts

[Considerations for provisioning IBM Storage Protect servers to handle backup of IBM Storage Scale file systems](#)

When backing up IBM Storage Scale file systems, the IBM Storage Protect server must be configured with adequate resources to handle a larger load of client data, sessions, and mount points than typically encountered with single-user file systems such as workstations.

[IBM Storage Protect data storage model](#)

When using IBM Storage Protect for backing up IBM Storage Scale file systems, the IBM Storage Protect data storage architecture and its implications need to be considered.

[How to identify backup and migration candidates](#)

When using IBM Storage Protect for backing up IBM Storage Scale file systems, there are several choices for the method used to identify backup and migration candidates.

[Comparison of snapshot based backups and backups from live system](#)

Backing up large file systems can take many hours or even days. When using the IBM Storage Scale command **mmbackup**, time is needed for the following steps.

[File system and fileset backups with IBM Storage Protect](#)

mmbackup supports backup from independent filesets in addition to backup of the whole file system. Independent filesets have similar capabilities as a file system in terms of quota management, dependent filesets, snapshots, and having their own inode space.

[Considerations for using fileset backup with IBM Storage Protect](#)

IBM Storage Protect stores backup data that is indexed by file or directory object path names in its database. However, file system objects are identified for IBM Storage Protect by their object ID extended attributes. In both cases, IBM Storage Protect is not aware of the fileset entity. Therefore, the fileset configuration of an IBM Storage Scale cluster is neither backed up nor it can be restored by IBM Storage Protect alone, unless separate action is taken.

Planning for quality of service for I/O operations

With the quality of service (QoS) for I/O operations capability, you can run a GPFS maintenance command without the risk of it dominating file system I/O performance and significantly delaying other tasks.

The GPFS *maintenance commands* are a group of about 20 potentially long-running, I/O-intensive commands that can generate hundreds or thousands of requests for I/O operations per second. This high demand for I/O can greatly slow down normal tasks that are competing for the same I/O resources. With QoS, you can assign an instance of a maintenance command to a QoS class that has a lesser I/O priority. Although the instance of the maintenance command now takes longer to run to completion, normal tasks have greater access to I/O resources and run more quickly. The maintenance commands are listed in the help topic for the *mmchqos* command in the *IBM Storage Scale: Command and Programming Reference Guide*.

Note the following requirements and limitations:

- QoS requires the file system to be at V4.2.0.0 or later.
- QoS works with asynchronous I/O, memory-mapped I/O, cached I/O, and buffered I/O. However, with direct I/O, QoS counts the I/O operations per second (IOPS) but does not regulate them.

- When you change allocations, mount the file system, or re-enable QoS, a brief delay due to reconfiguration occurs before QoS starts applying allocations.

To set up QoS, you allocate shares of I/O operations per second (IOPS) to two QoS classes, `maintenance` and `other`. Each storage pool has these two classes. The `maintenance` class is the class in which the GPFS maintenance commands run by default. Typically, you assign a smaller share of IOPS to this class so that the commands in it do not overwhelm the file system with I/O requests. The `other` class is the class to which all other processes belong. Typically, you assign an unlimited number of IOPS to this class so that normal processes have greater access to I/O resources and finish more quickly.

When a process issues an I/O request, QoS first finds the QoS class to which the process belongs. It then finds whether the class has any I/O operations available for consumption. If so, QoS allows the I/O operation. If not, then QoS queues the request until more I/O operations become available for the class. For a particular storage pool, all instances of maintenance commands that are running concurrently compete for the IOPS that you allocated to the `maintenance` class of that storage pool. Similarly, all processes that belong to the `other` class compete for the IOPS that you allocated to the `other` class of that storage pool. If no processes in the `maintenance` class are running, processes in the `other` class can also compete for the `maintenance` IOPS.

If you do not configure a storage pool for QoS, then processes compete for I/O operations in that storage pool in the normal way without interference.

By default, for both the `maintenance` class and the `other` class, QoS divides the IOPS allocation equally among the nodes of the cluster that have the file system mounted. For example, you might assign 1000 IOPS to the `maintenance` class of storage pool SP1. If your allocation applies to all the nodes in the cluster (the default setting) and the cluster has five nodes, then the `maintenance` class of each node gets an allocation of 200 IOPS for storage pool SP1. At run time, on each node all processes that belong to the `maintenance` class and that access storage pool SP1 compete for the 200 IOPS.

You can also divide IOPS among a list of nodes, a user-defined node class, or the nodes of a remote cluster that is served by the file system.

To get started with QoS, follow the steps that are outlined in the topic *Setting the quality of service (QoS) for I/O operations* in *IBM Storage Scale: Administration Guide*.

Use the `mmchqos` command to allocate IOPS to the QoS classes `maintenance` and `other`. You can set up allocations for one pool or several pools in one call to the command. You can designate a single node, several nodes, or all the nodes in the current cluster to participate in the QoS operations. You can also designate the nodes in a remote cluster to participate. To preserve a particular configuration, you can set up a stanza file that the `mmchqos` command reads when it runs.

You can reset or change IOPS allocations at any time without unmounting the file system. You can also disable QoS at any time without losing your IOPS allocations. When you reenables QoS, it resumes applying the allocations.

Remember the following points:

- You can allocate shares of IOPS separately for each storage pool.
- QoS divides each IOPS allocation equally among the specified nodes that mounted the file system.
- Allocations persist across unmounting and remounting the file system.
- QoS stops applying allocations when you unmount the file system and resumes when you remount it.

Use the `mm1sqos` command to display the consumption of IOPS. The command displays data for subperiods within the period that you specify. You can configure the type of information that is displayed. You can display information for specified storage pools or for all storage pools. You can display information for various time periods, from 1 second up to about 15 minutes. You can display I/O performance for each QoS class separately or summed together. You can also display the I/O performance of each node separately or summed across the participating nodes.

Planning for extended attributes

IBM Storage Scale supports extended attributes for files, directories, block devices, pipes, and named sockets in a file system.

Extended file attributes are key-value pairs that can be set programmatically by the file system, by other middleware such as the Data Management API, by the operating system, or by users. The name of an extended attribute consists of a namespace name followed by a dot followed by an attribute name, as in the following extended attribute names:

```
gpfs.Encryption
system.posix_acl_access
```

The public namespaces are *system*, *security*, *trusted*, and *user*. IBM Storage Scale has a few private namespaces, which include *gpfs* and *dmapl*.

You can access extended attributes in the public namespaces through operating system commands and functions. Linux has commands such as `setfattr` and `getfattr` and functions such as `setxattr()` and `getxattr()` for accessing extended attributes.

- In Linux, the following facts are relevant:
 - You must have root-level access to work with extended attributes in the *system*, *security*, and *trusted* namespaces.
 - To set an extended attribute for a file or file object (directory, block device, pipe, or named socket) you must have write access to the entity.
 - The **cp --preserve=xattr** Linux command copies either the POSIX or the NFSv4 ACL extended attributes when an IBM Storage Scale file is copied. Also some Linux system calls are extended to handle attributes for POSIX and NFSV4 ACLs when the system calls are applied to files in IBM Storage Scale file systems. For more information, see *Managing GPFS access control lists* in *IBM Storage Scale: Administration Guide*.
- In Windows, the following facts are relevant:
 - Log on as an administrator to work with extended attributes in the *system*, *security*, and *trusted* namespaces.
 - Windows has file permissions for reading and writing extended attributes that are separate from its file permissions for reading or writing file data.
- In AIX, the following facts are relevant:
 - IBM Storage Scale does not support AIX commands or functions that access extended attributes.

The native operating system commands and functions that are referred to in the preceding paragraph cannot access the private namespaces of IBM Storage Scale. You can access the extended attributes in the IBM Storage Scale private namespaces only through IBM Storage Scale resources, which are described in the following list. These resources can also access the extended attributes in the public namespaces. IBM Storage Scale supports these resources on all platforms, Linux, Windows, and AIX:

- The rules for automating file management include a set of functions for accessing, setting, and evaluating extended attributes. For more information, see *Extended attribute functions* in the *IBM Storage Scale: Administration Guide*.
- The **mmchattr** and **mmlsattr** commands provide parameters to create, set, get, or delete extended attributes. For more information, see *mmchattr command* and *mmlsattr command* in the *IBM Storage Scale: Command and Programming Reference Guide*.

Note: The **mmchattr** command does not set some DMAPi extended attributes because of their side effects. However, you can set these extended attributes through DMAPi. For more information, see *IBM Storage Scale Data Management API for GPFS information* in the *IBM Storage Scale: Command and Programming Reference Guide*.

- The GPFS programming interfaces include a set of subroutines for getting and setting extended attributes. For more information, see *GPFS programming interfaces* in the *IBM Storage Scale: Command and Programming Reference Guide*.

Fast extended attributes are extended attributes that are stored in the inode of the file or file object so that they can be read or written quickly. Some features of IBM Storage Scale require support for fast extended attributes, such as DMAPI, file encryption, and file heat. To work with fast extended attributes, you use the same methods that you use with regular extended attributes.

The following features in IBM Storage Scale also handle extended attributes:

- Limitations:
 - Extended attributes are not encrypted when a file is encrypted.
 - Extended attributes cannot be added, deleted, or modified in an immutable or an appendOnly file. This restriction does not apply to extended attributes in the **dmapi** or **gpfs** namespaces.
- Size:
 - To determine the maximum allowed size of extended attributes, see Q6.11 in [IBM Storage Scale FAQ in IBM Documentation](#).
 - Fast extended attributes are stored first in the inode of a file, then, if more space is needed, in a maximum of one overflow block. The overflow block is the same size as the current metadata block size. The maximum space in the overflow block for storing extended attributes is the minimum of 64 KB or the metadata block size.
- Administration commands:
 - The **mmafmconfig** command has parameters that enable or disable extended attributes or sparse file support in the AFM cache.
 - The **mmclone** command does not copy extended attributes.
 - See the administration commands in the item "Backup and restore" that follows.
- Encryption:
 - In encryption, file extension keys (FEKs) are stored in the **gpfs**.Encryption extended attribute. It is a good idea to create a file system with an inode size of 4 K or larger to accommodate FEKs that are encrypted several times. If you encounter an error message that says that the extended attributes are too large, you must either change the encryption policy so that the file key is wrapped fewer times, reduce the number of keys that are used to wrap a file key, or create a file system that has a larger inode size. For more information, see the following links:
 - *Encryption policies* in the *IBM Storage Scale: Administration Guide*
 - *Preparation for encryption* in the *IBM Storage Scale: Administration Guide*
 - Error message 6027-3470 [E] in the *IBM Storage Scale: Problem Determination Guide*
- Backup and restore:
 - Scale Out Backup and Restore (SOBAR) preserves and restores extended attributes. For more information, see *Scale Out Backup and Restore (SOBAR)* in the *IBM Storage Scale: Administration Guide*.
 - The **mmbackup** and **mmrestorefs** commands preserve extended attributes. The **mmrestorefs** command has options to restore encryption attributes and to not restore external attributes. The IBM Storage Protect product has a setting in the **dsm.sys** configuration file that omits backing up a file if only an extended attribute is changed (SKIPACLUPDATECHECK). For more information, see *Options in the IBM Storage Protect configuration file dsm.opt* in the *IBM Storage Scale: Administration Guide*.
 - The **mmrestorefs** command has options to restore encryption attributes and to not restore external attributes.
- Disaster recovery
 - In outband disaster recovery, if you are transferring files that contain extended attributes, such as object data, you must use a method that transfers extended attributes, such as IBM Storage Protect, cross-cluster mount, or AFM.
- NFS v4

- NFS v4 uses extended attributes to store some of its file attributes. For more information, see *Linux ACLs and extended attributes* in the *IBM Storage Scale: Administration Guide*.

Fast extended attribute support is now the default on newly created file systems. You can verify that a file system had fast extended attributes that are enabled by issuing the following command:

```
mmfsfs <Device> --fastea
```

Where, *<Device>* is the name of the file system, for example, *gpfs1*. If the command output shows that fast extended attributes are not enabled, then run the following command to enable fast extended attributes:

```
mmigratefs <Device> --fastea
```

Where, *<Device>* is the name of the file system. For more information, see [“Completing the upgrade to a new level of IBM Storage Scale” on page 753](#).

Planning for systemd

IBM Storage Scale supports systemd version 219 and later on Linux operating systems.

IBM Storage Scale automatically installs and configures itself as a systemd service in systems that have systemd version 219 or later installed. The service unit files are installed in the systemd `lib` directory, as shown in the following list:

- Red Hat Enterprise Linux and SUSE Linux Enterprise Server:

```
/usr/lib/systemd/system
```

- Ubuntu:

```
/lib/systemd/system
```

IBM Storage Scale includes the following systemd services. For references to IBM Storage Scale daemons, see the descriptions that immediately follow this list:

gpfs

Starts or stops the GPFS daemon (**mmfsd**).

gpfs-online

Marker service which switches to active after GPFS has reached a working state. It uses the GPFS state as shown in the **mmhealth node show** command output for its status.

gpfs-wait-mount.service

Marker service which switches to active after the GPFS filesystems have been mounted. It uses the filesystem state as shown in the **mmhealth node show** command output for its status. The status for the **mmhealth** command only looks at filesystems which have automount value set to yes.

mmccrmonitor

Starts or stops the **mmccrmonitor** daemon. This service is enabled to start at boot time. It runs all the time. If it is killed by some means other than the systemd **systemctl** command, restart it manually with the **systemctl** command.

mmautoload

Starts the **mmautoload** daemon after the node is rebooted and also starts the GPFS daemon if autostart is specified. This service also shuts down the GPFS daemon when the operating system shuts down.

The systemd manager invokes this service when the node is shut down or rebooted. This service is enabled to start at boot time after **mmccrmonitor** is started. It cannot be started or stopped with the systemd **systemctl** command.

mmsdrserv

Starts or stops the **mmsdrserv** daemon. This service is started after the node is rebooted and is stopped after the GPFS daemon is started. It is also stopped as needed by other GPFS systemd services.

mmsysmon

Starts or stops system health monitoring in a cluster in which the Clustered Configuration Repository (CCR) is enabled. This service is enabled to start when the node is booted and is meant to run continuously until the node is shut down. It also can be started by other services.

lxtrace

Does trace operations. This service is controlled by the **mmtracectl** command. Do not start or stop this service with the **systemctl** command.

scaleadmd

Starts, stops, restarts, and displays the status of the administration daemon for the IBM Storage Scale native REST API feature. The following **systemctl** commands are supported:

- To start the **scaleadmd** service, issue the following command:

```
systemctl start scaleadmd
```

- To verify the status of the **scaleadmd** service, issue the following command:

```
systemctl status scaleadmd
```

A sample is as follows:

```
scaleadmd.service - Storage Scale Admin Daemon Service
Loaded: loaded (/usr/lib/systemd/system/scaleadmd.service; enabled; vendor preset:
disabled)
Active: active (running) since Thu 2025-04-17 15:20:38 CEST; 6 days ago
Process: 498376 ExecStartPre=/bin/sh -c targetLink=$(/bin/readlink -f /var/mmfs/run) ;
Main PID: 498382 (scaleadmd)
Tasks: 13
Memory: 102.2M
CGroup: /system.slice/scaleadmd.service
498382 /usr/lib/systemd/system/scaleadmd
[-n "TargetLink" -a "${targetLink##*/}" = "mmfs" ] &6
{ /bin/m>
Apr 17 15:20:38 mos-dev-1.novalocal systemd[1]: Starting Storage Scale Admin Daemon
Service...
Apr 17 15:20:38 mos-dev-1.novalocal systemd[1]: Started Storage Scale Admin Daemon
Service. Apr 17 15:20:38 mos-dev-1.novalocal scaleadmd[498382]: REST server stopped with
error: unable to start admin daemon rest server: node identity not configured
Apr 17 15:20:38 mos-dev-1.novalocal scaleadmd[498382]: TCP gRPC server stopped with
error: unable to start admin daemon internal server: node identity not configur>
Apr 17
15:20:38 mos-dev-1.novalocal scaleadmd[498382]: Starting GRPC server on unix:///var/
mmfs/run/scaleadmd.grpc.sock, boot mode 1
Apr 17
15:40:32 mos-dev-1.novalocal scaleadmd[498382]: REST server stopped with error: unable to
start admin daemon rest server: node identity not configured
Apr 17 15:40:44 mos-dev-1.novalocal scaleadmd[498382]: Starting GRPC server on https://
[::]:50052, boot mode 2
Apr 17 15:40:44 mos-dev-1.novalocal scaleadmd[498382]: Starting REST API server on
https://0.0.0.0:46443, boot mode
```

- To restart the **scaleadmd** service, issue the following command:

```
systemctl restart scaleadmd
```

- To stop the **scaleadmd** service, issue the following command:

```
systemctl stop scaleadmd
```

- To verify the messages from the **systemd** service for the administration daemon, issue the following command:

```
journalctl -xe > /tmp/journal.out
```

After command is successful, the `journal.out` output file is generated with the **systemd** log entries. Verify the **scaleadmd** service log entries in the generated output file. For example:

```
Apr 24 08:52:56 mos-dev-1.novalocal systemd[1]: Stopping Storage Scale Admin Daemon Service...
-- Subject: Unit scaleadmd.service has begun shutting down
-- Defined-By: systemd
-- Support: https://access.redhat.com/support
-- Unit scaleadmd.service has begun shutting down.
Apr 24 08:52:56 mos-dev-1.novalocal systemd[1]: scaleadmd.service: Succeeded.
-- Subject: Unit succeeded
-- Defined-By: systemd
-- Support: https://access.redhat.com/support
-- The unit scaleadmd.service has successfully entered the 'dead' state.
Apr 24 08:52:56 mos-dev-1.novalocal systemd[1]: Stopped Storage Scale Admin Daemon Service.
-- Subject: Unit scaleadmd.service has finished shutting down
-- Defined-By: systemd
-- Support: https://access.redhat.com/support
- Unit scaleadmd.service has finished shutting down.
```

For more information about native REST API, see *IBM Storage Scale native REST API* in the *IBM Storage Scale: Concepts, Planning, and Installation Guide*.

Descriptions of the IBM Storage Scale daemons:

- The GPFS daemon (**mmfsd**) is the main IBM Storage Scale daemon that runs on the node.
- The **mmccrmonitor** daemon starts and stops the **mmfsd** daemon if **mmfsd** is not running.
- The **mmfsd** daemon provides access to configuration data to the rest of the nodes in the cluster when the **mmfsd** daemon is not running.
- The **mmauto** daemon sets up GPFS resources when the operating system is started.
- The administration daemon (**scaleadmd**) is the alternate control plane management that is included for the IBM Storage Scale native REST API feature.

Planning for protocols

This topic describes the planning considerations for using protocols.

Some protocol use considerations:

- At the time of release, several features in GPFS are not explicitly tested with protocol functionality. Explicitly not tested features are Local Read Only Cache, Multi-cluster, Encryption, File Placement Optimizer, and Hierarchical Storage Management. They are expected to work with protocols and are tested with protocols over time. However, if you use one of these features before IBM claims support of it, ensure that it is tested with the expected workloads before putting into production.

For protocol support on remote clusters, see *Using NFS/SMB protocol over remote cluster mounts* topic in the *IBM Storage Scale: Administration Guide*.

- Use of Clustered NFS (CNFS) is not compatible with use of Clustered Export Service (CES). You must choose one or the other. CNFS (which uses kernel NFS, a different NFS stack from the stack used by the CES infrastructure) continues to be supported. However, if you choose CNFS, you cannot take advantage of the integration of SMB server functionality. If you choose to migrate from CNFS to CES, the CES infrastructure does not support a complete equivalent of CNFS group feature to control failover of IP addresses.
- For more information about specific limitations about protocols and their integration with GPFS, see the [IBM Storage Scale FAQ in IBM Documentation](#) and [IBM Storage Scale in IBM Documentation](#).

Scaling considerations

When planning for scaling the protocols function, consider the maximum supported or maximum recommended number of protocol nodes and client connections.

Scaling considerations for the protocols function include the following:

- The number of protocol nodes.
 - If you are using SMB in any combination of other protocols you can configure only up to 16 protocol nodes. This is a hard limit and SMB cannot be enabled if there are more protocol nodes. If only NFS is enabled, you can have 32 nodes configured as protocol nodes.
 - If you are using CES S3 you can configure up to 10 protocol nodes.
- The number of client connections.

A maximum of 3000 SMB connections is recommended per protocol node with a maximum of 20000 SMB connections per cluster. A maximum of 4000 NFS connections per protocol node is recommended. A maximum of 3000 S3 connections per protocol nodes is recommended. The maximum number of connections depends on the amount of memory configured and sufficient CPU. A minimum of 64 GB of memory is recommended only for NFS use cases. If you have multiple protocols enabled or if you have SMB enabled we recommend 128 GB of memory on the system.
- S3 protocol supports the following scaling limits:
 - Up to 10 TB single object size
 - Up to 5000 S3 accounts
 - Up to 5000 S3 buckets
 - Up to 100 million objects per S3 bucket

Authentication considerations

To enable read and write access to directories and files for the users on the IBM Storage Scale system, you must configure user authentication on the system. Only one user authentication method, and only one instance of that method, can be supported.

The following matrix gives a quick overview of the supported authentication configurations with the IBM Storage Scale system for both file and object access.

- ✓: Supported
- X: Not supported
- NA: Not applicable

Table 32. General authentication support matrix								
Authentication method	ID-mapping method	File						S3
		SMB	SMB with Kerberos	NFSV3	NFSV3 with Kerberos	NFSV4	NFSV4 with Kerberos	
User-defined	User-defined	NA	NA	NA	NA	NA	NA	NA
LDAP with TLS	LDAP	✓	NA	✓	NA	✓	NA	✓
LDAP with Kerberos	LDAP	✓	✓	✓	✓	✓	✓	X
LDAP with Kerberos and TLS	LDAP	✓	✓	✓	✓	✓	✓	X
LDAP without TLS and without Kerberos	LDAP	✓	NA	✓	NA	✓	NA	✓

Table 32. General authentication support matrix (continued)

Authentication method	ID-mapping method	File						S3
		SMB	SMB with Kerberos	NFSV3	NFSV3 with Kerberos	NFSV4	NFSV4 with Kerberos	
LDAP with SSL		NA	NA	NA	NA	NA	NA	X
AD	Automatic	✓	✓	X	X	X	X	NA
AD	RFC2307	✓	✓	✓	✓	✓	✓	NA
AD	LDAP	✓	✓	✓	X	X	X	NA
AD with SSL		NA	NA	NA	NA	NA	NA	NA
AD with TLS		NA	NA	NA	NA	NA	NA	NA
Network Information Service (NIS)	NIS	NA	NA	✓	NA	✓	NA	NA
Local	None	NA	NA	NA	NA	NA	NA	NA
Local (OpenStack Keystone)	None	NA	NA	NA	NA	NA	NA	NA
Local (OpenStack Keystone) with SSL	None	NA	NA	NA	NA	NA	NA	NA

Note:

- Network Information Service (NIS) authentication is not supported on:
 - RHEL 9 and newer releases
 - SLES 16 and newer releases
- In the user-defined mode, the customer is free to choose the authentication and ID-mapping methods and manage on their own. That is, the authentication needs to be configured by the administrator outside of the IBM Storage Scale commands and ensure that it is common and consistent across the cluster.
- If LDAP-based authentication is used, ACL management for SMB is not supported.

Unified identity in protocols: In this case, we need to ensure that the users get the same user UID and GID across NFS and SMB.

The following diagram shows the high-level overview of the authentication configuration.




Figure 39. High-level overview of protocol user authentication

The authentication requests that are received from the client systems are handled by the corresponding services in the IBM Storage Scale system. For example, if a user needs to access the NFS data, the NFS services resolves the access request by interacting with the corresponding authentication and ID-mapping servers.

For more information about how to configure authentication, see *Managing protocol user authentication* in the *IBM Storage Scale: Administration Guide*.

For more planning information, for example, prerequisites, see *Configuring authentication and ID mapping for file access* in the *IBM Storage Scale: Administration Guide*.

Impacts of authentication on enabling and disabling protocols

The following are the recommendations for enabling and disabling protocols.

- If authentication is already configured in the system, then the protocols that are enabled in the system to access data cannot be disabled, unless the system administrator cleans up the authentication configuration.
- If authentication is configured, a disabled protocol cannot be enabled. Hence, it is vital to plan properly to decide the protocols that you want to avail now and in near future and enable them correctly in advance.

- If authentication method chosen is active directory (AD), then authentication cannot be configured unless the SMB protocol is enabled and started.
- Protocols cannot be disabled unless authentication is unconfigured. Disabling of a protocol is a critical action and it might result in loss of access to the data hosted by that protocol or other protocols. After you disable the protocol, the authentication can be reconfigured with the same authentication servers and types so that the other enabled protocols can start functioning again.

Note: Instead of disabling protocols, the system administrators might stop a protocol if they do not need it. Stopping a protocol is the recommended action instead of disabling the protocol. Stopping a protocol is as good as removing the protocol from the cluster but it does not have an impact on the access to the data. Users can retain the access when the protocol is started again. Use the **mmces stop -nfs -a** command to stop the NFS protocol on the CES nodes. For more details about these options, see **mmces command** in *IBM Storage Scale: Command and Programming Reference Guide*. If a service is stopped and not disabled, then it starts automatically post the system reboot.

- The restrictions for enabling or disabling protocols are not applicable if user-defined authentication method is used for file or object access.
- The authentication configuration needs to be removed to enable a new protocol in the system. If other protocols are already enabled in the system, the system administrator must reconfigure the authentication with the same authentication servers and types. Because of this authentication reconfiguration, the other enabled protocols can start functioning again.

You can also remove ID mappings, along with authentication, if you want to completely remove the authentication configuration. This authentication configuration removal results in permanent loss of access to the data.

Related concepts

Authentication and ID mapping for file access

The system supports an external authentication service to authenticate users on the system. Before you configure an authentication method, ensure that the external authentication service is set up correctly.

Deleting authentication and ID mapping

You can choose to delete the ID mappings that are generated and stored on IBM Storage Scale. The authentication method that is configured in the system can be deleted too. This deletion might result in a complete loss of data access. Before you delete an ID mapping, determine how to maintain data access if needed.

Authentication and ID mapping for file access

The system supports an external authentication service to authenticate users on the system. Before you configure an authentication method, ensure that the external authentication service is set up correctly.

The following steps are involved in user authentication for file access:

1. User tries to connect to the IBM Storage Scale system by using their credentials.
2. The IBM Storage Scale system contacts the authentication server to validate the user.
3. The IBM Storage Scale system contacts the ID map server that provides UIDs and GIDs of the user and user group to verify the identity of the user.
4. If the user credentials are valid, the user gains access to the system.

The following diagram shows the high-level flow of authentication for File-based protocols.




Figure 40. High-level flow of authentication for File protocols

ID mapping

The authentication of the user or groups of users is also associated with the identification of their unique identifiers. To support data access to Microsoft Windows clients (SMB protocol) and to allow interoperability, that is, to share data among UNIX and Windows clients (SMB and NFS protocols), the IBM Storage Scale system must map Windows SID to UNIX UID/GID. This process is referred to as ID mapping and the map is referred to as ID map. The ID mapping can be done either internally in the IBM Storage Scale system or in an external authentication server.

ID mapping is part of the user identification process in user authentication. The purpose of identification is to identify users and infrastructure components. Identification methods include unique user identifiers (IDs), keys, or fingerprints such as a public Secure Shell (SSH) key, and digital certificates such as a certificate of a web server.

UNIX based systems such as the IBM Storage Scale system use user names and user identifiers (UIDs) to represent users of the system. The user name is typically a human-readable sequence of alphanumeric characters and the UID is a positive integer value. When a user logs on to a UNIX system, the operating system looks up the UID and then uses this UID for further representation of the user. User names, UIDs, and the mapping of user names to UIDs are stored locally in the `/etc/passwd` file or on an external directory service such as Active Directory (AD), Lightweight Directory Access Protocol (LDAP), Keystone, or Network Information Service (NIS).

UNIX systems implement groups to maintain sets of users that have the same group permissions to access certain system resources. Similar to user names and UIDs, a UNIX system also maintains group names and group identifiers (GID). A UNIX user can be a member of one or more groups, where one group is the primary or default group. Group names, GIDs, the mapping of group names to GIDs, and the memberships of users to groups are stored locally in the `/etc/group` file or on an external directory service such as Active Directory, LDAP, Keystone, or NIS. The primary group of a user is stored in `/etc/passwd` or in an external directory service.

Windows systems reference all operating system entities as resources. For example, users, groups, computers, and so on are Windows resources. Each resource is represented by a security identifier (SID). Resource names and SIDs are stored locally in the Windows registry or in an external directory service.

such as Active Directory or LDAP. The following methods are used to map Windows SID to UNIX UID and GID:

- External ID mapping methods
 - RFC2307 when AD-based authentication is used
 - LDAP when LDAP-based authentication is used
- Internal ID mapping method
 - Automatic ID mapping when AD-based authentication is used

External ID mapping

A UID or GID of a user or group is created and stored in an external server such as Microsoft Active Directory, NIS server, or LDAP server.

External ID mapping is useful when user UID or group GID is preexisting in the environment. For example, if NFS client with UID and GID as 1000 exists in the environment, and you want a certain share to be accessed by both SMB and NFS client, then you can use an external ID mapping server, assign UID/GID 1000 to the SMB user, and thus, allow both SMB and NFS client to access same data.

Note: The external server administrator is responsible for creating or populating the UID/GID for the user/group in their respective servers.

The IBM Storage Scale system supports the following servers for external ID mapping:

- LDAP server, where the UID or GID is stored in a dedicated field in the user or group object on the LDAP server.
- AD server with RFC2307 schema extension defined. The UID or GID of a user or group that is defined in AD server is stored in a dedicated field of the user or group object.

The UID/GID defined in external server can be used by the IBM Storage Scale system.

LDAP ID mapping is supported only when the IBM Storage Scale is configured with LDAP authentication.

With the external ID mapping methods, a user is always mapped to a *uid* and a group is always mapped to a *gid*. This matches the user and group management in POSIX systems. It carries the limitation that a file or a directory can never be owned by a group. In Microsoft Windows and SMB semantics, it is possible for objects to be owned by a group. However, in IBM Storage Scale and an external ID mapping method, that cannot be supported.

Internal ID mapping

The UID or GID of a user or group is created automatically by the IBM Storage Scale system and stored in the internal repositories.

When an external ID mapping server is not present in the environment or cannot be used, the IBM Storage Scale system uses its internal ID mapping method to create the UID/GID.

IBM Storage Scale supports an Automatic ID mapping method if AD-based authentication is used. The Automatic ID mapping method uses a reserved ID range to allocate an ID based on the following logic: A user or group in AD is identified by SID, which includes a component that is called RID. Whenever a user or group from an AD domain accesses IBM Storage Scale, a range is allocated per AD domain. The UID or GID is then allocated depending upon this range and the RID of the user/group.

Internal ID mapping cannot be used when the user UID or group GID is preexisting in the environment. However, while using internal ID mapping, if an NFS client wants to access data, then the NFS client must have a UID or GID that is identical to the one created by the IBM Storage Scale system.

With the internal ID mapping method, both a *uid* and a *gid* are assigned to the same object at the same time. Querying both even reports user and group information for the same object. That allows to follow SMB semantics more closely, as users and groups are treated very similar in many aspects. This also enables to have an object owned by a group.

Other supported authentication elements for file access

The system supports the following authentication elements as well:

- **Netgroups:** Groups of hosts are used to restrict access for mounting NFS exports on a set of hosts, and deny mounting on the remainder of the hosts. The IBM Storage Scale system supports only the netgroups that are stored in NIS and in Lightweight Directory Access Protocol (LDAP).
- **Kerberos:** Kerberos is a network authentication protocol that provides secured communication by ensuring passwords are not sent over the network to the system. The system supports Kerberos with both AD and LDAP-based authentication. When you configure AD-based authentication for any ID mapping method (automatic, RFC2307, LDAP), the Kerberos is enabled for the SMB access by default. However, Kerberos for NFS access is supported only for RFC2307 ID mapping method in AD-based authentication, and to enable NFS Kerberos access for AD-based authentication with RFC2307 ID mapping; you need to use the `--enable-nfs-kerberos` and `--unixmap-domains` options in the **mmuserauth** command.

Kerberos is optional for both SMB and NFS access in LDAP. To enable Kerberos with LDAP, you need to integrate the system with MIT KDC.

- **Transport Level Security (TLS):** The TLS protocol is primarily used to increase the security and integrity of data that is sent over the network. These protocols are based on public key cryptography and use digital certificates based on X.509 for identification.

Caching user and user group information

Every UID and GID, whether generated automatically or through RFC2307 or NIS, is cached during a login attempt. For each successful user login attempt, the UID and GID are stored for seven days. For each failed user login attempt, the UID and GID are cached for two minutes.

When using an external ID mapping server, it is recommended that you not to change any previously created UID or GID. If you must change the UID or GID, plan this activity carefully, preferably before any data associated with the UID or GID exists on the IBM Storage Scale cluster. After the change, ensure that the cached entries are flushed out.

Caching user and user group information while using LDAP-based authentication

In LDAP-based user authentication, the **mmuserauth service create** command populates the LDAP bind user and bind password in its configuration file that is located at: `/etc/pam_ldap.conf`. Using this bind user and password, the *username* is looked up in the LDAP server with the configured *login_attribute*. This is the *uid* by default. As a result, *uid=<username>* is looked up in LDAP to fetch the DN for this user. If found, the DN and the *password* are used to perform an LDAP bind operation. The authentication fails if either the *username* is not found or if the bind operation fails.

Note: If LDAP is used for UNIX style login attempts, the *username* is also compared with what is returned from the LDAP server.

For SMB authentication, the ability to perform the LDAP bind with 'username' and 'password' are not available, since the password is in clear text. Therefore, the SMB client passes an NT hash for the password. The Samba server compares the NT hash from the client to the associated NT hash fetched from LDAP server. For this, the **mmuserauth service create** command configures the Samba registry with the LDAP bind user and password. The LDAP bind user and password are used to look up the Samba-related information from the LDAP server, such as the SID and the NT hash. If the NT hashes match for the bind user, then system access is granted.

User credentials, group ID mapping, and group membership caching: User credentials are not cached in the IBM Storage Scale system. The user and group ID mapping, and the group membership cache information, are stored in the SSSD component. The cache, local to each node, is for 90 minutes. There is no IBM Storage Scale native command to purge this cache. However, the `sss_cache` command (from the `sss-common` package) can be used to refresh the cache.

SMB data caching: There is a seven-day-cache period for the user/group SID within the Samba component of the IBM Storage Scale system. This cache is local to each node. There is no IBM Storage

Scale native command to manage this cache. However, the 'net cache flush' command may be used, on a per node basis, to purge the cache. A negative cache entry persists for 2 minutes.

Netgroup caching: The netgroup information from LDAP is also cached for 90 minutes with the SSSD component. The `sss_cache` command can be used to refresh the cache.

Caching user credentials while using NIS-based authentication for file access

Note: Network Information Service (NIS) authentication is not supported on:

- RHEL 9 and newer releases
- SLES 16 and newer releases

In IBM Storage Scale, an NIS server may be used only for users and groups, UID/GID, and group membership lookups for NFS access. In addition, the netgroups defined in NIS are honored as NFS client attributes within the NFS export configuration. Users and groups, UID and GID information, names, and group membership are cached within the SSSD component. This cache, local to each node, is for 90 minutes. There is no IBM Storage Scale native command to purge this cache.

The **sss_cache** command from the `sssd-common` package may be used to refresh the cache. The `sss_cache` utility only marks the entries in the respective database as expired. If the NIS server is online and available, the expired entries are refreshed immediately. If the NIS server is not available, the old, cached entries that are still marked expired become valid. These values are refreshed once the server is available.

The netgroup information from the NIS is also cached for 90 minutes within the SSSD component. The `sss_cache` command can be used to refresh the cache.

Caching user and group ID mapping and group membership details while using AD-based authentication

The Samba component in the IBM Storage Scale system uses the SMB protocol NTLM and Kerberos authentication for user authentication. User and group SID to UID and GID mapping information is cached within the Samba component. This cache, local to each node, is maintained for seven days. There is no IBM Storage Scale native command to manage this cache. However, the **net cache flush** command may be used to refresh the cache. The negative cache entry persists for 2 minutes.

The group membership cache, local to each IBM Storage Scale protocol node, lies within the Samba component. For an authenticated user, its group membership cache is valid for the lifetime of that session. The group membership is only refreshed on a new authentication request. Therefore, if an SMB tree connect is requested on an existing session, the group cache is not refreshed. So, if there are group membership changes made on the AD server, all of the existing sessions for that user must be disconnected in order to obtain a new authentication request and subsequent cache refresh for the user.

If there is no new authentication request made on behalf of the user, a simple user and group information look-up is requested. For example, if you issue the **id** command on the protocol node, the winbind component in IBM Storage Scale searches the database to check whether that user's cached information exists from a previous authentication request. If an appropriate entry is found, it is returned. Otherwise, the winbind fetches the mapping information from the AD server and caches it for five minutes.

Note: If a user was authenticated in the past, its group membership cache information is within the database and valid for the session lifetime. Winbind keeps referring to this information and will never try to fetch new information from the AD server. To force winbind to fetch new information, you need to make an authentication request, on behalf of the user, to the node in the same way as you would connect from a CIFS client.

Related concepts

Impacts of authentication on enabling and disabling protocols

The following are the recommendations for enabling and disabling protocols.

Deleting authentication and ID mapping

You can choose to delete the ID mappings that are generated and stored on IBM Storage Scale. The authentication method that is configured in the system can be deleted too. This deletion might result in a

complete loss of data access. Before you delete an ID mapping, determine how to maintain data access if needed.

Deleting authentication and ID mapping

You can choose to delete the ID mappings that are generated and stored on IBM Storage Scale. The authentication method that is configured in the system can be deleted too. This deletion might result in a complete loss of data access. Before you delete an ID mapping, determine how to maintain data access if needed.

It is not possible to delete both an authentication method and the current ID mappings at the same time. If you need to delete the current ID mappings, you need to first delete the authentication method that is configured in the system. If only the authentication settings are deleted, you can return to the existing setting by reconfiguring authentication. You cannot recover a deleted ID mapping. Typically, IDs are deleted in test systems or as part of a service operation.

For more information about how to delete authentication and ID mappings, see *Deleting authentication and ID mapping configuration* in *IBM Storage Scale: Administration Guide*.

Related concepts

[Impacts of authentication on enabling and disabling protocols](#)

The following are the recommendations for enabling and disabling protocols.

[Authentication and ID mapping for file access](#)

The system supports an external authentication service to authenticate users on the system. Before you configure an authentication method, ensure that the external authentication service is set up correctly.

Planning for NFS

You must make a number of decisions before you deploy NFS in an IBM Storage Scale environment. An outline of these decision points is as follows.

The IBM Storage Scale for NFS architecture is shown in [Figure 41 on page 390](#).




Figure 41. IBM Storage Scale for NFS architecture

Each IBM Storage Scale protocol node that has an enabled NFS, runs the NFS service, other protocol services (SMB and S3, if enabled), and the IBM Storage Scale client. NFS users make requests through an NFS client to perform an NFS operation on a mounted NFS file system, for example, to read or to write a file. The request is routed to an NFS protocol node. The NFS server on the protocol node handles the request against the backend IBM Storage Scale storage file system. The request completion status is then returned to the user through the NFS client.

File system considerations for the NFS protocol

Ensure that all IBM Storage Scale file systems, which export data via NFS, are mounted with the `syncnfs` option to prevent clients from running into data integrity issues during failover.

Linux operating systems

On Linux operating systems, the IBM Storage Scale `syncnfs` setting is the default mount option.

AIX operating system

On AIX, the IBM Storage Scale `nosyncnfs` setting is the default mount option.

Enter the following command to change the setting to `syncnfs`:

```
mmchfs fileSystem -o syncnfs
```

Enter the following command to verify:

```
mmlsfs fileSystem -o
```

For more information about mount options, see *Mount options specific to IBM Storage Scale* in *IBM Storage Scale: Command and Programming Reference Guide*.

For fileset considerations for the NFS protocol, see [“Fileset considerations for creating protocol data exports”](#) on page 403.

NFS tested clients

IBM Storage Scale tested some clients that can be used for accessing IBM Storage Scale data by using NFS protocols.

IBM support can help resolve NFS client/server issues when these issues can be re-created on one of the NFS clients that are listed here as tested.

Table 33. Tested NFS clients					
Operating System	PPC LE	PPC BE	X86_64	X86_32	s390x
RHEL6	✓	X	✓	X	
RHEL7	✓	X	✓	X	✓
RHEL8	✓		✓		✓
SLES12		X	✓		✓
SLES15			✓		✓
AIX6		X			
AIX7		✓			
z/OS® V2					✓
z/OS V3					✓
Ubuntu 20.04	✓		✓		

You must always comply with the following requisites:

- Install the latest fix level for the operating system.
- Run an operating system still supported by the vendor.
- Keep IBM Storage Scale on the latest fix level.

The following clients are not explicitly supported:

- The blank cells in [Table 33 on page 391](#) denote an unsupported environment.
- All versions of Windows NFS client.

CES NFS limitations

Describes the limitations of IBM Storage Scale CES NFS protocol.

IBM CES NFS stack limitations

CES NFS limitations are described here:

- Changes to the IBM CES NFS global configuration are not dynamic. NFS services automatically restart during the execution of the **mmnfs export load** and **mmnfs config change** commands. During this time, an NFS client with a soft mount might lose connectivity. This might result in an application failure on the client node. An NFS client with a hard mount might "stall" during the NFS restart.

- Whenever NFS is restarted, a grace period will ensue. The NFS grace period is user configurable, and the default NFS grace period is 90 seconds. If NFS global configuration changes are performed sequentially, then NFS services will be restarted multiple times, leading to a cumulative extended grace period. This might prevent NFS clients from reclaiming their locks, possibly leading to an application failure on the client node.
- A maximum of 4000 NFS connections per protocol node is recommended.
- The maximum number of NFS exports supported per protocol cluster is 1000.
- Exporting symbolic links is not supported in CES NFS.
- If NFS starts before a filesystem is mounted, exports from the filesystem might not be available. Therefore, to make the exports function work, you need to restart the `nfs-ganesha` service.

NFS protocol node limitations

When mounting an NFSv3 file system on a protocol node, the Linux kernel **lockd** daemon registers with the `rpcbind`, preventing the CES NFS lock service from taking effect. If you need to mount an NFSv3 file system on a CES NFS protocol node, use the `-o nolock` mount option to prevent invoking the Linux kernel **lockd** daemon.

NFS Ganesha protocol server's client limitation

- Microsoft Windows are not supported and tested client for CES NFS Ganesha protocol server. If you want to use Windows as client, CES SMB is the recommended protocol.
- RHEL 7 clients using IPv6 cannot communicate with the quota server on CES nodes. (Red Hat Bugzilla 2217437)

Limitations while using nested exports with NFS

Creating nested exports (such as `/path/to/folder` and `/path/to/folder/subfolder`) is not recommended as this might lead to serious issues in data consistency. Remove the higher-level export that prevents the NFSv4 client from descending through the NFSv4 virtual filesystem path. In case nested exports cannot be avoided, ensure that the export with the common path, called as the top-level export, has all the permissions for this NFSv4 client. Also, NFSv4 client that mounts the parent (`/path/to/folder`) export does not see the child export subtree (`/path/to/folder/inside/subfolder`) unless the same client is explicitly allowed to access the child export as well. If the parent folder and its subfolder are exported, access to the parent folder must be granted to the clients of the subfolder.

NFS export considerations for versions prior to NFSv4

For NFS exported file systems, the version of NFS you are running with may have an impact on the number of inodes you need to cache, as set by both the `maxStatCache` and `maxFilesToCache` parameters on the `mmchconfig` command. The performance of the `ls` command in NFSv3 in part depends on the caching ability of the underlying file system. Setting the cache large enough will prevent rereading inodes to complete an `ls` command, but will put more of a CPU load on the token manager.

Also, the clocks of all nodes in your GPFS cluster must be synchronized. If this is not done, NFS access to the data, as well as other GPFS file system operations, may be disrupted.

NFSv4 export considerations

For information on NFSv4, refer to *NFS Version 4 Protocol* and other information found in the [Network File System Version 4 \(nfsv4\)](http://datatracker.ietf.org/wg/nfsv4/documents) section of the IETF Datatracker website (datatracker.ietf.org/wg/nfsv4/documents).

To export a GPFS file system using NFSv4, there are two file system settings that must be in effect. These attributes can be queried using the `mmllsfs` command, and set using the `mmcrfs` and `mmchfs` commands.

1. The `-D nfs4` flag is required. Conventional NFS access would not be blocked by concurrent file system reads or writes (this is the POSIX semantic). NFSv4 however, not only allows for its requests to

block if conflicting activity is happening, it insists on it. Since this is an NFSv4 specific requirement, it must be set before exporting a file system.

flag	value	description
-D	nfs4	File locking semantics in effect

2. The `-k nfs4` flag is required. To export a file system by using NFSv4, NFSv4 ACLs must be enabled. Since NFSv4 ACLs are vastly different and affect several characteristics of the file system objects (directories and individual files), they must be explicitly enabled. This is done by specifying `-k nfs4`.

flag	value	description
-k	nfs4	ACL semantics in effect

3. For NFS users with more than 16 groups, set `MANAGE_GIDS=TRUE`, else user can not get access to NFS exports.

Considerations for NFS clients

When you use NFS clients in an IBM Storage Scale environment, the following considerations apply.

- If you mount the same NFS export on one client from two different IBM Storage Scale NFS protocol nodes, data corruption might occur.
- IBM Storage Scale allows concurrent access to the same file data by using SMB, NFS, and native POSIX access. For concurrent access to the same data, some limitations apply. For more information, see *Multiprotocol export considerations* in *IBM Storage Scale: Administration Guide*.
- The NFS protocol version that is used as the default on a client operating system might differ from what you expect. If you are using a client that mounts NFSv3 by default, and you want to mount NFSv4, then you must explicitly specify the relevant NFSv4.0 or NFSv4.1 version in the **mount** command. For more information, see the **mount** command for your client operating system. If you are planning to use NFSv4.1 version, ensure that the client supports the version.
- It is recommended to use the CES IP address of the IBM Storage Scale system to mount the NFS export on an NFS client.
 - You can use the **mmces address list** command to view the CES IPs.
 - You can use the **mmclscluster --ces** command to determine which nodes host which CES IPs.
- For items that might affect NFS access, see the *Authentication limitation* topic in the *IBM Storage Scale: Administration Guide*.

CES NFS scaling considerations

When planning for scaling the protocols function, consider the maximum supported or maximum recommended number of protocol nodes and client connections. For detailed information, see [“Scaling considerations”](#) on page 380.

Planning for SMB

Several steps need to be taken before the SMB service is implemented in IBM Storage Scale.

The SMB support for IBM Storage Scale allows clients to access the GPFS file system by using SMB clients. The SMB service on the protocol node provides file serving to SMB clients. Requests are routed to an SMB protocol node, typically by using DNS round robin. The SMB server on the protocol node handles these requests by issuing calls to the IBM Storage Scale file system.

Note: The SMB service is based on a clustering component CTDB. You must have the CTDB network traffic on some dedicated private network. The CTDB traffic is using the same network interface as the GPFS daemon. It can be configured or changed through the **--daemon-interface** option in the **mmaddnode** and **mmchnode** commands. The recommendation is to ensure that this is on a private network that is not visible outside the cluster.

SMB connections

Each IBM Storage Scale protocol node is capable of handling many SMB connections.

The number of concurrent SMB connections that a protocol node can handle is dependent on the following factors:

- Number of processors
- Amount of memory installed in the protocol node
- The IO workload. This depends on the following factors:
 - Number of concurrent SMB connections
 - Lifetime and frequency of SMB connections
 - Overhead due to metadata operations
 - Frequency of operations on each SMB connection
 - Concurrent access to the same files
- The storage configuration and advanced functions that are configured to run while serving high number of SMB connections

SMB fail-over scenarios and upgrade

When you are planning an IBM Storage Scale system configuration that has a high number of SMB connections, you must consider the impact that a fail-over can have on the performance of the system.

When an IBM Storage Scale protocol node fails, the IP addresses that are hosted by that protocol node are moved to another IBM Storage Scale protocol node. SMB clients must reconnect to one of the remaining CES nodes. After the fail-over is processed, the same IP address can be used. The remaining protocol nodes handle all the SMB connections.

Therefore, when you plan an IBM Storage Scale system that has a high number of SMB connections, some buffer, in terms of number of SMB connections, must be factored into the overall system configuration. This contingency buffer prevents a system overload during a fail-over scenario, thus reducing any adverse effects on the system performance.

A similar consideration applies to SMB upgrades. The SMB upgrade happens in two phases. During the first phase of the upgrade process, the first half of the nodes are updated. The remaining nodes handle the SMB connections of the updated node. When the first half of the nodes are updated, the upgrade moves to the second phase. In the second phase, SMB is shut down completely. This practice is followed to update the SMB code on all of the remaining protocol nodes concurrently. However, it leads to a brief outage of the SMB service. For more information, see [“Upgrading SMB packages” on page 697](#).

SMB limitations

This topic describes the limitations of IBM Storage Scale SMB protocol.

IBM Storage Scale allows concurrent access to the same data file through SMB and NFS, and through native POSIX access.

For more details on these options, see the *mmsmb* command in the *IBM Storage Scale: Command and Programming Reference Guide*.

For more information on SMB client access and limitations, see *Multiprotocol export considerations* in the *IBM Storage Scale: Administration Guide*.

SMB server limitations

The SMB server does not indicate mount points to the SMB client. This can have the effect that SMB clients query free space in the wrong directory, if different filesets are linked to subfolders in an SMB share. When an SMB CREATE request for creating or replacing a file fails, an empty file can be left behind.

SMB share limitations

Consider all the limitations and support restrictions before you create an SMB share.

SMB share limitations include the following:

- NTFS alternate data streams are not supported. For example, named streams generated by a Mac OS X operating system cannot be stored directly.

Enabling `vfs_fruit` will partially lift this limitation. For more information, see the *Support of `vfs_fruit` for the SMB protocol* topic in the *IBM Storage Scale: Administration Guide*.

- The encryption status of files cannot be queried or changed from SMB clients. Use the **CLI** commands instead.
- When propagation of opportunistic locks across protocols is enabled (SMB option `gpfs:leases`), then Level 2 oplocks are not granted and Exclusive or batch oplocks are not broken down to Level 2 oplocks and are revoked from the system.
- Concurrent access to the same files or directories from many SMB connections on multiple nodes can result in scalability problems.
- Symbolic links cannot be created or changed from SMB clients and are not reported as symbolic links.
- Symbolic links created via NFS or directly in the file system will be respected as long as they point to a target under the same shared directory.
- Windows Internet Name Service (WINS) is not supported.
- Retrieving Quota information using `NT_TRANSACT_QUERY_QUOTA` is not supported.
- Setting Quota information using `NT_TRANSACT_SET_QUOTA` is not supported.
- Setting the maximum number of connections to a share is not supported. The MMC GUI allows specifying this parameter, but it cannot be set on an IBM Storage Scale cluster.
- UNIX Extensions are not supported.
- You cannot create a shadow copy of a shared folder using a remote procedure call from a shadow copy client. Backup utilities, such as Microsoft Volume Shadow Copy Service, cannot create a shadow copy of a shared folder using a procedure call.
- The Branch Cache hash operations using `SRV_READ_HASH_IOCTL` are not supported.
- Leases are not supported.
- Only the SMB2 and SMB3 protocol versions are supported.
- No support of dynamic ACLs and SACLs.
- SMB Direct (SMB over RDMA) is not supported.
- Multi-channel for SMB is not supported.
- Durable handles are not supported.
- Resilient handles are not supported.
- Persistent handles are not supported.
- Storing arbitrary data in Security Descriptors is not supported.
- Pre-allocating space larger than the file size is not supported. A check for the available disk space is used instead when processing an allocation request from an SMB client.
- The Witness notification protocol is not supported.
- Continuous Availability (CA) SMB exports are not supported.
- Scale out SMB exports are not supported.
- Creating snapshots from an SMB client is not supported.
- Application instance ID is not supported.
- `FSCTL_SET_ZERO_DATA` is only supported on sparse files.
- Setting compression from an SMB client or querying the compression state from an SMB client is not supported.

- All CES nodes need to be the same hardware architecture (x86 versus Power) and the same endianness (little endian versus big endian).
- Rolling code upgrade is not supported for SMB. The SMB service needs to be stopped on all CES nodes for the upgrade to ensure that the same SMB software level is used on all CES nodes.
- When upgrading IBM Storage Scale from version 4.2.3.19 to version 5.0.5, the max protocol level for SMB support does not automatically change to the latest. This has to be done manually.
- The AAPL create context used by Mac OS clients is not supported. When enabling VFS fruit AAPL create contexts are supported.
- The number of open files for each SMB connection is limited to 16384.
- IBM Storage Scale can host a maximum of 1,000 SMB shares. There must be less than 3,000 SMB connections per protocol node and less than 20,000 SMB connections across all protocol nodes.
- No support of SID history.
- The SMB service requires file names and directory names to use the UTF-8 encoding. Files and directories with different encoding cannot be accessed through the SMB protocol.
- The NFSv4 SYNCHRONIZE bit is no longer automatically set on file ACLs. Windows clients expect SYNCHRONIZE on ACLs to correctly allow rename.
- 8.3 short name creation is not supported
- IBM Storage Scale can act as a MS-DFS target or can redirect an entire share. But it does not allow any redirects from within an exported tree.

SMB node limitations

IBM Storage Scale allows a maximum of 16 SMB server nodes in a CES cluster.

SMB user limitations

When a new security feature is added for CVE-2020-25717, Samba by default blocks access for users having UNIX uid below 1000. This was done to enhance the samba security. Additionally, a new **min domain uid** parameter (default 1000) has been added and no matter how we obtain the UNIX uid to use in the process token (we may eventually read /etc/passwd or similar), by default no UNIX uid below this value will be accepted.

If user wants the UIDs below 1000 to work, they must set **min domain uid** to the lowest UNIX user uid that they contain.

Important: At present, this **min domain uid** parameter is not supported through **mmsmb** command. Therefore, the lowest UNIX user uid can be set by using samba net command.

The following sample command shows usage of samba net command to set the **min domain uid** parameter value:

```
/usr/lpp/mmfs/bin/net conf setparm global 'min domain uid' 500
```

Here, 500 is the minimum Unix uid of the user that need SMB access.

For information on CVE-2020-25717 security feature, see [Samba security CVE-2020-25717](#).

SMB encryption on Power systems

SMB encryption on Power systems requires Openssl 1.0.2 or newer for crypto-hardware acceleration, which is provided by RHEL 7.4 or newer. Also, the hardware crypto-acceleration is only available on POWER8 or newer systems.

SMB scaling considerations

When planning for scaling the protocols function, consider the maximum supported or maximum recommended number of protocol nodes and client connections. For detailed information, see [“Scaling considerations” on page 380](#).

SMB data migration to IBM Storage Scale

This topic documents a procedure for migrating data from an SMB system to IBM Storage Scale while preserving SMB metadata like ACLs and SMB file attributes.

Prerequisites for SMB data migration

This topic describes the prerequisites that must be met before you go ahead with SMB data migration

The recommended tool for migration is Robocopy, which should be running from one or more Windows client machines. The data is read through the SMB protocol from the source machine and written through the SMB protocol to IBM Storage Scale file system.

The prerequisites are the following:

- Authentication method has to be Active Directory with the source SMB server and target IBM Storage Scale cluster using the same Active Directory environment.
- Windows client machines for the data migration can access the source SMB server and the IBM Storage Scale cluster
- A user ID used for the migration needs to be defined and that user must have permissions to read all data from the source SMB server.
- SID history needs to be removed or converted on the source data.

Setting up the IBM Storage Scale environment for capacity and performance

The first step for data migration is to set up the IBM Storage Scale nodes if you have not already done so. For more information, see [Chapter 3, “Establishing and starting your IBM Storage Scale cluster,” on page 467](#).

To achieve optimum performance while migrating small files, ensure the following:

- Use many Windows clients to scale out the number of SMB sessions
- Use Robocopy with /MT option
- Disable the syncops: onclose option while creating the SMB exports, if applicable

Configuring IBM Storage Scale authentication

The authentication method to be used is Active Directory (AD) authentication. For more information about setting up file authentication, see the **mmuserauth service create** section of the **mmuserauth** command in the *IBM Storage Scale: Command and Programming Reference Guide*.

Note: Robocopy reads SIDs from the source SMB server, sends them to IBM Storage Scale. IBM Storage Scale creates new UNIX UIDs and GIDs based on the ID mapping that is configured on the IBM Storage Scale CES cluster.

Creating SMB exports

Create SMB exports for the target of the data migration. For more information on creating SMB exports, see the **mm smb** command in the *IBM Storage Scale: Command and Programming Reference Guide*. Ensure to configure the migration user ID as "admin user" for all SMB exports. See the **mm smb export change SMBexport --option 'admin users=domain\user' options**.

Note: The user identifier used for the migration should have rights to set file ownership and ACLs.

Accessing the source SMB server and the IBM Storage Scale file system

Do the following:

1. From the Windows clients, access the source SMB server and the IBM Storage Scale cluster.
2. Map network drives for the source SMB server and the IBM Storage Scale cluster.

If data for multiple SMB exports are transferred, multiple Windows clients can be used, where each one accesses one pair of source and target SMB exports. This allows to speed up the transfer.

Running the Robocopy commands

For each SMB export that you want to migrate, run the **Robocopy** command as follows:

```
Robocopy \\<source-smb-server>\<smbexport> \\<target-scale-cluster>\<smbexport> /copy:ATSO /secfix /Z /E /MT:32 /R:5 /W:3 /sl /log:logfile .
```

The Robocopy options are as follows:

- /COPY:ATSO: Copy file information.
 - A: Attributes
 - T: Timestamps
 - S: Security - NTFS ACLs
 - O: Owner information
- /secfix: Fixes the security on all files
- /Z: Ensures Robocopy can resume the transfer of a large file in mid-file instead of restarting.
- /E : Copy all subdirectories including the empty ones.
- /MT : Create multi-threaded copies with N threads. N must be an integer between 1 and 128. The default value for N is 8.
- /R: Specifies the number of retries on failed copies. The default value is 1,000,000 (one million retries).
- /W: Specifies the wait time between retries, in seconds. The default value is 30 (wait time 30 seconds).
- /sl : Copies the symbolic link instead of the target.
- /log: Writes the status output to the log file (overwrites the existing log file)

If the above command fails with errors such as "access denied", issue the following command:

```
Robocopy \\<source-smb-server>\<smbexport> \\<target-scale-cluster>\<smbexport> /copy:DATSO /secfix /Z /E /IS /MT:32 /R:5 /W:3 /sl /log:logfile /V
```

- /IS: Includes the same files. The same files have identical name, size, times, and all attributes.
- /copy:DATSO: Copies file contents.

For example,

```
Robocopy \\10.11.12.15\smbshare1 \\10.90.40.16\smbshare1 /copy:DATSO /secfix /Z /E /IS /MT:32 /R:5 /W:3 /sl /log:C:\robocopylog.txt /V
```

Cleanup after migration

Do the following steps after migration:

1. Unmount all SMB exports from the Windows clients.
2. Remove the "admin users" entry as it is no longer required and can be harmful if used by accident.

SMB best practices

Some SMB best practices are highlighted in the following sections.

IBM Storage Scale configuration

Learn about the factors that determine the maximum number of active SMB connections.

The two important factors that determine the maximum number of active SMB connections per protocol node are:

- The protocol node hardware
- The underlying storage systems.

Storage planning and configuration

The storage configuration must be carefully planned based on the overall workload characteristics that are intended for IBM Storage Scale.

The factors to be considered in planning the storage configuration include:

- Expected response time, during peak and off peak hours.
- Workload profiling to understand the data access requirements, both in terms of throughput (MiB per second) and I/O operations per second (IOPS).
- Plan for adequate storage resources not just for “normal” network file serving but also for advanced functions.

The storage configuration includes the following details:

- Number of storage nodes.
- Storage disk subsystems.
- Number and type of disk drives.
- Total capacity
- Projected IO throughput and IOPS

A file system that resides on fewer disks, or disks that are mapped on to RAID arrays that comprise slower lower speed disk drives, can result in a lower number of active concurrent SMB connections per protocol node.

SMB Options to enable or disable interoperability

Read about options to enable or disable interoperability, that is, the options that control whether file access metadata is forwarded into the file system to make it available to other clients (POSIX, NFS, etc.) as well.

Leases, locking, and sharemodes

If your files are only accessed via SMB clients, the overhead that is incurred while ensuring consistency across multiple NAS protocols can be avoided if you disable the following interoperability options. This is recommended if the data in the file system is only accessed through the SMB protocol and not through other NAS protocols, such as NFS, FTP, HTTPS, or by POSIX clients.

gpfs:leases

Understand the scenarios in which gpfs:leases are enabled or disabled.

gpfs:leases are enabled by default when an SMB export is created. When enabled, it specifies that clients, which access the file over other NAS protocols, can break the opportunistic lock of an SMB client. This feature informs the SMB client when another client is accessing the same file at the same time through a non-SMB protocol.

Disabling this feature provides a slight increase in performance each time that a file is opened. You can only disable this option if the files are exclusively accessed via SMB. You risk data corruption otherwise. If only the SMB layer manages oplocks, the overhead of keeping them in sync with the core file system is saved.

gpfs:leases can be disabled for a particular SMB export by specifying the

```
--option "posix:locking"=no
```

option on the "create" or "change" subcommands in the *mmsmb* command topic in the *IBM Storage Scale: Command and Programming Reference Guide*.

posix:locking

Understand the scenarios in which *posix:locking* is enabled or disabled.

posix:locking is enabled by default when an SMB export is created. When enabled, it determines whether a byte range file control lock is already present on the requested portion of the file when a byte range lock is granted to an SMB client.

Clients accessing the same file that uses another NAS protocol, such as NFS, are able to determine whether an SMB export sets a lock on that file.

If an export is accessed only by an SMB export, disable the inter-protocol level byte-range locking to enhance the serving performance of the SMB export.

Locking is disabled for a particular SMB export at the Inter-protocol level by specifying the

```
--option "posix:locking"=no
```

option on the "create" or "change" subcommands in the *mmsmb* command topic in the *IBM Storage Scale: Command and Programming Reference Guide*.

gpfs:sharemodes

Read about share modes and how they impact SMB export.

The SMB protocol allows an application simultaneous access to a file by defining share modes when it is first opened. The share modes can be in any of the following combinations.

- SHARE_READ
- SHARE_WRITE
- SHARE_DELETE

If no sharemode is specified, all attempts of simultaneous access by another application or client to open a file, in a manner that conflicts with the existing open mode, is denied. Access is denied even if the user has the appropriate permissions that are granted by share and file system access control lists.

The "sharemodes" option is enabled by default when an SMB export is created. When enabled, the share modes that are specified by SMB clients are respected by other NAS protocols. When disabled, it specifies that the share modes are applicable only to access by SMB clients. Clients that use all other NAS protocols are granted or denied access to a file without regard to any share mode defined by an SMB client.

If the export is not being accessed by clients that use other network file protocols (such as NFS), then it is highly recommended that

```
--option gpfs:sharemodes=no
```

option is specified on the **mmsmb export add** or **mmsmb export change** commands.

For more information, see [Upgrading SMB packages](#).

Note: If your environment requires data sharing over multiple –protocols and these options cannot be disabled, you might not be able to achieve the maximum active SMB connections per node. In that case, consider adding more protocol nodes and increased storage bandwidth.

Creating home directory exports using substitution variables

Guidance to set up SMB exports for efficient home directories

Having many Windows users all concurrently accessing the same SMB export can lead to performance bottlenecks because Windows clients automatically open the root folder of an export when they connect. In a home directory environment, it is recommended that substitution variables be used

when creating SMB exports for home directories. For example, home directory exports can be created by using the `%U` substitution variable that represents the user name on the **mmsmb export add** <export> <path with substitution variable> " command (**mmsmb export add home /ibm/gpfs0/.../%U**). For more information, see the *mmsmb* command topic in the *IBM Storage Scale: Command and Programming Reference Guide*.

Sharing files and directories among SMB clients - concurrent access to files

Information on performance optimization when performance slows down due to extensive sharing.

If your environment calls for extensive file and directory sharing among many users, such as a large set of department documents, you can experience slowdown in performance. In this type of environment, it is possible to improve the performance based on the user needs. Consider the following options to optimize performance:

- Use the SMB export coherency options that are described in the following section.
- Limit all sharing through a single protocol node, or as few protocol nodes as possible. Restricting SMB connections that export data to a single protocol node helps reduce internal communication among the protocol nodes.
- If possible, distribute workload in sub directories to reduce number of SMB connections that access the same directory at the same time.

SMB export coherency options: fileid:algorithm

IBM Storage Scale provides a coherency, that is, `fileid:algorithm`, option to control data consistency needs for an SMB export. This option applies when an export is being accessed only by SMB clients. When the default value of **fsname** is changed, it helps to improve performance. However, extreme caution must be taken to determine right settings for your data as it impacts data integrity.

The applications must ensure that files or directories are not modified by multiple processes at the same time. For example, reading and writing of the same file does not happen simultaneously by different processes or alternatively, when the application is coordinating all file accesses to avoid conflicts.

The coherency (**fileid:algorithm**) option can be changed for a particular SMB export by specifying the **--option fileid:algorithm = {fsname | hostname | fsname_nodirs | fsname_norootdir}** option on the **mmsmb export add** or **mmsmb export change** commands.

Table 34. Coherency option and description	
Coherency option	Description
fsname:norootdir	Disables synchronization of directory locks for the root directory of the specified export, but keeps lock coherency for all files and directories within and underneath the export root. This option is useful for scenario where large sets of connections are accessing different sub directories with the same export. A typical example is an export that is used for home directories (<code>/ibm/gpfs0/homeroot</code>) which then contains a subdirectory for each user.
fsname:nodirs	Disables synchronization of directory locks across the cluster nodes, but leaves lock coherency that is enabled for files. This option is useful if data sharing is not dependent of the changes to the directory attributes like time stamps, consistent view of the directory contents.
fsname	Enables cross-node lock coherency for both directories and files. This is the default setting.
hostname	Disables cross-node lock coherency for both directories and files. It needs to be used only with applications that guarantee data consistency and all other options to enhance performance are exhausted.

Scheduling advanced functions for data management

Information and tips on data management tasks timings and schedule planning.

In most environments, it is typical to have an off-peak window of time at some point during the day. This time can be used to conduct data management tasks such as nightly backup, snapshots, and asynchronous replication. You need to ensure that you have some time of lower SMB file activity on the server and that this time window is sufficient for the planned advanced functions to complete.

When running advanced functions that require a file system policy scan such as backup, asynchronous replication, policy invocations, Advanced File Management (AFM), or Active Cloud Engine (ACE) cache pre-population, schedule them sufficiently apart to allow adequate time to complete the policy scan. The planned time gap helps to avoid two overlapping policy scans.

In a typical system, a couple of hours gap between two advanced functions is sufficient. However, you need to review the logs of each function to ensure that the scan completes before the next scheduled advanced function starts. If necessary, make adjustments like increasing the time gap, adding more protocol nodes, more disks for metadata, or adding Solid State Disks (SSDs) for metadata for IBM Storage Scale gateway configuration.

As you plan data management tasks for your environment, you need to ensure adequate resources are available to complete all data management tasks. If these tasks do not complete in the expected time window or they impact overall performance of the system during peak hours, consider adding additional resources such as dedicated protocol nodes for backup or extra storage resources to enhance storage response time. This can also eliminate bottlenecks.

Planning for fail-over scenarios and upgrade

Information on performance impact and management during failover scenario of IBM Storage Scale protocol node.

During the planning of a system that has a high number of active concurrent SMB connections, the potential performance impact during fail-over scenarios needs to be considered.

If an IBM Storage Scale protocol node fails, the IP addresses that are hosted by that protocol node are relocated to other IBM Storage Scale protocol nodes, and SMB client reconnections are redistributed among the remaining IBM Storage Scale nodes.

When a high number of active concurrent SMB connections are planned, the overall system configuration needs to factor some buffer space in terms of maximum active concurrent SMB connections. This accounts for the potential performance implications during these fail-over scenarios.

During the IBM Storage Scale software upgrade process, IP addresses are frequently relocated. Depending on the configuration, multiple protocol nodes are suspended concurrently to minimize the upgrade time. This leaves fewer protocol nodes to serve various protocol clients including SMB. Therefore, the maximum number of active SMB connections cannot be sustained during the IBM Storage Scale software upgrade process. You need to plan for upgrade during the off-peak hours or schedule a maintenance window to minimize the impact on clients that access IBM Storage Scale.

For more information on upgrade planning, see [“SMB fail-over scenarios and upgrade” on page 394](#).

Analysis of performance concerns and fine-tuning

Information on addressing performance problems due to the high number of active concurrent SMB connections on each protocol node of IBM Storage Scale.

If IBM Storage Scale experiences performance problems, the following actions can be taken:

- You can use the IBM Storage Scale center GUI or CLI to understand which physical resource such as CPU, memory, networking, disk in the system are highly used resources. This information helps to gain an insight on the physical system resource that is possibly inhibiting or limiting performance. For more information, see the *mmpfmon* command topic in the *IBM Storage Scale: Command and Programming Reference Guide*.

- Ensure that the IBM Storage Scale protocol nodes are configured with the maximum number of processors, memory, and networking adapters.
- Add more IBM Storage Scale protocol nodes to your system.
- Move certain advanced functions, such as TSM and AFM, to periods of time when SMB file activities that are related to the server is lower.
- Reduce the frequency at which snapshots are being created and deleted either at the file system or the file set level or both, especially during the periods of highest SMB user activity.
- Check that the file system cluster has appropriate values for pagepool, number of worker threads, and number of receiver threads. For more information, see the *Configuring and tuning your system for GPFS* topic in the *IBM Storage Scale: Administration Guide*.
- Investigate and fine-tune the performance of the underlying disk storage systems that contain the file systems on which the SMB export resides. The checklist items are described in the following list:
 - Ensure that the file system disks of a storage system are distributed between the pair of IBM Storage Scale storage nodes to which that storage system is attached. One half of the file system disks in a storage system should have one of the storage nodes that are identified as the primary NSD server. The other half of the file system disks in the storage system should have the other IBM Storage Scale storage node in the pair that is assigned as the primary NSD server. The **mmlsdisk** CLI command option shows the IBM Storage Scale storage nodes that are the primary and secondary NSD server for each file system disk.
 - If an IBM Storage Scale metadata replication or data replication is being used, ensure that you assign file system disks to failure groups that balance the I/O and data across a set of file system disks, RAID arrays, and disk storage systems. The **mmlsdisk** CLI command shows the failure group to which each file system disk is assigned. The **mmchdisk** CLI command can be used to change the failure group to which each file system disk is assigned.
 - If the underlying disk storage systems on which the file system reside is becoming a performance bottleneck, consider adding more physical resources to those disk storage systems. More resources include more cache memory, more disk drives, more RAID arrays and more file system disks for the file systems that contain the SMB exports.
 - If the existing disk storage systems on which the file system resides reach their limit in terms of either capacity or performance or both, then consider adding more disk storage systems and extending the file system on which the SMB exports resides. Capacity and performance improvement can be done by adding new file system disks that are residing on the new disk storage systems to the existing disc storage system.

Considerations for SMB clients

Recommendations to consider before you mount SMB share on Linux systems.

Linux systems can mount SMB shares from protocol nodes. For more information, see your Linux distribution documentation and the `mount.cifs` man page. Workloads running on Linux SMB clients that often query `stat()` metadata can run into performance limitations. The performance limitation is due to the increase in processor usage as a result of the high number of metadata queries in the SMB protocol and the IBM Storage Scale SMB implementation.

The Linux SMB mount offers the "actimeo" mount option (see "mount.cifs" man page) for metadata caching on the Linux SMB client. The "actimeo" mount option reduces the metadata queries to the IBM Storage Scale SMB server. For workloads with a high number of metadata queries, you can set this parameter to improve the overall performance.

Fileset considerations for creating protocol data exports

You can create exports on the entire file system, on sub-directories of a file system, or on filesets.

A fileset is a file system object that enables you to manage data at a finer granularity than the file system. You can perform administrative operations such as defining quotas, creating snapshots, and defining file

placement policies and rules, and specifying inode space values at the fileset level, especially when the fileset is independent.

In IBM Storage Scale, you can create exports even without filesets. Depending on your data management strategy, choose either of the following ways to create exports:

Create exports on the entire file system or on sub-directories of the file system

In this option, the export represents a large space. You can create independent filesets over the directories in this space and have finer control over the export directory paths. Universities and organizations that require a departmental multi-tenancy solution can choose this option.

Review the following example to better understand this option.

As a storage administrator of an organization, you want to create separate storage space for every department and user of the organization:

1. Export the root directory of the file system.

```
mmnfs export add /gpfs/fs0
```

Note: You can create a sub-directory in the root directory and export it. For example: `mmnfs export add /gpfs/fs0/home`

For more information, see ***mmnfs command*** in *IBM Storage Scale: Command and Programming Reference Guide*.

2. Create independent filesets in the root directory, linked to the subdirectory `/gpfs/fs0/home`.

In the following example, it is assumed that there is a user `user1` that is a part of the group `group/HR`.

```
mmcrfileset fs0 hr_fileset --inode-space=new
mmlinkfileset fs0 hr_fileset -J /gpfs/fs0/home/hr
mmcrfileset fs0 user1_fileset --inode-space=new
mmlinkfileset fs0 user1_fileset -J /gpfs/fs0/home/user1
```

For more information, see the following commands in the *IBM Storage Scale: Command and Programming Reference Guide*.

- **`mmcrfileset` command**
- **`mmlinkfileset` command**

3. Similarly, create independent filesets for other departments and users.


You can assign quota to each user and department via the independent filesets.

The NFS and SMB clients can now mount the export and access their directories.

Access to the data in the export directories is controlled via the group and user ACLs. In this case, only the users who are in group HR, which has group ACLs to read/write into the `hr` directory, can access the directory. The user `user1` who is in the group `group/HR` can perform read/write to the `user1` directory and the `hr` directory.

Create exports on the entire file system or on sub-directories of the file system

 = Where export occurs



Create exports on independent filesets

In this option, the independent filesets represent discrete projects. One or more exports can be created on each fileset. You can apply the Information Lifecycle Management (ILM) policies over the filesets to automate the placement and management of file data. You can protect the export data by granting access permissions to specific workstations or IP addresses. Also, data in the exports can be preserved by the independent fileset's snapshot policy.

Review the following example to better understand this option.

You are a storage administrator of a private cloud hosting storage system that stores webcam data. You want to ensure that a webcam has access only to its storage directory. Also, you want the data analyzers to access all data so that they can look for activity and generate analytical reports.

1. Create an independent fileset web_cam_data.

```
mmcrfileset fs0 web_cam_data --inode-space=new
mmlinkfileset fs0 web_cam_data -J /gpfs/fs0/web_cam_data
```

Data from all webcams is stored in /gpfs/fs0/web_cam_data.

2. Create exports for both webcams.


```
mkdir /gpfs/fs0/web_cam_data/camera1
mkdir /gpfs/fs0/web_cam_data/camera2

mmnfs export add "/gpfs/fs0/web_cam_data/camera1" \
-c
"198.51.100.2(Access_Type=RW);203.0.113.2(Access_Type=R0);203.0.113.3(Access_Type=R0)"
mmnfs export add "/gpfs/fs0/web_cam_data/camera2" \
-c "198.51.100.3(Access_Type=RW);203.0.113.2(Access_Type=R0);203.0.113.3(Access_Type=R0)"
```

The webcam1 (IP: 198.51.100.2) mounts and records data to the camera1 export and the webcam2 (IP: 198.51.100.3) mounts and records data to the camera2 export. The data analyzers (IP: 203.0.113.2 and 203.0.113.3) are given 'Read Only' type access to both exports. Thus, the data is accessible only from the specified IP addresses.

Create exports on independent filesets

 = Where export occurs



Planning for S3

Before you deploy the S3 service on IBM Storage Scale, ensure that all prerequisites are met.

After the installation and enablement of the S3 service, for secure communication between the S3 client and the S3 service, see the *Setting up self-signed SSL/TLS certificates for secure communication between the S3 client and the S3 service* in the *IBM Storage Scale: Administration Guide*.

For planning for S3 over RDMA, see the *S3 over RDMA overview and deployment* section *IBM Storage Scale: Concepts, Planning, and Installation Guide*.

Configuration considerations for CES clusters

The configuration considerations for a CES cluster, on which the S3 protocol is enabled, are as follows:

- Up to 10-node CES cluster based on IBM Storage Scale 5.2.2
- For each CES node:
 - Operating System: RHEL 8.x or RHEL 9.x
 - Architecture: x86_64, Power ppc64le, Z - s390x (RH 9.x)
 - Memory: Up to 64 GB RAM for NooBaa service configured with up to eight endpoint processes
 - Depending on the workload requirements you can increase or decrease the number of S3 NooBaa protocol stack 'endpoint' processes by using the **mms3 config** command. Ensure that endpoint processes are sufficient to manage workloads.

Note: By default, CES S3 is configured with two endpoint processes. Increase the endpoints to have 8–16 endpoint processes based on workload and available memory of the system.

Configuration considerations for each NooBaa endpoint fork process

The configuration considerations for each NooBaa endpoint fork process are as follows:

- Minimal configuration: 4 GB RAM and 2 vCPU
- Recommended configuration: 4 GB RAM and 3 vCPU

Planning for CES HDFS

For CES HDFS planning information, see *Planning and Limitations and Recommendations* under [CES HDFS](#) in *Big data and analytics support documentation*.

Planning for cloud services

Proper planning is one of the most important activities that you need to carry out for successful implementation of cloud services on the IBM Storage Scale cluster.

Hardware requirements for cloud services

You must meet certain hardware requirements to be able to install and use cloud services on the IBM Storage Scale cluster.

For running cloud services, adhere to the recommended CPU and memory requirements for a CES node even when installing cloud services on an NSD node. For more information, see the [IBM Spectrum Scale FAQ in IBM Documentation](#).

The hardware requirements for cloud services are:

- Any standard x86 64-bit servers or Power Linux nodes that run supported Linux distributions.
- The minimum size that is required for the `/var/MCStore` folder is 12 GB.

Note: For better performance, it is recommended to have a minimum of 2 CPU socket server of the latest Intel variety with at least 128 GB of memory.

A high CPU count promotes better cloud tiering throughput because although object storage can be slow in I/O operations per thread, object storage can support many threads. Use sixteen or more CPUs when you select your hardware.

Cloud tiering services demand a large amount of memory, which is why the minimum recommended memory size is 128 GB. Memory size requirements increase if the number of files increases, as you add files on the cloud means you must increase the memory that is on your system. For larger deployments, it is recommended that you use 10 - 20 times as much memory that is required so the cloud services can cache its directory database data.

Software requirements for cloud services

You must meet certain software requirements to be able to install and use cloud services on the IBM Storage Scale cluster.

Note: cloud services are shipped along with IBM Storage Scale Advanced Edition, IBM Storage Scale Data Management Edition, or IBM Storage Scale Developer Edition or IBM Storage Scale Erasure Code Edition. Therefore, IBM Storage Scale Advanced Edition, IBM Storage Scale Data Management Edition, or IBM Storage Scale Developer Edition or IBM Storage Scale Erasure Code Edition is required for the installation.

The software requirements for cloud services are

- For the complete list of OS support, see [Table 35 on page 408](#).
- For Kernel versions, see [IBM Storage Scale FAQ in IBM Documentation](#).
- The following packages must be installed on the Transparent cloud tiering core server nodes:
 - Redhat-release-server
 - Unzip
 - Tar
 - bc - An arbitrary precision calculator language
 - SQLite 3.7 and above
 - GPFS base RPMs

- GPFS RPMs required for a CES node
- GPFS Advanced Edition RPMs
- For bulk delete functionality to work in Swift, verify that the Bulk middleware to be in pipeline, as follows:

```
vim /etc/swift/proxy-server.conf
[pipeline:main] pipeline = healthcheck cache bulk authtoken keystone proxy-server
```

- To run client-assist, clients need all the packages listed for server nodes except SQLite.

<i>Table 35. OS support matrix</i>		
OS	TCT server	TCT client
RHEL 7.1 (x86)	yes	yes
RHEL 7.2 (x86)	yes	yes
RHEL 7.3 (x86)	yes	yes
RHEL 7.4 (x86)	yes	yes
RHEL 7.5 (x86)	yes	yes
RHEL 7.1 (Power8 and Power9 LE)	yes	yes
RHEL 7.2 (Power8 and Power9 LE)	yes	yes
RHEL 7.3 (Power8 and Power9 LE)	yes	yes
RHEL 7.4 (Power8 and Power9 LE)	yes	yes
RHEL 7.5 (Power8 and Power9 LE)	yes	yes
RHEL 8.0 (Power8 and Power9 LE)	yes	yes
RHEL 7.5 (s390x)	yes	yes
SLES 11.2	no	yes
SLES 12.0	no	yes
SLES12 SP3 (s390x)	no	yes
Ubuntu 14.x and 16.4	no	yes
Ubuntu 16.04 and 18.04 (s390x)	no	yes

For the latest support matrix, see *Transparent cloud tiering questions* in [IBM Spectrum Scale FAQ](#).

Network considerations for cloud services

This topic describes the network considerations that you need to follow for installing cloud services on your IBM Storage Scale cluster.

Note:

- Network Time Protocol (NTP) must be correctly set on all nodes where the cloud services are installed. If correct date and time are not set, some of the cloud storage providers will refuse to authenticate your requests.

- cloud services always use the GPFS cluster (daemon) network and not the CES network.

Cloud storage supports storing cold data on a public cloud object storage service, such as IBM Softlayer or Amazon S3. When using a public cloud, the connection is over a WAN.

cloud services for both transparent cloud tiering and Cloud data Sharing can consume a considerable amount of WAN bandwidth. Customers need to ensure that there is sufficient WAN bandwidth for migrations and recalls along with any other WAN usage that they might require. When cloud services run, an administrator can monitor WAN usage and other running processes in the IBM Storage Scale performance monitor. For a list of metrics, see the *cloud services* section in the *List of Metrics* topic in the *IBM Storage Scale: Administration Guide*.

To avoid any conflict of cloud services I/O traffic and internal IBM Storage Scale cluster I/O traffic, it is recommended that the cloud services node is configured with two or more NICs. One NIC card is designated as the target for all inter-cluster I/O, and the other one is designated for cloud services I/O traffic.

After the NIC card is set up and connected to the WAN, it is recommended to run any external network speed test, such as speedtest.net, to ensure that the I/O bandwidth to the external network meets your expectations.

If customers choose to configure a private cloud instance, they need to ensure that there is sufficient bandwidth between the IBM Storage Scale cluster and the private cloud to support migrating and recalling files. It is preferred to have a 10 GB network for higher throughput.

Cluster node considerations for cloud services

This topic describes the cluster considerations that need to be followed before you install cloud services on the IBM Storage Scale cluster.

You can install the core cloud services packages on a maximum of 4 CES or NSD (or a combination of both) nodes that you want as members of your cloud services node class. We support up to 4 node classes. The client package is installed on all other nodes within the same cluster. Cloud data transfer (migration, recall, import, or export) occurs only from the cloud services nodes. However, you can initiate a cloud data transfer from either a server or a client node.

You can install the transparent cloud tiering server package on a non-gateway node. In this client-assisted recall configuration, the non-gateway node can recall from the same node instead of routing the commands through the gateway nodes. Except for recall, all other commands are routed through the gateway nodes. Thus, file-read requests originating on the client nodes are served immediately from the same node.

Use of client-assisted recall configuration is recommended in the following scenarios:

- The non-gateway node has direct network access to the cloud object storage.
- High CPU and memory resources are available on the node as the client-assist package starts a Java virtual machine per node-class managed
- If many recalls are initiated through the non-gateway nodes.

Do not use a mixed node cluster configuration in which there are nodes on the cluster that do not allow the installation of clients. For example, Windows nodes do not allow the installation of clients because there is no Transparent cloud tiering Windows client. These configurations:

- Do not support transparent data recalling
- Cannot respond to requests that initiate from nodes on which a client is not installed
- Cannot efficiently perform deletions

cloud services is supported for use only in IBM Storage Scale clusters and on any associated remotely mounted clients, with Intel x86 Linux and Power LE Linux nodes. Use of Windows, PowerPC Big Endian, Linux on Z or AIX nodes that access file systems, where cloud services is used (either within the cluster or via a remotely mounted client), is untested and unsupported. Inclusion of such nodes is allowed only if there is no possibility that they will ever access cloud services; for example, such nodes are tested and

supported for PowerPC Big Endian nodes in Elastic Storage Server (ESS) BE & LE servers, where no user applications are supported on Elastic Storage Server (ESS).

IBM Cloud Object Storage considerations

The following information describes about the points that you need to consider before you use IBM Cloud Object Storage as the object storage provider.

Note: IBM Cloud Object Storage 3.7.3.2 or above is required for the cloud services functions. If you are running an older version of IBM Cloud Object Storage, contact IBM support for upgrading your version.

Before you begin, ensure that an access pool is created and available on the IBM Cloud Object Storage system. The access pool must have 'Cloud Storage Object' configured as the API type.

Before you configure cloud services with IBM Cloud Object Storage as the object storage provider, ensure that the following settings are done through IBM Cloud Object Storage **dsNet Manager** GUI. If these settings are not correctly set, the cloud account configuration or data migration fail.

- During the cloud services setup process, two vaults are created on the IBM Cloud Object Storage system. One of the vaults is used for storing data and the other one is used for storing metadata. The data vault that is created by cloud services contains the “container-prefix” appended with a unique ID.

Note: IBM Cloud Object Storage endpoint URL (“cloud_url” command-line option) must not have this container or vault name.

- In order for cloud services to be able to create the vault, the user (whose access key is specified in the **mmcloudgateway account create** configuration command) must have the "Vault Provisioner" role that is assigned through the **dsNet Manager** UI.

To create vaults, go to **dsNet Manager** UI > **Security** tab and click **user** > **Roles** > **Add "Vault Provisioner"** role.

Make sure that either “Create Only” or “Create and Delete” permission is selected under **Administration** > **Provisioning API configuration**. Doing this enables cloud services to create vaults by using the IBM Cloud Object Storage vault provisioning APIs. It is not necessary to allow cloud services privileges to create the vault. You can create the vault separately by using IBM Cloud Object Storage management services.

Note: Delete access is not required for the “Vault Provisioner”.

Note: To specify the container name of an existing container, use the **--data-name/--meta-container** parameter of the **mmcloudgateway** command.

- For IBM Cloud Object Storage deployed on public cloud and with container mode enabled, you must contact the cloud object storage admin to obtain the accessor IP address or hostname, S3 credentials, and provisioning code (location) for the container vault.
- To create vaults through the provisioning API, IBM Cloud Object Storage uses provisioning templates. You can provide a provisioning template to the cloud services in two ways:

Using default template

A default vault template can be set on the **dsNet Manager** UI. Go to **dsNet Manager** UI > **Configure** tab > **Storage pools** > **Create Vault template**. Then, **dsNet system** > **Default Vault template** and then select the newly created template. The vault that is created by cloud services uses the default template.

Note: The default template on the IBM Cloud Object Storage system must have index that is enabled.

Using Provisioning code

If the default vault template is not set or not preferred, then a provisioning code can be used. Ensure that during the creation of the vault template, a unique user-defined provisioning code is specified. The provisioning code is different from the name.

To look up the provisioning code, go to **dsNet Manager** UI > **Configure** tab > **Storage pools** > **Vault Templates** > select the template, and look for “Provisioning Code”). Use the **--location**

option in the **mmcloudgateway account create** command to specify the provisioning code. Using this mechanism, the vault template that is configured for cloud services is used for vault provisioning.

Note: If there is no default provisioning template set on the IBM Cloud Object Storage system and a provisioning code is not specified to the **mmcloudgateway account create** command, the command fails. If the provisioning code specified is not found on the IBM Cloud Object Storage system, the command fails.

The following settings are recommended when you create a vault template on IBM Cloud Object Storage dedicated for transparent cloud tiering.

<i>Table 36. Recommended settings when you create a vault template on IBM Cloud Object Storage</i>		
Configuration	Recommended Values	Comment
Width	See IBM Cloud Object Storage documented configuration for production.	
Threshold	See IBM Cloud Object Storage documented configuration for production.	
WriteThreshold	See IBM Cloud Object Storage documented configuration for production.	
Alert Level	See IBM Cloud Object Storage documented configuration for production.	
Alert Level	See IBM Cloud Object Storage documented configuration for production.	
SecureSlice Technology	Disabled	When using cloud data sharing services, the user might consider enabling SecureSlice Technology encryption. If using the transparent cloud tiering services, encryption is not needed and is redundant since data is encrypted by the transparent cloud tiering service before the data is stored on object storage.
SecureSliceAlgorithm	Not applicable since SecureSlice is disabled.	
Versioning	Disabled	transparent cloud tiering has built-in versioning capability, hence IBM Cloud Object Storage versioning can be unavailable. For the Cloud Data Sharing service, versioning might or might not be turned off depending on the needs for retaining versioning on the data.
DeleteRestricted	Yes/No	The gateway does not attempt to delete the vaults, so this setting can be set to yes or no.
Name Index	Disabled	Disabling this setting can result in improved vault performance.

Table 36. Recommended settings when you create a vault template on IBM Cloud Object Storage (continued)

Configuration	Recommended Values	Comment
Recovery Listing	Enabled	For performance reasons, the vault that is used for storing data has Name Index disabled and for searchability reasons, the other vault has index that is enabled. On the second provisioning template, Name Index is enabled and the rest of the settings are the same as above.

Essentially, IBM Cloud Object Storage needs two provisioning templates. One of them is used for storing data and the other one is used for metadata or book-keeping. This vault provisioning template must be set as default (Click the **Configure** tab and scroll down to see the option to set the default template). Pass the provisioning code ('demo ' in the example) of the first vault provisioning template to the **mmcloudgateway** command during account creation by using the `--location` parameter.

Firewall recommendations for cloud services

This topic describes the firewall recommendations that you need to follow to be able to implement cloud services on your cluster.

Port that can be used is 8085 TCP

- Enables connections to cloud services nodes from all IBM Storage Scale nodes on this port. All communications from non-cluster nodes on this port can be blocked.
- cloud services nodes are required to communicate with the configured Object storage provider. Typically, this communication occurs over HTTPS (443) or HTTP (80). Contact your Object storage provider for more details.
- The internal port that is used by cloud services can be changed from 8085 to any other port by using the **mmcloudgateway config** command.

Table 37. Port requirements

Port number	Protocol	Service name	Component involved in communication
8085	TCP	Transparent cloud tiering	Intra cluster
Object storage provider dependent	TCP	Transparent cloud tiering	Transparent cloud tiering connection to Object storage provider on the external network. Typically HTTPS (443) or HTTP (80)

Note: For firewall recommendations for other components such as performance monitoring tool, protocol access, and GUI, see the *Securing the IBM Storage Scale system using firewall* topic in the *IBM Storage Scale: Administration Guide*.

Performance considerations

While default configurations on the IBM Storage Scale and cloud services work for the migration and recall of data, some adjustments might be required for optimized performance. This topic addresses some of such adjustments that will likely have the most significant effect on performance.

IBM Storage Scale cluster hardware

Memory, CPU, storage, and network play a significant role in determining how well the cloud tiering is going to perform.

Memory: cloud services requires at least 2 GB RAM for the heap memory, and the IBM Storage Scale page pool should also be reasonably high. It is recommended to have a minimum of 64 GB RAM for page pool on the Gateway nodes.

CPU: cloud services can run more threads to transfer data if the system has more CPU cores available. For better performance, it is recommended to have a minimum of 24 or 32 CPU cores (the more, the better) on the cloud services gateway nodes.

Storage: The read/write performance from the IBM Storage Scale file system is one of the major factors that determine the migrate/recall performance. To achieve desirable performance, the read/write throughput should be more than the available network bandwidth to the cloud. Some of the general recommendations are:

- Have flash disks for storing metadata on the file system.
- If possible, have a separate flash-based file system to keep the cloud services database. This will ensure that the internal cloud service operations that can span the entire database (like reconcile) run in a reasonable amount of time.

Note: Adding flash storage helps particularly with scaling up the number of files that are supported by cloud services and is apparent when you get to the hundreds of millions of files migrated to the cloud. So, this recommendation on flash storage only applies if you are scaling to those kinds of numbers.

Network: It is recommended to use 10 GB Ethernet for the network interfaces that are used to transmit data to the cloud. In case of on-premise object store, the object store also should have 10 GB Ethernet. It is recommended to have a dedicated external links to the cloud storage and not to use these interfaces for intra-cluster communication.

Software: cloud services uses the Data Management Application programming Interface (DMAPI) of the IBM Storage Scale to read and write the data from the file system. As cloud services can migrate data only as fast as it might read from the file system, it is important to ensure that the read/write performance of the file system is good. Set the following key configuration parameters by using the **mmchconfig** command:

- `pagepool` – 16 GB
- `maxGeneralThreads` – 1280
- `dmapiWorkerThreads` – 64
- `maxFilesToCache` – 100k
- `workerThreads` – 1024

Another important IBM Storage Scale configuration is the file system block size. It is recommended to have the block size greater than the cloud services slice size (default 512 KB). This will ensure that you do not read from the disk more than once to fill in a single slice.

cloud services configurations

cloud services has configurable parameters, which can be adjusted based on the environment:

- `connector.server.migrate-threadpool-size` – This determines the number of internal threads that accept the migrate requests from the command line. The default value is 32, and is ideal for most of the situations where you have up to 32 CPU cores. If you have a gateway node with more than 32

CPU cores, it is advisable to increase the number to match the number of CPU cores that you have, so cloud services can process more migrate requests at once.

Note: This recommendation is for Intel servers. For Power LE servers, you can efficiently run two threads per CPU core, so adjust your thread count accordingly.

- `connector.server.recall-threadpool-size` – Normally you do not adjust this parameter. This determines the number of internal threads that accept recall requests from the command line. If you are sure of your workload, you can adjust the migrate and recall thread pool size. For example, if you have a very cold archive application and expect more migrates and a very limited number of recalls, then you could increase the migrate thread pool size and reduce the recall thread pool size.

For more information, see the *Tuning Cloud services parameters* topic in the *IBM Storage Scale: Administration Guide*.

Recommendations on number of Gateway nodes: While increasing the number of gateway nodes to migrate data can increase the total performance, remember that you must find an optimum limit. For example, if the network link from each node to the cloud is 10 Gb and overall read throughput of the file system is 2 GB/s on cluster, just two gateway nodes could achieve the maximum throughput possible. Depending on the file system read/write throughput and the network link, you must make an informed decision on what should be the optimal number of gateway nodes.

Recommendation on number of Cloud Storage Access Points (CSAPs): If you are going to a public cloud service such as the public IBM Cloud Object Storage (COS) service, there is no need to add in multiple access points since these services have a built-in load balancer. For on-premise object storage, you may choose to add in multiple Cloud Storage Access Points (CSAP) if you are not using a load balancer as follows:

- If there are multiple gateway nodes in the cluster, using more than one CSAP for a Cloud Account is a good idea. The number of requests that can be processed by a single Cloud Storage Access Point is limited, and multiple nodes hitting the same access point can cause the performance to drop. A general guideline is to have about as many CSAPs as the number of gateway nodes.

Mixing migrate and recall workloads

cloud services internally shares threads that carry out an operation on each slice of data, as well as threads that carry out the interactions with cloud object storage. Hence, running a huge workload of migrates and recalls simultaneously, would affect the total turn-around of each activity, compared to running migrates and recalls separately. cloud services expects migrates to work in batches, where policy identifies a set of cold files to be moved to a cloud storage tier. Recalls, however, are expected to be accessed on demand and are given higher priority than migrates. It is wise to schedule migration policies where possible during off-peak hours to avoid recall performance impacts.

Note: Bulk policy-based recall operations are supported, but not expected to be the common method for recalling files. Since recalls are generally expected to be on demand, recalls are prioritized ahead of migrates.

Distributing files using policies

IBM Storage Scale policies can be used to distribute the migrations and recalls to all the nodes in the node class. To effectively use the policy, refer to the following points:

- The **mmapplypolicy** takes two parameters for the number of parallel threads processing the files (-m) and number of files processed per thread (-b).
 - -m has a default value of 24. If you have more CPU cores available, it is advisable to adjust the -m value accordingly
 - -b has a default value of 100. You can increase the number if you want to process more files in one iteration.
 - Plan the -m and -b carefully, so that all the nodes are given equal number of files. For example, if you have 5000 files to migrate and 4 nodes with cloud services installed, the default values of -m 24 and -b 100 will try to distribute 2400 files to each node. There will not be enough files to be distributed to all nodes, and two nodes will be idle.

- -N parameter of the **mmapplypolicy** will accept the node class name, so the policy will distribute the files to all the nodes in the node class.

Using multiple file system or file sets for the same node group

A customer can configure multiple file systems and multiple file sets for a single node group, and run migrates in parallel for all of them. If they are looking for high performance, it is recommended that they configure the policy in a manner that, migrates run on a single file system or file set at a time. cloud services design uses a single server process, and resources such as memory, networking and threading are shared between the containers. Running parallel migrates can split the resource, and hence the effective throughput of a single container can be impacted.

Migration policy runs should not be done during the weekly and daily maintenance windows as migrations conflict with some activities that lock the Cloud directory for long periods of time.

Container spillover

For achieving better performance and ease of maintenance, it is recommended to keep no more than 100 million entries per container. After that, the customer must create a new container for the same file system/fileset and continue migrations. This will ensure that a new database is created for the new container, and the container cloud database size does not grow too large. Operation such as reconcile for a 100 million file database can take a few hours. Growing the size much beyond that results in maintenance operations taking too long since the duration of these operations relative to database size is non-linear.

Note:

The automatic container spillover (auto-spillover) is enabled by default during container creation. Auto-spillover is applicable for containers used with tiering service only. Auto-spillover is not applicable for containers used with sharing service. After reaching the default 100 million entries in a container, as an administrator you need not manually create the new container for the same file system anymore. When number of entries in an Active container reaches the given threshold value, a new spillover container is automatically created by the Cloud service during the next reconcile operation on container. This container is made an Active container for the configured file system or filesset, for subsequent data migrations. An index value is used as a suffix for new container name. For example, when a container named `testContainer` reaches spillover threshold value, a new container named `testContainer001` is created automatically by the Cloud service. The next spillover index can be seen in the output of **mmcloudgateway containerpairset list**.

Container level auto-spillover settings can be tuned using **mmcloudgateway containerpairset create / update**. More details of auto-spillover CLI options are available under **mmcloudgateway help**. You can create a new container for the given file system at any time, even before reaching configured threshold entries per container. Auto-spillover occurs only when the Cloud service finds an active container crossing the configured threshold entries.

Object storage configurations

Use of a load balancer

Object stores such as IBM Cloud Object Storage (formerly known as Cleversafe®) gives you an option to have more than one node that can accept the requests from the client. While using a single CSAP, cloud services can send the requests to only one cloud end-point URL. You can overcome this restriction by using a load balancer to distribute the requests to multiple endpoints.

- NGINX(<https://www.nginx.com/>) - IBM Cloud Object Storage recommends the use of NGINX with COS, and has a solution guide, which can be used to configure NGINX with COS. For most customers, this load balancer does not perform as well as the native load balancer provided by cloud services.
- HAProxy (<http://www.haproxy.org/>) can be configured on a system, and used to send the requests to more than one end point
- You can use a simple DNS-based load balancer where each DNS lookup for a host name returns an IP address for a different end-point. The host name can be used to create an account with the cloud services.

If you would prefer that cloud services distribute the requests to multiple cloud storage endpoint URLs, you can configure more than one CSAP, and cloud services uses a very basic load balancing algorithm to distribute load across those CSAPs.

IBM Cloud Object Storage

For on-premise configurations, it is recommended to have multiple accessors setup, so that a load balancer or cloud services itself will have no single point of failure in its cloud services with COS. For public IBM COS service, there is a built-in load balancer – no additional CSAP or load balancer work is necessary. For public IBM COS service, there is a built-in load balancer, so no additional CSAP or load balancer work is necessary.

Amazon S3

Use the appropriate supported Amazon S3 region while configuring the cloud account. Load balancing or multiple CSAPs are not necessary as Amazon S3 has a built-in load balancer.

Security considerations

You can integrate cloud services with IBM Security Key Lifecycle Manager (ISKLM) to provide security to the data that is stored on the cloud storage tier, or you can use the native key manager provided with the cloud services.

transparent cloud tiering: For information on integration of ISKLM with transparent cloud tiering, see *Configuring cloud services with SKLM in IBM Storage Scale: Administration Guide*.

Note: Ensure that you back up your security keys by using the **mmcloudgateway service backupConfig** command. Data encrypted by using the cloud services cannot be recovered, if security keys are lost.

cloud data sharing: cloud data sharing currently supports importing any cloud data, assuming it is not read in an encrypted format. cloud data sharing supports encrypted data only if it was exported by cloud data sharing and the encryption keys are shared between the importer and the exporter. For this reason, it is recommended that the native cloud services and cloud data sharing encryption are used to provide encryption at rest. A secure connection should be used to transfer data between IBM Storage Scale and the cloud. Thus, data is stored encrypted on IBM Storage Scale, and stored encrypted on the cloud, and is secured by the connection when transferring between the two systems.

Planning for maintenance activities

cloud services has five key maintenance activities that need to take place. Some are for keeping normal day-to-day operations going and some are contingencies for service restoration in case of a disaster.

The following two activities are not automated and must be addressed by the administrator as follows:

- Backing up the key library manager - a new copy needs to be made every time there is any change to a key in the key manager. For more information, see the *Backing up the Cloud services configuration* topic in the *IBM Storage Scale: Administration Guide*.

Note: If the key manager is lost, there is no backdoor or secret IBM internally known way to recover the data from cloud storage. So, it is important to have a backup copy of the key manager.

- cloud services leverages SOBAR as a backup mechanism of transparent cloud tiering metadata that can be used to restore the transparent cloud tiering service in case of a failure. A sample script is provided in the transparent cloud tiering directory that can be deployed to run the backups. For more information, see the *Scale out backup and restore (SOBAR) for Cloud services* topic in the *IBM Storage Scale: Administration Guide*.

The following three activities are provided for by an automated cloud services maintenance service:

- Background removal of deleted files from the object storage. This is recommended to be done daily.
- Backing up the cloud services full database to the cloud. This is recommended to be done weekly.

- Reconciling the transparent cloud tiering database. This is recommended to be done every four weeks.

The transparent cloud tiering maintenance service comes with default daily (at night) and weekly (on the weekend) maintenance windows where these activities run. You can use the **mmcloudgateway maintenance status** command to query what your current maintenance window is. It is important to avoid running migration during these activities so you need to make sure your migration policies run outside of those windows. There are scaling considerations when setting up your maintenance windows. They must be long enough so that maintenance activities fit inside the windows. Consider the following guidelines when you plan these activities:

- Reconcile is by far the longest activity and the main one to consider when you plan your service windows. A reconcile for every 100 million file container (which is the default spillover value) takes a few hours if you run metadata with flash. If you run it from a disk, it takes more like 6 to 12 hours.
- Each transparent cloud tiering node can run one service activity at a time during the maintenance window. For example, if you have three transparent cloud tiering nodes, you might be running three maintenance activities in parallel at the same time.
- Keep in mind that maintenance activities run to completion even after the maintenance window has passed. The longest duration outside the maintenance window would be an activity that was scheduled and started just prior to the maintenance window closing. That activity would run almost entirely after the maintenance window had completed.

Here is an explanation of how the maintenance window size affects the ability to scale number of transparent cloud tiering files. If you have a weekly maintenance window of 24 hours, each node would be able to process two-to-three 100 million file containers per week. Since reconcile maintenance needs to be done every 4 weeks, it would follow that a 24-hour maintenance window can accommodate eight-to-twelve 100 million file containers per node every week. With a full 4-node setup, this adds up to 32 - 48 containers (roughly, that is support for 3-5 billion files) per 24-hour maintenance window for a single node group.

It follows that if you want to use transparent cloud tiering with more than this, you are likely going to want to consider putting the IBM Storage Scale metadata and the transparent cloud tiering database in flash storage. This will greatly increase the number of files the maintenance window can handle.

For setting up maintenance activities, see the *Setting up maintenance tasks* topic in the *IBM Storage Scale: Administration Guide*.

Backup considerations for Transparent cloud tiering

You must adhere to some guidelines for backing up files that need to be migrated to a cloud storage tier.

In many cases, Transparent cloud tiering (cooler data) might be run along with file system backups (for hot and warm data) for data protection. Transparent cloud tiering should not be used as a replacement for a valid backup strategy. Files that are migrated with Transparent cloud tiering might be lost if accidentally deleted or if there is a complete site failure. File system backups can protect from these types of events.

When you work with Transparent cloud tiering file system backups, it is important that a file is backed up before it is sent to a cloud storage. A file that is not backed up is recalled from the cloud storage tier when a backup program is run. This can cause unnecessary recall traffic, and also cause the backup to take longer to run. Recalling data also involves cost if you are using a public cloud provider.

To ensure that files are not migrated before they are backed up, it is best to define a migration policy that explicitly excludes files that are recently modified. For example, if nightly backups are run, at minimum, a migration policy should exclude files that are modified within the past day, so they can be backed up before they are migrated. In most cases, it is best to wait longer to avoid recalls if a backup window is missed for some reason. For example, if you run a nightly backup, it is advisable to exclude files that are modified in the past 3 days. If there is an issue with running backups for a long period, it might be necessary to disable a migration policy until the problem with backups is resolved. This avoids many recalls during the next successful backup.

The following is a sample migration policy and it should be tuned as per the backup frequency that is configured for an IBM Storage Scale file system.

Note: The amount of time to wait to migrate a file should be greater than the frequency of file system backups.

For information on backing up configuration data, see *Backing up the Cloud services configuration* topic in *IBM Storage Scale: Administration Guide*.

For information on backing up the cloud database, see *Backing up the Cloud services database to the cloud* topic in *IBM Storage Scale: Administration Guide*.

You can exclude recently modified files with a policy statement such as the following:

```
WHERE (DAYS(CURRENT_TIMESTAMP) - DAYS(MODIFICATION_TIME)) > 5
```

This statement can also be added for manual application of the policy for a directory.

```
/* Define a cloud pool, using the 'mmcloudgateway' command to perform migrations to the pool */
RULE EXTERNAL POOL 'cloudpool' EXEC '/usr/lpp/mmfs/bin/mmcloudgateway files' OPTS '-F'

/* Define a migrate rule, in this example we will migrate from the system pool */
RULE 'MigrateToCloud' MIGRATE FROM POOL 'system'
/* Define the threshold to use, in this example we will migrate when the pool is 85% full,
and attempt to migrate files until the pool is 75% full */
THRESHOLD (85, 75)
/* Next, define a weight for the file to determine which files to migrate first */
/* Here we will use the last access time for the file by subtracting it from the current time,
giving a higher weight to older files */
WEIGHT(CURRENT_TIMESTAMP - ACCESS_TIME)
/* Define that we want to move to the cloud pool */
TO POOL ('cloudpool')
/* We don't want to migrate files that are smaller than 8k in size, because there are little to
no space savings due to file overhead */
WHERE (KB_ALLOCATED > 8)
/* Do not migrate cloud services internal files, which reside in the .mcstore directory */
AND NOT(PATH_NAME LIKE '%/.mcstore/%')
/* Ignore files that have been modified in the past 5 days. This needs to be customized based on
the backup frequency. */
AND (DAYS(CURRENT_TIMESTAMP) - DAYS(MODIFICATION_TIME)) > 5
```

Quota support for tiering

transparent cloud tiering supports container-level quotas in ways that are described here.

IBM Storage Scale does not account for data that transparent cloud tiering has migrated to the Cloud for quota support.

- Files in non-resident states do not count against quota as defined by IBM Storage Scale.

Object storage generally supports quotas and might be used to provide quota support.

- For example, IBM Cloud Object Storage supports hard quotas at the granularity of a container.

General guidelines for quota support:

- Use fileset support – this provides the level of granularity you need for cloud storage quotas to be useful.
- Set the desired quota for each fileset by using the object storage quota setting for the container.
- On container spillover, the size for each container in the fileset needs to be accounted for.
 - Since all but the most recent container will no longer be growing, you can set the quota for the most recent (and active) container based on the remaining capacity after subtracting for the current size of the older containers.
 - If older containers have a significant number of files deleted from them, the quota for the active container might need to be periodically adjusted upwards.

- A quota support script is available in the samples directory. that provides a report of space taken in the cloud for a given filesystem. The report does not include space allocated for files that are deleted or are overwritten with a new version that are being retained for the specified retention period.

Note: For more information on setting quota at a container level, see documentation for the specific cloud account that you use in your environment.

Client-assisted recalls

Non-gateway nodes can have an extended client which can service the recalls locally.

You must install a transparent cloud tiering server package on a non-gateway node.

To enable recalls from a non-gateway node, run **mmcloudgateway clientassist enable** on the non-gateway node where the package is installed. If a node is client-assist enabled, the transparent cloud tiering server nodes might still be called upon for transparent recalls if the node gets overloaded with transparent recall requests with a full queue.

All recalls originating from this node are serviced from this node only, without routing through the gateway node.

To disable client assisted recalls, run **mmcloudgateway clientassist disable**.

Planning for IBM Storage Scale on public clouds

Learn about the planning considerations for using **cloudkit** on public clouds.

After going through the planning considerations, see the following sections for installation, configuration, cloudkit command options, and cloudkit issues:

- *Installing IBM Storage Scale on public cloud with the cloudkit* in the IBM Storage Scale: Concepts, Planning, and Installation Guide.
- *cloud kit command* in the IBM Storage Scale : Command and Programming Reference Guide.
- *Cloudkit issues* in the IBM Storage Scale : Problem Determination Guide

Note: Currently, the **cloudkit** supports to deploy an IBM Storage Scale cluster on the public cloud of either Amazon Web Services (AWS) or Google Cloud Platform (GCP). For more information, see *Supported Features of cloudkit* in IBM Storage Scale: Concepts, Planning, and Installation Guide.

IBM Storage Scale on cloud: Support statement

IBM Storage Scale deployments and functions are supported on a variety of clouds, directly through IBM, and also through Business Partners who have crafted specific solutions.

Recommended: Use of the IBM Storage Scale **cloudkit** for deployments is highly recommended, because each option and configuration parameter has been vetted against best practices, thus resulting in a known and workable cloud configuration. If the **cloudkit** does not support your use cases, refer to the alternate methods for provisioning.

Alternatively: Manually provisioning cloud resources and deploying IBM Storage Scale without using the **cloudkit** is also supported on the following clouds.

- Amazon Web Services (AWS), Azure, IBM Cloud, Google Cloud Platform (GCP)

Restriction: IBM Storage Scale consists of many features, some of which are highly useful and relevant for cloud environments, while other features are not recommended in a cloud environment. A few examples:

- SMB, NFS, S3 protocols using Cluster Export Services (CES stack). The CES stack relies upon existence of an assigned IP pool (non-DHCP) that can freely rotate across nodes. The IBM Storage Scale CES stack controls assignment and rotation in cases of node failure. In a cloud environment, IPs can only be assigned and moved by using cloud-specific APIs; which means that the IBM Storage Scale CES stack is currently incompatible with cloud APIs without further development.

- AFM home (using NFS protocol). See the previous CES stack explanation, because of the CES NFS protocol is used for AFM home.
- Cloud block disk services are a feature of every cloud and provide block access to storage. IBM Storage Scale uses disk services to create NSDs for building the IBM Storage Scale file system. For compatibility to be ensured, these disk services must be fully vetted with IBM Storage Scale. An inclusion list of services specifically compatible with IBM Storage Scale is shown in the supported features matrix for each cloud.
- Custom Kubernetes services are a feature of every cloud. IBM Storage Scale supports manually installed Vanilla Kubernetes with its CSI driver. Custom services such as EKS and GKS are not supported by the CSI driver.
- OpenShift is available in every cloud as either a self-managed instantiation or a bundled service. Refer to the IBM Storage Scale Container Native Storage Access (CNSA) documentation for a list of supported clouds for OpenShift.
- Virtualization technologies on the cloud: VMware/KVM. The IBM Storage Scale FAQ contains a compatibility matrix of features and functions supported in VMware/KVM. This same matrix applies to VMware/KVM on the cloud.
- For other exceptions, as well as lists of supported features, see the supported features matrix of IBM Storage Scale for each cloud.

If you choose not to use the **cloudkit**, it is your responsibility to make sure that you adhere to all restrictions and best practices in the rest of the IBM Storage Scale documentation.

If you are using a business partner’s cloud deployment of IBM Storage Scale, contact the business partner for support statement details.

Planning for deployment patterns on the cloud

Learn about the different cloud deployment patterns that IBM Storage Scale supports.




Figure 42. Cloud deployment modes

Cloudkit offers various deployment modes that help the users to select from among various storage configuration capabilities (like persistence storage, scratch storage, single AZ, multi-AZ), based on their workload preferences for IBM Storage Scale deployment on cloud.

Performance Persistence Storage mode

In this mode, persistence storage is used so that both IBM Storage Scale NSDs and data persist even after a shutdown. The supported Instance profile option is available for users to select their preferred instance profile, which is then deployed on a single available zone. Cloudkit automatically

calculates the number of disks per instance that is needed to saturate instance bandwidth, and also calculates the number of instances that are needed to fulfill the file system capacity. *This mode is recommended for novice users.*

Performance Advanced Storage mode

In this mode, persistence storage is used so that both IBM Storage Scale NSDs and data persist even after a shutdown. The supported Instance profile option is available for users to select their preferred instance profile, which is then deployed on a single available zone. Cloudkit provides to the users the option to select disk and instance count. *This mode is meant of storage capacity rather than performance; and it is recommended mostly for advanced users.*

Balanced mode

In addition to include all features that are offered for advanced users, this mode supports deployment of IBM Storage Scale on multiple (3) availability zones, and supports configuration of the IBM Storage Scale file system in such a way that data and metadata get replicated across availability zones. *This mode is recommended for high-availability workloads.*

Performance Scratch Storage mode

This mode uses single availability zone and locally attached SSD disks with placement group. This mode enables workloads to achieve the low-latency network performance that is necessary for the tightly coupled node-to-node communication that is typical of high-performance computing (HPC) applications. The data that is stored in this mode is volatile and can be lost if the instance is stopped or ended. Therefore, it is recommended to make frequent backups. This mode must not be used for long-term storage or if the data is not backed up elsewhere. *This mode is recommended for high-transient workloads.*

Planning for resources creation

Cloudkit offers a composable deployment architecture that provides flexibility for cloud administrators to pick and choose resource creation through **cloudkit**. These resources can be created either from scratch or based on an existing resource of the IBM Storage Scale deployment.

Cloud administrators can create the following types or resources by using the **cloudkit** command:

- Repository
- Network
- Jump host
- Domain name system (DNS)
- Image
- Storage cluster
- Compute cluster
- Combined cluster

Preparing the installer node

The installer node refers to the node where the cloudkit is being run.

The installer node can be of one of the following types:

- **On-premise host:** A physical host or virtual machine that is located on the premises. A reliable and secure network connection must exist between the installer node and the cloud environment. The cloudkit uses this connection to communicate with the cloud infrastructure to provision and configure the required resources.
- **Cloud VM:** A virtual machine that is located on the cloud.

The cloudkit execution is currently only supported on Red Hat Enterprise Linux (RHEL). Therefore, the installer node must run any one of these operating systems. For information about the supported RHEL versions, see [Table 50 on page 604](#).

It is essential that this machine has access to the base Red Hat repositories.

Note: The installer node stores critical information regarding the IBM Storage Scale cluster on the cloud, therefore, it is crucial to ensure its security. In addition, this critical information that is stored in the user's home directory (/<home_directory>/scale-cloudkit/data) must be backed up to ensure that it is not lost.

Extracting the cloudkit binary

The cloudkit rpm binary is bundled as a part of IBM Storage Scale self-extraction (SE) package. For more information, see [“Extracting the IBM Storage Scale software on Linux nodes” on page 473](#).

By default, the **cloudkit rpm** is extracted from the SE package into the /usr/lpp/mmfs/<release_version>/cloudkit_rpms directory.

Installation of cloudkit package for Red Hat Enterprise Linux (RHEL)

To install IBM Storage Scale cloudkit packages, change the directory to where the installation package is extracted. For example, run this command:

```
cd /usr/lpp/mmfs/<x.x.x.x>/cloudkit_rpms
```

Then, run the following command:

```
yum -y install gpfs.cloud-deployer*.rpm
```

Now, the cloudkit is extracted into the /usr/lpp/mmfs/<release_version>/cloudkit directory.

Note: The **cloudkit rpm** was introduced with IBM Storage Scale 6.0.0 release. For versions earlier than IBM Storage Scale 6.0.0, the cloudkit binary is directly extracted from the SE package and placed in the /usr/lpp/mmfs/<release_version>/cloudkit directory.

Preparing the cloudkit environment file

The **cloudkit** encrypts all data that is stored inside it by using a key that is stored in a passphrase file (cloudkit_config.ini), which is provided by the user; the owner of the passphrase file is granted read and write permission. The cloudkit_config.ini file must be created in the installer node and follow the next format:

Note: Ensure that this passphrase file is backed-up to a secure location.

```
[default]
cloudkit_key = <PassPhrase>
```

During the cloudkit initialization, by default, feature_gate = Stable gets set. This is recommended for production clusters and lists only the generally available features. To enable *technical preview features* on *test clusters*, you need to change this setting to: feature_gate = TechPreviewNoUpgrade, and then run the **cloudkit init** command.

```
[default]
cloudkit_key = <PassPhrase>
feature_gate = TechPreviewNoUpgrade
```



Warning: Enabling the TechPreviewNoUpgrade feature set on your cluster prevents future upgrades and it cannot be undone. You should *not* enable this feature set on production clusters.

Example:

In this example, the passphrase file is created at /usr/lpp/mmfs/<release_version>/cloudkit/cloudkit_config.ini with the following content:

```
# cat .env
[default]
cloudkit_key = XXXXXXXX
```

This passphrase file path is required during **cloudkit init**:

1. To export, use `export env file`, with the syntax `export CLOUDKIT_PASSPHRASE_FILE_PATH<ENV_FILE_PATH>`:

```
export CLOUDKIT_PASSPHRASE_FILE_PATH=/usr/lpp/mmfs/<release_version>/cloudkit/  
cloudkit_config.ini
```

2. To configure, run the **cloudkit init** command:

```
# ./cloudkit init  
I: Logging at /root/scale-cloudkit/logs/cloudkit-06-12-2024_0-2-45.log  
? Passphrase file path for encrypting DB contents: /usr/lpp/mmfs/6.0.1.0/cloudkit/  
cloudkit_config.ini
```

Planning for AWS cloud

Learn about prerequisites, configurations, and other important planning information that you must consider before the deployment of an IBM Storage Scale cluster on AWS.

Consider the following steps when you are planning for an IBM Storage Scale cluster on the public cloud of Amazon Web Services (AWS):

1. Preparing the AWS environment to deploy the IBM Storage Scale cluster.
2. Planning the virtual network cloud (VPC) architecture.
3. Planning for DNS.
4. Planning for bastion.
5. Creating an IBM Storage Scale Amazon Machine Image (AMI) in advance.
6. Planning for the IBM Storage Scale deployment architecture on AWS.
7. Planning IBM Storage Scale cluster deployment profiles.
8. Determining your performance, scalability, data availability, and data protection requirements.
9. Planning cloudkit notifications.
10. Planning for encryption at rest.

Preparing the AWS environment

Complete the following steps to prepare the AWS environment:

1. Create an account in [AWS](#), if you do not already have an account.
2. Create the necessary IAM users along with the access key required for cloudkit to make programmatic requests to the AWS API. For more information, see [Managing access keys for IAM users](#).
3. Use the region selector in the navigation bar to choose the AWS region to deploy the IBM Storage Scale cluster.
4. Create a key pair in the preferred region. For more information, see [Amazon EC2 Key Pairs](#).
5. Verify that the wanted AWS region and availability zone support the EC2 instance types that are chosen to be provisioned for the IBM Storage Scale storage and compute nodes. You can verify the AWS EC2 instance type from **AWS console > EC2 Service Page > Instance Type**.
6. Verify the AWS EC2 quota limits and make sure that enough quotas exist for the EC2 instance types that you intend to deploy. Verify the AWS EC2 quota limits from **AWS console > EC2 Service Page > Limits**.

If necessary, request a service limit increase for the EC2 instance types that you intend to deploy from [AWS website](#). To request a service limit increase, in the AWS Support Center, choose **Create Case > Service Limit Increase > EC2 instances**, and then complete the fields in the **Limit Increase** form.

7. Prepare the installer node where the cloudkit is meant to run. For more information, see [“Preparing the installer node” on page 421](#).

Planning the virtual private cloud (VPC) architecture for AWS

When resources are being deployed in the cloud, the cloudkit can either create a new virtual private cloud (VPC) and provision the resources into it, or use a previously created VPC.

When the cloudkit creates a new VPC, the cloudkit designs your network infrastructure from scratch. It chooses the subnets, network address ranges, and security groups that best suit the IBM Storage Scale deployment.

A VPC can be created by using the **cloudkit create network** command, which creates only the VPC. Or you can use the **cloudkit create cluster** command.

The cloudkit can also deploy resources in a previously created Virtual Private Cloud (VPC). For information about creating a VPC, see [AWS documentation](#)

```
? VPC Mode:[Use arrows to move, type to filter]
> New
Existing
```

When the IBM Storage Scale cluster is to be deployed in an existing VPC, the following requirements must be met:

- This is a mandatory step. DHCP options set or DNS configured.
- This is a mandatory step. Private subnets (with allocatable IP address) in the availability zone (minimum 1 private subnet per availability zone).
- This is an optional step. A public subnet (a subnet with internet gateway attached) is only required if either the following cases are met:
 1. Where the cluster is planned to be accessed through a jump host. For more information, see [“Other considerations”](#) on page 448.
 2. If it is necessary to precreate an IBM Storage Scale AMI. For more information, see [“Precreating an IBM Storage Scale AMI”](#) on page 425.
- This is an optional step. A NAT gateway attached to the private subnet is only required when the IBM Storage Scale instances need access to the internet (for operating system updates and security patches, and so on).

Planning for DNS

You can create a DNS domain by using the **cloudkit create dns** command.

This command can create a new DNS zone or use an existing DNS zone:

```
? Do you wish to use an existing DNS zone: (y/N)
```

To consult the available deployment options, issue the **cloudkit create cluster** command:

Remember: Creation of a cloud DNS is mandatory and must be created before you run the **cloudkit create cluster** command.

```
E: No existing DNS zones found. In order to configure, use 'cloudkit create dns'.
```

Planning for bastion

Use the **cloudkit create jumphost** command to create a jump host or bastion in the public subnet.

To consult the available deployment options, issue the **cloudkit create cluster** command:

```
? Bastion OS: RHEL-9.2.0_HVM-20230905-x86_64-38-Hourly2-GP2 | ami-05a5f6298acdb05b6 | ec2-user
? Bastion instance type: t3.small | vCPU(1) | RAM (2.0 GiB) | CPU Credits/hr (12)
? Key pair to be used for launching Bastion/Jumphost host instance(s): prvnkeypair
? Bastion/Jumphost SSH private key file path (will be used only for configuration): /home/
```

```
user1/.ssh/id_rsa
? Bastion CIDR allow list: xxx.xxx.xxx.xx/32
```

The cloudkit provides the following options for VM instances types for a jump host:

> t3.micro	vCPU(1)	RAM (1.0 GiB)	CPU Credits/hr (6)
t3.small	vCPU(1)	RAM (2.0 GiB)	CPU Credits/hr (12)
t3.medium	vCPU(2)	RAM (4.0 GiB)	CPU Credits/hr (24)
t3.large	vCPU(2)	RAM (8.0 GiB)	CPU Credits/hr (36)
m4.large	vCPU(2)	RAM (8.0 GiB)	Dedicated EBS Bandwidth (450 Mbps)
m4.xlarge	vCPU(4)	RAM (16 GiB)	Dedicated EBS Bandwidth (750 Mbps)
m4.2xlarge	vCPU(8)	RAM (32 GiB)	Dedicated EBS Bandwidth (1000 Mbps)
m4.4xlarge	vCPU(16)	RAM (64 GiB)	Dedicated EBS Bandwidth (2000 Mbps)

Precreating an IBM Storage Scale AMI

When the cloudkit is used to create an IBM Storage Scale cluster by using the `cloudkit create cluster` command, it can either automatically create a stock Amazon Machine Image (AMI) or the customer can provide a previously created custom image.

The cloudkit can automatically create an image by using Red Hat Enterprise Linux (RHEL) 8.10 or 9.4. When the option to use a stock image is chosen, the cloudkit automatically performs all required actions to create an IBM Storage Scale VM image and create the IBM Storage Scale cluster.

A custom image can be created by using the `cloudkit create image` command. The `create image` command can accept inputs in the form of an existing image or provides the ability to create an image from a Red Hat OS image.

A custom existing image consists of:

1. A Red Hat version that is supported by IBM Storage Scale.
2. Optionally, for any customer applications that are already preinstalled, the `cloudkit create image` command installs all required IBM Storage Scale packages on top of the existing image.

Planning for the IBM Storage Scale deployment architecture on AWS

Before the IBM Storage Scale cluster can be created on the cloud, the deployment architecture must be planned.

The cloudkit offers the following deployment models for IBM Storage Scale clusters:

- **Combined-compute-storage**

A unified IBM Storage Scale cluster with both storage and compute nodes. It is recommended that any customer workload runs only on the compute nodes.

- **Compute-only**

An IBM Storage Scale cluster that only consists of compute nodes. This cluster does not have any local file system and therefore must remote mount the file system from an IBM Storage Scale storage cluster.

- **Storage-only**

This IBM Storage Scale cluster consists of only storage nodes, which have access to storage where the file system is created. This file system is remote mounted to any number of compute-only IBM Storage Scale clusters.

When you run the `cloudkit create cluster` command, the following prompt asks the user to select from among the different deployment models that are offered:

```
? IBM Spectrum Scale deployment model:[Use arrows to move, type to filter, ? for more help]
> Storage-only
Compute-only
Combined-compute-storage
```

Deployment purpose

The deployment purpose is applicable only to storage clusters.

- **Non-production**

This profile lists instance types that are experimental (best suited for proof of concepts) and may require tuning for running scale-suited workloads.

- **Production**

This profile lists instance types that are well suited for scale workloads. Optimal tuning parameters are calculated and set during the deployment.

When the next prompt appears, select one of the supported deployment purposes.

```
? IBM Storage Scale deployment model: Storage-only
? Deployment purpose: [Use arrows to move, type to filter]
Non-Production
> Production
```

IBM Storage Scale cluster deployment profiles

Cloudkit offers the following deployment profiles:

- **Throughput-Performance-Persistent-Storage**

This profile uses a single availability zone with persistent storage, which means that the file system device retains data after the instance is stopped. In this mode, the cloudkit calculates the number of storage nodes based on the provided file system capacity. This mode uses gp3 as disk type.

Important: When you choose this profile, the file system is configured in such manner that the data is not replicated, only the metadata gets replicated across the instances.

- **Throughput-Performance-Scratch-Storage**

This profile uses a single availability zone and a placement group with `cluster` policy, which means that it packs instances close together inside an availability zone. This strategy enables workloads to achieve the low-latency network performance necessary for tightly coupled node-to-node communication that is typical of high-performance computing (HPC) applications, with instance or temporary storage, meaning that the file system device loses data after the instance is stopped.

In this mode, the number of storage instances is limited to 10; compute instances are limited to 65.

Important: This profile uses local NVMe SSD or instance storage, which offers high performance and low latency for data-intensive workloads. However, the data that is stored in this mode is volatile and can be lost if the instance is stopped or terminated. Therefore, it is recommended to take frequent backups.

This profile must not be used for long-term storage or if the data is not backed up elsewhere.

- **Throughput-Advance-Persistent-Storage**

This profile uses a single availability zone with persistent storage, which means that the file system device retains data after the instance is stopped. This mode is meant for storage capacity rather performance; the number of storage instances is limited to 64, compute instances is limited to 65. This profile offers these disk types to choose from: gp2, gp3, and io1.

Important: When you choose this profile, the file system is configured in such manner that the data is not replicated, only the metadata gets replicated across the instances.

- **Balanced**

This profile uses multiple (3) availability zones to deploy the IBM Storage Scale cluster into. In this mode, the number of storage instances is limited to 64. It offers a choice of disk types gp2, gp3, and io1.

Important: The IBM Storage Scale file system is configured in such a way that data and metadata get replicated across availability zones.

Instances are spread across the first two availability zones that are specified in the selection, and the tie-breaker instance gets provisioned in the three availability zones that are specified during the selection.

```
? Tuning profile: [Use arrows to move, type to filter, ? for more help]
Throughput-Performance-Scratch-Storage
> Throughput-Performance-Persistent-Storage
Throughput-Advance-Persistent-Storage
Balanced
```

Determining your performance, scalability, data availability, and data protection requirements

Before deploying the IBM Storage Scale cluster, it is important to understand your requirements in terms of performance, scalability, data availability, and data protection. These criteria determine what AWS instance types to use for the storage nodes and compute nodes, also what elastic block store types should be used.

The cloudkit provides the following choices for VM instances types.

For IBM Storage Scale compute nodes, all the VM instance types are supported, according to their availability per region.

For IBM Storage Scale storage nodes, the support depends on the profile, as described in the following list.

- For the **Throughput-Performance-Persistent-Storage** and **Balanced** profiles, the next VM instance types are supported.

m6in.2xlarge to 25 Gbps	8 vCPU	32 GiB RAM	Network Bandwidth Up to 40 Gbps	EBS Bandwidth Up
m6in.4xlarge to 25 Gbps	16 vCPU	64 GiB RAM	Network Bandwidth Up to 50 Gbps	EBS Bandwidth Up
m6in.8xlarge Gbps	32 vCPU	128 GiB RAM	Network Bandwidth 50 Gbps	EBS Bandwidth 25
m6in.12xlarge 37.5 Gbps	48 vCPU	192 GiB RAM	Network Bandwidth 75 Gbps	EBS Bandwidth
m6in.16xlarge Gbps	64 vCPU	256 GiB RAM	Network Bandwidth 100 Gbps	EBS Bandwidth 50
m6in.24xlarge Gbps	96 vCPU	384 GiB RAM	Network Bandwidth 150 Gbps	EBS Bandwidth 75
m6in.32xlarge Gbps	128 vCPU	512 GiB RAM	Network Bandwidth 200 Gbps	EBS Bandwidth 100
c6in.2xlarge to 25 Gbps	8 vCPU	16 GiB RAM	Network Bandwidth Up to 40 Gbps	EBS Bandwidth Up
c6in.4xlarge to 25 Gbps	16 vCPU	32 GiB RAM	Network Bandwidth Up to 50 Gbps	EBS Bandwidth Up
c6in.8xlarge Gbps	32 vCPU	64 GiB RAM	Network Bandwidth 50 Gbps	EBS Bandwidth 25
c6in.12xlarge 37.5 Gbps	48 vCPU	96 GiB RAM	Network Bandwidth 75 Gbps	EBS Bandwidth
c6in.16xlarge Gbps	64 vCPU	128 GiB RAM	Network Bandwidth 100 Gbps	EBS Bandwidth 50
c6in.24xlarge Gbps	96 vCPU	192 GiB RAM	Network Bandwidth 150 Gbps	EBS Bandwidth 75
c6in.32xlarge Gbps	128 vCPU	256 GiB RAM	Network Bandwidth 200 Gbps	EBS Bandwidth 100

- For the **Throughput-Performance-Scratch-Storage** profile, the next VM instance types are supported.

i3en.2xlarge Bandwidth Up to 25 Gbps	8 vCPU	64 GiB RAM	Instance Storage 2 x 2500 NVMe SSD	Network
i3en.3xlarge Bandwidth Up to 25 Gbps	12 vCPU	96 GiB RAM	Instance Storage 1 x 7500 NVMe SSD	Network
i3en.6xlarge Bandwidth 25 Gbps	24 vCPU	192 GiB RAM	Instance Storage 2 x 7500 NVMe SSD	Network
i3en.12xlarge Bandwidth 100 Gbps	48 vCPU	384 GiB RAM	Instance Storage 4 x 7500 NVMe SSD	Network
i3en.24xlarge Bandwidth 100 Gbps	96 vCPU	768 GiB RAM	Instance Storage 8 x 7500 NVMe SSD	Network

- For the **Throughput-Advance-Persistent-Storage** profile, the next VM instance types are supported.

m6in.2xlarge	8 vCPU	32 GiB RAM	Network Bandwidth Up to 40 Gbps	EBS Bandwidth Up to 25 Gbps
m6in.4xlarge	16 vCPU	64 GiB RAM	Network Bandwidth Up to 50 Gbps	EBS Bandwidth Up to 25 Gbps
m6in.8xlarge	32 vCPU	128 GiB RAM	Network Bandwidth 50 Gbps	EBS Bandwidth 25 Gbps
m6in.12xlarge	48 vCPU	192 GiB RAM	Network Bandwidth 75 Gbps	EBS Bandwidth 37.5 Gbps
m6in.16xlarge	64 vCPU	256 GiB RAM	Network Bandwidth 100 Gbps	EBS Bandwidth 50 Gbps
m6in.24xlarge	96 vCPU	384 GiB RAM	Network Bandwidth 150 Gbps	EBS Bandwidth 75 Gbps
m6in.32xlarge	128 vCPU	512 GiB RAM	Network Bandwidth 200 Gbps	EBS Bandwidth 100 Gbps
c6in.2xlarge	8 vCPU	16 GiB RAM	Network Bandwidth Up to 40 Gbps	EBS Bandwidth Up to 25 Gbps
c6in.4xlarge	16 vCPU	32 GiB RAM	Network Bandwidth Up to 50 Gbps	EBS Bandwidth Up to 25 Gbps
c6in.8xlarge	32 vCPU	64 GiB RAM	Network Bandwidth 50 Gbps	EBS Bandwidth 25 Gbps
c6in.12xlarge	48 vCPU	96 GiB RAM	Network Bandwidth 75 Gbps	EBS Bandwidth 37.5 Gbps
c6in.16xlarge	64 vCPU	128 GiB RAM	Network Bandwidth 100 Gbps	EBS Bandwidth 50 Gbps
c6in.24xlarge	96 vCPU	192 GiB RAM	Network Bandwidth 150 Gbps	EBS Bandwidth 75 Gbps
c6in.32xlarge	128 vCPU	256 GiB RAM	Network Bandwidth 200 Gbps	EBS Bandwidth 100 Gbps

For more information on choosing instance types, see [choosing instance types](#) in AWS documentation.

The cloudkit provides the following choices for the disks (AWS Elastic Block storage) that can be attached to the IBM Storage Scale storage nodes:

- gp2
- gp3
- io1

For more information on choosing appropriate EBS volumes, see [Choosing the appropriate EBS volumes](#) in Amazon Web Service documentation.

Planning cloudkit notifications

The following prompt offers the option of subscribing to an Amazon Simple Notification Service (Amazon SNS). A confirmation message is sent to the provided email address; to receive the cloudkit notifications, the user needs to confirm the subscription.

```
? Operator Email (Optional): cloudkit@example.com
```

Planning for encryption at rest

The cloudkit offers an easy and simplified way to enable EBS encryption used by IBM Storage Scale instances. For more information, see [Amazon EBS encryption](#) in AWS documentation.

During the interactive method of **cloudkit create cluster**, the following encryption-related questions are shown:

```
? EBS Disk type: gp2
? Do you wish to encrypt boot and data volumes: Yes
? Block device encryption key (Key ID/key ARN/Alias Name/Alias ARN) (Optional): xxxxxx
```

The **? Do you wish to encrypt boot and data volumes:** prompt offers you the next options.

- Specify No. The root and data corresponding EBS volumes remain unencrypted.

- Specify Yes and do not provide a block device encryption key. The root and data corresponding to EBS volumes are encrypted by using the default key that is provided by AWS.
- Specify Yes and provide a block device encryption key. The root and data corresponding to EBS volumes are encrypted by using the key that you provided.

Note:

- If the key used for encryption IBM Storage Scale EBS volumes is deleted, data cannot be retrieved.
- Invalid key or user who is configured to run cloudkit lacks permissions to read the key, which leads to failures.

Limitations

- `Throughput-Performance-Scratch-Storage` is not encrypted through cloudkit. Since the disks provided in scratch profile instance type, Instance storage, are hardware encrypted internally using an XTS-AES-256 block cipher. Hence, it does not require external encryption.
- Encryption cannot be turned on or off after the EBS volume is created. This limitation means that unencrypted volumes that are created through cloudkit cannot be encrypted later and vice versa.
- Key rotation and life-cycle management are the responsibility of the users.
- If the key used for encrypting IBM Storage Scale is deleted, data cannot be retrieved.

Planning for GCP cloud

Learn about prerequisites, configurations, and other important planning information you must consider before you deploy an IBM Storage Scale cluster on Google Cloud Platform (GCP).

Consider the following steps when you are planning for an IBM Storage Scale cluster on the public cloud of GCP:

1. Preparing the GCP environment to deploy the IBM Storage Scale cluster.
2. Planning the virtual network cloud (VPC) architecture.
3. Planning for DNS.
4. Planning for bastion.
5. Creating an IBM Storage Scale virtual machine (VM) image in advance.
6. Planning for the IBM Storage Scale deployment architecture on GCP.
7. Planning for IBM Storage Scale cluster deployment profiles.
8. Determining your performance, scalability, data availability, and data protection requirements.
9. Planning for encryption at rest.

Preparing the GCP environment

Complete the following steps to prepare the GCP environment:

1. Create a service account with sufficient privileges and quota to provision all the required resources along with the credentials to access the GCP API. For information about creating a service account from the GCP console, see [Create a service account](#) in GCP documentation.
2. Make sure to have the next roles configured. These roles are required for the GCP service account to run **cloudkit**.

```
Artifact Registry Administrator
Cloud KMS CryptoKey Encrypter/Decrypter
Compute Instance Admin (v1)
Compute Network Admin
Compute Security Admin
DNS Administrator
Service Account User
Storage Admin
Browser
```

Note: Roles can be added or removed by following the procedures describe in [GCP documentation](#).

3. Verify the GCP quota limits and make sure that enough quotas exist for the GCP instance types that you intend to deploy. Consult the quota limits of the GCP instance from **Google Cloud console > Quota page**.

If necessary for the GCP instance types that you intend to deploy, request a service limit increase. For more information, see [Manage your quota using the Google Cloud console](#) in GCP documentation.

4. Prepare the installer node where the cloudkit is meant to run. For more information, see [“Preparing the installer node” on page 421](#).

Planning the virtual private cloud (VPC) architecture for GCP

When resources are being deployed in the cloud, the cloudkit can either create a new virtual private cloud (VPC) and provision the resources into it or use previously created VPC.

When the cloudkit creates a new VPC, the cloudkit designs your network infrastructure from scratch. It chooses the subnets, network address ranges, and security groups that best suit the IBM Storage Scale deployment.

A VPC can be created by using the **cloudkit create network** command, which creates only the VPC. Or you can use the **cloudkit create cluster** command.

The following deployment options are available to users:

```
? VPC Mode:[Use arrows to move, type to filter]
> New
Existing
```

- **Deploy IBM Storage Scale into a new VPC with a single availability zone**

This option builds a new GCP environment that consists of the VPC, subnets, firewall rules, bastion hosts, and other infrastructure components; then, it deploys IBM Storage Scale into this new VPC with a single availability zone.

- **Deploy IBM Storage Scale into an existing VPC**

This option provisions IBM Storage Scale in your existing GCP infrastructure.

When the IBM Storage Scale cluster is to be deployed in an existing VPC, the following requirements must be met:

- This is a mandatory step. Private subnets (with allocatable IP address) in the availability region (minimum 1 private subnet per availability region).
- This is an optional step. A public subnet (a subnet with internet gateway attached) is only required if either the following cases are met:
 1. Where the cluster is planned to be accessed via a jump host. For more information, see [“Other considerations” on page 448](#).
 2. If there is a need to precreate an IBM Storage Scale VM image. For more information, see [“Planing for GCP by precreating a VM image of IBM Storage Scale” on page 431](#).
- This is an optional step. A cloud NAT attached to the private subnet is only required when the IBM Storage Scale instances need access to the internet (for operating system updates and security patches, and so on).

Planning for DNS

You can create a DNS domain by using the **cloudkit create dns** command.

This command can create a new DNS zone or use an existing DNS zone:

```
? Do you wish to use an existing DNS zone: (y/N)
```

To consult the available deployment options, issue the **cloudkit create cluster** command:

Remember: Creation of a cloud DNS is mandatory and must be precreated before you run the **cloudkit create cluster** command.

```
E: No existing DNS zones found. In order to configure, use 'cloudkit create dns'.
```

Planning for bastion

Use the **cloudkit create jumphost** command to create a jump host or bastion in the public subnet.

To consult the available deployment options, issue the **cloudkit create cluster** command:

```
? Bastion OS: CentOS Stream 9 (latest minor version) | centos-stream-9-v20240709 | gcpuser
? Bastion instance type: n1-standard-1 | vCPU(1) | RAM (3.75 GB)
? Do you wish to continue by creation of jumphost ssh key pair: Yes
I: Bastion instance(s) private key path: /home/user1 /bastion_ssh_keys/id_rsa
? Bastion CIDR allow list: xxx.xxx.xxx.xx/32
```

The cloudkit provides the following options for VM instances types for a jump host:

```
> n1-standard-1 | vCPU(1) | RAM (3.75 GB)
n1-standard-2 | vCPU(2) | RAM (7.50 GB)
n1-highmem-2 | vCPU(2) | RAM (13 GB)
n1-highcpu-4 | vCPU(4) | RAM (3.60 GB)
```

Planing for GCP by precreating a VM image of IBM Storage Scale

When the cloudkit is used to create an IBM Storage Scale cluster by using the **cloudkit create cluster** command, it can either automatically create a stock VM image or the customer can provide a previously created custom image.

The cloudkit can automatically create an image by using Red Hat Enterprise Linux (RHEL) 8 (latest minor version) and RHEL 9 (latest minor version). When the option to use a stock image is chosen, it will automatically perform all required actions to create an IBM Storage Scale VM image and create the IBM Storage Scale cluster.

A custom image can be created by using the **cloudkit create image** command. The **create image** command can accept inputs in the form of an existing image or provides the ability to create an image from a Red Hat operating system image.

A custom existing image consists of:

- A Red Hat version that is supported by IBM Storage Scale.
- This is an optional step. For any customer applications that are already preinstalled, the **cloudkit create image** command installs all required IBM Storage Scale packages on top of the existing image.

Planning for the IBM Storage Scale deployment architecture on GCP

Before the IBM Storage Scale cluster can be created on the cloud, the deployment architecture must be planned.

The cloudkit offers the following deployment models for IBM Storage Scale clusters:

- **Combined-compute-storage**

A unified IBM Storage Scale cluster with both storage and compute nodes. It is recommended that any customer workload runs only on the compute nodes.

- **Compute-only**

An IBM Storage Scale cluster that only consists of compute nodes. This cluster does not have any local file system and therefore must remote mount the file system from an IBM Storage Scale storage cluster.

- **Storage-only**

This IBM Storage Scale cluster only consists of storage nodes that have access to storage where the file system is created. This file system is remote mounted to any number of compute-only IBM Storage Scale clusters.

When you run the **cloudkit create cluster** command, the following prompt asks the user to select from among the different deployment models that are offered:

```
? IBM Spectrum Scale deployment model:[Use arrows to move, type to filter, ? for more help]
> Storage-only
Compute-only
Combined-compute-storage
```

Deployment purpose

The deployment purpose is applicable only to storage clusters.

- **Non-Production**

This profile lists instance types that are experimental (best suited for proof of concepts) and may require tuning for running scale-suited workloads.

- **Production**

This profile lists instance types that are well suited for scale workloads. Optimal tuning parameters are calculated and set during the deployment.

When the next prompt appears, select one of the supported deployment purposes.

```
? IBM Storage Scale deployment model: Storage-only
? Deployment purpose: [Use arrows to move, type to filter]
Non-Production
> Production
```

IBM Storage Scale cluster deployment profiles

Cloudkit offers the following deployment profiles:

- **Throughput-Performance-Persistent-Storage**

This profile uses a single availability zone with persistent storage, which means that the file system device retains data after the instance is stopped. In this mode, the number of storage instances is limited to 64; compute instances is limited to 65. This profile offers these disk types to choose from: pd-balanced, pd-standard, and pd-ssd.

Important: When you choose this profile, the file system is configured in such manner that the data is not replicated, only the metadata gets replicated across the instances.

- **Throughput-Performance-Scratch-Storage**

This profile uses a single availability zone and a placement group with `cluster` policy, which means that it packs instances close together inside an availability zone. This strategy enables workloads to achieve the low-latency network performance necessary for tightly coupled node-to-node communication that is typical of high-performance computing (HPC) applications, with instance or temporary storage, meaning that the file system device loses data after the instance is stopped.

In the mode, the number of storage instances is limited to 10; compute instances are limited to 65.

Important: This profile uses local NVMe SSD or instance storage, which offers high performance and low latency for data-intensive workloads. However, the data that is stored in this mode is volatile and can be lost if the instance is stopped or terminated. Therefore, it is recommended to take frequent backups.

This profile must not be used for long-term storage or if the data is not backed up elsewhere.

- **Balanced**

This profile uses multiple (3) availability zones to deploy the IBM Storage Scale cluster into. In this mode, the number of storage instances is limited to 64. It offers a choice of disk types pd-balanced, pd-standard, and pd-ssd.

Important: The IBM Storage Scale file system is configured in such a way that data and metadata will be replicated across availability zones.

Instances are spread across the first two availability zones they are specified in the selection, and the tie-breaker instance is provisioned in the three availability zones that are specified during the selection.

```
? Tuning profile: [Use arrows to move, type to filter, ? for more help]
Throughput-Performance-Scratch-Storage
> Throughput-Performance-Persistent-Storage
Balanced
```

Determining your performance, scalability, data availability, and data protection requirements

Before the deployment of an IBM Storage Scale cluster, it is important to understand your requirements of performance, scalability, data availability, and data protection. These criteria determine what GCP instance types to use for the storage nodes and computes, also what elastic block store types should be used.

The cloudkit provides the following choices for VM instances types.

For IBM Storage Scale compute nodes, all the VM instance types are supported, according to their availability per region.

For IBM Storage Scale storage nodes, the support depends on the profile, as described in the following list.

- For **Throughput-Performance-Persistent-Storage** and **Balanced** profiles, the next VM instance types are supported.

> n2-standard-2	vCPU(2)	RAM (8.0 GB)	Egress Network Bandwidth (Up to 10 Gbps)
n2-standard-4	vCPU(4)	RAM (16.0 GB)	Egress Network Bandwidth (Up to 10 Gbps)
n2-standard-8	vCPU(8)	RAM (32.0 GB)	Egress Network Bandwidth (Up to 16 Gbps)
n2-standard-16	vCPU(16)	RAM (64.0 GB)	Egress Network Bandwidth (Up to 32 Gbps)
n2-standard-32	vCPU(32)	RAM (128.0 GB)	Egress Network Bandwidth (Up to 50 Gbps)
n2-standard-48	vCPU(48)	RAM (192.0 GB)	Egress Network Bandwidth (Up to 50 Gbps)

- For **Throughput-Performance-Scratch-Storage** profile, the next VM instance types are supported.

> n2-standard-32	vCPU(32)	RAM (128.0 GB)	Egress Network Bandwidth (Up to 50 Gbps)
n2-standard-48	vCPU(48)	RAM (192.0 GB)	Egress Network Bandwidth (Up to 50 Gbps)

For more information on choosing instance types, see [Machine families resource and comparison guide](#) in GCP documentation.

The cloudkit provides the following choices for the disks that can be attached to the IBM Storage Scale storage nodes:

```
? Disk type: [Use arrows to move, type to filter, ? for more help]
> pd-balanced
pd-standard
pd-ssd
```

Planning for encryption at rest

The cloudkit offers an easy and simplified way to enable disk encryption used by IBM Storage Scale instances.

GCP provides built-in disk encryption options that you can enable during disk creation. Data is encrypted by default. The following encryption methods are the most common:

- Google-managed encryption key. No configuration is required.
- Customer-managed encryption key (CMEK). Managed through Google Cloud key management service.
- Customer-supplied encryption key (CSEK). Managed outside of Google Cloud.

During the interactive method of **cloudkit create cluster**, the following encryption-related questions are shown:

```
? Data stored in boot and data disk(s) are encrypted automatically. Select an encryption key
management solution: [Use arrows to move, type to filter]
> Google-managed-encryption-key
  Customer-managed-encryption-key
```

You must select an encryption key management solution.

Note:

- If the key used for encrypting IBM Storage Scale disk volumes is deleted, data cannot be retrieved.
- An invalid key or user, one that was not configured to run **cloudkit**, lacks permissions to read the key, which leads to failures.
- CSEK (managed outside of Google Cloud) is not supported.

The Google-managed-encryption-key solution does not require any further configuration.

If you choose the Customer-managed-encryption-key solution, make the following configurations:

- In these prompts, input a KMS key ring and a KMS key name:

```
? Data stored in boot and data disk(s) are encrypted automatically. Select an encryption key
management solution: Customer-managed-encryption-key
? Google Cloud key ring: cloudkit-uscentral
? Customer-managed encryption key (CMEK): cloudkit-key
```

- To be able to read the key, make sure that the service user has a viewer role for the `cloudkms.cryptoKeyEncrypterDecrypter` and the Cloud KMS.
- Apart from the service user, the default compute engine service user also needs to be provided with `cloudkms.cryptoKeyEncrypterDecrypter`.

Note: If you choose this solution, consider these aspects of its implementation:

- When the cluster deployment configuration is storage only, both data and boot volumes are encrypted.
- When the cluster deployment configuration is compute only, only boot volumes are encrypted.

Limitations

- Both boot and data-related volumes are encrypted.
- All data and boot volumes are encrypted by using a single key.
- Scratch or instance storage (`local-ssd`) disks are hardware-encrypted internally. Hence, they do not require external encryption.
- CSEK is out of scope.
- The KMS key ring must be local to the region of operation.

Planning for Microsoft Azure cloud

Learn about prerequisites, configurations, and other important planning information you must consider before the deployment of an IBM Storage Scale cluster on Azure.

Consider the following steps when you are planning for an IBM Storage Scale cluster on the public cloud of Azure:

1. Preparing the Azure environment to deploy the IBM Storage Scale cluster.
2. Planning the virtual network (VNET) architecture.

3. Planning for domain name system (DNS).
4. Planning for bastion.
5. Creating an IBM Storage Scale virtual machine (VM) image in advance.
6. Planning for the IBM Storage Scale deployment architecture on Azure.
7. Planning for IBM Storage Scale cluster deployment profiles.
8. Determining your performance, scalability, data availability, and data protection requirements.
9. Planning for encryption at rest.

Preparing the Azure environment

Complete the following steps to prepare the Azure environment:

1. Create a service principle by using the **az ad sp create-for-rbac** command.

To consult the credentials, issue the following command:

```
az ad sp create-for-rbac --query "{ client_id: appId, client_secret: password, tenant_id: tenant }"
```

The following output is an example of the previous command:

```
{
  "client_id": "00000000-0000-0000-0000-34XXXXXXX",
  "client_secret": "00000000-0000-0000-0000-50XXXXXXX",
  "tenant_id": "00000000-0000-0000-0000-40XXXXXXX",
}
```

To authenticate on Azure, you also need to obtain your Azure subscription ID. Issue the **az account show** command to consult this ID:

```
az account show --query "{ subscription_id: id }"
```

The following output is an example of the previous command:

```
{
  "subscription_id": "00000000-aea2-4177-a0f1-7117aXXXX"
}
```

Remember: Make a note of the next values because they are required during the configuration operations of **cloudkit**:

- client_id
 - client_secret
 - tenant_id
 - subscription_id
2. Make sure to configure at least the Contributor role. These roles are required for the Azure service principle to run **cloudkit**.
 3. Verify the Azure quota limits and make sure that enough quotas exist for the Azure instance types that you intend to deploy. To consult the quota limits of the Azure instance, see [Azure Quotas documentation](#).
 4. Prepare the installer node where the cloudkit is meant to run. For more information, see [“Preparing the installer node” on page 421](#).
 5. Optionally, create a resource group that holds all Azure resources that are specific for IBM Storage Scale cluster. For more information, see [Create resource groups](#).

Planning the virtual network (VNET) architecture for Azure

When resources are being deployed in the cloud, the cloudkit can either create a new virtual network (VNET) and provision the resources into it or use a previously created VNET.

When the cloudkit creates a new VNET, the cloudkit designs your network infrastructure from scratch. It chooses the subnets, network address ranges, and security groups that best suit the IBM Storage Scale deployment.

A VNET can be created by using the **cloudkit create network** command, which creates only the VNET. Or you can use the **cloudkit create cluster** command.

The following deployment options are available to users while they are using the **cloudkit create cluster**:

```
? VPC Mode:[Use arrows to move, type to filter]
> New
Existing
```

- **Deploy IBM Storage Scale into a new VNET with a single availability zone**

This option builds a new Azure environment that consists of the VNET, private and public subnets, network security groups, application security groups, security rules, bastion hosts, and other infrastructure components; then, it deploys IBM Storage Scale into this new VNET with a single availability zone.

- **Deploy IBM Storage Scale into an existing VNET**

This option provisions IBM Storage Scale in your existing Azure infrastructure.

When you are deploying the IBM Storage Scale cluster in an existing VNET, the following requirements must be met:

- This is a mandatory step. Private subnets (with allocatable IP address range) in the availability region (minimum 1 private subnet per availability region).
- This is an optional step. A public subnet is only required if either of the following cases is met:
 1. Where the cluster is planned to be accessed through a jump host. For more information, see [“Other considerations”](#) on page 448.
 2. If there is a need to precreate an IBM Storage Scale VM image. For more information, see [“Planning for Microsoft Azure by precreating a VM image of IBM Storage Scale”](#) on page 437.
- This is an optional step. A cloud NAT attached to the private subnet is only required when the IBM Storage Scale instances need access to the internet (for operating system updates and security patches, and so on).
- If you configured an Azure Private DNS, you must provide the Azure Private DNS information during the cluster deployment..

```
? Do you wish to use Azure Cloud DNS: Yes
```

- Make sure to select the appropriate resource group that is aimed to deploy the IBM Storage Scale cluster.

Planning for domain name system (DNS)

To facilitate hostname resolution, use the **./cloudkit create dns** command to create or associate a DNS.

This command can create a new DNS zone or use an existing DNS zone:

```
? Do you wish to use an existing DNS zone: (y/N)
```

The following deployment options are available to users when they are using the **cloudkit create cluster** command:

Remember: Creation of a cloud DNS is mandatory and must be precreated before you run the **cloudkit create cluster** command.

```
E: No existing DNS zones found. In order to configure, use 'cloudkit create dns'.
```

Important: Make sure to select the same resource group that you selected when you created the VNET.

Planning for bastion

Use the **cloudkit create jumphost** command to create a jump host or bastion in the public subnet.

The following deployment options are available to users when they are using the **cloudkit create cluster** command:

```
? Bastion OS: RHEL-HA | RedHat | 8_8-gen2 | 8.8.2023121916 | azureuser
? Bastion instance type: Standard_B2s | vCPU(2) | RAM (4.0 GB)
? Do you wish to continue by creation of jumphost ssh key pair: Yes
I: Bastion instance(s) private key path: /user/clouduser/bastion_ssh_keys/id_rsa
? Bastion CIDR allow list: xxx.xxx.xxx.20/32
```

The cloudkit provides the following options for VM instances types for a jump host:

Standard_B2s	vCPU(2)	RAM (4.0 GB)
Standard_B2s_v2	vCPU(2)	RAM (8.0 GB)
Standard_B4s_v2	vCPU(4)	RAM (16.0 GB)

Important: Make sure to select the same resource group that you selected when you created the VNET.

Planning for Microsoft Azure by precreating a VM image of IBM Storage Scale

When the cloudkit is used to create an IBM Storage Scale cluster by using the **cloudkit create cluster** command, it can either automatically create a stock VM image or the customer can provide a previously created custom image.

The cloudkit can automatically create an image by using Red Hat Enterprise Linux (RHEL) 8.8 or RHEL 9.2. When the option to use a stock image is chosen, the cloudkit automatically performs all required actions to create an IBM Storage Scale and create the IBM Storage Scale cluster.

A custom image can be created by using the **cloudkit create image** command. The **create image** command can accept inputs in the form of an existing image or provides the ability to create an image from a Red Hat operating system image.

A custom existing image consists of:

- A Red Hat version that is supported by IBM Storage Scale.
- This is an optional step. For any customer applications already preinstalled, the **cloudkit create image** command installs all required IBM Storage Scale packages on top of the existing image.
- Up to two virtual machine images.
- An image that is a logical volume management (LVM) partition.

Planning for the IBM Storage Scale deployment architecture on Azure

Before the IBM Storage Scale cluster can be created on the cloud, the deployment architecture must be planned.

The cloudkit offers the following deployment models for IBM Storage Scale clusters:

- **Combined-compute-storage**

A unified IBM Storage Scale cluster with both storage and compute nodes. It is recommended that any customer workload runs only on the compute nodes.

- **Compute-only**

An IBM Storage Scale cluster that consists only of compute nodes. This cluster does not have any local file system and therefore must remote mount the file system from an IBM Storage Scale storage cluster.

- **Storage-only**

This IBM Storage Scale cluster consists only of storage nodes that have access to storage where the file system is created. This file system is remotely mounted to any number of compute-only IBM Storage Scale clusters.

When you run the **cloudkit create cluster** command, the following prompt asks the user to select from among the different deployment models that are offered:

```
? IBM Spectrum Scale deployment model:[Use arrows to move, type to filter, ? for more help]
> Storage-only
Compute-only
Combined-compute-storage
```

Deployment purpose

The deployment purpose is applicable only to storage clusters.

- **Non-Production**

This profile lists instance types that are experimental (best suited for proof of concepts) and may require tuning for running scale-suited workloads.

- **Production**

This profile lists instance types that are well suited for scale workloads. Optimal tuning parameters are calculated and set during the deployment.

When the next prompt appears, select one of the supported deployment purposes.

```
? IBM Storage Scale deployment model: Storage-only
? Deployment purpose: [Use arrows to move, type to filter]
Non-Production
> Production
```

IBM Storage Scale cluster resource group

Azure resource group plays an important role during the deployment of an IBM Storage Scale cluster because cluster deployment depends on resources such as VNET and DNS, which are tied to a single resource group. When you are deploying an IBM Storage Scale cluster, select the same resource group to access all the dependent group resources.

Cloudkit provides the option to choose among existing resource group to deploy IBM Storage scale cluster on Azure.

When you run the **cloudkit create cluster** command, following cloudkit create cluster prompt asks the user to select from among the different available resource group.

```
? Resource Group: scale-cluster
```

Important: Make sure to select the same resource group that you selected when you created the VNET.

IBM Storage Scale cluster deployment profiles

Cloudkit offers the following deployment profiles:

- **Throughput-Performance-Persistent-Storage**

This profile uses a single availability zone with persistent storage, which means that the file system device retains data after the instance is stopped. In this mode, cloudkit calculates the number of storage nodes based on the provided file system capacity. The number of storage instances is limited to 64; compute instances are limited to 65.

Important: When you choose this profile, the file system is configured in such manner that the data is not replicated, only the metadata gets replicated across the instances.

- **Throughput-Performance-Scratch-Storage**

This profile uses a single availability zone and a proximity placement group, which means that it packs instances close together inside an availability zone. This strategy enables workloads to achieve the low-latency network performance necessary for tightly coupled node-to-node communication that is typical of high-performance computing (HPC) applications, with instance or temporary storage, meaning that the file system device loses data after the instance is stopped.

In this mode, the number of storage instances is limited to 10 and compute instances is limited to 65.

Important: This profile uses temporary storage (NVMe), which offers high performance and low latency for data-intensive workloads. However, the data that is stored in this mode is volatile and can be lost if the instance is stopped or terminated. Therefore, it is recommended to take frequent backups.

This profile must not be used for long-term storage or if the data is not backed up elsewhere.

- **Throughput-Advance-Persistent-Storage**

This profile uses a single availability zone with persistent storage, which means that the file system device retains data after the instance is stopped. This mode is meant for storage capacity rather performance; the number of storage instances are limited to 64; compute instances are limited to 65. This profile offers these disk types to choose from: Standard_LRS, StandardSSD_LRS, Premium_LRS, and PremiumV2_LRS.

Important: When you choose this profile, the file system is configured in such manner that the data is not replicated, only the metadata gets replicated across the instances.

- **Balanced**

This profile uses multiple (3) availability zones to deploy the IBM Storage Scale cluster into. In this mode, the number of storage instances is limited to 64. It offers a choice of disk types Standard_LRS, StandardSSD_LRS, Premium_LRS, and PremiumV2_LRS.

Important: The IBM Storage Scale file system is configured in such a way that data and metadata are replicated across availability zones.

Instances are spread across the first two availability zones that are specified in the selection, and the tie-breaker instance gets provisioned in the three availability zones that are specified during the selection.

```
? Tuning profile: [Use arrows to move, type to filter, ? for more help]
  Throughput-Performance-Scratch-Storage
> Throughput-Performance-Persistent-Storage
  Throughput-Advance-Persistent-Storage
  Balanced
```

Determining your performance, scalability, data availability, and data protection requirements

Before the deployment of an IBM Storage Scale cluster, it is important to understand your requirements in terms of performance, scalability, data availability, and data protection. These criteria determine what Azure instance types to use for the storage nodes and compute nodes, also what elastic block store types should be used.

The cloudfit provides the following choices for VM instances types:

For IBM Storage Scale compute nodes, all the VM instance types are supported, according to their availability per region.

For IBM Storage Scale storage nodes, the support depends on the profile, as described in the following list.

- For **Throughput-Performance-Persistent-Storage** and **Balanced** profiles, the next VM instance types are supported.

Standard_F8s_v2	vCPU(8)	RAM (16.0 GB)	Max network bandwidth (12500 Mbps)
Standard_F16s_v2	vCPU(16)	RAM (32.0 GB)	Max network bandwidth (12500 Mbps)
Standard_F32s_v2	vCPU(32)	RAM (64.0 GB)	Max network bandwidth (16000 Mbps)

- For **Throughput-Performance-Scratch-Storage** profile, the next VM instance types are supported.

Standard_L8s_v3	vCPU(8)	RAM (64)	Instance Storage (1 x 1.92 TB NVMe Disk)
Expected network bandwidth (12500 Mbps)			
Standard_L16s_v3	vCPU(16)	RAM (128)	Instance Storage (2 x 1.92 TB NVMe Disk)
Expected network bandwidth (12500 Mbps)			
Standard_L32s_v3	vCPU(32)	RAM (256)	Instance Storage (4 x 1.92 TB NVMe Disk)
Expected network bandwidth (16000 Mbps)			
Standard_L48s_v3	vCPU(48)	RAM (384)	Instance Storage (6 x 1.92 TB NVMe Disk)
Expected network bandwidth (24000 Mbps)			
Standard_L64s_v3	vCPU(64)	RAM (512)	Instance Storage (8 x 1.92 TB NVMe Disk)
Expected network bandwidth (30000 Mbps)			
Standard_L80s_v3	vCPU(80)	RAM (640)	Instance Storage (10 x 1.92 TB NVMe Disk)
Expected network bandwidth (32000 Mbps)			

- For **Throughput-Advance-Persistent-Storage** profile, the next VM instance types are supported.

Standard_F8s_v2	vCPU(8)	RAM (16.0 GB)	Max network bandwidth (12500 Mbps)
Standard_F16s_v2	vCPU(16)	RAM (32.0 GB)	Max network bandwidth (12500 Mbps)
Standard_F32s_v2	vCPU(32)	RAM (64.0 GB)	Max network bandwidth (16000 Mbps)

For more information on choosing instance types, see [Virtual Machine series](#) in Azure documentation.

The cloudkit provides the following choices for the disks that can be attached to the IBM Storage Scale storage nodes:

```
? Disk type: [Use arrows to move, type to filter, ? for more help]
> Standard_LRS
StandardSSD_LRS
Premium_LRS
PremiumV2_LRS
```

Planning for encryption at rest

The cloudkit offers an easy and simplified way to enable disk encryption used by IBM Storage Scale instances.

Azure provides built-in disk encryption options that you can enable during disk creation. Data is encrypted by default. The following encryption methods are the most common:

- Platform-managed encryption key. No configuration is required. By default disks are automatically encrypted-at-rest with platform-managed keys
- Customer-managed encryption key (CMEK). Managed through Azure Key Vault.

To enable encryption at rest, you need Key Vault Crypto Service Encryption permissions, which rely on encryption set. Key vault and key creation steps are out of scope.

When you are creating the vault, make sure that the following are enabled:

- Azure virtual machines for deployment
- Azure Resource Manager for template deployment
- Azure Disk Encryption for volume encryption

During the interactive method of **cloudkit create cluster**, the following encryption-related questions are shown:

```
? Data stored in boot and data disk(s) are encrypted automatically. Select an encryption key
management solution: [Use arrows to move, type to filter]
> Platform-managed-encryption-key
Customer-managed-encryption-key
```

You must select a solution for encryption key management.

Note:

- If the key used for encrypting IBM Storage Scale disk volumes is deleted, data cannot be retrieved.
- An invalid key or user, one that was not configured to run **cloudkit**, lacks permissions to read the key, which causes failures.

The Platform-managed-encryption-key solution does not require any further configuration.

If you choose the Customer-managed-encryption-key solution, make the following configurations:

- In these prompts, input a key name for the Azure Key Vault:

```
? Data stored in boot and data disk(s) are encrypted automatically. Select an encryption key
management solution: Customer-managed-encryption-key
? Cloud Vault Name:  scalettest1
? Customer-managed encryption key (CMEK):  scalekey1
```

Note: If you choose this solution, consider these aspects of its implementation:

- When the cluster deployment configuration is storage only, both data and boot volumes are encrypted.
- When the cluster deployment configuration is compute only, only the boot volumes are encrypted.

Limitations

- Both boot and data-related volumes are encrypted.
- All data and boot volumes are encrypted by using a single key.
- Temporary disks or scratch storage disks are not encrypted when Customer-managed-encryption-key is used.
- The Azure Key Vault must use the same resource group that is used for IBM Storage Scale cluster creation.
- You cannot use the **cloudkit** command in the Azure Cloud VM.

Planning for IBM Cloud

Learn about prerequisites, configurations, and other important planning information that you must consider before the deployment of an IBM Storage Scale cluster on IBM Cloud.

Consider the following steps when you are planning for an IBM Storage Scale cluster on the public cloud of IBM Cloud:

1. Checking the limitations.
2. Preparing the IBM Cloud environment to deploy the IBM Storage Scale cluster.
3. Planning the virtual network cloud (VPC) architecture.
4. Planning for DNS.
5. Planning for bastion.
6. Creating an IBM Storage Scale custom image in advance.
7. Planning for the IBM Storage Scale deployment architecture on IBM Cloud.
8. Determining the deployment purpose.
9. Selecting the deployment profiles of the IBM Storage Scale cluster.
10. Determining your performance, scalability, data availability, and data protection.

Checking the limitations

These are the applicable limitations for the deployment of IBM Storage Scale on IBM Cloud:

- If the key that is used to encrypt IBM Storage Scale is deleted, data cannot be retrieved.
- Only the general purpose type of disk is supported with IBM Cloud.

Preparing the IBM Cloud environment to deploy the IBM Storage Scale cluster

Complete the following steps to prepare the IBM Cloud environment.

1. If you do not already have one, create an IBM Cloud account.
2. Create the necessary IAM API keys that are required for cloudkit to make programmatic requests to IBM Cloud. For more information, see [Managing user API keys](#) in the IBM Cloud documentation.
3. Configure the required permission to the IAM API key by following the procedure from *Configuring permissions to use cloudkit on public clouds* in the *IBM Storage Scale: Administration Guide*.
4. In the preferred region, create a key pair.
5. Prepare the installer node where cloudkit is meant to run. For more information, see *Preparing the installer node* in the *IBM Storage Scale: Concepts, Planning, and Installation Guide*

Planning the virtual network cloud (VPC) architecture

When resources are being deployed in the cloud, cloudkit can either create a new virtual private cloud (VPC) and provision the resources into it, or use a previously created VPC. When cloudkit creates a new VPC, cloudkit designs your network infrastructure from scratch; it chooses the subnets, network address ranges, and security groups that best suit the IBM Storage Scale deployment.

Cloudkit offers the next commands to create a VPC:

- `Cloudkit create network` command, which creates only the VPC.
- `Cloudkit create cluster` command, which creates the cluster and its VPC.

The following deployment options are available to users:

```
? VPC Mode:  [Use arrows to move, type to filter, ? for more help]
> New
  Existing
```

If the IBM Storage Scale cluster is to be deployed in an existing VPC, the following requirements must be met:

- Dynamic Host Configuration Protocol (DHCP) options set or DNS configured.
- In the availability zone, subnets with allocatable IP address; the minimum is one private subnet per availability zone.

Optionally, meet the next requirements for the deployment of an IBM Storage Scale cluster in an existing VPC:

- A public subnet (a subnet with internet gateway attached) is only required if either the following cases are met:
 - Where the cluster is planned to be accessed through a jump host. For more information, see [“Other considerations”](#) on page 448.
 - If it is necessary to precreate an IBM Storage Scale image. For more information, see [“Precreating an IBM Storage Scale AMI”](#) on page 425.
- A network address translation (NAT) gateway that is attached to the private subnet is only required when the IBM Storage Scale instances need access to the internet (for operating system updates and security patches, for example).

Planning for DNS

Use `cloudkit create dns` command to create a DNS domain. During the DNS creation process, you must create a service instance.

To create a new DNS zone or use an existing DNS zone, issue the next command:

```
? Do you wish to use an existing DNS zone:  (y/N)
```

Remember: Creation of a cloud DNS is mandatory and must be created before you run `cloudkit create cluster` command.

To consult the available deployment options, issue `cloudkit create cluster` command:

```
E:
    No existing DNS zones found. In order to configure, use 'cloudkit create
    dns'
```

If no existing zones are found, this command fails and displays an error message.

Planning for bastion or jump host

Use `cloudkit create jumphost` command to create a jump host or bastion in the public subnet.

To consult the available deployment options, issue `cloudkit create cluster` command:

```
? Bastion OS:  ibm-redhat-9-6-minimal-amd64-6 | r006-2bc93792-b76e-45ca-aa8a-d7f44029aae4 |
x86_64 | vpcuser
? Bastion instance type:  bx2-2x8 | 2 vCPU | 8 GB RAM | Up to 10 Gbps bandwidth
? Bastion/Jumphost SSH private key file path (will be used only for configuration):
<private_key_path>
? JumpHost CIDR allow list: xxx.xxx.xxx.xxx/0
```

For a jump host or bastion, cloudkit provides the following options for VM instances types:

```
? Bastion instance type: [Use arrows to move, type to filter]
> bx2-2x8 | 2 vCPU | 8 GB RAM | Up to 10 Gbps bandwidth
  bx2-4x16 | 4 vCPU | 16 GB RAM | Up to 10 Gbps bandwidth
  bx2-8x32 | 8 vCPU | 32 GB RAM | Up to 16 Gbps bandwidth
```

Creating an IBM Storage Scale custom image in advance

With cloudkit, you can either automatically create a stock image or provide a custom image.

Automatically created image by using `cloudkit create cluster` command

Red Hat Enterprise Linux (RHEL) 8.10 or 9.6 is required for cloudkit to automatically create a stock image. When you select this option, cloudkit automatically performs all required actions to create an IBM Storage Scale VM image and create the IBM Storage Scale cluster.

New or existing custom image by using `cloudkit create image` command

A custom image can be created by using `cloudkit create image` command. This command supports two options to create a custom image: it accepts inputs in the form of an existing image and it provides the ability to create an image from a RHEL image.

An existing custom image consists of the following components:

1. A RHEL version that is supported by IBM Storage Scale.
2. Optionally, for any customer applications that are already preinstalled, `cloudkit create image` command also installs all required IBM Storage Scale packages on top of the existing image.

Planning for the IBM Storage Scale deployment architecture on IBM Cloud

Before the IBM Storage Scale cluster can be created on the cloud, plan the deployment architecture.

Cloudkit offers the following deployment models for IBM Storage Scale.

Combined-compute-storage

A unified IBM Storage Scale cluster with both storage and compute nodes. It is recommended that any customer workload runs only on the compute nodes.

Compute-only

An IBM Storage Scale cluster that only consists of compute nodes. This cluster does not have any local file system and therefore must remotely mount the file system from an IBM Storage Scale storage cluster.

Storage-only

This IBM Storage Scale cluster consists of only storage nodes, which have access to storage where the file system is created. This file system is remotely mounted to any number of compute-only IBM Storage Scale clusters.

When you run `cloudkit create cluster` command, the following prompt asks you to select from among the different deployment models that are offered:

```
? IBM Storage Scale deployment model: [Use arrows to move, type to filter, ? for more help]
> Storage-only
  Compute-only
  Combined-compute-storage
```

Determining the deployment purpose

The deployment purpose is applicable only to storage clusters.

Non-Production

This profile lists instance types that are experimental, which are best suited for proof of concepts. A non-production deployment may require tuning to run scale-suited workloads.

Production

This profile lists instance types that are well suited for scale workloads. Optimal tuning parameters are calculated and set during the deployment.

When the next prompt appears, select one of the supported deployment purposes.

```
? Deployment purpose: [Use arrows to move, type to filter]
  Non-Production
> Production
```

Selecting the deployment profiles for an IBM Storage Scale cluster

Cloudkit offers the following deployment profiles.

Throughput-Performance-Persistent-Storage

This profile uses a single availability zone with persistent storage, which means that the file system device retains data after the instance is stopped. In this mode, cloudkit calculates the number of storage nodes based on the provided file system capacity. This mode uses `general-purpose` as disk type.

Important: When you choose this profile, the file system is configured in such manner that the data is not replicated, only the metadata gets replicated across the instances.

Throughput-Performance-Scratch-Storage

This profile uses a single availability zone and a placement group with `cluster` policy, which means that it packs instances close together inside an availability zone. This strategy enables workloads to achieve the low-latency network performance that is necessary for tightly coupled node-to-node communication that is typical of high-performance computing (HPC) applications, with instance or temporary storage; meaning that the file system device loses data after the instance is stopped.

In this mode, the number of storage instances is limited to 10, compute instances are limited to 65.

Important: This profile uses local NVMe SSD or instance storage, which offers high performance and low latency for data-intensive workloads. However, the data that is stored in this mode is volatile and

can be lost if the instance is stopped or terminated. Therefore, it is recommended to take frequent backups.

This profile must not be used for long-term storage or if the data is not backed up elsewhere.

Throughput-Advance-Persistent-Storage

This profile uses a single availability zone with persistent storage, which means that the file system device retains data after the instance is stopped. This mode is meant for storage capacity rather performance; the number of storage instances is limited to 64, compute instances is limited to 65.

Important: When you choose this profile, the file system is configured in such manner that the data is not replicated, only the metadata gets replicated across the instances.

Balanced

This profile uses multiple (3) availability zones to deploy the IBM Storage Scale cluster into. In this mode, the number of storage instances is limited to 64.

Important: The IBM Storage Scale file system is configured in such a way that data and metadata get replicated across availability zones.

```
? Tuning profile: [Use arrows to move, type to filter, ? for more help]
> Throughput-Performance-Persistent-Storage
  Throughput-Advance-Persistent-Storage
  Balanced
  Throughput-Performance-Scratch-Storage
```

Determining performance, scalability, data availability, and data protection

Before you deploy the IBM Storage Scale cluster, it is important to understand your requirements in terms of performance, scalability, data availability, and data protection. These criteria determine what IBM Cloud instance types to use for the storage nodes and compute nodes, also what elastic block store types should be used.

Cloudkit provides the following choices for VM instances types.

- For IBM Storage Scale compute nodes, all the VM instance types are supported, according to their availability per region.
- For IBM Storage Scale storage nodes, the support depends on the profile, as described in the following list.
 - For the Throughput-Performance-Persistent-Storage, Throughput-Advance-Persistent-Storage, and Balanced profiles, the next VM instance types are supported.

- cx2-16x32	16 vCPU	32 GiB RAM	Max network bandwidth 32 Gbps
- cx2-32x64	32 vCPU	64 GiB RAM	Max network bandwidth 64 Gbps
- cx2-48x96	48 vCPU	96 GiB RAM	Max network bandwidth 80 Gbps
- cx2-64x128	64 vCPU	128 GiB RAM	Max network bandwidth 80 Gbps
- cx2-96x192	96 vCPU	192 GiB RAM	Max network bandwidth 80 Gbps

- For the Throughput-Performance-Scratch-Storage profile, the next VM instance types are supported.

- ox2-8x64	8 vCPU	64 GiB RAM	Instance Storage 2 x 1200 GB	local SSD	Network
- ox2-16x128	16 vCPU	128 GiB RAM	Instance Storage 2 x 2400 GB	local SSD	Network
- ox2-32x256	32 vCPU	256 GiB RAM	Instance Storage 3 x 3200 GB	local SSD	Network
- ox2-64x512	64 vCPU	512 GiB RAM	Instance Storage 6 x 3200 GB	local SSD	Network
- ox2-96x768	96 vCPU	768 GiB RAM	Instance Storage 9 x 3200 GB	local SSD	Network

Installing IBM Storage Scale in a network restricted (air gap) setup with the cloudkit

Learn how to install IBM Storage Scale in air gap setup.

Planning

To be able to execute **cloudkit** when using the air gap installation option, you must first meet different prerequisites and complete the required configurations.

Two virtual machines are required to set up an air gap environment. These machines serve distinct roles, the cloudkit installer node and the staging node.

cloudkit installer node prerequisite

- The cloudkit installer node is the node responsible for executing the cloudkit binary. This node must be located in AWS and cannot be an on-premise node.
- The cloudkit installer node must be a Virtual Machine (VM) which should be spun in a private subnet (with no NAT attachment) in an AWS existing VPC.
- The VM must have Instance Metadata Service Version 1 (IMDSv1) enabled.
- The security group associated with the VM must have egress (outbound) open to allow traffic from all sources.
- Ensure this VPC has enough private IP addresses, as the same VPC will be utilized for IBM Storage Scale deployment.
- Ensure IBM Storage Scale AMI is already created.

Staging node prerequisite

- The staging node is the node which has public network associated either through internet gateway or NAT gateway.
- This node is used to download the SE package and prerequisite packages.
- This node must have the ability to reach (through s3 bucket, VPC peering or public subnet within the same VPC as cloudkit installer node) the cloudkit installer node to copy the packages.

Configure the VPC endpoints

Virtual Private Cloud (VPC) endpoints offer a private connection between a VPC and other AWS services (such as, EC2, etc), allowing traffic to be routed over the AWS network instead of over the internet.

In an air gap mode, the cloudkit requires EC2, STS, KMS, and SNS endpoints to be configured to the VPC. Complete the next steps to configure the endpoints:

1. Log in to the **AWS Management Console** and navigate to the **VPC dashboard**.
2. In the navigation pane, select **Endpoints** and click **Create Endpoint** button.
3. Choose the EC2, STS, KMS, and SNS services for creating endpoints.
4. Select the VPC and subnet which is used to spin the cloudkit installer VM to create the endpoint.
5. Select the cloudkit installer VM security group to use the endpoints.
6. Ensure that the DNS name and ipv4 settings are enabled.
7. Ensure that the VPC endpoint policy is set to allow Full access.
8. Review the endpoint settings and click **Create Endpoint**.

Configure the terraform provider(s) to work in isolated or air gap mode

Complete the following steps:

1. On the air gap cloudkit installer node, create the below directory layout.

Note: The versions included in the directory path shown below should be treated as the minimum version required.

```
mkdir -p $HOME/.terraform.d/plugins/registry.terraform.io/hashicorp/aws/4.61.0/linux_amd64/
mkdir -p $HOME/.terraform.d/plugins/registry.terraform.io/hashicorp/null/3.2.1/linux_amd64/
mkdir -p $HOME/.terraform.d/plugins/registry.terraform.io/hashicorp/template/2.2.0/
linux_amd64/
mkdir -p $HOME/.terraform.d/plugins/registry.terraform.io/hashicorp/time/0.9.1/linux_amd64/
mkdir -p $HOME/.terraform.d/plugins/registry.terraform.io/hashicorp/local/2.4.0/linux_amd64/
mkdir -p $HOME/.terraform.d/plugins/registry.terraform.io/hashicorp/tls/4.0.4/linux_amd64/
mkdir -p $HOME/.terraform.d/plugins/registry.terraform.io/integrations/github/5.20.0/
linux_amd64/
```

2. Copy the appropriate terraform providers to the appropriate directories.

As an example, you can use the staging VM to download the providers as shown below:

```
wget https://releases.hashicorp.com/terraform-provider-aws/4.61.0/terraform-provider-aws_4.61.0_linux_amd64.zip
unzip terraform-provider-aws_4.61.0_linux_amd64.zip
rm -rf terraform-provider-aws_4.61.0_linux_amd64.zip
wget https://releases.hashicorp.com/terraform-provider-null/3.2.1/terraform-provider-null_3.2.1_linux_amd64.zip
unzip terraform-provider-null_3.2.1_linux_amd64.zip
rm -rf terraform-provider-null_3.2.1_linux_amd64.zip
wget https://releases.hashicorp.com/terraform-provider-template/2.2.0/terraform-provider-template_2.2.0_linux_amd64.zip
unzip terraform-provider-template_2.2.0_linux_amd64.zip
rm -rf terraform-provider-template_2.2.0_linux_amd64.zip
wget https://releases.hashicorp.com/terraform-provider-time/0.9.1/terraform-provider-time_0.9.1_linux_amd64.zip
unzip terraform-provider-time_0.9.1_linux_amd64.zip
rm -rf terraform-provider-time_0.9.1_linux_amd64.zip
wget https://releases.hashicorp.com/terraform-provider-local/2.4.0/terraform-provider-local_2.4.0_linux_amd64.zip
unzip terraform-provider-local_2.4.0_linux_amd64.zip
rm -rf terraform-provider-local_2.4.0_linux_amd64.zip
wget https://releases.hashicorp.com/terraform-provider-tls/4.0.4/terraform-provider-tls_4.0.4_linux_amd64.zip
unzip terraform-provider-tls_4.0.4_linux_amd64.zip
rm -rf terraform-provider-tls_4.0.4_linux_amd64.zip
wget https://releases.hashicorp.com/terraform-provider-github/5.20.0/terraform-provider-github_5.20.0_linux_amd64.zip
unzip terraform-provider-github_5.20.0_linux_amd64.zip
rm -rf terraform-provider-github_5.20.0_linux_amd64.zip
```

Copy these providers from staging node to air gapped cloudkit installer node.

```
cd providers
scp terraform-provider-aws_v4.61.0_x5 <cloudkit-vm>:/root/.terraform.d/plugins/
registry.terraform.io/hashicorp/aws/4.61.0/linux_amd64/
scp terraform-provider-null_v3.2.1_x5 <cloudkit-vm>:/root/.terraform.d/plugins/
registry.terraform.io/hashicorp/null/3.2.1/linux_amd64/
scp terraform-provider-template_v2.2.0_x4 <cloudkit-vm>:/root/.terraform.d/plugins/
registry.terraform.io/hashicorp/template/2.2.0/linux_amd64/
scp terraform-provider-time_v0.9.1_x5 <cloudkit-vm>:/root/.terraform.d/plugins/
registry.terraform.io/hashicorp/time/0.9.1/linux_amd64/
scp terraform-provider-local_v2.4.0_x5 <cloudkit-vm>:/root/.terraform.d/plugins/
registry.terraform.io/hashicorp/local/2.4.0/linux_amd64/
scp terraform-provider-tls_v4.0.4_x5 <cloudkit-vm>:/root/.terraform.d/plugins/
registry.terraform.io/hashicorp/tls/4.0.4/linux_amd64/
scp terraform-provider-github_v5.20.0 <cloudkit-vm>:/root/.terraform.d/plugins/
registry.terraform.io/integrations/github/5.20.0/linux_amd64/
```

3. Create CLI configuration file (.terraformrc) in the user's home directory.

```
# cat .terraformrc
disable_checkpoint = true
disable_checkpoint_signature = false
provider_installation {
  filesystem_mirror {
    path = "/root/.terraform.d/plugins/"
    include = ["hashicorp/aws", "hashicorp/null", "hashicorp/template", "hashicorp/
```

```
time", "hashicorp/tls", "hashicorp/local", "integrations/github"]
    }
}
```

Extract the SE package on the air gapped cloudkit installer node

Copy the IBM Storage Scale SE package to the air gapped cloudkit installer node (you can copy it from the staging node and then further copy it to the cloudkit installer VM). Extract the SE package on the installer node to obtain the cloudkit binary and the associated scale components.

1. Install the rpm dependencies (createrepo_c, python3-pip, sqlite-3, and pip3 dependencies, such as jmespath and ansible).
2. Disable the strict host key checking.

```
echo "Host *" > ~/.ssh/config
echo " StrictHostKeyChecking no" >> ~/.ssh/config
chmod 400 ~/.ssh/config
```

Verify if all the dependencies are met or not using:

```
./cloudkit init --airgap
```

Note: To install the ansible or toolkit dependencies, you can go to [/usr/lpp/mmfs/5.1.8.0/ansible-toolkit/ansible/ansible_library/rhel9](#) and execute **`pip3 install ansible-5.9.0.tar.gz -f ./ --no-index --no-compile --disable-pip-version-check --no-cache-dir`**.

cloudkit using air gap option

Perform the following steps to execute the cloudkit using the air gap option:

1. To create a cluster in the air gap based mode, use the option below.

```
./cloudkit create cluster --airgap
```

2. To grant a filesystem access to a compute cluster in the air gap based mode, use the option below.

```
./cloudkit grant filesystem --airgap
```

3. To revoke a filesystem access provided to the compute cluster in the air gap based mode, use the option below.

```
./cloudkit revoke filesystem --airgap
```

Limitations

- Create repository and image functionalities are not supported in air gap mode.
- Grant and revoke filesystem functionalities are supported to compute cluster only.

Other considerations

Consider security, monitoring, and maintenance aspects while you are planning for the deployment of IBM Storage Scale on public clouds.

Planning for accessing a new IBM Storage Scale cluster

When accessing the IBM Storage Scale cluster from an on-premise environment, it is recommended to use a reliable and secure network connection between on-premise and cloud environment.

The cloudkit allows following way to connect to the IBM Storage Scale cluster on the cloud.

- Jump host or bastion host
- Direct
- Cloud VM

A direct connect entails either of these options:

- Directly connected leased, dedicated, and private connection
- A VPN encrypted connection established between the on-premise data center and cloud VPC.

If a jump host is preferred, it can be used as an intermediary server to provide secure access to other resources within the cloud network.

With a jump host, users first connect to it and then use it to connect to other resources in the network. This can be useful in scenarios where direct connectivity is not possible or desired.

If Cloud VM is preferred, it connects from a Cloud VM deployed on same cloud.

```
? Connectivity method to cloud:  [Use arrows to move, type to filter]
> JumpHost | Recommended when cloudkit is executed from an on-premise/local VM and your network
cannot reach your cloud private network (Provisions an intermediary jump host)
  Direct | Recommended when cloudkit is executed from an on-premise/local VM and your network
is directly connected to your cloud network (e.g. AWS Direct Connect)
  Cloud-VM | Recommended when cloudkit is executed from a cloud VM and you want to provision
the scale VM(s) using the same VPC as the cloud VM
```

Planning for security

While cloudkit configures the scaled instances with limited permissions during deployment, it is important to note that security is ultimately the responsibility of the user. It is recommended to ensure adequate security measures are in place to protect the resources and data stored in the cloud environment. This includes using strong authentication and access controls, encrypting sensitive data, regularly monitoring, and auditing access and activity logs, and keeping software and applications up to date with the latest security patches.

To ensure the security of your bastion private SSH key, it should be protected with file permissions of 400 or 600. This will restrict access to only the key's owner and prevent unauthorized access.

Monitoring and maintenance

Regularly monitor the performance of your filesystem using IBM Storage Scale GUI / system health components and perform routine maintenance to ensure it continues to meet your requirements.

Infracost API key generation

Infracost is a tool that helps organizations estimate the cost of cloud infrastructure. To use the infracost bundled with cloudkit, generate an API key from the infracost website using the following process:

1. Sign up for an account on the infracost website.
2. Navigate to the **API Keys** section of your account settings.
3. Click the **New API Key** button to generate a new API key.
4. Copy the API key to your clipboard and use it when running infracost commands.

Infracost provides a quick and easy way to estimate the cost of cloud infrastructure before you deploy it, which can help you avoid unexpected charges and better plan your budget.

Configuring the IBM Storage Scale environment in the AWS Cloud

Setting up a VPC with the default parameters builds the IBM Storage Scale environment in the AWS Cloud with the following properties.

Note: You can choose to create a new VPC for the IBM Storage Scale deployment or use your existing VPC on AWS. The template that deploys the setup into an existing VPC skips the first four components of this list:

- A VPC that spans either a single Availability Zone or three Availability Zones. A single Availability Zone configuration that includes one public and one private subnet. A multi Availability Zone configuration includes three public and three private subnets that are spread across each Availability Zone.
- An internet gateway to allow access to the internet.
- The managed NAT gateways in the public subnets allow outbound internet access for resources in the private subnets.
- A Bastion host in a public subnet provides Secure Shell (SSH) access to the IBM Storage Scale cluster. The Bastion host instance is managed by an auto scaling group of one, which ensures that there is always at least one host available.
- An AWS Identity and Access Management instance role with fine-grained permissions for access to AWS services necessary for the deployment process.
- Appropriate security groups for each instance or function to restrict access to only necessary protocols and ports. The AWS offering opens only ports for SSH and the IBM Storage Scale daemon.




Figure 43. Single availability zone architecture diagram




Figure 44. Multi availability zone architecture diagram

Each IBM Storage Scale storage node, also referred to as the IBM Network Shared Disk storage server or NSD server, has:

- A100-GB Amazon Elastic Block Store (Amazon EBS) volume for the root device.
- By default, one 500-GB EBS volume is attached for use as an NSD storage per NSD server. You can change the number and size of the NSD storage EBS volumes that is attached per NSD server during the deployment.

Each IBM Storage Scale compute instance has:

- A 100-GB EBS volume for the root device.
- If a multi Availability Zone configuration is defined, an extra IBM Storage Scale node is created to improve the resiliency of the IBM Storage Scale cluster. The extra node is created to avoid loss of file system access after a disk failure. This node has an extra 5 GB EBS volume that is attached to it, and fulfills the role of a tiebreaker for the node and the file system descriptor quorum decisions. For more information, see the *Synchronous mirroring with GPFS replication* section in the IBM Storage Scale documentation.

Note: In IBM Storage Scale terminology, the disk is called a descriptor quorum disk, and it is added to avoid a file system descriptor quorum loss.




Figure 45. Synchronous mirroring by using GPFS replication

Planning for AFM

The following topics assist you in planning for AFM.

Requirements for UID and GID on the cache and home clusters

This topic describes the UID and GID requirements on the cache and home clusters.

User IDs and group IDs must be managed the same way across the cache and home clusters. However, for AFM relationships, which are using the native GPFS protocol, where user IDs are different on the home and the cache, you can remap the IDs by using the GPFS UID remapping. For more information about ID mapping, see *Configuring ID mappings in IMU* in *IBM Storage Scale: Administration Guide*.

Recommended workerThreads on a cache cluster

Specifies the number of **workerThreads** on a cache cluster.

It is recommended to increase the **workerThreads** to twice the default value while you use GPFS backend. For NFS protocol, the number of **workerThreads** can be equal to the default value. The default value is 48.

Table 38. workerThreads value per backend	
Backend	Minimum workerThreads value
NFS	48 (default)
GPFS	96 (2*default)

Inode limits to set at cache and home

To control the number of inodes that are used by a cache or home fileset, you can specify a limit on the number of inodes for that fileset during fileset creation time. For more information about the

--inode-limit option, see *mmcrfileset command* in *IBM Storage Scale: Command and Programming Reference Guide*.

If a home is a dependent fileset in the file system, the inode limit used are the ones for the independent fileset where the dependent fileset is located. If the dependent fileset is located in the "root" fileset then the file system limits are used. However, if the dependent fileset is located in an independent fileset that is not the "root" fileset then the inode limits for that independent fileset are used.

Note: AFM does not perform eviction of file inodes, even if the inode limit is reached. The cache eviction is applicable only for file data block quotas.

Planning for AFM gateway nodes

An AFM cache cluster must include minimum one node running the Linux operating system. This node can be assigned the gateway nodes role for the AFM cache filesets. Every AFM fileset is assigned to a unique AFM gateway, which is called as a primary gateway. This gateway handles all the communication between the AFM fileset and the target on the home site for synchronization.

You can assign and remove a gateway node role by using the following commands:

- To assign gateway node role to the cluster member node, issue the following command:

```
mmchnode -N node1,node2,... --gateway
```

- To remove the gateway node role from the cluster member node, issue the following command:

```
mmchnode -N node1,node2,... --nogateway
```

The AFM gateway node is licensed as a server node.

Multiple gateway nodes

An IBM Storage Scale cluster can have one or more gateway nodes. Multiple gateway nodes are helpful if you want to assign different filesets to a different gateway node and configure parallel data transfer. The assignment of the gateway node to the AFM cache fileset is based on the Hashing Algorithm. You can also assign a gateway node to an AFM cache fileset.

For more information about Parallel Data Transfer, see [“Parallel data transfers” on page 77](#).

AFM is enabled with IBM Storage Scale Data Access Edition.

For more information about the licensing, see [“IBM Storage Scale product editions” on page 298](#).

General recommendations for AFM gateway node configuration

To configure an AFM gateway node, the general recommendations are as follows:

- The memory and CPU requirement for a gateway node depends on the number of assigned AFM, AFM-DR fileset, or AFM to cloud object storage fileset and files in the filesets.
- Minimum 128 GB of memory is needed for the gateway node.
- An AFM gateway node can be assigned up to 20 AFM, AFM-DR, or AFM to cloud object storage filesets.
- A gateway node is a dedicated node in the cluster without any other designations such as NSD server, CES protocol nodes, quorum, and manager.

The following table contains the cluster and fileset-level parameters:

Parameter	Value	Level
Pagepool	8G	Cluster
afmHardMemThreshold	40G	Cluster
afmDIO	2	Cluster

Parameter	Value	Level
maxFilestoCache	10000	Cluster

To set these parameters, use the **mmchconfig** command. For example,

```
# mmchconfig pagepool=8G
```

For more information about configuration parameters, see *Parameters for performance tuning and optimization* in the *IBM Storage Scale: Administration Guide*.

Note:

- All these recommendations are based on observations in a controlled test environment by running moderate or reasonable workloads. The parameters values might vary based on the setup or workload.
- These recommendations are not limitations.

General guidelines and recommendations for AFM SW and IW mode fileset

Use the IBM guidelines and the recommendations to configure and maintain AFM Single Writer (SW) and Individual Writer (IW) mode fileset.

The IBM guidelines for the deployment of IBM Storage Scale AFM SW/IW are as follows:

AFM filesets and files recommendation

- The number of AFM SW/IW mode filesets that can be configured in an IBM Storage Scale cluster is up to 100.
- The number of files in an IBM Storage Scale AFM SW/IW environment is approximately 100 million files in every fileset.

Gateway nodes recommendation

The number of gateway nodes that can be configured in an IBM Storage Scale cluster depends upon the number of AFM filesets. Keep AFM gateway nodes equal to 1/10th of the number of filesets. That is, for 20 AFM SW/IW filesets in a cache cluster, the proposed number of gateway nodes is '2'. So that, every gateway can be assigned up to 10 filesets. The gateway node must have a dedicated role and no other role must be assigned to the gateway node. Gateway nodes are needed at cache Site only. AFM does not need any Gateway node at the home or target site.

Memory and storage recommendation for a gateway node

Random Access Memory (RAM) on the gateway node needs 128 GB. The /var partition space storage needs to be provisioned for bigger storage (based on the number of filesets handled by gateway node * files in fileset * 255 bytes) space for internal usage during the recovery or resynchronization.

Configuration parameters recommendation for AFM

Tunable	Value
Pagepool	8G
afmHardMemThreshold	40G
afmNumFlushThreads	8
afmDIO	2
maxFilestoCache	10000
afmMaxParallelRecoveries	3

To set these preceding parameters, use the **mmchconfig** command. For example, **mmchconfig pagepool=4G**. For more information about configuration parameters, see *Parameters for performance tuning and optimization* in the *IBM Storage Scale: Administration Guide*.

Network and computing requirement

Plan necessary network bandwidth between cache and home clusters to replicate incoming changes based on a workload. This requirement varies based on the workload. The application is slower on AFM SW/IW filesets because of the message that waiting on the gateway node for the replication. This performance hit varies from 1.2x to 4x based on network and gateway node resources. If the gateway node handles many filesets, more computational power is needed for filtering and dependency checking logic.

Limitations

For more information about the limitations of AFM, see [AFM and AFM DR limitations](#).

Planning for AFM DR

The following topics assist you in planning for AFM DR.

Requirements for UID/GID on primary and secondary clusters

User IDs and group IDs must be managed the same way across the primary and secondary clusters.

However, for AFM DR relationships, which use native GPFS protocol, where user IDs are different on the primary and secondary clusters, the IDs may be mapped again by using GPFS UID remapping.

Recommended worker1Threads on primary cluster

IBM recommends that you increase the **worker1Threads** to twice the number of AFM DR filesets that are using GPFS backend, which exist on all file systems in the cache cluster.

For example, if a cluster has 50 primary fileset that is using NFS protocol and 60 primary using GPFS protocol, the **worker1Threads** must be at least 160 ($50 + 2 \times 60$). For NFS protocol, the number of **worker1Threads** can be equal to the number of primary filesets.

NFS setup on the secondary cluster

NFS setup on the secondary cluster follows the same rules as the NFS setup on the AFM home cluster.

General guidelines and recommendations for AFM-DR

AFM-DR is a disaster recovery solution, which is by nature a business-critical solution. Its success depends on administration discipline, including careful design, configuration, and testing. Considering this, IBM requires that the AFM Async DR (AFM-ADR) deployment must follow certain guidelines and recommendations for the deployment. AFM DR is enabled by default. You can follow the recommendations and guideline to configure AFM DR.

However, the general recommendations and prerequisites for the AFM-DR setup for general usage might differ for use-case specific configurations. The general recommendations help to understand the requirement better for the AFM-DR deployment.

Guidelines or recommendations for the deployment of IBM Storage Scale AFM-DR are as follows:

AFM-DR filesets and files recommendation

- The number of AFM DR filesets that can be configured in an IBM Storage Scale cluster is up to 100.
- The number of files in an IBM Storage Scale AFM-ADR environment is approximately 100 million files in every fileset.

Gateway nodes recommendation

- The number of gateway nodes that can be configured in an IBM Storage Scale cluster depends upon the number of AFM filesets.
- Keep AFM gateway nodes equal to 1/10th of the number of filesets. That is, for 20 ADR filesets in a primary cluster, the recommended number of gateway nodes is '2'. So that, every gateway can be assigned up to 10 filesets.

- The gateway node must have a dedicated role and no other role must be assigned to the gateway node.
- Gateway node(s) are needed at primary Site only. AFM or AFM-DR do not need any Gateway node at the secondary site.

Memory and storage recommendation for a gateway node

- Random access memory (RAM) on the gateway node should be 128 G.
- The /var partition space storage needs to be provisioned for bigger storage (based on the number of filesets handled by gateway node * files in fileset * 255 bytes) space for internal usage during the recovery or resynchronization.

Configuration parameters recommendation for AFM-ADR

Tunable	Recommended value
Pagepool	8G
afmHardMemThreshold	40G
afmNumFlushThreads	8
afmDIO	2
maxFilesToCache	10000
afmMaxParallelRecoveries	3

To set these parameters, use the **mmchconfig** command. For example,

```
# mmchconfig pagepool=4G
```

For more information about configuration parameters, see *Parameters for performance tuning and optimization* in the *IBM Storage Scale: Administration Guide*.

Network and computing requirement

- Plan necessary network bandwidth between primary and secondary clusters to replicate incoming changes based on a workload.
- This requirement varies based on the workload. The application is slower on AFM DR filesets because of the message queuing on the gateway node for the replication. This performance hit varies from 1.2x to 4x based on network and gateway node resources. If the gateway node handles many filesets, more computational power is required for filtering and dependency checking logic.

Data synchronization for the AFM fileset conversion

- You can convert an existing GPFS-independent fileset to an AFM-ADR fileset by using the `--inband` trucking method. You must not use the `--outband` (deprecated) option even though data was copied by using the external tool.
- If the secondary site has already data, AFM does not copy data during the initial resynchronization. You can still truck (out of band) data by using any other tool such as `rsync` and convert the fileset by using the `--inband` option. AFM expects that files are in sync with nanoseconds time difference. Use `rsync` version `>= 3.10` on both source and destination to copy the data and later use the `inband` conversion. The `inband` conversion skips the data, if data exists (and matches) on the secondary site.

Limitations

For more information about the limitations of AFM-DR, see [“AFM and AFM DR limitations” on page 128](#).

Planning for AFM to cloud object storage

The considerations for the planning AFM to cloud object storage relation are as follows:

The AFM to cloud object storage relation is between an AFM to cloud object storage fileset that is hosted on an IBM Storage Scale cluster and cloud object storage bucket, or a directory within an AFM to

cloud object storage fileset and a cloud bucket. For more information, see [“AFM to cloud object storage operation modes” on page 156](#). The important elements that need consideration in the IBM Storage Scale cluster are the host system, storage, operating system, and networking. For more information, see [“Planning for GPFS” on page 309](#).

Before you set the AFM to cloud object storage relation, consider the following points:

1. Cloud object storage such as Amazon S3, Microsoft Azure Blob, and IBM Cloud Object Storage must be in accordance with the workload requirement. Access keys and secret keys that are provided by these services need to be set up for AFM to cloud object storage filesets.

To transfer the data between AFM to cloud object storage and cloud object storage endpoints, select an appropriate protocol, for example, HTTP or HTTPS.

2. On an IBM Storage Scale cluster, file system space and inodes (number of inodes) need to be planned for an AFM to cloud object storage fileset first before provision. For more information, see [“File system creation considerations” on page 345](#).

To control the number of inodes that are used by the AFM to cloud object storage fileset, you can specify a limit on the number of inodes for that fileset. To set inodes limit, use the **mmchfileset** command after the **mmafmcconfig** command is run to set the relationship. For more information, see *mmchfileset command* in the *IBM Storage Scale: Command and Programming Reference Guide*.

3. Planning of gateway nodes

An IBM Storage Scale cluster can have one or more gateway nodes. Multiple gateway nodes are helpful, if you want to assign different filesets to a different gateway node. The multiple nodes on a gateway node are useful to load balance the workload and also helpful for recovery and upgrade scenarios.

Gateway nodes need to have access to cloud object storage endpoints. For more information about gateway node recommendation, see [“General recommendations for AFM gateway node configuration” on page 453](#).

4. Network considerations

Plan necessary network bandwidth between , and IBM Storage Scale clusters to replicate data and workload.

The networking element is the major contributor in the overall performance delivery of the IBM Storage Scale cluster where a few nodes act as Network Shared Disk (NSD) servers while most of the client nodes access the file system over a TCP/IP network. Gateway nodes in this host system connect to the outside world, which in this case is Amazon S3 cloud object storage, Microsoft Azure Blob, or IBM Cloud Object Storage.

All the firewalls are set up to allow the gateway nodes to discover and communicate with these cloud storage endpoints.

5. When you work with temporary credentials for accessing Cloud Object Storage such as Amazon S3, Microsoft Azure Blob, IBM Cloud, and other services, you must obtain the security token and create fileset by specifying **--user-keys** option of **mmafmcconfig** command.

The `mmuid2keys (/var/mmfs/etc/mmuid2keys)` file must be configured to provide credentials in the following format as shown:

```
Access_Key:Secret_Key:STS
```

6. Licensing requirement:

To use IBM Storage Scale multi-site replication feature, upgrade to advance/DM/ECE edition.

Firewall recommendations for AFM to cloud object storage

Consider these firewall recommendations for the AFM to cloud object storage setup.

To enable communication between a cloud object storage server and IBM Storage Scale AFM to cloud object storage, firewall must keep open a specific port or ports on all the gateway nodes. AFM to cloud

object storage communicates to a cloud object storage server on a specified port such as 80 and 443. The specified port is provided as a part of an endpoint for a fileset. Different filesets can have a different port. For this communication, a specific port needs to be opened or enabled for the AFM to cloud object storage.

Planning for IBM Storage Scale native REST API

Learn about the hardware and software requirements, prerequisites, and other important planning information that you must consider before the deployment of the IBM Storage Scale native REST API feature.

Supported platforms

IBM Storage Scale native REST API is supported on all Linux operating systems that are supported by IBM Storage Scale, except where noted in the Feature exceptions for Linux table in the IBM Storage Scale FAQ. For more information, see [IBM Storage Scale FAQ](#).

This feature is not supported on Windows and AIX operating system nodes within a cluster.

Note: The entire IBM Storage Scale cluster must consist of nodes with supported operating system and architecture. Do not intermix unsupported nodes with native REST API clusters.

Support for IPv6

The native REST API feature supports both IPv4 and IPv6 network configurations. This feature allows the administrator to configure the following components with IPv6 protocol support:

- Native REST API administration daemon **grpc** interface.
- Native REST API administration daemon REST interface
- **mm**-commands administrative node-to-node interface
- I/O daemon interface

Networking prerequisites

- Ensure that a local IPv6 network and a valid IPv6 address are configured on a local network interface. The address must not be a link-local IPv6 address. For example:

```
ip addr
6: ens12: <BROADCAST,MULTICAST,UP,LOWER_UP> mtu 1450 qdisc fq_codel state UP group
default qlen 1000

    link/ether fa:16:3e:15:44:f7 brd ff:ff:ff:ff:ff:ff
    altname enp0s12
    inet6 fd9f:f0d0:1002:11::1c3/64 scope global
        valid_lft forever preferred_lft forever
```

- Ensure that both hostname and reverse IP address lookups can be performed. To enable consistent resolution, configure the IPv6 address and hostname mappings in the `/etc/hosts` file on all nodes in the cluster. Each node's `/etc/hosts` file should include the IPv6 address and hostname pairs for all nodes in the cluster. You can verify the configuration by pinging the host's IPv6 address locally and remotely. For example:

```
ping mos-dev-3-ipv6-ens12

PING mos-dev-3-ipv6-ens12(mos-dev-3-ipv6-ens12 (fd9f:f0d0:1002:11::1c3)) 56 data bytes
64 bytes from mos-dev-3-ipv6-ens12 (fd9f:f0d0:1002:11::1c3): icmp_seq=1 ttl=64 time=0.048
ms
64 bytes from mos-dev-3-ipv6-ens12 (fd9f:f0d0:1002:11::1c3): icmp_seq=2 ttl=64 time=0.055
ms
```

```
ping -6 fd9f:f0d0:1002:11::1c3
PING fd9f:f0d0:1002:11::1c3(fd9f:f0d0:1002:11::1c3) 56 data bytes
64 bytes from fd9f:f0d0:1002:11::1c3: icmp_seq=1 ttl=64 time=0.052 ms
64 bytes from fd9f:f0d0:1002:11::1c3: icmp_seq=2 ttl=64 time=0.057 ms
64 bytes from fd9f:f0d0:1002:11::1c3: icmp_seq=3 ttl=64 time=0.054 ms
```

Configuring the administration daemon for IPv6

The IBM Storage Scale native REST API administration daemon is started through the **systemd** interface at the end of the package installation process. The **scaleadmd** service starts with default settings, which include network addresses and port numbers for the **grpc** server and REST API interface. By default, the service uses the IPv4 *any* interface notation for the following settings:

```
remote_grpc_server.host = "0.0.0.0"
remote_grpc_server.port = "50052"
```

To configure the administration daemon **grpc** server, REST API interface or both for a specific IPv6 address and a single network interface, specify the IPv6 address on the **host** option to either one or both components. To update the configuration values, use the configuration APIs or CLI. For more information, see *scalectl config command* in the *IBM Storage Scale: Command and Programming Reference Guide*. In the following example, "fd9f:f0d0:1002:11::1c3" is the IPV6 address:

```
remote_grpc_server.host = "[fd9f:f0d0:1002:11::1c3]"
remote_grpc_server.port = "50052"
rest_server.host = "[fd9f:f0d0:1002:11::1c3]"
rest_server.port = "46443"
```

Note: Administration daemons use the `remote_grpc_server` definition to determine the address and port for communication with other administration daemons. This configuration is required when you run **scalectl cluster create** or **scalectl node add** commands. You must specify a host name or IP address for the node. If the administration daemon is configured with an IPv6 address, specify the corresponding admin host name that maps to the IPv6 address in the **scalectl cluster create** or **scalectl node add** CLI or REST API methods. When specifying the IPv6 address directly, enclose it in square brackets.

Restart the IBM Storage Scale native REST API administration daemon service with the **systemd** interface to start listening with the new IPv6 addresses and interfaces.

```
systemctl restart scaleadmd
```

The output from the **systemctl status scaleadmd** command must display the messages that state the IPv6 address for the **grpc** and REST API services are using.

You can verify that the IBM Storage Scale native REST API administration daemon is running with the IPv6 protocol by issuing the following command:

```
netstat -tulnp6 | grep scaleadmd

tcp6      0      0 fd9f:f0d0:1002:11:50052 :::*          LISTEN      2039323/./
scaleadmd

tcp6      0      0 fd9f:f0d0:1002:11:46443 :::*          LISTEN      2039323/./
scaleadmd
```

scalectl client considerations

The IBM Storage Scale native REST API **scalectl** CLI can send requests over HTTPS. Use the **--url** option to specify the endpoint for secure communication. When specifying an IPv6 address, enclose the address in square brackets (**[IPv6]:port**), as shown in the following example:

```
scalectl cluster list --domain StorageScaleDomain --json --url  
[fd9f:f0d0:1002:11::1c3]:46443 --insecure-skip-tls-verify
```

When you configure a local **.netrc** file, enclose the IPv6 address in square brackets as shown in the following example:

```
cat ~/.netrc  
machine [fd9f:f0d0:1002:11::1c3]  
  login root  
  password yourPassword
```

Creating nodeID certs for IPv6 support

When you create node ID certificates for the cluster, specify the IP address for each host during certificate generation. These addresses must match the ones used in the administration daemon **grpc** and REST API interface configurations. For IPv6 configurations, use the raw address (without brackets) during certificate creation. After the certificates are generated, import them on each node that you are creating or adding to the cluster.

Create the cluster with IPv6

When you configure the IBM Storage Scale native REST API administration daemon **grpc** interface to run or listen on an IPv6 address, use either the host name that maps to the IPv6 address or the address itself enclosed in brackets with the **-N** or **--admin-node-name** option. This configuration sets how nodes within the cluster can contact the administration daemon running on those IPv6 address. It also sets the IP for the classic **mm** commands to reach remote nodes. If you use a direct IPv6 address, enclose it in square brackets as shown in the following example:

```
scalectl cluster create -N [fd9f:f0d0:1002:11::1c3] --admin-node-name  
[fd9f:f0d0:1002:11::1c3]  
  or a host name mapped to the IPv6 address
```

```
scalectl cluster create -N ipv6-host name
```

Note: The **-N** option specifies the host name or IP address of the node for the GPFS I/O daemon to I/O daemon communication. If it is an IPv6 address, it must be wrapped in square brackets. If **--admin-node-name** is not specified, the value from **-N** is used for administration daemon communication as well.

To verify that the **mmsdrserv** process is using the correct IPv6 address when the I/O daemon is not started, run the following command:

```
netstat -tulnp6 | grep mmsdrserv
```

A sample output is as follows:

```
tcp6      0      0 :::1191          :::*
```

To verify the **mmfsd** process when the I/O daemon is running, run the following command:

```
netstat -tulnp6 | grep mmfs
```

A sample output is as follows:

```
tcp6      0      0 :::1191          :::*              LISTEN           3079124/  
mmfsd
```

```
netstat -tulnp6 | grep mmfs
```



```
tcp6      0      0 :::1191      :::*          LISTEN     3079124/
mmfsd
```

Add a node with IPv6

When you add a new node for IBM Storage Scale native REST API administration daemon grpc interface to run or listen on an IPv6 address, use the host name that maps to that IPv6 address or the direct IP address with the **-N** or **--admin-node-name** option. This configuration sets how nodes within the cluster can contact the administration daemon running on that IPv6 address. If you use a direct IPv6 address, enclose it in square brackets as shown in the following example:

```
scalectl node add -N [fd9f:f0d0:1002:11::1c3] --admin-node-name [fd9f:f0d0:1002:11::1c3]
or a host name mapped to the IPv6 address
```

```
scalectl node add -N ipv6-host name --admin-node-name [fd9f:f0d0:1002:11::1c3]
```

Note: The **-N** option specifies the host name or IP address of the node for the GPFS I/O daemon to I/O daemon communication. If it is an IPv6 address, it must be wrapped in square brackets. If **--admin-node-name** is not specified, the value from **-N** is used for administration daemon communication as well.

Planning for performance monitoring tool

The performance monitoring tool (perfmom) planning for the memory and filesystem provisioning on the nodes that are running the pmcollector service.

The pmcollector service maintains an in-memory database that stores performance metrics collected from IBM Storage Scale Systems. With the default perfmom configuration, a single collector instance can support up to 150 sensor nodes and 4 IBM Storage Scale Systems building blocks, assuming Fabric Hospital monitoring is enabled. For large-scale compute clusters, it may be beneficial to activate the majority of sensors only on I/O nodes, while enabling a reduced subset of sensors on compute nodes. This approach helps to optimize resource usage and improves scalability.

Some sensors are computationally intensive, and modifying the perfmom configuration to run these sensors more frequently can significantly reduce the number of nodes that a single collector can support. For more information on sensor impact, see *Impact of sensor configuration* in *IBM Storage Scale: Problem Determination Guide*. To further enhance scalability, the performance monitoring tool supports a multi-collector federation configuration, allowing multiple collectors to be deployed across the environment.

For environments with very large clusters, such as those used in supercomputing, customers are encouraged to contact IBM support for guidance on planning and tuning perfmom settings to meet specific scalability and performance requirements.

The performance monitoring tool is not supported for AIX and Window platforms. For more information, see feature exceptions in the [IBM Spectrum Scale FAQ Q2.1/A2.1 Supported OS and software versions](#).

Memory provisioning

The following table presents a rough estimation of memory that is consumed by pmcollector for the different cluster configurations. The memory limits that are specified here only considers the pmcollector memory requirements and needs to be added to the memory provisioned. For example, the GPFS pagepool or the IBM Storage Scale GUI.

Table 39. Monitoring data consumption limits by pmcollector memory	
Cluster configuration	Recommended pmcollector memory provisioning
Clusters up to 10 nodes	8 GB (Minimum) and 16 GB (recommended)
Clusters up to 100 nodes	32 GB
Clusters with more than 100 nodes	64 GB

The numbers are based on the default perfmon sensor configuration that is created by the **mmperfmon** config generate command. Changing the default settings or adding more sensors, for example, for AFM, protocols or other features, can increase the amount of required memory.

File space provisioning

The pmcollector service requires 20 GB file system space under /opt/IBM/zimon/ folder.

Performance monitoring limitations

The following section lists the limitations of the performance monitoring tool:

- Performance monitoring is not supported on the following operating systems:
 - x86_64/sles11
 - ppc64le/Ubuntu14.04
 - ppc64/aix
- The interface for NFS of the performance monitoring tool does not work on SLES 11 nodes.
- The performance monitoring tool does not support systems with mixed endianness. For example, big-endian nodes (s390x) and little-endian nodes (x86_64 or ppc64le) are not supported in the same cluster.

Important: When you are using the performance monitoring tool, ensure that the clocks of all the nodes in the cluster are synchronized. It is recommended to configure the Network Time Protocol (NTP) on all nodes.

Planning for UEFI secure boot on x86_64 and secure boot on Linux on Z

Starting with 5.1.9.0, IBM Storage Scale supports the UEFI secure boot feature on x86_64; 5.2.1.0 onward, secure boot is supported on Linux on Z. For IBM Storage Scale, using secure boot means that the kernel modules are cryptographically signed by IBM so that their integrity can be verified when the system starts.

To use signed kernel modules for IBM Storage Scale on x86_64, the following prerequisites must be met:

- UEFI and secure boot are enabled in the cluster nodes.

```
sudo mokutil --sb-state
```

A sample output is as follows:

```
SecureBoot enabled
```

- The `pesign`, `mokutil`, `openssl`, `kernel-devel`, and `keyutils` tools are installed.
- An RPM package with the signed kernel modules is available. The kernel version of these modules must match the kernel version that runs on the system.
- RHEL 9.2 or later is the operating system for x86_64.
- The public key that is provided by IBM for the verification of the signed kernel modules is available. The public key is either part of the self-encrypting package or is included as part of the signed kernel stand-alone package.

To use signed kernel modules for IBM Storage Scale on Linux on Z, the following prerequisites must be met:

- Secure boot is enabled in the cluster nodes.
- An RPM package with the signed kernel modules is available. The kernel version of these modules must match the kernel version that runs on the system.
- RHEL 9.4 or later is the operating system for Linux on Z.
- The public key is imported by using the IBM Z Hardware Management Console (HMC).

Firewall recommendations

The IBM Storage Scale system administrator is supposed to follow certain recommendations to set up firewall to secure the IBM Storage Scale system from unauthorized access.

For more information about firewall recommendations, see *Securing the IBM Storage Scale system using firewall* in *IBM Storage Scale: Administration Guide*.

Considerations for GPFS applications

Application design should take into consideration the exceptions to Open Group technical standards with regard to the `stat()` system call and NFSv4 ACLs. Also, a technique to determine whether a file system is controlled by GPFS is provided.

For more information, see the following topics in *IBM Storage Scale: Administration Guide*:

- *Exceptions to Open Group technical standards*
- *Determining if a file system is controlled by GPFS*
- *GPFS exceptions and limitation to NFSv4 ACLs*
- *Considerations for GPFS applications*

Security-Enhanced Linux support

IBM Storage Scale runs on Red Hat Enterprise Linux operating systems with Security-Enhanced Linux (SELinux).

IBM Storage Scale supports SELinux that is set to one of the available SELinux modes such as enforcing or permissive and the SELinux policy set to a targeted policy.

For more information about SELinux support, see the following topics:

- See *Monitoring file audit logging with rsyslog and SELinux* in the *IBM Storage Scale: Problem Determination Guide* for file audit logging.

Note: IBM Storage Scale daemons runs in unconfined domains.

Space requirements for call home data upload

Certain space requirements must be met by the system before uploading the call home data. The following section details the space requirements for call home.

The call home component uses the directory which is specified in the IBM Storage Scale settings variable `dataStructureDump`, for saving the temporary data. By default, this directory is `/tmp/mmfs`, but it can be changed by the customer by using the `mmchconfig` command. The current value can be read by running the following command:

```
mmdiag --config | grep "dataStructureDump"
```

The space requirements differ depending on the use case. The following two situations are possible:

- For uploading daily and weekly packages, and non-PMR or Salesforce case-related files you require space equivalent to two times the size of the file that is to be uploaded. For example, if you need to upload 1 GB of data, then there needs to be at least 2 GB of disk space for the file to be properly uploaded.
- For uploading the ticket-related files, a dynamical buffered chunking system is used. In this case, the minimum space that is required for an upload is 1/250 of the file size or 0.4% of the size of the file to be uploaded. At least 200 MB of space is recommended to minimize the chunking overhead and maximize the transfer speed.

Shared memory communications direct planning

Before you configure shared memory communications direct (SMC-D), ensure that all requirements are met.

SMC-D requirements are as follows:

- Hardware: IBM z16 or later
- Operating Systems. RHEL8: 8.9 or later, RHEL9: 9.3 or later, SLES: 15 SP5 or later
- All IBM Storage Scale nodes must be located on the same IBM Z Central Processor Complex (CPC)

Planning for data tiering

Plan your data tiering implementation by understanding requirements, validation checks, and filesystem configuration options.

Before planning for data tiering, ensure you have:

- IBM Storage Scale 6.0.1.0 or later installed
- Understanding of storage pool concepts
- Fast storage devices available (NVMe, SSD)
- Knowledge of your workload access patterns

Before implementing data tiering, you must plan your storage pool configuration and understand the requirements and limitations. Proper planning ensures optimal performance and cost efficiency.

Requirements for the fast tier pool

The data tiering feature requires a fast tier pool in the file system.

The following requirements must be met:

- Create the pool with the `fastTierPool=yes` attribute.
- Define the pool during file system creation or while adding disks.
- Use high-performance devices such as NVMe or SSD drives.
- Configure the pool as either `dataOnly` or `dataAndMetadata`.
- Ensure that only one fast tier pool exists in the file system.

Requirements for the file system

A file system that uses data tiering must contain the following storage components:

- System pool:

Required for metadata storage. The system pool can also act as the fast tier pool if it is configured as `dataAndMetadata`.

- Regular data pool

Required to store blocks that are outside the configured fast tier range. This can be the system pool or a separate data pool.

- Fast tier pool

Required for data tiering. Stores the initial file blocks that fall within the configured threshold.

Validation and limitations

These are the operational limitations:

- Existing pools cannot be modified into fast tier pools after creation.

- A fast tier pool designation cannot be modified.
- Only one fast tier pool is allowed per file system.
- Migration policy rules that use fast tier ranges fail validation if no fast tier pool exists.
- Data tiering settings apply only to regular files, not directories.
- File systems that use both the fastTierPool and performance pool are not supported.
- If the fast tier range changes after a snapshot is created, a file might be placed incorrectly when it is restored from the snapshot by using the `mmrestorefs` command. If you do not need to restore the fast tier range, specify the `--suppress-external-attributes` option when you run the `mmrestorefs` command.

Chapter 3. Establishing and starting your IBM Storage Scale cluster

There are several steps that you must perform to establish and start your IBM Storage Scale cluster. This topic provides the information that you need for performing those steps.

About this task

The installation toolkit, available on a few Linux distributions, automates the following many steps. For more information, see [“Installing IBM Storage Scale on Linux nodes with the installation toolkit”](#) on page 509.

You can install IBM Storage Scale and deploy protocols either manually or by using the installation toolkit. This topic provides the information that you need for establishing and starting your IBM Storage Scale cluster manually. If you have already installed IBM Storage Scale with the installation toolkit, then these steps are already completed.

Tip: Ensure that there is approximately 5 GB of free space on the nodes to download, extract, and install the packages.

Follow these steps to establish an IBM Storage Scale cluster:

1. Review supported hardware, software, and limits in the [IBM Storage Scale FAQ](#) in [IBM Documentation](#) for the latest recommendations on establishing an IBM Storage Scale cluster.
2. Install the IBM Storage Scale licensed program on your system:
 - For existing systems, see [Chapter 20, “Upgrading,”](#) on page 671.
 - For new systems:
 - For your Linux nodes, see [Chapter 4, “Installing IBM Storage Scale on Linux nodes and deploying protocols,”](#) on page 469.
 - For your AIX nodes, see [Chapter 5, “Installing IBM Storage Scale on AIX nodes,”](#) on page 575.
 - For your Windows nodes, see [Chapter 6, “Installing IBM Storage Scale on Windows nodes,”](#) on page 579.
3. Decide which nodes in your system you want to be the quorum nodes (see [“Quorum”](#) on page 314).
4. Create your GPFS cluster by issuing the `mmcrcluster` command. See [“GPFS cluster creation considerations”](#) on page 319.
5. Use the `mmchlicense` command to assign an appropriate GPFS license to each of the nodes in the cluster. For more information, see [“IBM Storage Scale license designation”](#) on page 300 .

If you use the installation toolkit to install IBM Storage Scale, then steps 2 to 5 in the following procedure are completed by the installation toolkit, and step 6 can be optionally completed by the installation toolkit.

After your IBM Storage Scale cluster is established:

1. Ensure you have configured and tuned your system according to the values suggested in the *Configuring and tuning your system for GPFS* topic in *IBM Storage Scale: Administration Guide*.
2. Start IBM Storage Scale by issuing the `mmstartup` command. For more information, see **`mmstartup` command** in *IBM Storage Scale: Command and Programming Reference Guide*.
3. Create new disks for use in your file systems by issuing the `mmcrnsd` command. See [“Network Shared Disk \(NSD\) creation considerations”](#) on page 329.
4. Create new file systems by issuing the `mmcrfs` command. See [“File system creation considerations”](#) on page 345.
5. Mount your file systems by issuing the `mmmout` command.

6. As an optional step, you can also create a temporary directory (/tmp/mmfs) to collect problem determination data. The /tmp/mmfs directory can be a symbolic link to another location if more space can be found there. If you decide to do so, the temporary directory should not be placed in an IBM Storage Scale file system, as it might not be available if IBM Storage Scale fails.

If a problem occurs, IBM Storage Scale might write 200 MB or more of problem determination data into /tmp/mmfs. These files must be manually removed when any problem determination is complete. This should be done promptly so that a NOSPACE condition is not encountered if another failure occurs. An alternate path can be specified by issuing the `mmchconfig dataStructureDump` command.

Chapter 4. Installing IBM Storage Scale on Linux nodes and deploying protocols

Use this information for installing IBM Storage Scale on Linux nodes and deploying protocols.

About this task

Before installing IBM Storage Scale, you must review the [“Planning for GPFS”](#) on page 309 and the [IBM Storage Scale FAQ](#) in IBM Documentation.

Installing IBM Storage Scale without ensuring that the prerequisites listed in [“Hardware requirements”](#) on page 309, [“Software requirements”](#) on page 310, and [“Installation prerequisites”](#) on page 470 are satisfied can lead to undesired results.

Note:

- CES supports NFS, SMB, S3, and HDFS protocols. CES HDFS follows the same installation methods and prerequisites as the other protocols as documented in this section. For specific information about installing CES HDFS, see *Installation* under [CES HDFS](#) in *Big data and analytics support documentation*.
- For information about installing IBM Storage Scale Erasure Code Edition, see IBM Storage Scale [Erasure Code Edition](#) documentation.

The following methods are available for installing IBM Storage Scale on Linux nodes.

Installation method	Description
Installation toolkit	On supported Linux distributions, you can install IBM Storage Scale and deploy protocols by using the installation toolkit. For more information, see “Installing IBM Storage Scale on Linux nodes with the installation toolkit” on page 509. For information on limitations if you are using the installation toolkit, see “Limitations of the installation toolkit” on page 518.
Manual	You can install the IBM Storage Scale packages manually by using commands, such as xpm or yum on Red Hat Enterprise Linux, xpm or zypper on SLES, and dpkg or apt on Ubuntu Linux. For more information, see “Manually installing the IBM Storage Scale software packages on Linux nodes” on page 477.

Deciding whether to install IBM Storage Scale and deploy protocols manually or with the installation toolkit

On supported Linux distributions, you can install IBM Storage Scale and deploy protocols either manually or using the installation toolkit.

Why would I want to use the installation toolkit?

While planning your installation, consider the advantages provided by the installation toolkit. The installation toolkit:

1. Simplifies IBM Storage Scale cluster creation.
2. Automatically creates NSD and file system stanza files.
3. Creates new NSDs, file systems, and adds new nodes.

4. Has a single package manager for all components of a release.
5. Deploys SMB, S3, NFS, and HDFS by automatically pulling in all prerequisites.
6. Configures file and object authentication during deployment.
7. Consistently sets and synchronizes time on all nodes.
8. Deploys the IBM Storage Scale GUI.
9. Configures Performance Monitoring consistently across all nodes.
10. Installs and configures file audit logging and call home functions.
11. Simplifies upgrade with a single command to upgrade all components on all nodes.

Before deciding which method you want to use, review the following topics:

- [“Installation prerequisites” on page 470](#)
- [“IBM Storage Scale packaging overview” on page 473](#)
- [“Understanding the installation toolkit options” on page 516](#)
- [“Limitations of the installation toolkit” on page 518](#)
- [“Using the installation toolkit to perform installation tasks: Explanations and examples” on page 532](#)
- [“Manually installing the IBM Storage Scale software packages on Linux nodes” on page 477](#)

Installation prerequisites

These are the several prerequisites for installing IBM Storage Scale including those for installing protocols and performance monitoring.

Required packages for supported Linux distributions

For the list of packages that must be installed on the system before you can install IBM Storage Scale, see [“Software requirements” on page 310](#).

For a list of supported operating systems, architectures, and kernels, see [IBM Storage Scale FAQ in IBM Documentation](#).

Cleanup required if you previously used the Object Redpaper configuration process

If you have already configured the system with OpenStack software to use GPFS (following instructions from the [Object Red paper](#)) be sure to follow the Object cleanup steps in [“Cleanup procedures required if reinstalling with the installation toolkit” on page 660](#) before you start any installation activity.

Same version required on protocol and quorum nodes

If you want to run protocols in a mixed-version cluster, the protocol nodes and all of the quorum nodes must be running on the same version.

Note: For the object protocol, the GPFS software is not required on the node that hosts the Keystone identity service if it is deployed on a separate node from the GPFS protocol nodes.

Additional GPFS requirements

- If you want to run protocols, Cluster Configuration Repository (CCR) must be available.
- `mmchconfig release=LATEST` must be run.
- A GPFS cluster and a file system are required. If this infrastructure does not already exist, you must install the GPFS software, create a GPFS cluster and create a file system. You can use the installation toolkit to do these tasks. For more information, see [“Using the installation toolkit to perform installation tasks: Explanations and examples” on page 532](#).
- The GPFS file system must be mounted on all GPFS protocol nodes.
- It is strongly recommended to configure the file system to only support NFSv4 ACLs. You can use the installation toolkit to do this task also if you use the installation toolkit to install GPFS. For more information, see [“Using the installation toolkit to perform installation tasks: Explanations and examples” on page 532](#).

Alternatively, you can use the `-k nfs4` parameter for `mmcrfs`. NFSv4 ACLs are a requirement for ACL usage with the SMB and NFS protocols. For more information, see `mmcrfs` examples in the *IBM Storage Scale: Command and Programming Reference Guide*.

- Quotas are not enabled by default in GPFS file systems but are recommended for use with SMB and NFS protocols. For more information about quota management, see *Enabling and disabling GPFS quota management* in *IBM Storage Scale: Administration Guide*.

Create a CES shared root directory and object fileset

The installation toolkit sets up a shared root directory for use by the CES protocol services. The CES shared root (`cesSharedRoot`) is required to store CES shared configuration data and support protocol recovery and other protocol-specific operations.

The starting point of this shared directory tree is called the CES shared root directory. The **`cesSharedRoot`** parameter defines the path to this directory, which is used only by protocol services and is not intended for user data. This directory tree can reside anywhere in a IBM Storage Scale file system. To keep protocol data separate from user data, use a dedicated file system. Once this option is set, changing it requires a restart of GPFS.

You can configure the CES shared root directory by using the installation toolkit, especially if you use the toolkit to install IBM Storage Scale and create file systems. For more information, see [“Using the installation toolkit to perform installation tasks: Explanations and examples”](#) on page 532.

If you prefer to configure the **`cesSharedRoot`** before launching the installation toolkit, follow these steps:

1. Create a file system or fileset to serve as the CES shared root directory. The recommended minimum size is 4 GB.
2. Use the following command to set the CES shared root path::

```
mmchconfig cesSharedRoot=path_to_the_filesystem/fileset_created_in_step_1
```

For object, the installation toolkit creates an independent fileset in the GPFS file system that you name.

SSH and network setup

The node used for initiating installation of SMB and NFS protocols by using the installation toolkit must be able to communicate through an internal or external network with all protocol nodes to be installed. All nodes also require SSH keys to be set up so that the installation toolkit can run remote commands without any prompts. Examples of prompts include a prompt for a remote node's password or a prompt for a yes or no question. No prompts should occur when using SSH among any cluster nodes to and from each other, and to and from the node designated for installation.

For information on ports required by IBM Storage Scale, see *Securing the IBM Storage Scale system using firewall* and *GPFS port usage* in *IBM Storage Scale: Administration Guide*.

For examples of how to open firewall ports on different operating systems, see *Examples of how to open firewall ports* in *IBM Storage Scale: Administration Guide*.

Repository setup

The installation toolkit contains all necessary code for installation. However, for manual installation or installation with the installation toolkit, there might be base operating system packages required as prerequisites. To satisfy any prerequisites, it is necessary for your nodes to have access to a DVD repository or an external repository accessible by network. Repositories containing IBM Storage Scale dependencies include the following x86_64 example:

```
rhel-x86_64-server-7 Red Hat Enterprise Linux Server
```

For information on setting up repositories on Red Hat Enterprise Linux nodes, see [Configuring Yum and Yum Repositories](https://access.redhat.com/documentation/en-US/Red_Hat_Enterprise_Linux/7/html/System_Administrators_Guide/sec-Configuring_Yum_and_Yum_Repositories.html) (https://access.redhat.com/documentation/en-US/Red_Hat_Enterprise_Linux/7/html/System_Administrators_Guide/sec-Configuring_Yum_and_Yum_Repositories.html).

For information on setting up repositories on SLES nodes, see [Managing Software Repositories and Services](https://documentation.suse.com/sles/15-SP1/html/SLES-all/cha-yast-software.html#sec-yast-software-instsource)(<https://documentation.suse.com/sles/15-SP1/html/SLES-all/cha-yast-software.html#sec-yast-software-instsource>).

For information on setting up repositories on Ubuntu nodes, see [Repositories/Ubuntu](https://help.ubuntu.com/community/Repositories/Ubuntu)(<https://help.ubuntu.com/community/Repositories/Ubuntu>).

Important: You must disable all configured EPEL repositories on all nodes added to the installation toolkit before proceeding with installation, deployment or upgrade.

Network time protocol (NTP) setup

You must configure time synchronization on all nodes in your cluster by using tools such as NTPD and Chrony. IBM Storage Scale necessitates all clocks to be in sync for installation, deployment, upgrade, and cluster operation. For information about using these tools with the operating system on your nodes, see the respective operating system documentation.

Collection of core dump data

For information about changing configuration to enable collection of core dump data from protocol nodes, see *Configuration changes required on protocol nodes to collect core dump data* in *IBM Storage Scale: Problem Determination Guide*.

Preparing the environment on Linux nodes

Before proceeding with installation, prepare your environment by following the suggestions in the following sections.

Add the GPFS bin directory to your shell PATH

Ensure that the PATH environment variable for the root user on each node includes `/usr/lpp/mmfs/bin`. (This is not required for the operation of IBM Storage Scale, but it can simplify administration.)

Other suggestions for cluster administration

GPFS commands operate on all nodes that are required to perform tasks. When you are administering a cluster, it can be useful to have a more general form of running commands on all nodes. A suggested way to do this is to use an OS utility like dsh or pdsh that can run the commands on all nodes in the cluster. For example, you can use dsh to check the kernel version of each node in your cluster:

```
# uname -opr
3.10.0-693.el7.x86_64 x86_64 GNU/Linux
```

When you have dsh set up, you can use it to install IBM Storage Scale on a large cluster. For details about setting up dsh or a similar utility, review the documentation for the utility.

Verify that that prerequisite software is installed

Before installing IBM Storage Scale, it is necessary to verify that you have the correct levels of the prerequisite software that is installed on each node in the cluster. If the correct level of prerequisite software is not installed, see the appropriate installation manual before proceeding with your IBM Storage Scale installation.

Installing IBM Storage Scale without ensuring that the prerequisites listed in [“Hardware requirements”](#) on page 309, [“Software requirements”](#) on page 310, and [“Installation prerequisites”](#) on page 470 are satisfied can lead to undesired results.

For the most up-to-date list of prerequisite software, see the [IBM Storage Scale FAQ](#) in [IBM Documentation](#).

The FAQ contains the latest information about the following:

- Supported Linux distributions and kernel levels

- Recommended or required RPM levels
- Software recommendations
- Configuration information

For more information, see [“GPFS and network communication” on page 17](#).

IBM Storage Scale packaging overview

The IBM Storage Scale self-extracting package consists of several components.

Components	Description
Installation toolkit and license	The components necessary to begin the protocol installation.
GPFS	Core GPFS is included in this self-extracting package. It also includes call home and file audit logging packages.
NFS (Ganesha)	The packages for Ganesha, the open source, user-space implementation of the NFS protocol. They can be installed, enabled, and configured by using the installation toolkit.
SMB (Samba)	The packages for Samba, the open-source implementation of the SMB protocol. They can be installed and enabled using the installation toolkit.
Performance monitoring tool	The packages necessary for performance monitoring by GPFS (including protocols). They can be installed and enabled by using the installation toolkit.

Preparing to install the IBM Storage Scale software on Linux nodes

Follow the steps in this topic in the specified order to prepare to install the IBM Storage Scale software.

About this task

Accepting the electronic license agreement on Linux nodes

The GPFS software license agreement is shipped with the GPFS software and is viewable electronically. When you extract the GPFS software, you are asked whether or not you accept the license.

About this task

The electronic license agreement must be accepted before the software installation can proceed. Read the software agreement carefully before accepting the license. See [“Extracting the IBM Storage Scale software on Linux nodes” on page 473](#).

Extracting the IBM Storage Scale software on Linux nodes

The IBM Storage Scale software is delivered in a self-extracting archive. This self-extracting archive can be downloaded from [Fix Central](#).

About this task

The self-extracting image contains the following:

- The IBM Storage Scale product installation SLES and Red Hat Enterprise Linux RPMs and Ubuntu Linux packages.
- The License Acceptance Process (LAP) tool.

The LAP tool is invoked for acceptance of the IBM Storage Scale license agreements. The license agreements must be accepted to obtain access to the IBM Storage Scale product installation images.

- A version of the Java Runtime Environment (JRE) is necessary to run the LAP tool.
1. Copy the self-extracting product image to a local directory, specifying the correct version of the product for your hardware platform and Linux distribution. For example:

```
cp /media/archive/Storage_Scale_Standard-6.0.1.x-x86_64-Linux-install/  
/tmp/Linux/Storage_Scale_Standard-6.0.1.x-x86_64-Linux-install
```

2. Verify that the self-extracting program has executable permissions, for example:

```
ls -l /tmp/Storage_Scale_Standard-6.0.1.x-x86_64-Linux-install
```

The system displays information similar to the following:

```
-rwxr-xr-x 1 root root 673341648 Nov 10 13:53 /tmp/Storage_Scale_Standard-6.0.1.x-x86_64-  
Linux-install
```

If it is not executable, you can make it executable by issuing the `chmod +x` command.

Note: To get a list of all packages in the self-extracting installation image, issue the following command (specifying the `--manifest` option):

```
./Storage_Scale_Standard-6.0.1.x-x86_64-Linux-install --manifest
```

A sample output is as follows:

```
manifest  
sw, rpm, gpfs.base-6.0.1-x.x86_64.rpm,  
Thu 16 Jun 2020 10:40:41 AM MST, md5sum: 99e940728f73649520533c06ae14b941  
sw, rpm, gpfs.adv-6.0.1-x.x86_64.rpm,  
Thu 16 Jun 2020 10:46:43 AM MST, md5sum: 8c7a635ee2381b075e7960813bef3c47  
sw, rpm, gpfs.gskit-8.0.55.x.x86_64.rpm,  
Thu 16 Jun 2020 10:45:15 AM MST, md5sum: 2e844aa5351dd87c234f5273ef9f3673  
sw, rpm, gpfs.gss.pmsensors-6.0.1-x.el7.x86_64.rpm,  
Wed 15 Jun 2020 09:56:29 AM MST, md5sum: 3f4c084e93f3bd0d7cdf4adb8fe65448  
sw, rpm, gpfs.gss.pmc collector-6.0.1-x.el7.x86_64.rpm,  
Wed 15 Jun 2020 09:56:30 AM MST, md5sum: 5fd3d21c935fea27d3b8c9b97b8c49e2  
sw, rpm, gpfs.gss.pmsensors-6.0.1-x.sles12.x86_64.rpm,  
Wed 15 Jun 2020 09:56:34 AM MST, md5sum: abab4976a83814877a140bab5e96c12f  
sw, rpm, gpfs.gss.pmc collector-6.0.1-x.sles12.x86_64.rpm,  
Wed 15 Jun 2020 09:56:31 AM MST, md5sum: 48ae6aedf41c6681c98bc167dc1babdf  
sw, rpm, gpfs.gpl-6.0.1-x.noarch.rpm,  
Thu 16 Jun 2020 10:52:11 AM MST, md5sum: 5989f71ace5dd4e7ca92f030de451819  
sw, rpm, gpfs.gui-6.0.1-x.noarch.rpm,  
Mon 13 Jun 2020 10:34:37 AM MST, md5sum: e562910893e468cf9790865093623129  
sw, rpm, gpfs.msg.en_US-6.0.1-x.noarch.rpm,  
Thu 16 Jun 2020 10:44:37 AM MST, md5sum: 6fe50b1bf9265d6dab322e0a472266ec
```

(The remaining packages in the installation image are displayed.)

3. Invoke the self-extracting image that you copied to a local directory by using a command such as `./Storage-Scale-Standard-6.0.1.x-x86_64-Linux-install` and accept the license agreement:
 - a. By default, the LAP tool, JRE, and IBM Storage Scale installation images are extracted to the target directory `/usr/lpp/mmfs/6.0.1.x`
 - b. The license agreement files on the media can be viewed in text-only mode or graphics mode:
 - Default mode is text-only. The output explains how to accept the agreement:

```
<...Last few lines of output...>  
Press Enter to continue viewing the license agreement, or  
enter "1" to accept the agreement, "2" to decline it, "3"  
to print it, "4" to read non-IBM terms, or "99" to go back  
to the previous screen.
```

- To view the files in graphics mode, invoke `Storage-Scale-Standard-6.0.1.x-x86_64-Linux-install`.

Using the graphics-mode installation requires a window manager to be configured.

- c. You can use the `--silent` option to accept the license agreement automatically.
- d. Use the `--help` option to obtain usage information from the self-extracting archive.

Upon license agreement acceptance, the IBM Storage Scale product installation images are placed in the extraction target directory (`/usr/lpp/mmfs/6.0.1.x`). This directory contains the IBM Storage Scale packages according to the target Linux distribution that are applicable for the IBM Storage Scale edition that you installed (Standard, Advanced, or Data Management). For more information, see [“Location of extracted packages”](#) on page 480.

Note: For this release, the IBM Global Security Kit (GSKit) version for RPMs and Ubuntu Linux packages must be at least 8.0.55.x or higher.

In this directory, there is a `license` subdirectory that contains license agreements in multiple languages. To view which languages are provided, issue the following command:

```
ls /usr/lpp/mmfs/6.0.1.x/license
```

The system displays information similar to the following:

```
Chinese_TW.txt  French.txt  Japanese.txt  notices.txt  Slovenian.txt
Chinese.txt     German.txt  Korean.txt    Polish.txt   Spanish.txt
Czech.txt      Greek.txt   Lithuanian.txt Portuguese.txt Status.dat
English.txt    Indonesian.txt Italian.txt  non_ibm_license.txt Russian.txt  Turkish.txt
```

The license agreement remains available in the extraction target directory under the `license` subdirectory for future access. The license files are written using operating system-specific code pages. This enables you to view the license in English and in the local language configured on your machine. The other languages are not guaranteed to be viewable.

Verifying signature of IBM Storage Scale packages

All IBM Storage Scale packages for Red Hat Enterprise Linux and SLES operating systems on all supported architectures are signed with a GPG (GNU Privacy Guard) key by IBM. The repository metadata is also signed by IBM. You can verify that an IBM Storage Scale package and repository metadata are signed by IBM as follows.

The public key is located in a file that is called `SpectrumScale_public_key.pgp` and this file is present in the IBM Storage Scale installation images that can be downloaded from IBM Fix Central. For the latest version of the public key, see [IBM Storage Scale FAQ in IBM Documentation](#).

About this task

Important: If you are using the installation toolkit, no additional steps are required. The installation toolkit checks, the signature of each package, and the repository metadata automatically before installation or upgrade.

For manual installation or upgrade, if you do not want to verify that the packages are signed, no additional steps are required. The signed packages functions the same as the unsigned packages. If you want to manually verify that the packages are signed by IBM, perform the following steps:

Procedure

1. Import the public key into the RPM database.

```
rpm --import SpectrumScale_public_key.pgp
```

Note: Some Red Hat packages are required for object installation. These packages are signed by Red Hat with a GPG key. This key is called `RPM-GPG-KEY-redhat-release` and it is located in

the `/usr/lpp/mmfs/<release>/Public_Keys/` directory. If you want to manually install object, along with importing the IBM public key, you can import the Red Hat key as follows:

```
rpm --import /usr/lpp/mmfs/<release>/Public_Keys/RPM-GPG-KEY-redhat-release
```

2. Confirm that the public key is imported into the RPM database.

```
rpm -q gpg-pubkey --qf '%{NAME}-%{VERSION}-%{RELEASE} \ n%{INSTALLTIME:date} \ n%{SUMMARY} \n\n' | grep SpectrumScale
```

3. Check the package's signature.

```
rpm -K PackageName
```

You can check the signature of more than one package by using wildcard characters. For example:

```
rpm -K *.rpm.
```

Extracting IBM Storage Scale patches (update SLES and Red Hat Enterprise Linux RPMs or Ubuntu Linux packages)

Typically, when you install an IBM Storage Scale system there are patches available. It is recommended that you always look for the latest patches when installing or updating an IBM Storage Scale node.

About this task

Note: For more information about restrictions pertaining to Ubuntu Linux, see the [IBM Storage Scale FAQ in IBM Documentation](#).

IBM Storage Scale patches (update SLES and Red Hat Enterprise Linux RPMs and Ubuntu Linux packages) are available from [Fix Central](#).

The updated SLES and Red Hat Enterprise Linux RPMs and Ubuntu Linux packages are distributed as self-extracting base software.

If you are applying a patch during the initial installation of IBM Storage Scale on a node, you only need to build the portability layer once after the base and update SLES or Red Hat Enterprise Linux RPMs or Ubuntu Linux packages are installed.

Related tasks

[“Building the GPFS portability layer on Linux nodes” on page 483](#)

Before starting GPFS, you must build and install the GPFS portability layer.

Installing the IBM Storage Scale man pages on Linux nodes

To use the IBM Storage Scale man pages, the `gpfs.docs` package (rpm or deb package) must be installed.

About this task

When you have installed the `gpfs.docs` package, the IBM Storage Scale man pages are located at `/usr/share/man/`.

You do not need to install the `gpfs.docs` package on all nodes if man pages are not required (for example, if local file system space on the node is minimal).

If you are using the installation toolkit on supported Linux distributions, it automatically installs the associated man pages when it runs **spectrumscale install** to create a new GPFS cluster.

For Linux on Z: Changing the kernel settings

For IBM Storage Scale to run on Linux on Z, the option "vmalloc=4096G" needs to be added to the kernel settings.

About this task

Before starting IBM Storage Scale, perform the following steps on each Linux on Z node.

If you are using Red Hat Enterprise Linux 8, Red Hat Enterprise Linux 9, or Red Hat Enterprise Linux 10:

1. Use the grubby utility to add vmalloc=4096G for all boot entries on your node as shown in the following example:

```
grubby --update-kernel=ALL --args="vmalloc=4096G"
```

2. Run the **zipl** command:

```
zipl -V
```

3. Reboot the node.

If you are using SLES:

1. Edit the /etc/default/grub file and add the following:

```
cat/etc/default/grub
GRUB_CMDLINE_LINUX_DEFAULT=" ... vmalloc=4096G "
```

2. Run the **grub2-mkconfig** command:

```
grub2-mkconfig -o /boot/grub2/grub.cfg
```

3. Reboot the node.

Note: For more information about installation and startup of IBM Storage Scale on IBM Z, see [Getting started with IBM Storage Scale for Linux on IBM Z](#).

Manually installing the IBM Storage Scale software packages on Linux nodes

Use this information to manually install the IBM Storage Scale software on Linux nodes.

Note: You can also use the installation toolkit to install and configure packages on supported Linux distributions. For more information, see [“Overview of the installation toolkit”](#) on page 509.

You can install the IBM Storage Scale packages manually by using tools such as **rpm** or **yum** on Red Hat Enterprise Linux, **rpm** or **zypper** on SLES, and **dpkg** or **apt** on Ubuntu Linux.

Required packages for SLES and Red Hat Enterprise Linux

The following packages are required for IBM Storage Scale Standard Edition or IBM Storage Scale Data Access Edition on SLES and Red Hat Enterprise Linux:

- gpfs.base*.rpm
- gpfs.gpl*.noarch.rpm
- gpfs.msg.en_US*.noarch.rpm
- gpfs.gskit*.rpm
- gpfs.license*.rpm

The gpfs.crypto*.rpm and gpfs.adv*.rpm packages are required for following editions respective to the operating system, in addition to the packages listed for IBM Storage Scale Standard Edition or IBM Storage Scale Data Access Edition:

- IBM Storage Scale Advanced Edition and IBM Storage Scale Data Management Edition on SLES
- IBM Storage Scale Advanced Edition, IBM Storage Scale Data Management Edition, and IBM Storage Scale Developer Edition on Red Hat Enterprise Linux

Required packages for Ubuntu Linux

The following packages are required for IBM Storage Scale Standard Edition or IBM Storage Scale Data Access Edition on Ubuntu Linux:

- gpfs.base*.deb
- gpfs.gpl*all.deb
- gpfs.msg.en-us*all.deb
- gpfs.gskit*.deb
- gpfs.license*.deb

The following packages are required for IBM Storage Scale Advanced Edition and IBM Storage Scale Data Management Edition on Ubuntu Linux, in addition to the packages listed for IBM Storage Scale Standard Edition or IBM Storage Scale Data Access Edition:

- gpfs.crypto*.deb
- gpfs.adv*.deb

Note: For this release, the IBM Global Security Kit (GSKit) version for RPMs and Ubuntu Linux packages must be at least 8.0.55.x or later.

Optional packages for SLES and Red Hat Enterprise Linux

- gpfs.docs*noarch.rpm
- gpfs.java*.rpm
- gpfs.compression*.rpm
- gpfs.gui*.rpm
- gpfs.librdkafka*.rpm
- gpfs.afm.cos*.rpm
- gpfs.scaleapi*.rpms
- gpfs.scalectl*.rpms

For more information about manually installing management GUI packages, see [“Manually installing IBM Storage Scale management GUI” on page 497](#).

The following performance monitoring packages are also optional for SLES and Red Hat Enterprise Linux.

- gpfs.gss.pmcollector*.rpm
- gpfs.gss.pmsensors*.rpm

The following protocols packages are available for Red Hat Enterprise Linux 8.x:

- gpfs.nfs-ganesha*.el8.*.rpm
- gpfs.nfs-ganesha-gpfs*.el8.*.rpm
- gpfs.nfs-ganesha-utils*.el8.*.rpm
- gpfs.smb*gpfs*.el8.*.rpm
- gpfs.mms3*.el8.*.rpm
- noobaa-core*.el8.*.rpm
- gpfs.pm-ganesha*.el8.*.rpm

The following protocols packages are available for Red Hat Enterprise Linux 9.x:

- gpfs.nfs-ganesha*.el9.*.rpm

- gpfs.nfs-ganesha-gpfs*.el9.*.rpm
- gpfs.nfs-ganesha-utils*.el9.*.rpm
- gpfs.smb*_gpfs_*.el9.*.rpm
- gpfs.pm-ganesha*.el9.*.rpm
- gpfs.mms3*.el9.*.rpm
- noobaa-core-*.el9.*.rpm

The following protocols packages are available for SLES 15:

- gpfs.nfs-ganesha*.sles15.x86_64.rpm
- gpfs.nfs-ganesha-gpfs*.sles15.x86_64.rpm
- gpfs.nfs-ganesha-utils*.sles15.x86_64.rpm
- gpfs.smb*_gpfs_*.sles15.x86_64.rpm
- gpfs.pm-ganesha*.sles15.x86_64.rpm

For information on manually installing packages including protocols packages on Red Hat Enterprise Linux, SLES, and Ubuntu nodes, see [“Manually installing IBM Storage Scale and deploying protocols on Linux nodes”](#) on page 486.

For information about manually installing performance monitoring packages, see [“Manually installing the performance monitoring tool”](#) on page 495.

Optional packages for Ubuntu Linux

The following packages are optional for Ubuntu Linux.

- gpfs.docs*all.deb
- gpfs.librdkafka*.deb
- gpfs.afm.cos*.deb
- gpfs.compression*.deb
- gpfs.scaleapi*.deb
- gpfs.scalectl*.deb

The following performance monitoring packages are optional for Ubuntu Linux .

- gpfs.gss.pmcollector*20*.deb
- gpfs.gss.pmsensors*20*.deb

The following protocols packages are available for Ubuntu 22.04:

- gpfs.nfs-ganesha*~jammy_amd64.deb
- gpfs.nfs-ganesha_gpfs*~jammy_amd64.deb
- gpfs.nfs-ganesha-dbgsym*~jammy_amd64.deb
- gpfs.python-nfs-ganesha*~jammy_all.deb
- gpfs.nfs-ganesha-doc*all.~jammy_deb
- gpfs.smb*gpfs*~jammy_amd64.deb
- gpfs.smb-dbg_*gpfs*~jammy_amd64.deb
- gpfs.pm-ganesha*~jammy_amd64.deb

The following protocols packages are available for Ubuntu 24.04:

- gpfs.nfs-ganesha*~noble_amd64.deb
- gpfs.nfs-ganesha_gpfs*~noble_amd64.deb
- gpfs.nfs-ganesha-dbgsym*~noble_amd64.deb
- gpfs.python-nfs-ganesha*~noble_all.deb

- `gpfs.nfs-ganesha-doc*all.~noble_deb`
- `gpfs.smb*gpfs*~noble_amd64.deb`
- `gpfs.smb-dbg_*gpfs*~noble_amd64.deb`
- `gpfs.pm-ganesha*~noble_amd64.deb`

The following packages are required (and provided) only on the Elastic Storage Server (ESS):

- `gpfs.gnr.base*.ppc64.rpm`
- `gpfs.gnr*.ppc64.rpm`
- `gpfs.gss.firmware*.ppc64.rpm`

For more information about Elastic Storage Server (ESS), see [Elastic Storage Server \(ESS\) documentation](#).

Location of extracted packages

The installation images are extracted to following component-specific directories.

GPFS packages

- Red Hat Enterprise Linux: `/usr/lpp/mmfs/6.0.x.x/gpfs_rpms`
 - File audit logging and watch folder packages: `/usr/lpp/mmfs/6.0.x.x/gpfs_rpms/rhel`
- Ubuntu: `/usr/lpp/mmfs/6.0.x.x/gpfs_debs`
 - File audit logging and watch folder packages: `/usr/lpp/mmfs/6.0.x.x/gpfs_debs/ubuntu`
- SLES: `/usr/lpp/mmfs/6.0.x.x/gpfs_rpms`

NFS packages

- Red Hat Enterprise Linux 8.x: `/usr/lpp/mmfs/6.0.x.x/ganesha_rpms/rhel8`
- Red Hat Enterprise Linux 9.x: `/usr/lpp/mmfs/6.0.x.x/ganesha_rpms/rhel9`
- Ubuntu 22.04: `/usr/lpp/mmfs/6.0.x.x/ganesha_debs/ubuntu/ubuntu22`
- Ubuntu 24.04: `/usr/lpp/mmfs/6.0.x.x/ganesha_debs/ubuntu/ubuntu24`
- SLES 15: `/usr/lpp/mmfs/6.0.x.x/ganesha_rpms/sles15`

SMB packages

- Red Hat Enterprise Linux 8.x: `/usr/lpp/mmfs/6.0.x.x/smb_rpms/rhel8`
- Red Hat Enterprise Linux 9.x: `/usr/lpp/mmfs/6.0.x.x/smb_rpms/rhel9`
- Ubuntu 22.04: `/usr/lpp/mmfs/6.0.x.x/smb_debs/ubuntu/ubuntu22`
- Ubuntu 24.04: `/usr/lpp/mmfs/6.0.x.x/smb_debs/ubuntu/ubuntu24`
- SLES 15: `/usr/lpp/mmfs/6.0.x.x/smb_debs/sles15`

S3 packages

- Red Hat Enterprise Linux 8.x: `/usr/lpp/mmfs/6.0.x.x/s3_rpms/rhel8`
- Red Hat Enterprise Linux 9.x: `/usr/lpp/mmfs/6.0.x.x/s3_rpms/rhel9`
- Red Hat Enterprise Linux 10.x: `/usr/lpp/mmfs/6.0.x.x/s3_rpms/rhel10`

Use RHEL 9 NooBaa RPMs on RHEL 10 clusters.

Performance monitoring packages

- Red Hat Enterprise Linux 8.x: `/usr/lpp/mmfs/6.0.x.x/zimon_rpms/rhel8`
- Red Hat Enterprise Linux 9.x: `/usr/lpp/mmfs/6.0.x.x/zimon_rpms/rhel9`
- Ubuntu 22.04: `/usr/lpp/mmfs/6.0.x.x/zimon_debs/zimon/ubuntu22`
- Ubuntu 24.04: `/usr/lpp/mmfs/6.0.x.x/zimon_debs/zimon/ubuntu24`
- SLES 15: `/usr/lpp/mmfs/6.0.x.x/zimon_rpms/sles15`

Note: The `pm-ganesha` package is extracted in the following directories:

- Red Hat Enterprise Linux 8.x: /usr/lpp/mmfs/6.0.x.x/zimon_rpms/rhel8
- Red Hat Enterprise Linux 9.x: /usr/lpp/mmfs/6.0.x.x/zimon_rpms/rhel9
- SLES: /usr/lpp/mmfs/6.0.x.x/zimon_rpms/sles15
- Ubuntu: /usr/lpp/mmfs/6.0.x.x/zimon_debs/ubuntu

IBM Storage Scale native REST API packages

- Red Hat Enterprise Linux: /usr/lpp/mmfs/6.0.x.x/scaleapi*_rpms
- Ubuntu: /usr/lpp/mmfs/6.0.x.x/scaleapi*_debs

Note: Starting with the 6.0.1.0 version, the `scalectl` binary is also provided as a separate package, `gpfs.scalectl`. This change enables installation of `scalectl` on remote systems outside the cluster without requiring the full set of dependencies associated with `gpfs.scaleapi`. The `gpfs.scaleapi` package continues to provide the core REST API components and depends on `gpfs.scalectl`, so `scalectl` remains available on IBM Storage Scale nodes.

Installation of packages for SLES

To install all required GPFS SLES RPMs for the IBM Storage Scale Standard Edition or IBM Storage Scale Data Access Edition, change the directory to where the installation image is extracted. For example, issue this command:

```
cd /usr/lpp/mmfs/6.0.x.x/gpfs_rpms
```

Then, issue this command for Standard Edition:

```
zypper install -y gpfs.base*.rpm gpfs.gpl*.rpm gpfs.license.std*.rpm gpfs.gskit*.rpm
gpfs.msg*.rpm
gpfs.docs*.rpm
```

Or issue the following command for Data Access Edition:

```
zypper install -y gpfs.base*.rpm gpfs.gpl*.rpm gpfs.license.da*.rpm gpfs.gskit*.rpm
gpfs.msg*.rpm
gpfs.docs*.rpm
```

To install all required GPFS SLES RPMs for the IBM Storage Scale Advanced Edition, IBM Storage Scale Data Management Edition, or IBM Storage Scale Developer Edition, change to the directory where the installation image is extracted. For example, enter the following command:

```
cd /usr/lpp/mmfs/6.0.x.x/gpfs_rpms
```

Then, issue the following command for Advanced Edition:

```
zypper install -y gpfs.base*.rpm gpfs.gpl*.rpm gpfs.license.adv*.rpm gpfs.gskit*.rpm
gpfs.docs*.rpm gpfs.msg*.rpm gpfs.adv*.rpm gpfs.crypto*.rpm
```

Or issue the following command for Data Management Edition:

```
zypper install -y gpfs.base*.rpm gpfs.gpl*.rpm gpfs.license.dm*.rpm gpfs.gskit*.rpm
gpfs.docs*.rpm gpfs.msg*.rpm gpfs.adv*.rpm gpfs.crypto*.rpm
```

Or issue the following command for Developer Edition:

```
zypper install -y gpfs.base*.rpm gpfs.gpl*.rpm gpfs.license.dev*.rpm gpfs.gskit*.rpm
gpfs.docs*.rpm gpfs.msg*.rpm gpfs.adv*.rpm gpfs.crypto*.rpm
```

To install IBM Storage Scale native REST API packages, change the directory to where the installation package is extracted. For example, issue this command:

```
cd /usr/lpp/mmfs/<x.x.x.x>/scaleapi_rpms
```

Then, issue the following command:

```
rpm -Fvh ./gpfs.scale*.rpm
```

Installation of packages for Red Hat Enterprise Linux

To install all required GPFS Red Hat Enterprise Linux RPMs for the IBM Storage Scale Standard Edition or IBM Storage Scale Data Access Edition, change the directory to where the installation image is extracted. For example, issue this command:

```
cd /usr/lpp/mmfs/6.0.x.x/gpfs_rpms
```

Then, issue this command for Standard Edition:

```
yum -y install gpfs.base*.rpm gpfs.gpl*.rpm gpfs.license.std*.rpm gpfs.gskit*.rpm gpfs.msg*.rpm gpfs.docs*.rpm
```

Or issue the following command for Data Access Edition:

```
yum -y install gpfs.base*.rpm gpfs.gpl*.rpm gpfs.license.da*.rpm gpfs.gskit*.rpm gpfs.msg*.rpm gpfs.docs*.rpm
```

To install all required GPFS Red Hat Enterprise Linux RPMs for the IBM Storage Scale Advanced Edition, IBM Storage Scale Data Management Edition, or IBM Storage Scale Developer Edition, change to the directory where the installation image is extracted. For example, enter the following command:

```
cd /usr/lpp/mmfs/6.0.x.x/gpfs_rpms
```

Then, issue the following command for Advanced Edition:

```
yum -y install gpfs.base*.rpm gpfs.gpl*.rpm gpfs.license.adv*.rpm gpfs.gskit*.rpm gpfs.docs*.rpm gpfs.msg*.rpm gpfs.adv*.rpm gpfs.crypto*.rpm
```

Or issue the following command for Data Management Edition:

```
yum -y install gpfs.base*.rpm gpfs.gpl*.rpm gpfs.license.dm*.rpm gpfs.gskit*.rpm gpfs.docs*.rpm gpfs.msg*.rpm gpfs.adv*.rpm gpfs.crypto*.rpm
```

Or issue the following command for Developer Edition:

```
yum -y install gpfs.base*.rpm gpfs.gpl*.rpm gpfs.license.dev*.rpm gpfs.gskit*.rpm gpfs.docs*.rpm gpfs.msg*.rpm gpfs.adv*.rpm gpfs.crypto*.rpm
```

Note: IBM Storage Scale Developer Edition is available only on Red Hat Enterprise Linux on x86_64.

To install IBM Storage Scale native REST API packages, change the directory to where the installation package is extracted. For example, issue this command:

```
cd /usr/lpp/mmfs/<x.x.x.x>/scaleapi_rpms
```

Then, issue the following command:

```
rpm -Fvh ./gpfs.scale*.rpm
```

Installation of packages for Ubuntu Linux

To install all GPFS Ubuntu Linux packages for the IBM Storage Scale Standard edition, change the directory to where the installation image is extracted. For example, issue this command:

```
cd /usr/lpp/mmfs/6.0.x.x/gpfs_debs
```

Then, issue this command for Standard Edition:

```
apt install ./gpfs.base*.deb ./gpfs.gpl*.deb ./gpfs.license.std*.deb ./gpfs.gskit*.deb ./gpfs.msg*.deb ./gpfs.docs*.deb
```

Or issue the following command for Data Access Edition:

```
apt install ./gpfs.base*.deb ./gpfs.gpl*.deb ./gpfs.license.da*.deb ./gpfs.gskit*.deb ./gpfs.msg*.deb ./gpfs.docs*.deb
```

To install all GPFS Ubuntu Linux packages for IBM Storage Scale Advanced Edition and IBM Storage Scale Data Management Edition, change the directory to where the installation image is extracted. For example, issue this command:

```
cd /usr/lpp/mmfs/6.0.1.x/gpfs_debs
```

Then, issue this command for Advanced Edition:

```
apt install ./gpfs.base*.deb ./gpfs.gpl*.deb ./gpfs.license.adv*.deb ./gpfs.gskit*.deb ./gpfs.msg*.deb ./gpfs.docs*.deb ./gpfs.adv*.deb ./gpfs.crypto*.deb
```

Or issue this command for Data Management Edition:

```
apt install ./gpfs.base*.deb ./gpfs.gpl*.deb ./gpfs.license.dm*.deb ./gpfs.gskit*.deb ./gpfs.msg*.deb ./gpfs.docs*.deb ./gpfs.adv*.deb ./gpfs.crypto*.deb
```

To install IBM Storage Scale native REST API packages, change the directory to where the installation package is extracted. For example, issue this command:

```
cd /usr/lpp/mmfs/<x.x.x.x>/scaleapi_debs
```

Then, issue the following command:

```
apt-get install --only-upgrade ./gpfs.scale*.deb
```

Related concepts

[“Manually installing IBM Storage Scale management GUI” on page 497](#)

The management GUI provides an easy way for the users to configure, manage, and monitor the IBM Storage Scale system.

[“Installing call home” on page 645](#)

The call home feature collects files, logs, traces, and details of certain system health events from different nodes and services. These details are shared with the IBM support center for monitoring and problem determination.

[“Manually installing the performance monitoring tool” on page 495](#)

The performance monitoring tool is automatically installed by the installation toolkit. You can also install the performance monitoring tool manually.

Building the GPFS portability layer on Linux nodes

Before starting GPFS, you must build and install the GPFS portability layer.

About this task

The GPFS portability layer is a loadable kernel module that allows the GPFS daemon to interact with the operating system.

Note: The GPFS kernel module must be updated every time the Linux kernel is updated. Updating the GPFS kernel module after a Linux kernel update requires rebuilding and installing a new version of the module.

Tip: You can configure a cluster to rebuild the GPL automatically whenever a new level of the Linux kernel is installed or whenever a new level of IBM Storage Scale is installed. This feature is available only on the Linux operating system. For more information, see the description of the **autoBuildGPL** attribute in the topic **mmchconfig** command in the *IBM Storage Scale: Command and Programming Reference Guide*.

To build the GPFS portability layer on Linux nodes, perform the following steps:

1. Check the following before building the portability layer:

- Updates to the portability layer at the [IBM Support Portal: Downloads for General Parallel File System](http://www.ibm.com/support/entry/portal/Downloads/Software/Cluster_software/General_Parallel_File_System) (www.ibm.com/support/entry/portal/Downloads/Software/Cluster_software/General_Parallel_File_System).
- The latest kernel levels supported in the [IBM Storage Scale FAQ](#) in IBM Documentation.
- The kernel development files and compiler utilities that are required to build the portability layer in “Software requirements” on page 310.

2. Build your GPFS portability layer in one of the following ways:

- Using the **mmbuildgpl** command (recommended) for nodes running GPFS 4.1.0.4 or later. For more information, see **mmbuildgpl** command in *IBM Storage Scale: Command and Programming Reference Guide*.
- Using the Autoconfig tool.
- Using the directions in **/usr/lpp/mmfs/src/README**.

Using the mmbuildgpl command to build the GPFS portability layer on Linux nodes

Starting with GPFS 4.1.0.4, you can use the **mmbuildgpl** command to simplify the build process.

About this task

To build the GPFS portability layer by using **mmbuildgpl**, enter the following command:

```
/usr/lpp/mmfs/bin/mmbuildgpl
```

Each kernel module is specific to a Linux version and platform. If you have multiple nodes running exactly the same operating system level on the same platform, and only some of these nodes have a compiler available, you can build the kernel module on one node, then create an installable package that contains the binary module for ease of distribution.

If you choose to generate an installable package for portability layer binaries, perform the following additional step:

```
/usr/lpp/mmfs/bin/mmbuildgpl --build-package
```

When the command finishes, it displays the location of the generated package as in the following examples:

```
Wrote: /root/rpmbuild/RPMS/x86_64/gpfs.gplbin-4.18.0-305.12.1.el8_4.x86_64-6.0.1-x.x86_64.rpm
```

or

```
Wrote: /tmp/deb/gpfs.gplbin-5.4.0-81-generic_6.0.1-x_amd64.deb
```

You can then copy the generated package to other machines for deployment. By default, the generated package can be deployed only to machines whose architecture, distribution level, Linux kernel, and IBM Storage Scale maintenance level are identical with those of the machine on which the **gpfs.gplbin** package was built. However, you can install the generated package on a machine with a different Linux kernel by setting the **MM_INSTALL_ONLY** environment variable before you install the generated package. If you install the **gpfs.gplbin** package, you do not need to install the **gpfs.gpl** package.

Note: During the package generation, temporary files are written to the **/tmp/rpm** or **/tmp/deb** directory, so be sure there is sufficient space available. By default, the generated package goes to **/usr/src/packages/RPMS/<arch>** for SUSE Linux Enterprise Server, **/usr/src/redhat/RPMS/<arch>** for Red Hat Enterprise Linux, and **/tmp/deb** for Ubuntu Linux.

Important:

- The GPFS portability layer is specific to both the current kernel and the GPFS version. If either the kernel or the GPFS version changes, a new GPFS portability layer needs to be built.
- Although operating system kernels might upgrade to a new version, they are not active until after a reboot. Thus, a GPFS portability layer for this new kernel must be built after a reboot of the operating system.
- Before you install a new GPFS portability layer, make sure to uninstall the prior version of the GPFS portability layer first.

Using the Autoconfig tool to build the GPFS portability layer on Linux nodes

To simplify the build process, GPFS provides an automatic configuration tool called Autoconfig.

About this task

The following example shows the commands required to build the GPFS portability layer by using the Autoconfig tool:

```
cd /usr/lpp/mmfs/src
make Autoconfig
make World
make InstallImages
```

Each kernel module is specific to a Linux version and platform. If you have multiple nodes running exactly the same operating system level on the same platform, or if you have a compiler available on only one node, you can build the kernel module on one node, then create an installable package that contains the binary module for ease of distribution.

If you choose to generate an installable package for portability layer binaries, perform the following additional step:

make rpm for SLES and RHEL Linux

make deb for Ubuntu Linux

When the command finishes, it displays the location of the generated package as in the following examples:

```
Wrote: /root/rpmbuild/RPMS/x86_64/gpfs.gplbin-4.18.0-305.12.1.el8_4.x86_64-6.0.1-x.x86_64.rpm
```

or

```
Wrote: /tmp/deb/gpfs.gplbin-5.4.0-81-generic_6.0.1-x_amd64.deb
```

You can then copy the generated package to other machines for deployment. By default, the generated package can be deployed only to machines whose architecture, distribution level, Linux kernel, and IBM Storage Scale maintenance level are identical with those of the machine on which the `gpfs.gplbin` package was built. However, you can install the generated package on a machine with a different Linux kernel by setting the `MM_INSTALL_ONLY` environment variable before you install the generated package. If you install the `gpfs.gplbin` package, you do not need to install the `gpfs.gpl` package.

Note: During the package generation, temporary files are written to the `/tmp/rpm` or `/tmp/deb` directory, so be sure there is sufficient space available. By default, the generated package goes to `/usr/src/packages/RPMS/<arch>` for SUSE Linux Enterprise Server, `/usr/src/redhat/RPMS/<arch>` for Red Hat Enterprise Linux, and `/tmp/deb` for Ubuntu Linux.

Important:

- The GPFS portability layer is specific to both the current kernel and the GPFS version. If either the kernel or the GPFS version changes, a new GPFS portability layer needs to be built.
- Although operating system kernels might upgrade to a new version, they are not active until after a reboot. Thus, a GPFS portability layer for this new kernel must be built after a reboot of the operating system.

- Before you install a new GPFS portability layer, make sure to uninstall the prior version of the GPFS portability layer first.

Manually installing IBM Storage Scale and deploying protocols on Linux nodes

The following tasks describe how to manually install IBM Storage Scale and deploy protocols on the systems running on supported Linux distributions.

Before you begin

Note: CES supports HDFS protocols. CES HDFS follows the same installation methods and prerequisites as the other protocols as documented in this section. For specific information about installing CES HDFS, see *Installation* under [CES HDFS](#) in *Big data and analytics support documentation*.

These prerequisites must be met before installing IBM Storage Scale on Linux systems.

- Repository must be set up:
 - For information on setting up repository on Red Hat Enterprise Linux, see *Red Hat Enterprise Linux System Administrator's Guide*.
 - For information on setting up repository on SLES, see *SLES Deployment Guide*.
 - For information on setting up repository on Ubuntu, see *Ubuntu Server Guide*.
- All prerequisite packages must be installed. For a list of prerequisite packages, see [“Software requirements”](#) on page 310 and [“Installation prerequisites”](#) on page 470.
- Passwordless SSH must be set up between all nodes in the cluster. For more information, see *Problems due to missing prerequisites* in *IBM Storage Scale: Problem Determination Guide*.
- Firewall configuration must be according to your requirements. It is recommended that firewalls are in place to secure all nodes. For more information, see *Securing the IBM Storage Scale system using firewall* in *IBM Storage Scale: Administration Guide*.

To check the status of the firewall on Red Hat Enterprise Linux, issue the following command:

```
systemctl status firewalld
```

To check the status of the firewall on SLES, issue the following command:

```
sudo /sbin/rcSuSEfirewall2 status
```

To check the status of the firewall on Ubuntu, issue the following command:

```
sudo ufw status
```

For examples of how to open firewall ports on different operating systems, see *Examples of how to open firewall ports* in *IBM Storage Scale: Administration Guide*.

- Every node must have a non-loopback IP address assigned. In some scenarios, a freshly installed node might have its host name pointing to 127.0.0.1 in `/etc/hosts`. 127.0.0.1 is a loopback IP address and it is not sufficient for multi-node IBM Storage Scale cluster creation. In these cases, each node needs a static IP with connectivity to the other nodes.
- Optionally, the bash profile can be updated to allow easier access to IBM Storage Scale commands.
 - Verify that the PATH environment variable for the root user on each node includes `/usr/lpp/mmfs/bin` and the required tool directories. This allows a user to execute IBM Storage Scale commands without having to first change directory to `/usr/lpp/mmfs/bin`.
 - Export the `WCOLL` variable used by **mmdsh**. In the following example, `/nodes` is a manually created file, listing line by line, each node within the cluster using the nodes' FQDN. Once IBM Storage Scale is installed, this configuration allows the **mmdsh** command to execute commands on multiple nodes simultaneously.

Example:

```
# cat ~/.bash_profile
# .bash_profile
# Get the aliases and functions
if [ -f ~/.bashrc ]; then
. ~/.bashrc
fi

#####
# User specific environment and startup programs
#####
#####
# Specifics for GPFS testing
#####
export PATH=$PATH:$HOME/bin:/usr/lpp/mmfs/bin
export WCOLL=/nodes
```

Log out and then log in again for the changes in the bash profile to take effect.

Installing IBM Storage Scale packages on Linux systems

Manually install IBM Storage Scale packages on systems that are running on supported Linux distributions as follows:

About this task

Do these steps on all nodes on which you want to install IBM Storage Scale, one by one unless noted otherwise.

Procedure

1. Download and extract IBM Storage Scale packages, and then accept the licensing agreement. For more information, see [“Extracting the IBM Storage Scale software on Linux nodes”](#) on page 473 and [“Accepting the electronic license agreement on Linux nodes”](#) on page 473.

For information on the location of extracted packages, see [“Location of extracted packages”](#) on page 480.

2. Install the IBM Storage Scale packages by issuing one of the following commands depending on the operating system:

Note: The following command example shows the installation of IBM Storage Scale Advanced Edition packages.

- On Red Hat Enterprise Linux, and SLES, issue the following command:

```
# cd /usr/lpp/mmfs/x.x.x.x/gpfs_rpms
# rpm -ivh gpfs.base*.rpm gpfs.gpl*.rpm gpfs.license*.rpm gpfs.gskit*.rpm
gpfs.msg*.rpm gpfs.compression*.rpm gpfs.adv*.rpm gpfs.crypto*.rpm
```

- On Ubuntu, issue the following command:

```
# cd /usr/lpp/mmfs/x.x.x.x/gpfs_debs
# dpkg -i gpfs.base*.deb gpfs.gpl*.deb gpfs.license.ad*.deb gpfs.gskit*.deb
gpfs.msg*.deb gpfs.compression*.deb gpfs.adv*.deb gpfs.crypto*.deb
```

A sample output is as follows:

```
Preparing... ##### [100%]
1:gpfs.base ##### [ 11%]
2:gpfs.compression ##### [ 27%]
3:gpfs.adv ##### [ 33%]
4:gpfs.crypto ##### [ 44%]
5:gpfs.gpl ##### [ 56%]
6:gpfs.license.adv ##### [ 67%]
7:gpfs.msg.en_US ##### [ 83%]
8:gpfs.gskit ##### [ 100%]
```

3. Build the portability layer by using the following command:

```
# /usr/lpp/mmfs/bin/mmbuildgpl
```

A sample output is as follows. The output varies depending on the operating system:

```
-----
mmbuildgpl: Building GPL module begins at Thu Mar 24 15:18:23 CST 2016.
-----
Verifying Kernel Header...
  kernel version = 31228004 (3.12.28-4-default, 3.12.28-4)
  module include dir = /lib/modules/3.12.28-4-default/build/include
  module build dir   = /lib/modules/3.12.28-4-default/build
  kernel source dir  = /usr/src/linux-3.12.28-4/include
  Found valid kernel header file under /lib/modules/3.12.28-4-default/build/include
Verifying Compiler...
  make is present at /usr/bin/make
  cpp is present at /usr/bin/cpp
  gcc is present at /usr/bin/gcc
  g++ is present at /usr/bin/g++
  ld is present at /usr/bin/ld
Verifying Additional System Headers...
  Verifying linux-glibc-devel is installed ...
    Command: /bin/rpm -q linux-glibc-devel
    The required package linux-glibc-devel is installed
make World ...
make InstallImages ...
-----
mmbuildgpl: Building GPL module completed successfully at Thu Mar 24 15:18:31 CST 2016.
-----
```

If the portability layer cannot be built due to missing prerequisite packages, you might need to install prerequisite packages. For a list of prerequisite packages, see [“Software requirements” on page 310](#).

If the repository setup and the kernel configuration are correct, you can install the prerequisite packages by using one of the following commands, depending on the operating system:

Red Hat Enterprise Linux

```
# /usr/bin/yum install -y package_name1 package_name2 ... package_nameN
```

SLES

```
# /usr/bin/zypper install -y package_name1 package_name2 ... package_nameN
```

Ubuntu

```
# /usr/bin/apt-get install -y package_name1 package_name2 ... package_nameN
```

After IBM Storage Scale is built on all nodes, you can create the cluster. It is recommended to have an odd number of quorum nodes and that your NSD nodes be designated as quorum nodes.

4. Create the cluster by using the following command from one of the nodes:

```
# /usr/lpp/mmfs/bin/mmcrcluster -N NodesList --ccr-enable -r /usr/bin/ssh -R /usr/bin/scp -C cluster1.spectrum
```

In this command example, *NodesList* is a file that contains a list of nodes and node designations to be added to the cluster and its contents are as follows:

```
node1:quorum
node2
node3
node4:quorum-manager
node5:quorum-manager
```

Running the **mmcrcluster** command generates output similar to the following snippet:

```
mmcrcluster: Performing preliminary node verification ...
mmcrcluster: Processing quorum and other critical nodes ...
mmcrcluster: Processing the rest of the nodes ...
mmcrcluster: Finalizing the cluster data structures ...
mmcrcluster: Command successfully completed
mmcrcluster: Propagating the cluster configuration data to all
```

```

affected nodes. This is an asynchronous process.
# Thu Mar 24 15:33:06 CST 2016: mmcommon pushSdr_async: mmsdrfs propagation started
Thu Mar 24 15:33:10 CST 2016: mmcommon pushSdr_async: mmsdrfs propagation completed; mmdsh
rc=0

```

5. Accept proper role licenses for the nodes by using the following commands from one of the nodes:

a) Accept the server licenses for the applicable nodes.

```
# /usr/lpp/mmfs/bin/mmchlicense server --accept -N node1,node4,node5
```

A sample output is as follows:

```

The following nodes will be designated as possessing server licenses:
  node1
  node4
  node5
mmchlicense: Command successfully completed
mmchlicense: Warning: Not all nodes have proper GPFS license designations.
  Use the mmchlicense command to designate licenses as needed.
mmchlicense: Propagating the cluster configuration data to all
affected nodes. This is an asynchronous process.
# Thu Mar 24 15:37:59 CST 2016: mmcommon pushSdr_async: mmsdrfs propagation started
Thu Mar 24 15:38:01 CST 2016: mmcommon pushSdr_async: mmsdrfs propagation completed;
mmdsh rc=0

```

b) Accept the client licenses for the applicable nodes.

```
# /usr/lpp/mmfs/bin/mmchlicense client --accept -N node2,node3
```

A sample output is as follows:

```

The following nodes will be designated as possessing client licenses:
  node2
  node3
mmchlicense: Command successfully completed
mmchlicense: Propagating the cluster configuration data to all
affected nodes. This is an asynchronous process.
# Thu Mar 24 15:38:26 CST 2016: mmcommon pushSdr_async: mmsdrfs propagation started
Thu Mar 24 15:38:28 CST 2016: mmcommon pushSdr_async: mmsdrfs propagation completed;
mmdsh rc=0

```

IBM Storage Scale cluster is now created. You can view the configuration information of the cluster by using the following command:

```
# /usr/lpp/mmfs/bin/mmlscluster
```

The system displays output similar to the following snippet:

```

GPFS cluster information
=====
GPFS cluster name:      cluster1.spectrum
GPFS cluster id:       993377111835434248
GPFS UID domain:       cluster1.spectrum
Remote shell command:  /usr/bin/ssh
Remote file copy command: /usr/bin/scp
Repository type:       CCR
GPFS cluster configuration servers:
-----
Primary server:        node1 (not in use)
Secondary server:      (none)
Node Daemon node name IP address   Admin node name Designation
-----
1  node1              203.0.113.11  node1          quorum
2  node4              203.0.113.14  node4          quorum-manager
3  node5              203.0.113.15  node5          quorum-manager
4  node2              203.0.113.12  node2
5  node3              203.0.113.13  node3

```

6. Start the GPFS daemons and the cluster by using the following command from one of the nodes:

```
# /usr/lpp/mmfs/bin/mmstartup -N NodesList
```

Where *NodesList* is the list of nodes on which the daemons and the cluster are to be started.

You can use the **mmgetstate -N *NodesList*** command to verify that the GPFS software is running on these nodes.

Creating NSDs and file systems as part of installing IBM Storage Scale on Linux systems

Create NSDs and file systems for installing IBM Storage Scale on supported Linux distributions as follows:

Before you begin

For more information about NSD creation considerations, see [“Network Shared Disk \(NSD\) creation considerations” on page 329](#).

Procedure

1. Create NSDs as follows:

- a) Identify available physical disks in the `/dev` directory.
- b) Create the NSD stanza file to be used with the **mmcrnsd** command.

For example, consider you have 2 disks, `/dev/sdc` and `/dev/sdd`, and your node name is `exnode1` and it is running Linux. In this case, your NSD stanza file contents are similar to the following:

```
%nsd:
    device=/dev/sdc
    nsd=nsd1
    servers=exnode1
    usage=dataAndMetadata
    failureGroup=-1
    pool=system

%nsd:
    device=/dev/sdd
    nsd=nsd2
    servers=exnode1
    usage=dataAndMetadata
    failureGroup=-1
```

Note: The server name used in the NSD stanza file must be resolvable by the system.

- c) Create the NSDs by using the following command:

```
# mmcrnsd -F NSD_Stanza_Filename
```

2. Create a GPFS file system using the following command:

```
# mmcrfs fs1 -F NSD_Stanza_Filename -k nfs4
```



Attention: When creating NSD stanza files ensure that you plan out the required file system layout. This includes number of file systems, pools in each file system, failure group layout, metadata versus data usage types, as well as other **mmcrnsd** and **mmcrnsd** stanzafile parameters. Additionally, a stanzafile created for NSD creation by using the **mmcrnsd** command, can also be used for file system creation through the **mmcrfs** command. If you are using an NSD stanzafile for file system creation then you must be aware that a single file system is created. If you require multiple file systems, then you must create a separate stanzafile for each file system.

Installing Cluster Export Services as part of installing IBM Storage Scale on Linux systems

Cluster Export Services (CES) provides highly available file and object services to a GPFS cluster by using Network File System (NFS), Hadoop Distributed File System (HDFS), Object, or Server Message Block (SMB) protocols.

For more information, see *Cluster Export Services overview* in *IBM Storage Scale: Concepts, Planning, and Installation Guide*.

CES requires a shared root directory (**cesSharedRoot**) to install and configure CES protocol services on each node. This directory is used by protocols such as NFS and Object to store CES shared configuration data and support protocol recovery and other protocol-specific operations. This directory tree can reside anywhere on a IBM Storage Scale file system. However, it is recommended to use a dedicated file system to clearly separate protocol data from user data. For more information, see *Creation of a file system or fileset or path for a CES shared root, and creation of an object fileset* in *IBM Storage Scale: Concepts, Planning, and Installation Guide* and *Setting up Cluster Export Services shared root file system* in *IBM Storage Scale: Administration Guide*.

Important:

- If you plan to install and configure CES, all protocols available on your platform must be installed. Installing a subset of available protocols is not supported.

Procedure

1. Unmount GPFS file systems and stop GPFS on all nodes by using the following command:

```
# mmshutdown -a
```

2. Configure the CES shared root file system by using the following command:

```
# mmchconfig cesSharedRoot=/gpfs/fs0
```

Note: It is recommended to configure an independent file system as the CES shared root file system. For more information, see *Creation of a file system or fileset or path for a CES shared root, and creation of an object fileset* in *IBM Storage Scale: Concepts, Planning, and Installation Guide* and *Setting up Cluster Export Services shared root file system* in *IBM Storage Scale: Administration Guide*.

3. Start GPFS on all nodes in the cluster by using the following command:

```
# mmstartup -a
```

4. Enable CES on the required nodes by using the following command:

```
# mmchnode --ces-enable -N prnode1,prnode2,prnode3
```

5. Add to the protocol nodes that the IP addresses designated to be used for CES connectivity by using the following command:

```
# mmces address add --ces-node prnode1 --ces-ip CESIPAddress
```

6. Verify the CES configuration by using the following commands:

```
# mmlscluster --ces  
# mmces address list
```

Currently, there are no services that are enabled. You can verify by using the **mmces service list -a** command. A sample output is as follows:

```
No CES services are enabled.
```

7. Download and extract IBM Storage Scale protocol packages, and then accept the licensing agreement. For more information, see [“Extracting the IBM Storage Scale software on Linux nodes” on page 473](#) and [“Accepting the electronic license agreement on Linux nodes” on page 473](#).

For more information on the location of extracted packages, see [“Location of extracted packages”](#) on page 480.

8. Remove conflicting Samba packages that might have been installed on each CES node by issuing one of the following commands, depending on the operating system.

- On Red Hat Enterprise Linux and SLES, issue the following command:

```
# mmdsh -N all rpm -e samba-common --nodeps
# mmdsh -N all rpm -e samba-client
```

- On Ubuntu, issue the following command:

```
# mmdsh -N all dpkg -P samba-common
# mmdsh -N all dpkg -P samba-client
```

9. Install IBM Storage Scale SMB package by issuing one of the following commands, depending on the operating system.

- On Red Hat Enterprise Linux and SLES, issue the following command:

```
# rpm -ivh gpfs.smb*.rpm
```

- On Ubuntu, issue the following command:

```
# dpkg -i gpfs.smb*.deb
```

For a list of packages for the current IBM Storage Scale release, see [“Manually installing the IBM Storage Scale software packages on Linux nodes”](#) on page 477.

10. Install IBM Storage Scale NFS packages by issuing one of the following commands, depending on the operating system.

- On Red Hat Enterprise Linux and SLES, issue the following command:

```
# rpm -ivh NFS_Package_Name1 NFS_Package_Name2 ... NFS_Package_NameN
```

- On Ubuntu, issue the following command:

```
# dpkg -i NFS_Package_Name1 NFS_Package_Name2 ... NFS_Package_NameN
```

If you cannot install NFS package because of any error by using the **dpkg -i** command, issue the following command:

```
# apt install ./NFS_Package_Name1 ./NFS_Package_Name2 ... ./NFS_Package_NameN
```

This command installs the specified local package files.

Note: Before installing IBM Storage Scale 5.1.x on Ubuntu operating system, make sure that upstream NFS Ganesha package is not installed or configured in the protocol nodes. If the NFS Ganesha package is installed or configured, then uninstall the package by using the **apt purge nfs-ganesha** command.

For a list of packages for the current IBM Storage Scale release, see [“Manually installing the IBM Storage Scale software packages on Linux nodes”](#) on page 477.

11. Install IBM Storage Scale S3 packages by issuing the following commands:

Install the pre-requisite on all CES nodes.

```
rpm -ivh bzip2
```

Install the respective S3 related RPMs.

```
rpm -ivh noobaa-core-* gpfs.mms3*.rpm
```


What to do next

If you plan to use IPv6 addresses, you must enable the CES interface mode. For more information, see *Configuring CES protocol service IP addresses* in *IBM Storage Scale: Administration Guide*.

Enabling NFS, SMB, HDFS, and S3 on Linux systems

Use this information to enable NFS, SMB, HDFS, and S3 on Linux systems after installing the packages.

Before you begin

Before you begin enabling NFS, SMB, HDFS, and S3, you must have installed the corresponding IBM Storage Scale packages.

For HDFS, follow the instructions in [Enable and Configure CES HDFS](#) under *IBM Storage Scale support for Hadoop in Big data and analytics support*.

Procedure

- Enable NFS as follows:
 - a) If the kernel NFS service is running, mask and stop the kernel NFS service on all nodes where CES NFS needs to be installed by using the following commands:

```
# systemctl mask nfs-server.service
# systemctl daemon-reload
# systemctl stop nfs
```

- b) Enable the NFS services by using the following command:

```
# mmces service enable NFS
```

If you get the `/sbin/rpc.statd: not found [No such file or directory]` error, try the following workaround.

Run the following commands on every protocol node.

```
# ln -s /usr/sbin/rpc.statd /sbin/rpc.statd
# ln -s /sbin/rpcinfo /usr/sbin/rpcinfo
```

After you have run these commands, try enabling the NFS services again.

- Enable the SMB services by using the following command:

```
# mmces service enable SMB
```

- Enable the S3 services by using the following command:

```
mmces service enable S3
```

Verifying the IBM Storage Scale installation on Ubuntu Linux nodes

You can verify the installation of the IBM Storage Scale Ubuntu Linux packages on each node.

About this task

To check that the software is successfully installed, use the **dpkg** command:

```
dpkg -l | grep gpfs
```

The system returns output similar to the following if you have IBM Storage Scale Standard Edition or IBM Storage Scale Data Access Edition installed:

ii	gpfs.base	6.0.1-x	amd64	GPFS File Manager
ii	gpfs.compression	6.0.1-x	amd64	IBM Spectrum Scale Compression Libraries

ii gpfs.docs	6.0.1-x	all	GPFS Server Manpages and Documentation
ii gpfs.gpl	6.0.1-x	all	GPFS Open Source Modules
ii gpfs.gskit	8.0.55.x	amd64	GPFS GSKit Cryptography Runtime
ii gpfs.librdkafka	6.0.1-x	amd64	librdkafka shared library installation
ii gpfs.msg.en-us	6.0.1-x	all	GPFS Server Messages - U.S. English
ii gpfs.license.std	6.0.1-x	amd64	IBM Spectrum Scale Standard Edition ITLM files

or

ii gpfs.license.da	6.0.1-x	amd64	IBM Spectrum Scale Data Access Edition ITLM files

If you have IBM Storage Scale Advanced Edition or IBM Storage Scale Data Management Edition installed, you should also see the following lines in the output:

ii gpfs.adv	6.0.1-x	amd64	GPFS Advanced Features
ii gpfs.crypto	6.0.1-x	amd64	GPFS Cryptographic Subsystem
ii gpfs.license.adv	6.0.1-x	amd64	IBM Spectrum Scale Advanced Edition ITLM files

or

ii gpfs.license.dm files	6.0.1-x	amd64	IBM Spectrum Scale Data Management Edition ITLM files

Verifying the IBM Storage Scale installation on SLES and Red Hat Enterprise Linux nodes

You can verify the installation of the IBM Storage Scale SLES or Red Hat Enterprise Linux RPMs on each node.

About this task

To check that the software is successfully installed, use the `rpm` command:

```
rpm -qa | grep gpfs
```

The system returns output similar to the following if you have IBM Storage Scale Standard Edition or IBM Storage Scale Data Access Edition installed:

```
gpfs.docs-6.0.1-x
gpfs.base-6.0.1-x
gpfs.msg.en_US-6.0.1-x
gpfs.compression-6.0.1-x
gpfs.gpl-6.0.1-x
gpfs.gskit-8.0.55.x
gpfs.license.std-6.0.1-x
```

or

```
gpfs.license.da-6.0.1-x
```

If you have IBM Storage Scale Advanced Edition, IBM Storage Scale Data Management Edition, or IBM Storage Scale Developer Edition installed, you should also see the following in the output:

```
gpfs.crypto-6.0.1-x
gpfs.adv-6.0.1-x
gpfs.license.dm-6.0.1-x or gpfs.license.adv-6.0.1-x or gpfs.license.dev-6.0.1-x
```

Note: IBM Storage Scale Developer Edition is available on Red Hat Enterprise Linux on x86_64.

For installations that include IBM Storage Scale RAID, you should also see the following in the output:

```
gpfs.gn1-6.0.1-x
```

For more information about IBM Storage Scale RAID, see *IBM Storage Scale RAID: Administration in Elastic Storage Server (ESS) documentation*.

To run the system health check for verifying IBM Storage Scale installation, issue the following command:

```
mmhealth node show -N all
```

Manually installing the performance monitoring tool

The performance monitoring tool is automatically installed by the installation toolkit. You can also install the performance monitoring tool manually.

- For more information about the performance monitoring tool, see *Configuring the performance monitoring tool* in *IBM Storage Scale: Administration Guide*.
- For information about packages required before installation, see [“Software requirements” on page 310](#).
- For information about the installation toolkit, see [“Understanding the installation toolkit options” on page 516](#).

A single collector can easily support up to 150 sensor nodes. The collector can be any node on the system. All sensors report to this node. Select any node in the system to be the collector node. For information on configuring the sensors and collectors, see the *Configuring the performance monitoring tool* section in the *IBM Storage Scale: Problem Determination Guide*. You can manually install the tool by doing the following steps:

1. Download the installation images and install the performance monitoring packages.

For information on the default directories in which these packages are extracted, see [“Location of extracted packages” on page 480](#).

Performance monitoring packages

```
gpfs.gss.pmsensors-version_release.os.target_arch.file_format  
gpfs.gss.pmcollector-version_release.os.target_arch.file_format
```

Then, use your operating system’s native package management mechanism to install the packages.

For example:

- To install performance monitoring sensors on a Red Hat Enterprise Linux 7.x x86_64 node, use the following command:

```
rpm -ivh gpfs.gss.pmsensors-x.x.x-el7.x86_64.rpm
```

- To install performance monitoring sensors on a Red Hat Enterprise Linux 8.x x86_64 node, use the following command:

```
rpm -ivh gpfs.gss.pmsensors-x.x.x-el8.x86_64.rpm
```

- To install performance monitoring sensors on a SLES 15 x86_64 node, use the following command:

```
rpm -ivh gpfs.gss.pmsensors-x.x.x-sles12.x86_64.rpm
```

- To install performance monitoring sensors on an Ubuntu amd64 node, use the following command:

```
dpkg -i gpfs.gss.pmsensors_x.x.x-U*amd64.deb
```

- To install a performance monitoring collector on a Red Hat Enterprise Linux 7.x PPC64LE node, use the following command:

```
rpm -ivh gpfs.gss.pmcollector-x.x.x-el7.ppc64le.rpm
```

2. To configure performance monitoring after the packages are installed on all the selected nodes, use the following commands:

```
mmperfmon config generate --collectors collectorNode1[,collectorNode2,...]  
mmchnode --perfmon -N sensorNode1[,sensorNode2,...]
```

3. To enable performance monitoring on a protocol node for NFS or SMB, complete the following steps:

- For NFS:

- a. Install the pm-ganesha package as follows:

- RHEL or SLES: `rpm -ivh gpfs.pm-ganesha_version-release.os.target_arch.rpm`
- Ubuntu: `dpkg -i gpfs.pm-ganesha_version-release.deb`

- b. Ensure that the `/opt/IBM/zimon/defaults/ZIMonSensors_nfs.cfg` sensor is created as shown:

```
sensors={
    name = "NFSIO"
    period = 10
    type = "Generic"
    restrict = "cesNodes"
}
```

- c. Run the following command to add the sensor:

```
mmperfmon config add --sensors /opt/IBM/zimon/defaults/ZIMonSensors_nfs.cfg
```

- d. Run the following code to ensure that the newly added sensor is made visible to the system:

```
mmhealth node show nfs --refresh -N cesNodes
```

- For SMB:

- a. Run the following command to add the sensor:

```
mmperfmon config add --sensors /opt/IBM/zimon/defaults/ZIMonSensors_smb.cfg
```

Note: No additional sensor package is needed for the SMB sensors.

- b. Run the following code to ensure that the newly added sensor is made visible to the system:

```
mmhealth node show smb --refresh -N cesNodes
```

4. If the protocol sensors are enabled on a GPFS-only node, you see an error message stating that the sensors are unavailable. However, the other sensors continue to run.
5. Start the sensors on each node using the **systemctl start pmsensors.service** command.
6. On the collector nodes, start the collector, using the **systemctl start pmcollector.service** command.
7. To ensure that sensors and collectors are restarted after the node reboots, you can enable them using the following commands:

Sensors

To enable sensors, use the **systemctl enable pmsensors.service** command.

To disable sensors, use the **systemctl disable pmsensors.service** command.

Collector

To enable the collector, use the **systemctl enable pmcollector.service** command.

To disable the collector, use the **systemctl disable pmcollector.service** command.

The collector node starts gathering all the requested metrics.

Important: Once the packages are installed, pmsensor and pmcollector services are activated automatically.

By default, the installation toolkit enables sensors on each protocol node (CES node) but not on the other GPFS nodes (non-CES nodes).

Metrics can be retrieved from any node in the system by using the **mmperfmon query** command. For more information, see *mmperfmon command* in *IBM Storage Scale: Command and Programming Reference Guide*.

For more information about the performance monitoring tool, see *Configuring the performance monitoring tool* in *IBM Storage Scale: Problem Determination Guide*.

Manually installing IBM Storage Scale management GUI

The management GUI provides an easy way for the users to configure, manage, and monitor the IBM Storage Scale system.

You can install the management GUI by using the following methods:

- Installing the management GUI by using the installation toolkit. For more information, see [“Installing IBM Storage Scale management GUI by using the installation toolkit”](#) on page 550.
- Manual installation of the management GUI. The following sections provide the details of how to manually install the management GUI.

Prerequisites

The prerequisites that are applicable for installing the IBM Storage Scale system through CLI are applicable for GUI installation as well. For more information, see [“Installation prerequisites”](#) on page 470.

The IBM Storage Scale GUI package is also part of the installation package. You need to extract this package to start the installation. The performance tool packages enable the performance monitoring tool that is integrated into the GUI. The following packages are important for performance monitoring tools in the GUI:

- The performance tool collector package. This package is placed only on the collector nodes. By default, every GUI node is also used as the collector node to receive performance details and display them in the GUI.
- The performance tool sensor package. This package is applicable for the sensor nodes, if not already installed.
- *iptables* is required for the installation of Linux operating systems. However, it might not be a prerequisite if the administrator configures the firewall for the specific GUI node. For more information, see *Firewall recommendations for IBM Storage Scale GUI* in *IBM Storage Scale: Administration Guide*.

Note: For GUI installations on RHEL9 you must install *nftables*. You can run the **dnf install nftables** command to install *nftables*.

However, to make the GUI functional in a system in which *iptables* already exist, you must first disable the existing *iptables* and then enable the *nftables*.

- To disable *iptables*, issue the following command:

```
systemctl disable --now iptables
```

- To enable *nftables*, issue the following command:

```
systemctl enable --now nftables
```

Note: The GUI must be a homogeneous stack. That is, all packages must be of the same release. For example, do not mix the 5.1.2 GUI rpm with a 5.1.3 base rpm. However, GUI PTFs and interim fixes (fixes) can usually be applied without having to install the corresponding PTF or interim fix of the base package. A GUI PTF and interim fix are helpful if you want to get rid of a GUI issue without changing anything on the base layer.

The following table lists the IBM Storage Scale GUI and performance tool package that are essential for different platforms.

Table 40. GUI packages essential for each platform

GUI Platform	Package name
Red Hat Enterprise Linux (RHEL) 7.x, 8.x and 9.x	gpfs.gui-6.0.1-x.noarch.rpm gpfs.java-6.0.1-x.x86_64.rpm gpfs.java-6.0.1-x.ppc64le.rpm gpfs.java-6.0.1-x.s390x.rpm
SUSE Linux Enterprise Server (SLES) 15	gpfs.gui-6.0.1-x.noarch.rpm gpfs.java-6.0.1-x.x86_64.rpm gpfs.java-6.0.1-x.ppc64le.rpm gpfs.java_6.0.1-x_s390x.rpm
Ubuntu 20	gpfs.gui_6.0.1-x_all.deb gpfs.java_6.0.1-x_amd64.deb gpfs.java_6.0.1-x_ppc64el.deb
Performance monitoring tool platform	Performance monitoring tool rpms
RHEL 9.x x86	gpfs.gss.pmcollector-6.0.1-x.el8.x86_64.rpm
RHEL 8.x x86	gpfs.gss.pmcollector-6.0.1-x.el8.x86_64.rpm gpfs.gss.pmsensors-6.0.1-x.el8.x86_64.rpm
RHEL 7.x x86	gpfs.gss.pmcollector-6.0.1-x.el7.x86_64.rpm gpfs.gss.pmsensors-6.0.1-x.el7.x86_64.rpm
RHEL 7 s390x	gpfs.gss.pmsensors-6.0.1-x.el7.s390x.rpm gpfs.gss.pmcollector-6.0.1-x.el7.s390x.rpm
RHEL 8.x ppc64 LE	gpfs.gss.pmcollector-6.0.1-x.el8.ppc64le.rpm gpfs.gss.pmsensors-6.0.1-x.el8.ppc64le.rpm
RHEL 7.x ppc64 LE	gpfs.gss.pmcollector-6.0.1-x.el7.ppc64le.rpm gpfs.gss.pmsensors-6.0.1-x.el7.ppc64le.rpm
RHEL 9.x ARM 64	gpfs.gui-6.0.1-x.noarch.rpm

Table 40. GUI packages essential for each platform (continued)	
GUI Platform	Package name
SLES 15 x86	gpfs.gss.pmcollector-6.0.1-x.SLES15.x86_64.rpm gpfs.gss.pmsensors-6.0.1-x.SLES15.X86_64.rpm
SLES 15 s390x	gpfs.gss.pmsensors-6.0.1-x.SLES15.s390x.rpm gpfs.gss.pmcollector-6.0.1-x.SLES15.s390x.rpm
Ubuntu 20.04 LTS sensor and collector packages	gpfs.gss.pmsensors_6.0.1-x.U20.04_amd64.deb gpfs.gss.pmcollector_6.0.1-x.U20.04_amd64.deb
Ubuntu 22.04.x ARM 64	gpfs.java-6.0.1-x.aarch64.rpm

Make sure that the performance tool collector runs on the same node as the GUI.

Yum repository setup

You can use yum repository to manually install the GUI rpm files. It is the preferred way of GUI installation as yum checks the dependencies and automatically installs missing platform dependencies like the postgres module, which is required but not included in the package.

Note: Refer to the PostgreSQL documentation for the latest supported version: <https://www.postgresql.org/docs/>

Installation steps

You can install the management GUI either by using the package manager (yum or zypper commands) or by installing the packages individually.

Installing management GUI by using package manager (yum or zypper commands)

It is recommended to use this method as the package manager checks the dependencies and automatically installs missing platform dependencies. Issue the following commands to install the management GUI:

Red Hat Enterprise Linux

```
yum install gpfs.gss.pmsensors-6.0.1-x.elx.<arch>.rpm
yum install gpfs.gss.pmcollector-6.0.1-x.elx.<arch>.rpm
yum install gpfs.java-6.0.1-x.<arch>.rpm
yum install gpfs.gui-6.0.1-x.noarch.rpm
```

SLES

```
zypper install gpfs.gss.pmsensors-6.0.1-x.SLES15.<arch>.rpm
zypper install gpfs.gss.pmcollector-6.0.1-x.SLES15.<arch>.rpm
zypper install gpfs.java-6.0.1-x.<arch>.rpm
zypper install gpfs.gui-6.0.1-x.noarch.rpm
```

Installing management GUI by using RPM

Issue the following commands for both RHEL and SLES platforms:

```
rpm -ivh gpfs.java-6.0.1-x.<arch>.rpm
rpm -ivh gpfs.gss.pmsensors-6.0.1-x.elx.<arch>.rpm
rpm -ivh gpfs.gss.pmcollector-6.0.1-x.elx.<arch>.rpm
rpm -ivh gpfs.gui-6.0.1-x.noarch.rpm
```

Installing the management GUI on Ubuntu by using dpkg and apt-get

Issue the following commands for Ubuntu platforms:

```
dpkg -i gpfs.java_6.0.1-x_<arch>.deb
dpkg -i gpfs.gss.pmsensors_6.0.1-x.<os>_<arch>.deb
dpkg -i gpfs.gss.pmcollector_6.0.1-x.<os>_<arch>.deb
apt-get install postgresql
dpkg -i gpfs.gui_6.0.1-x_all.deb
```

The sensor package must be installed on any additional node that you want to monitor. All sensors must point to the collector node.

Note: In IBM Storage Scale 5.1.4 you can disable the GUI service after you have installed it without restricting your access to the REST API service. You can configure `WEB_GUI_ENABLED = false` in the `gpfsGui.properties` file to access only the REST API service. The **WEB_GUI_ENABLED** parameter value is set to `true` by default.

Start the GUI

Start the GUI by issuing the **systemctl start gpfsGui** command.

Note: After installing the system and GUI package, you need to create the first GUI user to log in to the GUI. This user can create other GUI administrative users to perform system management and monitoring tasks. When you launch the GUI for the first time after the installation, the GUI welcome page provides options to create the first GUI user from the command line prompt by using the `/usr/lpp/mmfs/gui/cli/mkuser <user_name> -g SecurityAdmin` command.

Installing GUI in a FIPS-enabled environment

You can install GUI in a FIPS-enabled environment.

Important: To verify whether FIPS mode is enabled on the system, use the following command:

```
fips-mode-setup --check
```

Perform the following steps on nodes where you would like to install GUI:

1. Extract the IBM Storage Scale SE package.
2. Install PostgreSQL:

```
yum install postgresql-server postgresql-contrib
```

3. Navigate to the `gpfs_rpms` directory and install the Java RPM with flags `--nofiledigest` and `--nodigest`:

```
rpm -ivh --nofiledigest --nodigest gpfs.java-5.2.3-3.x86_64.rpm
```

4. Install the GUI RPM with flags `--nofiledigest` and `--nodigest`:

```
-
rpm -ivh --nofiledigest --nodigest gpfs.gui-5.2.3-3.noarch.rpm --nodeps
```

5. Restart the GUI service:

```
systemctl restart gpfsGui
```

Note: Once the installation is complete, deploy the cluster using the Ansible Toolkit.

Enabling performance tools in management GUI

The performance tool consists of sensors that are installed on all nodes that need to be monitored. It also consists of one or more collectors that receive data from the sensors. The GUI expects that a collector runs on a GUI node. The GUI queries the collectors for performance and capacity data. The following steps use the automated approach to configure and maintain performance data collection by using the **mmperfmon** CLI command. Manually editing the `/opt/IBM/zimon/ZIMonSensors.cfg` file is not compatible with this configuration mode.

1. Install the necessary software packages. Install the collector software package, `gpfs.gss.pmcollector`, on all GUI nodes. Install the sensor software packages, `gpfs.gss.pmsensors`, on all nodes, which are supposed to send the performance data.
2. Initialize the performance collection. Use the **mmperfmon config generate --collectors [node list]** command to create an initial performance collection setup on the selected nodes. The GUI nodes must be configured as collector nodes. Depending on the installation type, this configuration might be already completed before. However, verify the existing configuration.
3. Enable nodes for performance collection. You can enable nodes to collect performance data by issuing the **mmchnode --perfmon -N [SENSOR_NODE_LIST]** command. **[SENSOR_NODE_LIST]** is a comma-separated list of sensor nodes' hostnames or IP addresses and you can also use a node class. Depending on the type of installation, nodes might be configured for performance collection.
4. Review peer configuration for the collectors. The **mmperfmon config update** command updates the multiple collectors with the necessary configuration. The collector configuration is stored in the `/opt/IBM/zimon/ZIMonCollector.cfg` file. This file defines the collector peer configuration and the aggregation rules. If you are using only a single collector, you can skip this step. The GUI must have access to all data from each GUI node. For more information, see *Configuring the collector* in *IBM Storage Scale: Administration Guide*.
5. Review aggregation configuration for the collectors. The collector configuration is stored in the `/opt/IBM/zimon/ZIMonCollector.cfg` file. The performance collection tool is configured with predefined rules on how data is aggregated when it gets older. By default, four aggregation domains are created as shown:
 - A raw domain that stores the metrics uncompressed.
 - A first aggregation domain that aggregates data to 30-second averages.
 - A second aggregation domain that stores data in 15-minute averages.
 - A third aggregation domain that stores data in 6-hour averages.

Do not change the default **aggregation configuration** because, if you do, the already collected historical metric information might get lost. You cannot manually edit the `/opt/IBM/zimon/ZIMonCollector.cfg` file in the automated configuration mode.

In addition to the aggregation that is done by the performance collector, the GUI might request aggregated data based on the zoom level of the chart. For more information, see *Configuring the collector* in *IBM Storage Scale: Administration Guide*.

6. Configure the sensors. Several GUI pages display performance data that is collected with the help of performance monitoring tools. If data is not collected, the GUI shows the error messages like "No Data Available" or "Objects not found" in the performance charts. Installation by using the *spectrumscale installation toolkit* manages the default performance monitoring installation and configuration. The GUI help that is available on the various pages shows performance metric information. The GUI context-sensitive help also lists the sensor names.

The **Services > Performance Monitoring** page provides option to configure the sensor configuration and provides hints for collection periods and restriction of sensors to specific nodes.

You can also use the **mmperfmon config show** command in the CLI to verify the sensor configuration. Use the **mmperfmon config update** command to adjust the sensor configuration to match your needs. For more information, see *Configuring the sensor* in *IBM Storage Scale: Administration Guide*.

The local file `/opt/IBM/zimon/ZIMonSensors.cfg` can be different on every node and the system might change this path whenever a configuration change occurs. Therefore, this file must not be edited manually when you are using the automated configuration mode. During distribution of the sensor configuration, the restrict clause is evaluated and the period for all sensors is set to 0 in the `/opt/IBM/zimon/ZIMonSensors.cfg` file. The setting is defined for those nodes that did not match the restrict clause. You can check the local file to confirm that a restrict clause worked as intended.

Configuring capacity-related sensors to run on a single-node

Several capacity-related sensors must run only on a single node as they collect data for a clustered file system. For example, *GPFSDiskCap*, *GPFSFilesetQuota*, *GPFSFileset* and *GPFSPool*.

It is possible to automatically restrict these sensors to a single node. For new installations, capacity-related sensors are automatically configured to a single node where the capacity collection occurs. An updated cluster, which was installed before ESS 5.3.7 (IBM Storage Scale 5.0.5), might not be configured to use this feature automatically and must be reconfigured. To update the configuration, you can use the **mmperfmon config update SensorName.restrict=@CLUSTER_PERF_SENSOR** command, where **SensorName** values include *GPFSFilesetQuota*, *GPFSFileset*, *GPFSPool*, and *GPFSDiskCap*.

To collect file system and disk level capacity data on a single node that is selected by the system, run the following command to update the sensor configuration.

```
mmperfmon config update GPFSDiskCap.restrict=@CLUSTER_PERF_SENSOR
```

If the selected node is in the DEGRADED state, then the CLUSTER_PERF_SENSOR is automatically reconfigured to another node that is in the HEALTHY state. The performance monitoring service is restarted on the previous and currently selected nodes. For more information, see *Automatic assignment of single node sensors* in *IBM Storage Scale: Problem Determination Guide*.

Note: If the *GPFSDiskCap* sensor is frequently restarted, it can negatively impact the system performance. The *GPFSDiskCap* sensor can cause a similar impact on the system performance as the **mmddf** command. Therefore, to avoid using the **@CLUSTER_PERF_SENSOR** for any sensor in the **restrict** field of a single node sensor until the node stabilizes in the HEALTHY state, it is advisable to use a dedicated healthy node. If you manually configure the **restrict** field of the capacity sensors then you must ensure that all the file systems on the specified node are mounted to record file system-related data, like capacity.

Use the **Services > Performance Monitoring** page to select the appropriate data collection periods for these sensors.

For the *GPFSDiskCap* sensor, the recommended period is 86400, which means once per day. As the *GPFSDiskCap*.period sensor runs **mmddf** command to get the capacity data, it is not recommended to use a value less than 10800 (every 3 hours). To show fileset capacity information, it is necessary to enable quota for all file systems where fileset capacity must be monitored. For more information, see the **-q** option in the **mmchfs** command and **mmcheckquota** command.

To update the sensor configuration for triggering an hourly collection of capacity-based fileset capacity information, run the **mmperfmon** command as shown in the following example,

```
mmperfmon config update GPFSFilesetQuota.restrict=@CLUSTER_PERF_SENSOR gui_node
GPFSFilesetQuota.period=3600
```

Checking GUI and performance tool status

Issue the **systemctl status gpfsgui** command to know the GUI status as shown in the following example.

```
systemctl status gpfsgui.service
gpfsgui.service - IBM_GPFS_GUI Administration GUI
Loaded: loaded (/usr/lib/systemd/system/gpfsgui.service; disabled)
Active: active (running) since Fri 2015-04-17 09:50:03 CEST; 2h 37min ago
Process: 28141 ExecStopPost=/usr/lpp/mmfs/gui/bin/cfgmantraclient unregister (code=exited, status=0/SUCCESS)
```

```

Process: 29120 ExecStartPre=/usr/lpp/mmfs/gui/bin/check4pgsql (code=exited, status=0/SUCCESS)
Main PID: 29148 (java)
Status: "GSS/GPFS GUI started"
CGroup: /system.slice/gpfsgui.service
<<<-29148 /opt/ibm/wlp/java/jre/bin/java -XX:MaxPermSize=256m -Dcom.ibm.gpfs.platform=GPFS
-Dcom.ibm.gpfs.vendor=IBM -Djava.library.path=/opt/ibm/wlp/usr/servers/gpfsgui/lib/
-javaagent:/opt/ibm/wlp/bin/tools/ws-javaagent.jar -jar /opt/ibm/wlp/bin/tools/ws-server.jar
gpfsgui
--clean

Apr 17 09:50:03 server-21.localnet.com java[29148]: Available memory in the JVM: 484MB
Apr 17 09:50:03 server.localnet.com java[29148]: Max memory that the JVM will attempt to use:
512MB
Apr 17 09:50:03 server.localnet.com java[29148]: Number of processors available to JVM: 2
Apr 17 09:50:03 server.localnet.com java[29148]: Backend started.
Apr 17 09:50:03 server.localnet.com java[29148]: CLI started.
Apr 17 09:50:03 server.localnet.com java[29148]: Context initialized.
Apr 17 09:50:03 server.localnet.com systemd[1]: Started IBM_GPFS_GUI Administration GUI.
Apr 17 09:50:04 server.localnet.com java[29148]: [AUDIT ] CWWKZ0001I: Application /
started in 6.459 seconds.
Apr 17 09:50:04 server.localnet.com java[29148]: [AUDIT ] CWWKF0012I: The server
installed the following features: [jdbc-4.0, ssl-1.0, localConnector-1.0, appSecurity-2.0,
jsp-2.2, servlet-3.0, jndi-1.0, usr:FscUserRepo, distributedMap-1.0].
Apr 17 09:50:04 server-21.localnet.com java[29148]: [AUDIT ] CWWKF0011I: ==> When you see
the service was started anything should be OK !

```

Issue the **systemctl status pmcollector** and **systemctl status pmsensors** commands to know the status of the performance tool.

You can also check whether the performance tool backend can receive data by using the GUI. As another option, you can also use a command-line performance tool that is called **zc**, which is available in **/opt/IBM/zimon** folder. For example,

```

echo "get metrics mem_active, cpu_idle, gpfs_ns_read_ops last 10 bucket_size 1" | /opt/IBM/
zimon/zc
1:      node-11|Memory|mem_active
2:      node-12|Memory|mem_active
3:      node-13|Memory|mem_active
4:      node-11|CPU|cpu_idle
5:      node-12|CPU|cpu_idle
6:      node-13|CPU|cpu_idle
7:      node-11|GPFSNode|gpfs_ns_read_ops
8:      node-12|GPFSNode|gpfs_ns_read_ops
9:      node-13|GPFSNode|gpfs_ns_read_ops
Row      Timestamp      mem_active      mem_active      cpu_idle
cpu_idle      gpfs_ns_read_ops      gpfs_ns_read_ops      gpfs_ns_read_ops
1      2025-11-07 09:09:06      1721760 883860      1040216 95.980000      98.000000
28.790000      null      null      null
2      2025-11-07 09:09:07      1733912 883988      1025988 70.050000      98.990000
77.270000      null      null      null
3      2025-11-07 09:09:08      1734344 883988      1024680 25.760000      99.500000
98.480000      null      null      null
4      2025-11-07 09:09:09      1722124 884120      1024680 74.240000      95.480000
98.490000      null      null      null
5      2025-11-07 09:09:10      1722124 883972      1024680 98.010000      98.510000
99.000000      null      null      null
6      2025-11-07 09:09:11      1722124 883972      1024680 98.480000      98.990000
98.480000      null      null      null
7      2025-11-07 09:09:12      1722124 883972      1024680 98.990000      99.490000
98.500000      null      null      null
8      2025-11-07 09:09:13      1722052 883972      1024680 97.990000      99.000000
98.990000      null      null      null
9      2025-11-07 09:09:14      1722052 883972      1024680 99.490000      98.500000
98.500000      null      null      null
10     2025-11-07 09:09:15      1722048 883972      1024680 94.030000      99.490000
98.000000      null      null      null
.

```

Node classes used for the management GUI

The IBM Storage Scale management GUI automatically creates the following node classes during installation:

- **GUI_SERVERS**: Contains all nodes with a server license and all the GUI nodes
- **GUI_MGMT_SERVERS**: Contains all GUI nodes

Each node on which the GUI services are started is added to these node classes.

Nodes can also be removed from Node classes. For more information, see [“Removing nodes from management GUI-related node class” on page 667](#).

For more information about node classes, see *Specifying nodes as input to GPFS commands* in *IBM Storage Scale: Administration Guide*.

Related concepts

[“Uninstalling the IBM Storage Scale management GUI” on page 666](#)

Do the following to uninstall management GUI and remove the performance monitoring components that are installed for the GUI:

[“IBM Storage Scale GUI” on page 194](#)

You can configure and manage various features that are available with the IBM Storage Scale system by using the IBM Storage Scale management GUI.

Migrating iptables rules to nftables

You can translate and migrate existing iptables rules to nftables during installation of GUI on RHEL 9 systems.

Before you begin

Before proceeding with the migration process, you must ensure that the following activities are completed:

- The nftables and iptables packages are installed.
- iptables and ip6tables rules are configured in the system.

About this task

While installing GUI on RHEL 9 systems you must install nftables. However, you can retain the previously created iptables rules that are already existing by translating and migrating them to nftables.

Procedure

1. To create a backup of the iptables and ip6tables rules, issue the following commands:

```
# iptables-save >/root/iptables.dump
# ip6tables-save >/root/ip6tables.dump
```

2. To convert the dump files to nftables, issue the following commands:

```
# iptables-restore-translate -f /root/iptables.dump > /etc/nftables/ruleset-migrated-from-iptables.nft
# ip6tables-restore-translate -f /root/ip6tables.dump > /etc/nftables/ruleset-migrated-from-ip6tables.nft
```

3. To edit the /etc/sysconfig/nftables.conf file and add the migrated files, issue the following command:

```
include "/etc/nftables/ruleset-migrated-from-iptables.nft"
include "/etc/nftables/ruleset-migrated-from-ip6tables.nft"
```

4. To disable the existing iptables, if any, issue the following command:

```
# systemctl disable --now iptables
```

5. To enable and start the nftables service, issue the following command:

```
# systemctl enable --now nftables
```

6. To verify the rules that are migrated, issue the following command:

```
# nft list ruleset
```

Root privilege considerations for IBM Storage Scale management GUI

The IBM Storage Scale GUI WebSphere Java process runs as a user named "scalemgmt". This provides improved security because web applications running as a non-root user are less vulnerable to security threats. The *scalemgmt* user is set up as a system account with no login privileges.

The GUI and GUI user still need root privileges in the backend to perform the following tasks:

- Monitoring and managing the cluster by issuing the IBM Storage Scale CLI commands.
- Bind to the privileged ports such as 80 (http) and 443 (https) to work with the GUI.

The system automatically performs the required configuration. It does not require any configuration to be done by the user.

Enabling the scalemgmt user to monitor and manage the system through GUI

As root privileges are not available to the GUI user, the system enables the scalemgmt user to run the CLI commands through sudo. The GUI installation adds the file `/etc/sudoers.d/scalemgmt_sudoers`, which allows the scalemgmt user to run commands that match the wildcard `"/usr/lpp/mmfs/bin/mm"`.

During installation of the IBM Storage Scale management GUI, the following lines are appended to the `/etc/sudoers` file:

```
## Read drop-in files from /etc/sudoers.d
#include_dir /etc/sudoers.d
```

Note: Do not alter these lines if you are editing the `/etc/sudoers` file. If these lines are removed or moved from the default location, the GUI cannot run the CLI commands as a non-root user.

Binding to the GUI privileged ports

A non-root process cannot bind privileged ports such as 80 and 443. So, the GUI process now runs on port 47443 and 47080 and uses *iptables* rules to forward port 443 to 47443 and 80 to 47080. The GUI is only available in the browser on port 443 through HTTPS. However, versions prior to 4.2.3 used port 80 to send events to the GUI, so to stay backwards compatible in clusters that still contain nodes with older versions, the GUI still accepts events on port 80.

To reach the GUI from outside, not only port 443 and 80 but also port 47443 and 47080 must be opened to the outside. This is because the port forwarding redirects port 443 to 47443 and 80 to 47080 before evaluating the rules of the *iptables* INPUT chain.

The *iptables* rules necessary for the port forwarding are automatically checked every time the GUI is started through the **systemctl start gpfsgui** command. The user does not have to configure anything manually.

Manually installing IBM Storage Scale bridge for Grafana with the RPM

You can install the IBM Storage Scale bridge for Grafana feature by using a uniquely named set of RPM packages.

Note: It's recommended to use the installation toolkit for production deployments because it automates installation and configuration tasks.

On a GPFS cluster, the Grafana bridge should be installed on a pmcollector node whenever possible. Installing it on a non-gpfs pmcollector node outside the cluster is possible, but it requires additional configuration and might lead to performance issues, such as slower system response times. For security, the Grafana bridge uses HTTPS connections by default.

Manual installation should be used only in specific scenarios, such as:

- Deploying an additional Grafana bridge instance on the same node

- Installing the Grafana bridge on non-scale nodes. This requires additional configuration that can impact performance.
- Environments not managed by the installation toolkit

Before you begin, ensure the following:

- Grafana bridge must run on a pmcollector-node
- If installing on non-scale nodes, ensure that the minimum required Python version is installed.

1. Download the Grafana bridge package with the format: `gpfs.grafana-bridge-version_release.arch.file_format`

For example:

- `gpfs.grafana-bridge-9.0.1-1.noarch.rpm`
- `gpfs.grafana-bridge_9.0.1-1_all.deb`

2. Install the package using any of these commands:

- To install Grafana bridge on Red Hat® Enterprise Linux® 9.x x86_64 node, use the following command:

```
rpm -ivh gpfs.grafana-bridge-x.x.x-x.noarch.rpm
```

- To install Grafana bridge on Ubuntu amd64 node, use the following command:

```
dpkg -i gpfs.grafana-bridge_x.x.x-x_all.deb
```

3. Add the following API keys to enable authenticated communication with the internal IBM Storage Scale performance monitoring system:

- `mmperfmon config add --apiKey scale_grafana`
- `mmperfmon config show --apiKey scale_grafana`

4. Update the `/etc/grafanbridge/config.ini` file by adding the following section for the API key:

```
##### GPFS Server #####
[server]
# The host name or ip address of the pmcollector node
server = localhost

# The https port to use
serverPort = 9980

# The sleep time in seconds before attempting to get metaData from the server again.
# This most likely the case if the query was sent to server before it did finish the
# initial startup (maximum 3 attempts)
retryDelay = 60

# The name of REST HTTPS API key name
apiKeyName = scale_grafana

# The REST HTTPS API key value, f.e:
apiKeyValue = e40960c9-de0a-4c75-bc71-0bcae6db23b2
```

5. Reload the systemd configuration and enable the Grafana bridge service:

- `systemctl daemon-reload`
- `systemctl enable gpfs.grafana-bridge.service`

6. Start the Grafana bridge service:

```
systemctl start gpfs.grafana-bridge.service
```

7. Verify the Grafana bridge service status:

```
systemctl status gpfs.grafana-bridge.service
```

A sample output is as follows:

```
gpfs.grafana-bridge.service - IBM Storage Scale bridge for Grafana
Loaded: loaded (/etc/systemd/system/gpfs.grafana-bridge.service; enabled; preset:
disabled)
Active: active (running) since Mon 2026-03-09 11:57:56 CET; 15s ago
Main PID: 2701837 (python3.12)
Tasks: 15 (limit: 74669)
Memory: 42.8M
CPU: 2.211s
CGroup: /system.slice/gpfs.grafana-bridge.service
└─2701837 /usr/bin/python3.12 /opt/IBM/grafanabridge/zimonGrafanaIntf.py --
configFile /etc/grafanabridge/config.ini

Mar 09 11:57:56 scale-12 grafana-bridge[2701837]: Using mmpython
Mar 09 11:57:57 scale-12 grafana-bridge[2701837]: All dependencies from /opt/IBM/
grafanabridge/requirements.in are installed for interpreter /usr/lpp/mmfs/bin/mmpython
Mar 09 11:57:57 scale-12 grafana-bridge[2701837]: TLS certs already present at /etc/
grafanabridge/ssl/certs
Mar 09 11:57:58 scale-12 grafana-bridge[2701837]: 2026-03-09 11:57 -
MainThread - INFO - *** IBM Storage Scale bridge for
Grafana - Version: 9.0.1 ***
Mar 09 11:57:58 scale-12 grafana-bridge[2701837]: 2026-03-09 11:57 -
MainThread - INFO - Successfully retrieved MetaData
Mar 09 11:57:58 scale-12 grafana-bridge[2701837]: 2026-03-09 11:57 -
MainThread - INFO - Received sensors:CPU, CTDBStats,
CTDBStats, DiskFree, GPFSBufMgr, GPFSXPELNODE, GPFSFilesystem, GPFSFilesystemAPI,
GPFSNSDDisk, GPFSNSDFS, GPFSNSDPool, GPFSNode, GPFSNodeAPI, GPFSRPCS, GPFSSTCOM, GPFSVFSX,
GPFSWaiters, GPFSmmhealth, Load, Memory, Netstat, Network, SMBGlobalStats, SMBStats,
TopProc, IPFIX, GPFSDiskCap, GPFSXPEL, GPFSFileset, GPFSInodeCap, GPFSPool, GPFSPoolCap
Mar 09 11:57:58 scale-12 grafana-bridge[2701837]: 2026-03-09 11:57 -
MainThread - INFO - Initial cherryPy server engine start
have been invoked. Python version: 3.12.5 (main, Aug 23 2024, 00:00:00) [GCC 11.5.0 20240719
(Red Hat 11.5.0-2)], cherryPy version: 18.10.0.
Mar 09 11:57:58 scale-12 grafana-bridge[2701837]: 2026-03-09 11:57 -
MainThread - INFO - Registered applications:
Mar 09 11:57:58 scale-12 grafana-bridge[2701837]: OpenTSDB Api listening on Grafana queries
Mar 09 11:57:58 scale-12 grafana-bridge[2701837]: 2026-03-09 11:57 -
MainThread - INFO - server started
```

If the Grafana bridge does not start using systemd (systemctl), you can try starting it manually by using one of the following methods:

- Run the Python command:

```
/usr/bin/python3.12 /opt/IBM/grafanabridge/zimonGrafanaIntf.py --configFile /etc/
grafanabridge/config.ini
```

- Run the startup script:

```
/opt/IBM/grafanabridge/bridge_startup.sh
```

8. Optional: To activate the Prometheus Exporter, change the settings in the Prometheus Exporter API Connection Defaults section:

```
##### Prometheus Exporter API Connection Defaults #####
[prometheues_exporter_plugin]
# Port number the bridge listening on for Prometheus server requests;
prometheus = 9250
```

9. To remove the Grafana bridge package, run one of the following commands:

- `yum remove gpfs.grafana-bridge`
- `rpm -e gpfs.grafana-bridge`

Troubleshooting

This section provides information to help identify and resolve issues with the Grafana bridge.

If the Grafana bridge does not start or function correctly, follow these recommendations:

- Review the Grafana bridge log file located at:

```
/var/log/ibm_bridge_for_grafana/zserver.log
```

The default log level is 15, as defined in `/etc/grafanabridge/config.ini`.

```
##### Logging #####  
  
[logging]  
# Directory where the bridge can store logs  
logPath = /var/log/ibm_bridge_for_grafana  
  
# log level 5 (TRACE), 10 (DEBUG), 15 (MOREINFO), 20 (INFO), 30 (WARN),  
# 40 (ERROR) (Default: 15)  
logLevel = 15  
  
# Log file name (Default: zserver.log)  
# Comment out this setting if you want to print trace messages directly on the command line  
logFile = zserver.log
```

You can also set the log level from the command line:

```
python3 /opt/IBM/grafanabridge/zimonGrafanaIntf.py --configFile /etc/grafanabridge/config.ini  
--logLevel LOGLEVEL
```

You can increase the log level in `config.ini` or use the `--logLevel` option to enable more detailed tracing, and then restart the Grafana bridge.

- If the `apiKey` is not set in `/etc/grafanabridge/config.ini`, the Grafana bridge fails to start.

A sample output is as follows:

```
# /opt/IBM/grafanabridge/bridge_startup.sh  
  
Using mmpython  
All dependencies from /opt/IBM/grafanabridge/requirements.in are installed  
TLS certs already present at /etc/grafanabridge/ssl/certs  
Missing mandatory ApiKey settings, quitting
```

Resolution: Uncomment and set the `apiKey` field in `/etc/grafanabridge/config.ini` (refer to Step 4), and then restart the Grafana bridge by using `systemctl restart gpfs.grafana-bridge.service`

- If required dependencies from `/opt/IBM/grafanabridge/requirements.in` are missing, you get a sample output as follows:

```
mmpython not found, falling back to system Python  
Auto-selected python version: 3.12  
  
The following dependencies are missing:  
- CherryPy
```

Resolution: Install the required dependencies by running:

```
/usr/bin/python3.12 -m pip install -r /opt/IBM/grafanabridge/requirements.in
```

After applying the troubleshooting recommendations, verify successful startup of the Grafana bridge.

Run the following command to start the Grafana bridge:

```
/opt/IBM/grafanabridge/bridge_startup.sh
```

- If `mmpython` is available, you can see the output similar to the following:

```
Using mmpython  
All dependencies from /opt/IBM/grafanabridge/requirements.in are installed for  
interpreter /usr/lpp/mmfs/bin/mmpython  
TLS certs already present at /etc/grafanabridge/ssl/certs  
2026-02-26 07:48 - MainThread - INFO - *** IBM Storage Scale bridge for Grafana - Version:
```



```

9.0.1 ***
2026-02-26 07:48 - MainThread - INFO - Successfully retrieved MetaData
2026-02-26 07:48 - MainThread - INFO - Received sensors:CPU, CTDBDBStats, CTDBStats,
DiskFree, GPFSSBufMgr, GPFSEXPELNODE, GPFSSFilesystem, GPFSSFilesystemAPI, GPFSSNDDisk,
GPFSSNDFS, GPFSSNDPool, GPFSSNode, GPFSSNodeAPI, GPFSSRPCS, GPFSTSCOM, GPFSSVFSX, GPFSSWaiters,
GPFSSmmhealth, Load, Memory, Netstat, Network, SMBGlobalStats, SMBStats, TopProc, IPFIX,
GPFSDiskCap, GPFSEXPEL, GPFSSFileset, GPFSSInodeCap, GPFSPool, GPFSPoolCap
2026-02-26 07:48 - MainThread - INFO - Initial cherryPy server engine start have been
invoked. Python version: 3.12.5 (main, Aug 23 2024, 00:00:00) [GCC 11.5.0 20240719 (Red Hat
11.5.0-2)], cherryPy version: 18.10.0.

```

- If mmpython is not available, you can see the output similar to the following:

```

mmpython not found, falling back to system Python
Auto-selected python version: 3.12
All dependencies from /opt/IBM/grafanabridge/requirements.in are installed for
interpreter /usr/bin/python3.12
TLS certs already present at /etc/grafanabridge/ssl/certs
2026-02-26 07:40 - MainThread - INFO - *** IBM Storage Scale bridge for Grafana - Version:
9.0.1 ***
2026-02-26 07:40 - MainThread - INFO - Successfully retrieved MetaData
2026-02-26 07:40 - MainThread - INFO - Received sensors:CPU, CTDBDBStats, CTDBStats,
DiskFree, GPFSSBufMgr, GPFSEXPELNODE, GPFSSFilesystem, GPFSSFilesystemAPI, GPFSSNDDisk,
GPFSSNDFS, GPFSSNDPool, GPFSSNode, GPFSSNodeAPI, GPFSSRPCS, GPFSTSCOM, GPFSSVFSX, GPFSSWaiters,
GPFSSmmhealth, Load, Memory, Netstat, Network, SMBGlobalStats, SMBStats, TopProc, IPFIX,
GPFSDiskCap, GPFSEXPEL, GPFSSFileset, GPFSSInodeCap, GPFSPool, GPFSPoolCap
2026-02-26 07:40 - MainThread - INFO - Initial cherryPy server engine start have been
invoked. Python version: 3.12.5 (main, Aug 23 2024, 00:00:00) [GCC 11.5.0 20240719 (Red Hat
11.5.0-2)], cherryPy version: 18.10.0.
2026-02-26 07:40 - MainThread - INFO - Registered applications:
OpenTSDB Api listening on Grafana queries
2026-02-26 07:40 - MainThread - INFO - server started

```

Installing IBM Storage Scale on Linux nodes with the installation toolkit

You can install IBM Storage Scale by using the installation toolkit. You can also use the installation toolkit for tasks such as deploying protocols on existing clusters and adding nodes to an existing IBM Storage Scale cluster.

Overview of the installation toolkit

The installation toolkit automates the steps that are required to install IBM Storage Scale, deploy protocols, and upgrade to a later IBM Storage Scale release. For a release-wise list of features available in the installation toolkit, see [Table 41 on page 511](#).

When using the installation toolkit, you provide environmental information based on which the installation toolkit dynamically creates a cluster definition file. Thereafter, the installation toolkit installs, configures, and deploys the specified configuration.

The installation toolkit enables you to do the following tasks:

- Install and configure IBM Storage Scale.
- Install the IBM Storage Scale native REST API.
- Add IBM Storage Scale nodes to an existing cluster.
- Deploy and configure SMB, NFS, HDFS, S3, and performance monitoring tools.
- Perform verification before installing, deploying, or upgrading. It includes checking whether passwordless SSH is set up correctly.
- Enable and configure call home and file audit logging functions.
- Upgrade IBM Storage Scale.
- Install the IBM Storage Scale Grafana bridge on scale collector nodes.
- Install and configure NFS core dump on NFS protocol nodes.

Installation and configuration are driven through commands.

From the self-extracting package, the installation toolkit is extracted to this directory, by default:

```
/usr/lpp/mmfs/package_code_version/ansible-toolkit
```

Using the installation toolkit is driven through the **spectrumscale** command in this directory, and this directory can optionally be added to the path.

The installation toolkit operation consists of four phases:

1. User input by using **spectrumscale** commands:

- a. All user input is recorded into a cluster definition file in `/usr/lpp/mmfs/6.0.1.0/ansible-toolkit/ansible/vars`.
- b. Review the cluster definition file to make sure that it accurately reflects your cluster configuration.
- c. As you input your cluster configuration, you can have the installation toolkit act on parts of the cluster by not specifying nodes that might have incompatible operating systems, OS versions, or architectures.

2. A **spectrumscale install** phase:

- a. Installation acts upon all nodes that are defined in the cluster definition file.
- b. GPFS and performance monitoring packages are installed.
- c. File audit logging and AFM to cloud object storage packages might be installed.
- d. GPFS portability layer is created.
- e. GPFS is started.
- f. A GPFS cluster is created.
- g. Licenses are applied.
- h. GUI nodes might be created and the GUI might be started upon these nodes.
- i. Performance monitoring, GPFS ephemeral ports, and cluster profile might be configured.
- j. NSDs are created.
- k. File systems are created.

3. A **spectrumscale deploy** phase:

- a. Deployment acts upon all nodes that are defined into the cluster definition file.
- b. SMB, NFS, HDFS, S3 protocol packages are copied to all protocol nodes and installed.
- c. SMB, NFS, HDFS, S3 services might be started.
- d. File audit logging might be configured.
- e. Licenses are applied.
- f. GUI nodes might be created and the GUI might be started upon these nodes.
- g. Performance monitoring, call home, file audit logging, GPFS ephemeral ports, and cluster profile might be configured.

4. A **spectrumscale upgrade** phase:

Note: The upgrade phase does not allow new function to be enabled. In the upgrade phase, the required packages are upgraded, but adding functions must be done either before or after the upgrade.

- a. Upgrade acts upon all nodes input into the cluster definition file. However, you can exclude a subset of nodes from the upgrade configuration.
- b. All installed or deployed components are upgraded. During the upgrade phase, any missing packages might be installed and other packages that are already installed are upgraded.
- c. Upgrade can be done in the following ways:

- Online upgrade (one node at a time)

Online upgrades are sequential with multiple passes. For more information, see [“Upgrade process flow” on page 725](#).

- Offline upgrade

Offline upgrades can be done in parallel saving a lot of time in the upgrade window.

- Upgrade while excluding a subset of nodes

- d. Allows for prompting to be enabled on a node to pause to allow for application migration from the node before proceeding with upgrade.

For more information, see [“Upgrading IBM Storage Scale components with the installation toolkit” on page 723](#) and [“Upgrade process flow” on page 725](#).

For information about command options available with the **spectrumscale** command, see the spectrumscale command description in the *IBM Storage Scale: Command and Programming Reference Guide*.

The following table lists the features available in the installation toolkit in the reverse chronological order of releases.

<i>Table 41. Installation toolkit: List of features</i>	
Release	Features
5.2.3.x	• Deployment by using IBM Storage Scale native REST API. • Ansible toolkit code enhancement to work with the latest Ansible library. • Rolling upgrade of CES S3 protocol nodes. • Extended certification and support for operating systems.
5.2.2.x	• Extended operating system certification and support. • Code enhancement to work with the latest Ansible library. • Support to upgrade CES S3. • Certification with Python 3.12. • Support for Ubuntu 24.04 to install and deploy.
5.2.1.x	• Supports installation of S3 protocol packages on CES nodes.
5.2.0.x	• Extended operating system certification and support. • Toolkit code enhancement to work with the latest Ansible library with base operating system repositories. • ECE multi-vdisk support for file system creation. • aarch64 (ARM) architecture
5.1.9.x	• Toolkit support for Remote mount configuration. • Extended operating system certification and support. • ECE SED installation, upgrade, and multi-DA support for vdisk creation.
5.1.8.x	• Extended operating system certification and support. • Problem determination enhancement in the installation and upgrade process.

Table 41. Installation toolkit: List of features (continued)

Release	Features
5.1.7.x	• Support for prebuilt gplbin installation and upgrade support on RHEL and SLES Operating system. • Toolkit support for option to specify a port to use for both tscTcpPort (daemon) and mmsdrservPort (sdr/cc). • ECE install toolkit enhancement to support checking fault tolerance for node upgrade. • ECE install toolkit enhancement to check and show how many DAs, type and usable spaces are available.
5.1.6.x	• Precheck problem determination enhancement in the installation and upgrade process. • Ansible toolkit ECE validation enhancement for block size and with RAID code. • Ansible library support for Red Hat Enterprise Linux® 9.0. • Support for Red Hat Enterprise Linux 8.7, 9.0, and 9.1 operating system on x86_64, PPC64LE, and s390x architectures. For more information, see IBM Storage Scale FAQ in IBM Documentation.
5.1.5.x	• Ansible collection support in the toolkit. • Precheck problem determination enhancement. • Config populate enhancement.
5.1.4.x	• Support for Ubuntu 22.04 on x86_64. • Support for Red Hat Enterprise Linux 8.6 on x86_64, PPC64LE, and s390x. • Modified the command to enable workload prompt to allow administrators to stop and migrate workloads before a node is shut down for upgrade. For more information, see “Upgrading IBM Storage Scale components with the installation toolkit” on page 723.
5.1.3.x	• Support for parallel offline upgrade of all nodes in the cluster. • Support for Ansible 2.10.x.
5.1.2.x	• Support for Red Hat Enterprise Linux 8.5 on x86_64, PPC64LE, and s390x. • Support for user-defined profiles of GPFS configuration parameters. • Support for populating cluster state information from mixed CPU architecture nodes. The information is populated from nodes that have the same CPU architecture as the installer node. • Several optimizations in the upgrade path resulting in faster upgrades than in earlier releases.

Table 41. Installation toolkit: List of features (continued)

Release	Features
5.1.1.x	• Migration to the Ansible® automation platform: – Enables scaling up to more number of nodes – Avoids issues that arise due to using an agent-based tooling infrastructure such as Chef – Enables parity with widely-adopted, modern tooling infrastructure • Support for Red Hat Enterprise Linux 8.4 on x86_64, PPC64LE, and s390x • Support for Ubuntu 20.04 on PPC64LE • Support for multiple recovery groups in IBM Storage Scale Erasure Code Edition • Simplification of file system creation: The file system is created during the installation phase rather than the deployment phase. • Support for IPv6 addresses • Support for the CES interface mode • IBM Storage Scale deployment playbooks are now open sourced on GitHub. Users can access the playbooks from the external GitHub repository and implement in their environment on their own. For more information, see https://github.com/IBM/ibm-spectrum-scale-install-infra. Note: The installation toolkit (./spectrumscale command) is only available as part of the IBM Storage Scale installation packages that can be downloaded from IBM Fix Central. • Added an option to enable workload prompt to allow administrators to stop and migrate workloads before a node is shut down for upgrade. • Discontinued the following functions: – NTP configuration Time synchronization configuration across all nodes in a cluster is recommended. Do it manually by using the available method. – File and object authentication configuration File and object authentication configuration must be done by using the mmuserauth command. – NSD balance Balance the NSD preferred node between the primary and secondary nodes by using the ./spectrumscale nsd servers command.

Table 41. Installation toolkit: List of features (continued)

Release	Features
5.1.0.x	• Support for Ubuntu 20.04 on x86_64 • Support for Red Hat Enterprise Linux 8.3 on x86_64, PPC64LE, and s390x • Support for Red Hat Enterprise Linux 7.9 on x86_64, PPC64LE, and s390x • Support for NFS and SMB protocols, and CES on SLES 15 on x86_64 • Support for installing and upgrading AFM to cloud object storage (gpfs.afm.cos) package
5.0.5.x	• Support for Red Hat Enterprise Linux 8.2 on x86_64, PPC64LE, and s390x • Support for Red Hat Enterprise Linux 7.8 on x86_64, PPC64, PPC64LE, and s390x • Support for packages and repository metadata signed with a GPG (GNU Privacy Guard) key • Enhanced handling of host entries in the /etc/hosts file. Support for both FQDN and short name
5.0.4.x	• Support for Hadoop Distributed File System (HDFS) protocol • Support for ESS 3000 environments. • Support for Red Hat Enterprise Linux 8.0 and 8.1 on x86_64, PPC64LE, and s390x • Support for Red Hat Enterprise Linux 7.7 on x86_64, PPC64, PPC64LE, and s390x • Support for SLES 15 SP1 on x86_64 and s390x • Note: NFS and SMB are not supported on SLES 15 SP1. • Improvements in online and offline upgrade paths • Removed the installation GUI • Support for IBM Storage Scale Developer Edition

Table 41. Installation toolkit: List of features (continued)

Release	Features
5.0.3.x	- Support for SLES 15 on x86_64 and s390x Note: NFS and SMB are not supported on SLES 15. - Upgrade related enhancements: - Upgrade flow changes to minimize I/O disruptions - Enhanced upgrade pre-checks to determine the packages that must be upgraded. Compare the versions of the installed packages with the versions in the repository of the packages you want to upgrade to. In a mixed operating system cluster, the comparison is done with the package repository applicable for the operating system running on the respective nodes. - Mixed OS support for upgrade - Enhanced upgrade post-checks to ensure that all packages have been upgraded successfully - Enhanced dependency checks to ensure dependencies are met for each required package - IBM Storage Scale Erasure Code Edition - Ability to define a new setup type ece - Ability to designate a scale-out node - Ability to define recovery group, vdisk set, and file system - Support for installation and deployment of IBM Storage Scale Erasure Code Edition - Support for config populate function in an IBM Storage Scale Erasure Code Edition environment - Offline upgrade support for IBM Storage Scale Erasure Code Edition
5.0.2.x	- Support for IBM Z (RHEL 7.x, SLES12.x, Ubuntu 16.04 and Ubuntu 18.04 on s390x) - Support for RHEL 7.6 on x86_64, PPC64, PPC64LE, and s390x - Support for offline upgrade of nodes or components while they are stopped or down - Support for excluding nodes from an upgrade run - Support for rerunning an upgrade procedure after a failure - Support for watch folder - Configuration of message queue for file audit logging and watch folder - Enhancements in CES shared root creation and detection in config populate - Upgraded bundled Chef package

Table 41. Installation toolkit: List of features (continued)

Release	Features
5.0.1.x	• Support for Ubuntu 18.04 on x86_64 • Support for RHEL 7.5 on x86_64, PPC64, and PPC64LE • Support for Ubuntu 16.04.4 on x86_64 • Config populate support for call home and file audit logging • Performance monitoring configuration-related changes
5.0.0.x	• Extended operating system support: – Ubuntu 16.04.0, 16.04.1, 16.04.2, 16.04.3 on x86_64 – RHEL 7.4 on x86_64, PPC64, and PPC64LE – SLES 12 SP3 on x86_64 • Improved deployment integration with Elastic Storage Server: The installation toolkit includes the capability to detect ESS nodes (EMS and I/O) and it ensures validation of permitted operations when you are adding protocol or client nodes to a cluster containing ESS. • File audit logging installation and configuration • Call home configuration • Cumulative object upgrade support • Enhanced network connectivity pre-checks including passwordless SSH validation from the admin node • Updated file system default block size for more likely best performance defaults

Understanding the installation toolkit options

Use the following information to understand how to work with the installation toolkit, and the options that are available with it.

Installation

When you use the installation toolkit to install IBM Storage Scale, the procedure comprises two stages:

1. Using a series of **spectrumscale** commands to add node and NSD specifications and configuration properties to the cluster definition file.
2. Using the **spectrumscale install** command to install IBM Storage Scale on the nodes specified in the cluster definition file and to apply the configuration options to the cluster.

Deployment

When you use the installation toolkit to deploy protocols, the procedure comprises two similar stages:

1. Using a series of **spectrumscale** commands to specify environment details, to identify which protocols are to be enabled, and to define protocol-specific properties in the cluster definition file.
2. Using the **spectrumscale deploy** command to deploy protocols as specified in the cluster definition file.

If you already have an existing IBM Storage Scale cluster, with IBM Storage Scale started and at least one file system for the CES shared file system, you can just define protocol nodes in the cluster definition and then deploy protocols.

Upgrade

When you use the installation toolkit to upgrade, the procedure comprises two similar stages:

1. Using a series of **spectrumscale** commands to specify environment details, or using **./spectrumscale config populate** to populate cluster wide properties in the cluster definition file in preparation for upgrade.
2. Using the **spectrumscale upgrade** command to upgrade all nodes and features as specified in the cluster definition file.

See “Using the installation toolkit to perform installation tasks: Explanations and examples” on page 532 for information on how to use the command options listed in Table 42 on page 517 to perform the installation and deployment. After you review the examples, you can tailor them according to your needs and then proceed to implement them.

<i>Table 42. spectrumscale command options for installing IBM Storage Scale and deploying protocols</i>	
spectrumscale command option	Purpose
node add	Adds node specifications to the cluster definition file
node delete	Removes node specifications from the cluster definition file
nsd add	Adds NSD specifications to the cluster definition file
nsd clear	Clears all NSDs
nsd delete	Removes a single NSD
nsd list	Lists all NSDs currently in the configuration
nsd modify	Modifies an NSD
filesystem list	Lists all file systems that currently have NSDs assigned to them
filesystem modify	Changes the block size and mount point of a file system
config gpfs	Adds IBM Storage Scale specific properties to the cluster definition file
install	Installs IBM Storage Scale on the configured nodes, creates a cluster, and creates NSDs.
scaleadmd enable	
nodeid define	
config protocols	Provides details about the IBM Storage Scale environment to be used during protocol deployment
config perfmon	Configures performance monitoring settings
config populate	Populates cluster definition file with current cluster configuration
deploy	Creates file systems and deploys protocols on your configured nodes
callhome	Configures call home settings
fileauditlogging	Configures file audit logging settings
upgrade	Upgrades the various components of the installation.

Table 42. *spectrumscale* command options for installing IBM Storage Scale and deploying protocols (continued)

spectrumscale command option	Purpose
nfs_core_dump enable	Enables NFS core dump collection on NFS protocol nodes to assist in troubleshooting NFS-related issues.

The list of **spectrumscale** command options listed is not exhaustive. For all command options available with the **spectrumscale** command and for more information about these command options, see the **spectrumscale** command description in the *IBM Storage Scale: Command and Programming Reference Guide*.

Limitations of the installation toolkit

Before using the installation toolkit to install IBM Storage Scale and deploy protocols, review the following limitations and workarounds, if any.

Table 43. <i>Installation toolkit limitations</i>		
Function	Description	Workaround, if any
CES groups	The installation toolkit does not support the configuration of CES groups. This causes protocol deployments to fail in multi-network environments. Having CES groups might cause issues with upgrades as well because the installation toolkit is not aware of CES groups during the upgrade.	
Clusters larger than 64 nodes	The installation toolkit does not restrict the number of nodes in a cluster in which the toolkit can be used. However, it is designed as a single node server so as the number of nodes increases, the bandwidth to the installer node decreases, and the latency goes up. This implementation might cause issues in clusters with more than 64 nodes. Note: If you want to use the installation toolkit in a cluster larger than 64 nodes, contact scale@us.ibm.com .	
Clusters without passwordless SSH between all nodes	If clusters are set up in an AdminCentral=True configuration, which is a widely used configuration, the installation toolkit and protocols might not function correctly.	Set up passwordless SSH between all nodes in the cluster and to the nodes themselves using FQDN, IP address, and host name.
Compression	The installation toolkit does not configure file system compression.	After installation, configure the compression function manually. For more information, see <i>File compression</i> in <i>IBM Storage Scale: Administration Guide</i> .

Table 43. Installation toolkit limitations (continued)

Function	Description	Workaround, if any
Concurrent upgrade	The installation toolkit does not support concurrent upgrade. You must plan for an outage depending on your setup. This brief outage prevents mixed code versions from running at the same time and it is also needed in a manual upgrade. Although the upgrade is non-concurrent, data is typically still accessible during the upgrade window. Access might be lost and need to be reestablished multiple times due to how the upgrade procedure is run among nodes.	
Configuration change	The installation toolkit does not support changing the configuration of entities that exist such as file systems and NSDs. For example, changing the block size or the default replication factor for data blocks for an existing file system is not supported. However, the installation toolkit can be used to add nodes, NSDs, or file systems to an existing cluster. The installation does not support authentication reconfiguration. It does not use the authentication section of the cluster definition file during upgrade.	• If you want to change the configuration of existing file systems or NSDs, use the IBM Storage Scale commands after installation or deployment. • If you want to change the authentication method, see <i>Modifying authentication method</i> in <i>IBM Storage Scale: Administration Guide</i>.
Customer designation of sensor or collector nodes	The installation toolkit does not support customer designation of sensor or collector nodes. The installation toolkit automatically sets up sensor or collectors without allowing the user to choose which nodes have these functions.	1. Use the -x flag to reconfigure performance monitoring, if needed. 2. Follow up with manual configuration of performance monitoring.
Disabling or uninstalling protocols and uninstalling GPFS	The installation toolkit does not support disabling or uninstalling protocols and uninstalling GPFS on an existing GPFS cluster.	Use the manual procedures. • For information about removing exports, see <i>Managing protocol data exports</i> in <i>IBM Storage Scale: Administration Guide</i>. • For information about disabling protocol services, see <i>Managing protocol services</i> in <i>IBM Storage Scale: Administration Guide</i>. • For information about uninstalling GPFS, see Chapter 19, “Uninstalling IBM Storage Scale,” on page 659.

<i>Table 43. Installation toolkit limitations (continued)</i>		
Function	Description	Workaround, if any
ESS and IBM Storage Scale Erasure Code Edition mixed environment	The installation toolkit does not support an environment in which ESS and IBM Storage Scale Erasure Code Edition coexist.	Use the manual procedure to install, deploy, or upgrade IBM Storage Scale Erasure Code Edition in a cluster that contains ESS.
Encryption	The installation toolkit does not support encrypted file systems. Therefore, installation and deployment by using the installation toolkit do not work if the CES shared root file system or any other file system that the installation toolkit works with is encrypted.	
EPEL repositories	The installation toolkit does not support the configuration of protocols when EPEL repositories are configured.	Remove or disable EPEL repositories before you use the installation toolkit for installation, deployment, or upgrade.
File system DMAPI flag set to Yes (-z) installation and deployment	The file system DMAPI flag is used in IBM Storage Scale for IBM Storage Protect for Space Management and policy management, attaching with an IBM Storage Protect server, and with IBM Spectrum Archive. The installation toolkit does not provide an option to add a DMAPI flag to a new or existing file system. If you need add a DMAPI flag to a file system, use the installation toolkit to create the file system and later set the DMAPI flag manually.	For information on manual procedures for installation and deployment, see “Manually installing the IBM Storage Scale software packages on Linux nodes” on page 477.

Table 43. Installation toolkit limitations (continued)

Function	Description	Workaround, if any
File system DMAPI flag set to Yes (-z) upgrade	An upgrade by using the installation toolkit is affected by the presence of the DMAPI flag in a few ways: 1. If the CES shared root shares a file system that has the DMAPI flag set: • In this situation, the installation toolkit cannot unmount or mount the file system on each node it attempts to upgrade unless the user has preemptively removed the DMAPI flag and stopped all DMAPI services. For more information, see “Upgrading IBM Storage Scale on IBM Storage Archive Enterprise Edition (EE) nodes by using the installation toolkit” on page 742. • This scenario is different from scenario 2 because the CES shared root file system is essential for maintaining healthy CES status of the cluster. Without the CES shared root mounted, the CES component remains unhealthy and prevents upgrade from continuing. 2. If a non-CES shared root file system has the DMAPI flag set: • In this scenario, the installation toolkit cannot unmount the file system during upgrade of each node unless the user stopped all DMAPI-related services and unmounted the file system. • Thereafter, the installation toolkit cannot remount the file system after upgrade of each node. The file system is left unmounted and the user needs to bring it up manually on each node individually after the GPFS portion of the upgrade finished. Otherwise, the user can wait until the upgrade finishes and start the file system on all nodes at once.	• For information on using the installation toolkit for upgrade in clusters where the DMAPI flag is set to Yes for any of the file systems, see “Upgrading IBM Storage Scale on IBM Storage Archive Enterprise Edition (EE) nodes by using the installation toolkit” on page 742. • In the case of a failed upgrade that needs to be restarted, if some of the nodes have been restarted, then steps need to be taken to get all of the file systems mounted again and the cluster healthy before another attempt can be made. This might require stopping DMAPI services again if they auto-started after a node was upgraded.

Table 43. Installation toolkit limitations (continued)

Function	Description	Workaround, if any
FPO configuration for disks	The installation toolkit does not support the extra stanza file flags required for FPO setup.	Do one of the following: 1. Create NSDs using the installation toolkit. 2. Manually edit the NSD stanza file afterward. The installation toolkit places the NSD stanza file in the <code>/usr/lpp/mmfs</code> directory. 3. Use <code>mmchnsd</code> to do the changes. OR 1. Create the cluster by using the installation toolkit. 2. Deploy protocols on the protocol nodes by using the installation toolkit. 3. Manually create the NSDs for the FPO setup.
Federal Information Processing Standard (FIPS) enabled	The installation toolkit does not support environments in which FIPS is enabled.	
GPG (GNU Privacy Guard) signed package support	The installation toolkit support for packages that are signed with the GPG key has the following limitation: • Repository metadata signing is not supported on Ubuntu.	
Host-based SSH authentication	The installation toolkit does not support host-based SSH authentication. It supports only key-based SSH authentication.	Either set up key-based SSH authentication temporarily for use with the toolkit, or follow the manual steps in “Manually installing the IBM Storage Scale software packages on Linux nodes” on page 477.
Kafka packages	To make the IBM Storage Scale cluster ready for file audit logging and watch folder functions, the installation toolkit automatically installs the <code>gpfs.librdkafka</code> package for supported operating systems and hardware architectures. This might lead to errors if the prerequisites for <code>gpfs.librdkafka</code> are not installed.	Ensure that the prerequisite packages for <code>gpfs.librdkafka</code> are installed. For more information, see “Requirements, limitations, and support for file audit logging” on page 647.

Table 43. Installation toolkit limitations (continued)

Function	Description	Workaround, if any
Local host	- The installation toolkit does not support specifying the IP address that the local host resolves to, such as 127.0.0.1, as the installer node. - The installation toolkit does not support adding the local host or the IP address that the local host resolves to, such as 127.0.0.1, as a node in the cluster configuration.	
Multiple CES networks	The installation toolkit does not support deployments with multiple CES networks. Attempting deployment in this configuration has a high probability of failure. This is because when the CES address pool has multiple subnets, the command for adding CES address assigns an IP to a CES node that cannot handle that subnet of address, which causes the deployment failure.	
Multiple clusters	The installation toolkit does not support multiple clusters being defined in the cluster definition.	
NFS or SMB exports configuration	The installation toolkit does not configure any exports on SMB or NFS.	Use the manual procedure. For information about configuring Cluster Export Services and creating exports, see <i>Configuring Cluster Export Services</i> and <i>Managing protocol data exports</i> in <i>IBM Storage Scale: Administration Guide</i> .
Node function addition during upgrade	The installation toolkit does not support designating node functionality during upgrade.	To add a function to the cluster or a node, designate this new function using the installation toolkit and proceed with an installation or a deployment. Perform this action either before or after an upgrade.
Non-English languages in client programs such as PuTTY	The installation toolkit does not support setting the language in client programs such as PuTTY to any language other than English.	Set the language of your client program to English.

Table 43. Installation toolkit limitations (continued)

Function	Description	Workaround, if any
NSD SAN attachment during initial installation	The installation toolkit cannot be used for NSD SAN attachment during initial installation because when adding NSDs using the installation toolkit, a primary and an optional comma-separated list of secondary NSD servers must be designated.	1. Create NSDs using the installation toolkit. 2. Manually edit the NSD stanza file afterward. The NSD stanza files created by the installation toolkit are named in the <code>StanzaFile.xx</code> format and they are located in the <code>/usr/lpp/mmfs/</code> directory. 3. Use mmchnsd to do the changes.
Online upgrade of a 2-node cluster without tie-breaker disks	Online upgrade of a 2-node cluster that does not have tie-breaker disks configured is not supported. To do an online upgrade of a 2-node cluster by using the installation toolkit, tie-breaker disks must be configured in the cluster. A 1-node cluster can be upgraded only offline.	Do an offline upgrade. For more information, see “Performing offline upgrade or excluding nodes from upgrade by using installation toolkit” on page 733.
Package managers other than yum, zypper, or apt-get	The installation toolkit requires the use of yum (RHEL), zypper (SLES), and apt-get (Ubuntu) package managers to function.	
PPC and x86_64 or PPC and s390x or x86_64 and s390x mix	The installation toolkit does not support mixed CPU architecture configurations.	Use the installation toolkit on a subset of nodes that are supported and then manually install or deploy on the remaining nodes. For upgrading a mixed CPU architecture cluster, you can use the installation toolkit in two hops. For more information, see “Upgrading mixed CPU architecture cluster” on page 739.
Quorum or manager configuration after cluster installation	The installation toolkit allows a user to add <code>-m</code> (to specify a manager node) and <code>-q</code> (to specify a quorum node) flags to various nodes as they are added. If the proposed configuration does not match the existing configuration, the installation toolkit does nothing to change it.	Manually change node roles by using the mmchnode command.

Table 43. Installation toolkit limitations (continued)		
Function	Description	Workaround, if any
Remote mounted file systems	- NFS and SMB deployments on remote mounted file systems work using the installation toolkit. - The setup of a CES shared root file system when the file system is remotely mounted works using the installation toolkit.	
Repository proxy	The installation toolkit does not support proxy setups when working with repositories. <code>yum repolist</code> must not have any failed repos and it must be clean.	Ensure that there are no stale or failed repositories and that only the base OS repositories are enabled during any installation toolkit activities such as installation, deployment, or upgrade.
RPMs that have dependencies upon GPFS RPMs and GPFS settings	In an environment where some RPMs have dependencies on base GPFS RPMs or GPFS settings, the installation toolkit cannot be used for installation or upgrade.	
Running mmchconfig release=LATEST to complete an upgrade	The installation toolkit does not run mmchconfig release=LATEST after an upgrade. This is to give users time to verify an upgrade success and decide if the code level upgrade should be finalized.	Use mmchconfig release=LATEST after an upgrade using the installation toolkit to finalize the upgrade across the cluster.
Separate admin and daemon network	The installation toolkit does not support separate admin and daemon network.	
Sudo user	The installation toolkit does not function correctly unless run as root. Running as sudo or as another user does not work.	
Support for AIX, Debian, PowerKVM, Windows	The installation toolkit does not support AIX, Debian, PowerKVM, Windows operating systems. If these operating systems are installed on any cluster nodes, do not add these nodes to the installation toolkit.	Use the installation toolkit on a subset of nodes that are supported and then manually perform installation, deployment, or upgrade on the remaining nodes. For information about manual upgrade, see Chapter 20, “Upgrading,” on page 671.
Tie-Breaker NSD configuration	The installation toolkit does not configure tie-breaker disks.	Manually set the tie-breaker configuration as required using mmchconfig after completing installation using the toolkit.
transparent cloud tiering	The installation toolkit does not install, configure, or upgrade transparent cloud tiering.	Use the manual procedures. For more information, see Chapter 8, “Installing transparent cloud tiering (discontinued),” on page 627.

Table 43. Installation toolkit limitations (continued)		
Function	Description	Workaround, if any
Unique NSD device configuration	The installation toolkit relies upon a user having already configured and run the nsddevices sample script provided within a GPFS installation. The mmcrnsd and mmchnsd commands require running of the nsddevices script beforehand. Therefore, the installation toolkit will fail if this is not done by the user.	
Upgrade while skipping over versions	The installation toolkit does not support skipping over major or minor versions of IBM Storage Scale releases when doing an online upgrade. For example, if an IBM Storage Scale cluster is at version 4.1.1.x, you cannot use the installation toolkit for an online upgrade directly to version 6.0.x.	Do an offline upgrade from version 4.1.1.x to 6.0.x. For more information, see “Performing offline upgrade or excluding nodes from upgrade by using installation toolkit” on page 733.
Upgrade on Ubuntu 22.04	The installation toolkit cannot perform upgrades for IBM Storage Scale 5.2.2.1 when the operating system is Ubuntu 22.04.	
Using the latest Python version on a RHEL operating system	If you are using a RHEL operating system, downloading, installing, or setting Python >=3.12 as the default path causes issues. The latest version of Python may cause failures in the installation toolkit operations because RHEL operating systems do not usually work with the latest version of Python. For example, when the latest Python version was 3.12, RHEL 8.10 used Python 3.6.x and RHEL 9.4 used Python 3.9.x.	For the proper functioning of the installation toolkit, use the recommended version of Python, as described in “Software requirements” on page 310 and “Preparing to use the installation toolkit” on page 528.
Support for the aarch64 (ARM) architecture	The installation toolkit does not support the next features with aarch64: • GUI • Protocols (NFS, SMB, HDFS, CES S3)	

Related concepts

[“Limitations of config populate option of the installation toolkit ” on page 552](#)

Mixed operating system support with the installation toolkit

You can use the installation toolkit to install GPFS and deploy protocols in a cluster that contains nodes that are running on different operating systems. Several validations are done when you use the **spectrumscale** command in a mixed operating system cluster.

The following operating systems are supported with the installation toolkit in a mixed operating system cluster.

Table 44. Operating systems supported with the installation toolkit in a mixed cluster	
CPU architecture	Operating system
x86_64	• Red Hat Enterprise Linux 7.x • Red Hat Enterprise Linux 8.x • SLES 15 • Ubuntu 20.04 and 22.04 ¹
s390x	• Red Hat Enterprise Linux 7.x • Red Hat Enterprise Linux 8.x • SLES 15

For latest information about supported operating systems, see [IBM Storage Scale FAQ](#) in [IBM Documentation](#).

Important:

- For using the installation toolkit in a mixed operating system cluster, all protocol nodes must be running on the same operating system. However, nodes that are running on different minor versions of Red Hat Enterprise Linux 7.x or 8.x can be designated as protocol nodes in the same cluster.
- For using the installation toolkit in a mixed operating system cluster, all nodes must have the same CPU architecture.

The installation toolkit performs the required validations in a mixed operating system cluster to prevent users from attempting any configurations that are not supported. These validations include the following.

Table 45. Validations by the installation toolkit in a mixed operating system cluster	
Operation that is attempted with the installation toolkit	Validations done
Node addition by using the spectrumscale node add <i>NodeName</i> command	Check whether the operating system of the target node is supported. If it is not supported, an error occurs.
Protocol node designation by using the spectrumscale node add <i>NodeName</i> -p command	• Check whether all protocol nodes being added are running on the same operating system. If you attempt to add protocol nodes running on different operating systems, then that operation is not allowed. For example, if you have added a protocol node running on Red Hat Enterprise Linux 7.x and then you try to add another protocol node running on SLES 15, then that operation is not allowed.
Protocol enablement by using the spectrumscale enable <i>Protocol</i> command	Check whether the protocol being enabled is supported on the operating system that is running on the node. If the protocol is not supported, an error occurs and you cannot continue with the installation and deployment. For example, if you try to enable object on a node running on SLES 15, an error occurs.

Mixed operating system support during upgrade

During the upgrade precheck, the installation toolkit performs a comparison of the versions of the installed packages with the versions of packages in the repository. In a mixed operating system cluster, the comparison is done with the package repository applicable for the operating system running on the respective nodes.

Note: ¹ On Ubuntu nodes in a mixed operating system cluster, upgrades by using the installation toolkit are not supported. As a workaround, you can use the installation toolkit to upgrade a mixed operating system cluster with Ubuntu nodes in two hops. For more information, see [“Upgrading mixed operating system cluster with Ubuntu nodes”](#) on page 739.

Preparing to use the installation toolkit

Before you use the installation toolkit, complete the following preparatory steps on all the nodes on which you plan to use the installation toolkit.

During the precheck phase of the installation toolkit, check the following items:

- Passwordless SSH is set up.
- Prerequisite kernel packages are installed.
- Needed ports are open.
- Supported OS and architectures are discovered on nodes.
- Base OS repositories are set up.

Procedure

1. Ensure that the following requirements are met.

Operating systems supported by the installation toolkit

For information about supported operating systems, see [IBM Storage Scale FAQ in IBM Documentation](#). Also, check with the operating system vendor to verify that a specific version of the operating system is supported by the vendor.

For information about how the installation toolkit can be used in a cluster that has nodes with mixed operating systems, see [“Mixed operating system support with the installation toolkit”](#) on page 526.

Required packages for the installation toolkit

The installation toolkit requires the following packages:

- Python 3.10 or higher

Note:

- It is recommended that you install Python by using the OS package manager to avoid potential installation issues. For example, **yum install python3**.
- If you have Python 2.x and Python 3.x in your environment, ensure that Python 3.x is the default Python version. You can use the **python --version** command to check the default Python version. For example:

```
# python3 --version Python 3.10
```

- net-tools
- Ansible 2.9.27

Supported Ansible version

The installation toolkit requires, at minimum, Ansible version 2.9.27. The installation toolkit also works in Ansible 2.10.x environments.

Important: The installation toolkit is not compatible with `ansible-core` version 2.24.x and later. Make sure that the `ansible-core` version is restricted to releases that are earlier than 2.24. The version 2.24 of `ansible-core` introduced changes in the Ansible core modules and dependencies that may cause compatibility issues with the installation toolkit.

The following Ansible installation considerations apply, depending on your operating system.

Operating system	Ansible installation considerations
Red Hat Enterprise Linux 8.x, and 9.x	The installation toolkit installs the supported version of Ansible on the installer node when you run the <code>./spectrumscale setup -s InstallNodeIP</code> command. Notes for the Ansible toolkit: • For Red Hat Enterprise Linux, the Ansible toolkit automatically installs the latest version of Ansible (<code>ansible-core</code>) that is available in the Base OS repositories and their dependencies package. • The Ansible toolkit uses the <code>ansible.posix</code> and <code>community.general</code> modules that are included in the Ansible collection packages.
SLES 15	You can use <code>zypper install ansible</code> to install Ansible, but ensure that the Ansible version is 2.9.1 or later.
Ubuntu 22.04, and 24.04	On Ubuntu nodes, Ansible 2.9.6 might be installed by default. You can install Ansible 2.9.x from the apt repository by using the <code>apt install ansible-core</code> command or the <code>pip3 install ansible-core</code> command. You can use the default Ansible version.

- You can manually install `ansible-core` by using the following command.

```
# pip3 install ansible-core
```

- You can manually install a specific version of `ansible-core`, for example version 2.14, by using the following command:

```
# pip3 install ansible-core==2.14
```

For more information, see [Installing Ansible](#).

Note: If the `pip3` command does not work with the installed Python version, use `yum`, `zypper`, or `apt-get` commands.

Optional: Ansible configuration settings to expedite Ansible tasks implementation

When the cluster to be deployed has more than 10 nodes, Ansible tasks implementation can be expedited by customizing the value for the `forks` parameter in the Ansible configuration file. The file is named `ansible.cfg` and can be saved in your home directory (`~/.ansible.cfg`).

The default value for `forks` is 5. You can change this value based on the number of nodes that the cluster is composed of. In the next example of an Ansible configuration file, the number of forks is set to 50.

```
[defaults]
callback_whitelist = timer, profile_tasks
forks = 50
```

```
gathering = smart
gather_subset = hardware
[ssh_connection]
pipelining = True
log_path = /var/log/ansible.log
```

For more information about tuning the parameters in an Ansible configuration file, see [Ansible configuration settings](#) in Ansible documentation.

Root access for the installation toolkit

The installation toolkit must be run as the root user. On Ubuntu, root login might be disabled by default. To use the installation toolkit on Ubuntu, you must enable the root login.

Call home information required for the installation toolkit

The installation toolkit requires call home data to work properly. You must have the *customer name*, the *customer id*, the *customer email*, and the *customer country code* available before installing and configuring call home with the installation toolkit.

Disable auto updates on Ubuntu nodes

On Ubuntu nodes, auto updates or unattended upgrades must be disabled, and kernel auto updates upon node reboot must also be disabled. This must be done to ensure that any updates that are not supported do not get installed. For example, a minor Ubuntu 22.04.x update that is not supported by IBM Storage Scale and installation toolkit might get installed if auto updates are not disabled.

Uninstall RPM Package Manager (RPM) on Ubuntu nodes

To use the installation toolkit on Ubuntu nodes, you must uninstall the RPM Package Manager (RPM) from these nodes.

Cluster configuration repository (CCR) is enabled

Before using the installation toolkit, ensure that CCR is enabled. For new installations, CCR is enabled by default. You can check that CCR is enabled in the cluster by using the following command.

```
# mm1scluster

GPFS cluster information
=====
GPFS cluster name:      node1.example.com
GPFS cluster id:       123412789099501234
GPFS UID domain:       node1.example.com
Remote shell command:  sudo wrapper in use
Remote file copy command: sudo wrapper in use
Repository type:       CCR
```

Uninstall Upstream nfs-ganesha package

Before installing IBM Storage Scale 5.1.x on Ubuntu operating system, make sure that the upstream NFS Ganesha package is not installed or configured in the protocol nodes. If the NFS Ganesha package is installed or configured, then uninstall the package by using the **apt purge nfs-ganesha** command.

For information on prerequisites for protocols and performance monitoring, see [“Installation prerequisites”](#) on page 470.

2. Set up passwordless SSH as follows.

- From the admin node to the other nodes in the cluster.
- From protocol nodes to other nodes in the cluster.
- From every protocol node to rest of the protocol nodes in the cluster.

Note: Passwordless SSH must be set up by using the FQDN and the short name of the node.

The installation toolkit performs verification during the precheck phase to ensure that passwordless SSH is set up correctly. This verification includes:

- Check whether passwordless SSH is set up between all admin nodes and all the other nodes in the cluster. If this check fails, a fatal error occurs.
 - Check whether passwordless SSH is set up between all protocol nodes and all the other nodes in the cluster. If this check fails, a warning is displayed.
 - Check whether passwordless SSH is set up between all protocol nodes in the cluster. If this check fails, a fatal error occurs.
- Set up the base OS repositories on your nodes so that package dependencies can be satisfied.
 - These package managers must be set up depending on the operating system.
 - Red Hat Enterprise Linux: yum repository must be set up on all nodes in the cluster.

Note: On Red Hat Enterprise Linux 8.x, two package repositories are available: BaseOS and Application Stream (AppStream). To streamline the installation and the upgrade of IBM Storage Scale packages on Red Hat Enterprise Linux 8.x nodes, it is recommended to configure the BaseOS and the Application Stream repositories.
 - SLES: zypper repository must be set up on all nodes in the cluster.
 - Ubuntu: apt repository must be set up on all nodes in the cluster.

Note: Ensure that the **apt-get update** command is working without any errors.
 - Configure repositories depending on whether you have internet connection or not. However, ensure that base operating system packages are available and EPEL repositories are disabled.
 - Ensure that the ports that are needed for installation are open.

Note: The installation toolkit checks if the firewall daemon is running and displays a warning if it is running. If the required ports are open, you can ignore the warning.

For information about the ports that need to be open, see *Securing the IBM Storage Scale system using firewall* in *IBM Storage Scale: Administration Guide*.
 - Ensure that the required kernel packages are installed.

For more information, see [“Software requirements” on page 310](#).

Note: You might have to force the installation of a specific version of these packages because a package of a version newer than the corresponding kernel might get picked up by default.
 - Ensure that networking is set up in one of the following ways.
 - DNS is configured such that all host names, either short or long, are resolvable.
 - All host names are resolvable in the `/etc/hosts` file. The host entries in the `/etc/hosts` file are required to be in the following order:

```
<IP address> <FQDN> <Short name>
```

If the entries in the `/etc/hosts` file are not in this order, the installation toolkit displays a warning and continues with the rest of the procedure. The installation toolkit supports both FQDN and short name during installation, upgrade, or deployment. During fresh installations, the nodes are configured based on the entries in `/etc/hosts`. If `/etc/hosts` contains short names then the nodes are added with the short names, if the host names are reachable. For existing clusters, the installation toolkit configures the new node according to the configuration of the existing cluster. For example, if the existing cluster is configured with short names then the new node is added with the short name.

- Obtain the IBM Storage Scale self-extracting installation package from IBM Fix Central.
- Ensure that the `LC_ALL` and the `LANGUAGE` parameters are set to `en_US.UTF-8`.

You can use the **locale** command to check the current settings. If these parameters are set to a value different than `en_US.UTF-8`, use the following export statements before you issue the **./spectrumscale** commands.

```
export LC_ALL=en_US.UTF-8
export LANGUAGE=en_US.UTF-8
```

You can also add these export statements in the `bash_profile`.

9. Perform pre-built/Custom Portability Layer rpm support through Installation toolkit.

Note: Prebuilt gplbin package installation is currently supported on RHEL and SLES on x86_64 architecture only.

- a) Create a repository directory inside `/usr/lpp/mmfs/6.0.1.0/`.

```
mkdir directory_name
```

- b) Copy the prebuilt generated gplbin package to the above repository directory. For example, `scp gpfs.gplbin.rpm /usr/lpp/mmfs/6.0.1.0/directory_name`.

- c) Install createrepo utility if it does not already exist.

```
yum install createrepo
```

- d) Create the repository metadata using createrepo command.

```
createrepo .
```

- e) Pass the defined rpm directory name to the install toolkit configuration file.

```
./spectrumscale config gpfs --gplbin_dir directory_name
```

Note: Toolkit always validates the directory name and the repository data file exists inside the dedicated path `/usr/lpp/mmfs/6.0.1.0/`. It gives fatal error if the file and repository data does not exist.

- f) Run the install/deploy process by executing the following commands:

- i) Define the rpm directory name in the configuration file.

```
./spectrumscale config gpfs --gplbin_dir directory_name
```

- ii) List the defined directory from the configuration file.

```
./spectrumscale config gpfs --list
```

- iii) Clear the defined directory from the configuration file.

```
./spectrumscale config clear gpfs --gplbin_dir
```

Using the installation toolkit to perform installation tasks: Explanations and examples

Use these explanations and examples of installation toolkit options and tasks to perform installation and deployment by using the installation toolkit.

Note: CES supports HDFS protocols. For installing CES HDFS, the same prerequisites are applicable as for the other protocols as documented in “Preparing to use the installation toolkit” on page 528 and this section. For specific information about installing CES HDFS, see *Installation* under [CES HDFS](#) in *Big data and analytics support documentation*.

After determining how you want your system to be configured, and after you have reviewed the `spectrumscale` command description in the *Command reference* section *IBM Storage Scale: Command and Programming Reference Guide*, you can tailor these examples to do the following tasks:

- Set up your installer node.
- Add node and NSD specifications, file system information, and configuration properties to your cluster definition file, and then install IBM Storage Scale and configure your cluster according to the information in that file.
- Specify environment details, identify which protocols are to be enabled, and define protocol-specific properties in the cluster definition file, and then deploy protocols on your system.

For information on using the installation toolkit for the installation and the deployment of protocol and client nodes in a cluster containing ESS, see [“Upgrading protocol nodes in an IBM Storage Scale System cluster”](#) on page 555.

Before using the installation toolkit for installation and deployment, ensure that the preparatory steps are completed. For more information, see [“Preparing to use the installation toolkit”](#) on page 528.

Using the installation toolkit to perform installation and deployment involves doing these tasks.

1. [“Setting up the installer node”](#) on page 533
2. [“Defining the cluster topology for the installation toolkit”](#) on page 534
3. [“Setting configuration parameters before installation”](#) on page 538
4. [“Installing IBM Storage Scale and creating a cluster”](#) on page 539
5. [“Deploying protocols”](#) on page 541

For more installation toolkit related information, see the following topics.

- [“Installation of performance monitoring tool using the installation toolkit”](#) on page 544
- [“Enabling and configuring file audit logging using the installation toolkit”](#) on page 546
- [“Enabling and configuring call home using the installation toolkit”](#) on page 548
- [“Installing IBM Storage Scale management GUI by using the installation toolkit”](#) on page 550
- [“Logging and debugging for installation toolkit”](#) on page 545

Setting up the installer node

The first step in using the installation toolkit for installing IBM Storage Scale and deploying protocols is to configure your installer node.

A good candidate for your installer node is the GPFS admin node because a prerequisite of this node is that it must be able to communicate with all other nodes without the need to provide a password. Also, SSH connections between the admin node and all other nodes must be set up to suppress prompting. Therefore, no prompts for password must display when you are using SSH among any cluster nodes to and from each other, and to and from the server.

Note: Before you configure a SLES or an Ubuntu node as the installer node, ensure that the supported version of Ansible is installed on that node. On Red Hat Enterprise Linux nodes, the installation toolkit installs the supported version of Ansible. For more information, see [“Supported Ansible version”](#) on page 528.

Tip: Ensure that there is at least 5 GB of free space on the node to download, extract, and install the packages.

1. Change the directory to where the installation toolkit is extracted. The default extraction path for 6.0.1.x is as follows. This path varies depending on the version.

```
cd /usr/lpp/mmfs/6.0.1.x/ansible-toolkit
```

2. Configure the installer node by issuing the following command:

```
./spectrumscale setup -s InstallNodeIP -i SSHIdentity
```

The `-s` argument identifies the IP address that nodes use to retrieve their configuration. This IP address is associated with a device on the installer node. This validation is automatically done during the setup phase.

Note: If there are multiple subnets in your cluster, it is recommended to specify a private IP address with the `./spectrumscale setup -s` command.

Optionally, you can specify a private SSH key to be used to communicate with nodes in the cluster definition file, by using the `-i` argument.

In an Elastic Storage Server (ESS) cluster, if you want to use the installation toolkit to install IBM Storage Scale and deploy protocols, you must specify the setup type as `ess` while setting up the installer node as follows.

```
./spectrumscale setup -s InstallNodeIP -i SSHIdentity -st ess
```

For more information, see [“Upgrading protocol nodes in an IBM Storage Scale System cluster”](#) on page 555.

If you want to use the installation toolkit to install IBM Storage Scale Erasure Code Edition, you must specify the setup type as `ece` while setting up the installer node as follows.

```
./spectrumscale setup -s InstallNodeIP -i SSHIdentity -st ece
```

For more information, see IBM Storage Scale Erasure Code Edition documentation.

Defining the cluster topology for the installation toolkit

Use these instructions to set up the cluster definition file before installing IBM Storage Scale and deploying protocols.

1. [“Adding node definitions to the cluster definition file”](#) on page 534
2. [“Adding and configuring NSD server nodes in the cluster definition file”](#) on page 536
3. [“Defining file systems”](#) on page 537

Adding node definitions to the cluster definition file

You can add node definitions to the cluster definition file by using the **./spectrumscale node add** command.

Note: For information on adding protocol nodes to the cluster definition file, see [“Deploying protocols”](#) on page 541.

1. If the installation toolkit is being used from a location outside any of the nodes to be installed, an admin node is required. The admin node is used to run cluster-wide commands.

To specify an admin node in the cluster definition file, use the `-a` argument.

```
./spectrumscale node add gpfsnode1 -a
```

If no admin node is specified in the cluster definition file, the node on which the installation toolkit is running is automatically designated as the admin node. If GUI nodes are to be installed, each GUI node must also be marked as an admin node.

The role of an admin node is to serve as the coordinator of the installation, deployment, and upgrade when using the installation toolkit. This node also acts as a central repository for all IBM Storage Scale packages. For larger clusters, it is important to have an admin node with plenty of network bandwidth to all other nodes in the cluster.

2. To add client nodes to the cluster definition file, provide no arguments.

```
./spectrumscale node add gpfsnode1
```

3. To add manager nodes to the cluster definition file, use the `-m` argument.

```
./spectrumscale node add gpfsnode2 -m
```

If no manager nodes are added to the cluster definition, the installation toolkit automatically designates manager nodes using the following algorithm:

- a. First, all protocol nodes in the cluster definition are designated as manager nodes.

- b. If there are no protocol nodes, all NSD nodes in the cluster definition are designated as manager nodes.
 - c. If there are no NSD nodes, all nodes in the cluster definition are designated as manager nodes.
4. To add quorum nodes to the cluster definition, use the `-q` argument.

```
./spectrumscale node add gpfsnode3 -q
```

If no quorum nodes are added to the cluster definition, the installation toolkit automatically designates quorum nodes using the following algorithm:

- a. If the number of nodes in the cluster definition is less than 4, all nodes are designated as quorum nodes.
- b. If the number of nodes in the cluster definition is between 4 and 9 inclusive, 3 nodes are designated as quorum nodes.
- c. If the number of nodes in the cluster definition is between 10 and 18 inclusive, 5 nodes are designated as quorum nodes.
- d. If the number of nodes in the cluster definition is greater than 18, 7 nodes are designated as quorum nodes.

This algorithm preferentially selects NSD nodes as quorum nodes. If the number of NSD nodes is less than the number of quorum nodes to be designated then any other nodes are selected until the number of quorum nodes is satisfied.

5. To add NSD servers to the cluster definition, use the `-n` argument.

```
./spectrumscale node add gpfsnode4 -n
```

6. To add Graphical User Interface servers to the cluster definition, use the `-g` argument. A GUI server must also be an admin node.

```
./spectrumscale node add gpfsnode3 -g -a
```

If no nodes have been specified as management GUI servers, then the GUI is not installed. It is recommended to have at least 2 management GUI interface servers and a maximum of 3 for redundancy.

7. To add a call home node to the cluster definition, use the `-c` argument.

```
./spectrumscale node add CallHomeNode -c
```

If a call home node is not specified, the installation toolkit assigns one of the nodes in the cluster as the call home node.

8. To display a list of all nodes in the cluster definition file, use the **`./spectrumscale node list`** command. For example:

```
./spectrumscale node list
```

```
[ INFO ] List of nodes in current configuration:
[ INFO ] [Installer Node]
[ INFO ] 192.0.2.1
[ INFO ]
[ INFO ] [Cluster Name]
[ INFO ] gpfscluster01
[ INFO ]
[ INFO ] GPFS Node      Admin  Quorum  Manager  NSD Server  Protocol  GUI Server  OS
Arch
[ INFO ] gpfsnode1      X      X
x86_64
[ INFO ] gpfsnode2      X      X      X      X      X      X      X      X
x86_64
[ INFO ] gpfsnode3      X      X      X      X      X      X      X      X
x86_64
[ INFO ] gpfsnode4      X      X      X      X      X      X      X      X
x86_64
```

- If you want to use the installation toolkit to install IBM Storage Scale and deploy protocols in an Elastic Storage Server (ESS) cluster, you must add the EMS node of the ESS system as follows.

```
./spectrumscale node add -e EMSNode
```

For more information, see [“Upgrading protocol nodes in an IBM Storage Scale System cluster”](#) on page 555.

- For information on adding nodes to an existing installation, see [“Adding nodes, NSDs, or file systems to an existing cluster”](#) on page 568.

Adding and configuring NSD server nodes in the cluster definition file

Note: A CES shared root file system is required for protocols deployment with IBM Storage Scale.

1. To configure NSDs, you must first add your NSD server nodes to the configuration:

```
./spectrumscale node add -n nsdserver1  
./spectrumscale node add -n nsdserver2  
./spectrumscale node add -n nsdserver3
```

2. Once NSD server nodes are in the configuration, add NSDs to the configuration:

```
./spectrumscale nsd add /dev/sdb -p nsdserver1 -s nsdserver2,nsdserver3,...
```

- The installation toolkit supports standalone NSDs which connect to a single NSD server or shared NSDs which connect to both a primary and a secondary NSD server.
- When adding a standalone NSD, skip the secondary NSD server parameter.
- When adding a shared NSD, it is important to know the device name on the node which is to become the primary NSD server. It is not necessary to know the device name on the secondary NSD server because the device is looked up using its UUID.

Note: Although it is not necessary to know the device name on the secondary NSD server, it may be helpful to create a consistent mapping of device names if you are using multipath. For more information, see [“NSD disk discovery”](#) on page 24.

Here is an example of adding a shared NSD to the configuration by specifying the device name on the primary server along with the primary and secondary servers.

3. The name of the NSD is automatically generated based on the NSD server names. This can be changed after the NSD has been added by using the modify command and specifying a new name with the `-n` flag; the new name must be unique:

```
./spectrumscale nsd modify nsd_old_name -n nsd_new_name
```

4. It is possible to view all NSDs currently in the configuration using the list command:

```
./spectrumscale nsd list
```

5. To remove a single NSD from the configuration, supply the name of the NSD to the delete command:

```
./spectrumscale nsd delete nsd_name
```

6. To clear all NSDs and start from scratch, use the clear command:

```
./spectrumscale nsd clear
```

7. Where multiple devices are connected to the same pair of NSD servers, they can be added in bulk either by providing a list of all devices, or by using wild cards:

```
./spectrumscale nsd add -p nsdserver1 -s nsdserver2 /dev/dm-1 /dev/dm-2 /dev/dm-3
```

or

```
./spectrumscale nsd add -p nsdserver1 -s nsdserver2 "/dev/dm-*"
```

A connection is made to the primary server to expand any wild cards and check that all devices are present. When using wild cards, it is important to ensure that they are properly escaped, as otherwise they may be expanded locally by your shell. If any devices listed cannot be located on the primary server, a warning is displayed, but the command continues to add all other NSDs.

8. When adding NSDs, it is good practice to have them distributed such that each pair of NSD servers is equally loaded. This is usually done by using one server as a primary for half of the NSDs, and the other server as primary for the remainder.

```
./spectrumscale nsd add "/dev/dm-*" -p serverA -s serverB
```

```
[ INFO ] Connecting to serverA to check devices and expand wildcards.
[ INFO ] Adding NSD serverA_serverB_1 on serverA using device /dev/dm-0.
[ INFO ] Adding NSD serverA_serverB_2 on serverA using device /dev/dm-1.
$ ./spectrumscale nsd list
[ INFO ] Name                               FS      Size(GB) Usage   FG Pool   Device
Servers
[ INFO ] serverA_serverB_1 Default 13      Default 1 Default /dev/dm-0 serverA,serverB
[ INFO ] serverA_serverB_2 Default 1       Default 1 Default /dev/dm-1 serverA,serverB
```

9. Ordinarily a connection is made to the primary NSD server when adding an NSD. This is done to check device names and so that details such as the disk size can be determined, but is not vital. If it is not feasible to have a connection to the nodes while adding NSDs to the configuration, these connections can be disabled using the `--no-check` flag. Extra care is needed to manually check the configuration when using this flag.

```
./spectrumscale nsd add /dev/sda -p nsdserver1 -s nsdserver2 --no-check
```

10. You can set the failure group, file system, pool, and usage of an NSD in two ways:

- using the add command to set them for multiple new NSDs at once
- using the modify command to modify one NSD at a time

```
./spectrumscale nsd add "/dev/dm-*" -p nsdserver1 -s nsdserver2 \
-po pool1 -u dataOnly -fg 1 -fs filesystem_1
./spectrumscale nsd modify nsd_name -u metadataOnly -fs filesystem_1
```

For information on adding NSDs to an existing installation, see [“Adding nodes, NSDs, or file systems to an existing cluster”](#) on page 568.

Defining file systems

File systems are defined with the NSD configuration and they are only created at the time of installation, if there are NSDs assigned to them.

Note: A CES shared root file system is required for protocols deployment with IBM Storage Scale.

1. To specify a file system, use the **nsd add** or **nsd modify** command to set the file system property of the NSD:

```
./spectrumscale nsd add "/dev/dm-*" -p server1 -s server2 -fs filesystem_1
[ INFO ] The installer will create the new file system filesystem_1 if it does not already exist.
```

```
./spectrumscale nsd modify server1_server2_1 -fs filesystem_2
[ INFO ] The installer will create the new file system filesystem_2 if it does not already exist.
```

2. To list all file systems that currently have NSDs assigned to them, use the **list** command. This also displays file system properties including the block size and mount point:

```
./spectrumscale filesystem list
[ INFO ] Name      BlockSize  Mountpoint      NSDs Assigned
[ INFO ] filesystem_1 Default    /ibm/filesystem_1 3
[ INFO ] filesystem_2 Default    /ibm/filesystem_2 1
```

3. To alter the block size and mount point from their default values, use the modify command:

```
./spectrumscale filesystem modify filesystem_1 -B 1M -m /gpfs/gpfs0
```

Important: NSDs and file systems are created when the **./spectrumscale install** command is issued.

It is not possible to directly rename or delete a file system; this is instead done by reassigning the NSDs to a different file system using the **nsd modify** command.

At this point, the cluster definition file contains:

- Nodes and node types defined
- NSDs optionally defined
- File systems optionally defined

To proceed with the IBM Storage Scale installation, go to the next task: [“Installing IBM Storage Scale and creating a cluster”](#) on page 539.

For information on adding file systems to an existing installation, see [“Adding nodes, NSDs, or file systems to an existing cluster”](#) on page 568.

Setting configuration parameters before installation

After you define nodes in the cluster definition file, you can set certain configuration parameters in the cluster definition. Use the **./spectrumscale config gpfs** command for this purpose.

- To specify a cluster name in the cluster definition, use the **-c** argument.

```
./spectrumscale config gpfs -c gpfscluster01.my.domain.name.com
```

In this example, `gpfscluster01.my.domain.name.com` is the cluster name.

If no cluster name is specified, the GPFS admin node name is used as the cluster name. If the user-provided name contains periods, it is assumed to be a fully qualified domain name. If the cluster name is not a fully qualified domain name, the cluster name domain name is inherited from the admin node domain name.

- To specify a profile in the cluster definition file to be set on cluster creation, use the **-p** argument.

```
./spectrumscale config gpfs -p randomio
```

The valid values for **-p** option are `default` (for `gpfsProtocolDefaults` profile), `randomio` (for `gpfsProtocolRandomIO` profile), and *ProfileName* (for user-defined profile).

- The system-defined profiles are based on workload type: sequential I/O (`gpfsProtocolDefaults`) or random I/O (`gpfsProtocolRandomIO`). These defined profiles can be used to provide initial default tunables or settings for a cluster. If more tunable changes are required, see `mmchconfig` command and the `mmcrcluster` command in *IBM Storage Scale: Command and Programming Reference Guide*.
- You can specify a user-defined profile of attributes to be applied. The profile file specifies GPFS configuration parameters with values different than the documented defaults.

A user-defined profile must have the following properties:

- It must not begin with the string `gpfs`.
- It must have the `.profile` suffix.

- It must be located in the `/var/mmfs/etc/` directory.

For more information, see *mmcrcluster command* and *mmchconfig command* in *IBM Storage Scale: Command and Programming Reference Guide*.

If no profile is specified in the cluster definition, the `gpfsProtocolDefaults` profile is automatically set on cluster creation.

- To specify the remote shell binary to be used by GPFS, use the `-r` argument.

```
./spectrumscale config gpfs -r /usr/bin/ssh
```

If no remote shell is specified in the cluster definition, `/usr/bin/ssh` is used as the default.

- To specify the remote file copy binary to be used by GPFS, use the `-rc` argument.

```
./spectrumscale config gpfs -rc /usr/bin/scp
```

If no remote file copy binary is specified in the cluster definition, `/usr/bin/scp` is used as the default.

- To specify an ephemeral port range to be set on all GPFS nodes, use the `-e` argument.

```
./spectrumscale config gpfs -e 61000-62000
```

For information about the ephemeral port range, see *GPFS port usage* in *IBM Storage Scale: Administration Guide*.

If no port range is specified in the cluster definition, 60000-61000 is used as default.

- To view the current GPFS configuration settings, issue the following command.

```
./spectrumscale config gpfs --list
```

```
[ INFO ] No changes made. Current settings are as follows:
[ INFO ] GPFS cluster name is gpfscluster01
[ INFO ] GPFS profile is default
[ INFO ] Remote shell command is /usr/bin/ssh
[ INFO ] Remote file copy command is /usr/bin/scp
[ INFO ] GPFS Daemon communication port range is 61000-62000
```

Installing IBM Storage Scale and creating a cluster

After you set up the installer node, and define nodes, NSDs, and file systems in the cluster definition file, and set certain configuration parameters, you can install IBM Storage Scale according to the topology defined in the cluster definition file.

1. If call home is enabled in the cluster definition file, specify the minimum call home configuration parameters.

```
./spectrumscale callhome config -n CustName -i CustID -e CustEmail -cn CustCountry
```

For more information, see [“Enabling and configuring call home using the installation toolkit” on page 548](#).

2. Check environment before initiating the installation.

```
./spectrumscale install -pr
```

This step is not mandatory because running `./spectrumscale install` with no arguments also runs these checks before the installation.

3. Start the IBM Storage Scale installation and the creation of the cluster.

```
./spectrumscale install
```

Understanding what the installation toolkit does during the installation

- If the installation toolkit is being used to install IBM Storage Scale on all nodes, create a new cluster, and create NSDs, it automatically does the following outlined steps.
- If the installation toolkit is being used to add nodes to an existing GPFS cluster and create new NSDs, it automatically does the following steps.
- If all nodes in the cluster definition file are in a cluster, then the installation toolkit automatically does the following steps.

To add nodes to an existing GPFS cluster, at least one node in the cluster definition must belong to the cluster where the nodes not in a cluster are to be added. The cluster name in the cluster definition must also exactly match the cluster name outputted by **mm1sc1uster**.

When you run the **./spectrumscale install** command, the installation toolkit does these steps:

Install IBM Storage Scale on all nodes, create a new GPFS cluster, and create NSDs

- Run preinstall environment checks
- Install the IBM Storage Scale packages on all nodes
- Build the GPFS portability layer on all nodes
- Install and configure performance monitoring tools
- Create a GPFS cluster
- Configure licenses
- Set ephemeral port range
- Create NSDs, if any are defined in the cluster definition
- Create file systems, if any are defined in the cluster definition
- Run post-install environment checks

Add nodes to an existing GPFS cluster and create any new NSDs

- Run preinstall environment checks
- Install the IBM Storage Scale packages on nodes to be added to the cluster
- Install and configure performance monitoring tools on nodes to be added to the cluster
- Add nodes to the GPFS cluster
- Configure licenses
- Create NSDs, if any are defined in the cluster definition
- Create file systems, if any are defined in the cluster definition
- Run post-installation environment checks

Important: The installation toolkit does not alter anything on the existing nodes in the cluster. You can determine the existing nodes in the cluster by using the **mm1sc1uster** command.

The installation toolkit might change the performance monitoring collector configuration if you are adding the new node as a GUI node or an NSD node, due to collector node prioritization. However, if you do not want to change the collector configuration then you can use the **./spectrumscale config perfmon -r off** command to disable performance monitoring before initiating the installation procedure.

All nodes in the cluster definition are in a cluster

- Run preinstall environment checks
- Skip all steps until NSD creation
- Create NSDs (if any new NSDs are defined in the cluster definition)
- Run post-install environment checks

For more information, see **spectrumscale command** in *IBM Storage Scale: Command and Programming Reference Guide*.

What to do next

Upon completion of the installation, you have an active GPFS cluster. Within the cluster, NSDs might be created, file systems might be created, performance monitoring is configured, and all product licenses are accepted.

The installation can be rerun in the future to:

- Add NSD server nodes
- Add GPFS client nodes
- Add GUI nodes
- Add NSDs
- Define new file systems

Deploying protocols

Use this procedure to deploy protocols in an IBM Storage Scale cluster by using the installation toolkit.

Before you begin

Before you deploy protocols, make sure that the following requirements are met:

- An IBM Storage Scale cluster exists.
- IBM Storage Scale is started.
- Minimum one file system is available for the CES shared root file system.
- If the OS-level Python version is earlier than 3.8, upgrade Python, and then run the installation toolkit.

About this task

Protocol services are deployed on a subset of the cluster nodes that are designated as protocol nodes by using the **./spectrumscale node add node_name -p** command. Protocol nodes have an extra set of packages installed that allow them to run the NFS, SMB, and S3 protocol services.

Data is served through these protocols from a pool of addresses that are designated as Export IP addresses or CES public IP addresses. You can configure these addresses by using the **./spectrumscale config protocols -e IP1,IP2,IP3...** command or add these addresses manually by using the **mmces address add** command. The cluster manages the allocation of addresses in this pool, and IP addresses are automatically migrated to other available protocol nodes if any node fails.

Notes: Considerations for deploying protocols are as follows:

- All the protocol nodes must be running the supported operating systems, and the protocol nodes must have the same CPU architecture. Although the other nodes in the cluster might be on other platforms and operating systems.

For information about supported operating systems for protocol nodes and their required minimum kernel levels, see [IBM Storage Scale FAQ in IBM Documentation](#).

- The packages for all protocols are installed on every node that is designated as a protocol node. This installation is done even if a service is not enabled in your configuration.
- Services are enabled and disabled cluster wide; this means that every protocol node serves all enabled protocols.
- If SMB is enabled, the number of protocol nodes is limited to 16 nodes.
- If your protocol node has Red Hat Enterprise Linux 7.x installed, the NFS service might be already running on the node that can cause issues with the installation of IBM Storage Scale NFS packages. To avoid these issues, before you start the deployment, you must do the following steps:

1. Stop the NFS service by using the **systemctl stop nfs.service** command.
2. Disable the NFS service by using the **systemctl disable nfs.service** command.

This command ensures that this change is persistent after system reboot.

- The installation toolkit does not support adding protocol nodes to an existing ESS cluster earlier than ESS version 3.5.

Procedure

1. To use protocol services, define a shared file system for protocols.

If IBM Storage Scale is already configured, specify the shared file system manually or rerun **spectrumscale install** to assign an existing NSD. Make sure that the NSD servers are compatible with the installation toolkit and listed in the cluster definition file.

- a) Define a shared file system and the shared file system mount point or path.

```
./spectrumscale config protocols -f ceshared -m /gpfs/ceshared
```

- b) View the current settings by issuing this command:

```
./spectrumscale config protocols --list
```

```
[ INFO ] No changes made. Current settings are as follows:  
[ INFO ] Shared File System Name is ceshared  
[ INFO ] Shared File System Mountpoint or Path is /gpfs/ceshared
```

2. Add protocol nodes to the cluster definition file.

To deploy protocols on nodes in your cluster, add the protocol nodes to the cluster definition file as protocol nodes. Issue the following command to designate a node as a protocol node:

```
./spectrumscale node add NODE_IP -p
```

3. Enable NFS, SMB, and S3 protocol.

To enable or disable a set of protocols, use the **./spectrumscale enable** and **./spectrumscale disable** commands.

Note: The S3 protocol is available on IBM Storage Scale 5.2.1.0 and later.

- a) Enable the protocols.

```
./spectrumscale enable smb nfs s3
```

A sample output is as follows:

```
[ INFO ] Enabling SMB on all protocol nodes.  
[ INFO ] Enabling NFS on all protocol nodes.  
[ INFO ] Enabling S3 on all protocol nodes.
```

- b) View the current list of enabled protocols.

```
./spectrumscale node list
```

A sample output is as follows:

```
[ INFO ] List of nodes in current configuration:  
[ INFO ] [Installer Node]  
[ INFO ] 192.0.2.1  
[ INFO ]  
[ INFO ] [Cluster Name]  
[ INFO ] ESDev1  
[ INFO ]  
[ INFO ] [Protocols]  
[ INFO ] S3 : Enabled
```

```

[ INFO ] SMB : Enabled
[ INFO ] NFS : Enabled
[ INFO ]
[ INFO ] GPFS Node Admin Quorum Manager NSD Server Protocol GUI
Server OS Arch
[ INFO ] ESDev1-GPFS1 X X X
X rhel7 x86_64
[ INFO ] ESDev1-GPFS2 X
X rhel7 x86_64
[ INFO ] ESDev1-GPFS3 X
X rhel7 x86_64
[ INFO ] ESDev1-GPFS4 X X X
X rhel7 x86_64
[ INFO ] ESDev1-GPFS5 X X X
X rhel7 x86_64

```

4. Add CES IP addresses.

CES public IP addresses, also known as export IP addresses, provide data access through NFS, SMB, and S3 protocols. File and object clients use these IP addresses to access data in GPFS file systems. All protocols share export IP addresses, which are organized into a public IP address pool. The number of export IP addresses can be less than the number of protocol nodes. Assign a hostname to each export IP address, and configure reverse DNS lookup for every address.

a) Add export IP addresses to your cluster by using this command:

```
./spectrumscale config protocols -e EXPORT_IP_POOL
```

Where *EXPORT_IP_POOL* is a comma-separated list of IP addresses.

In the CES interface mode, you must specify the CES IP addresses with the installation toolkit in the Classless Inter-Domain Routing (CIDR) notation. In the CIDR notation, the IP address is followed by a forward slash and the prefix length.

```
IPAddress/PrefixLength
```

For example,

IPv6

```
2001:0DB8::/32
```

IPv4

```
192.0.2.0/20
```

Specify the prefix length for every CES IP address. Else, you cannot add the IP address by using the installation toolkit.

- When you are using IPv6 addresses, the prefix length must be in the range 1 - 124.
- When you are using IPv4 addresses, the prefix length must be in the range 1 - 30.

b) If you are using the CES interface mode, specify the interfaces.

```
./spectrumscale config protocols -i INTERFACES
```

Where *INTERFACES* is the comma-separated list of network interfaces, for example, eth0, eth1.

c) View the current configuration.

```
./spectrumscale node list
```

View the CES shared root and the IP address pool.

```
./spectrumscale config protocols -l
```

View the object configuration.

```
./spectrumscale config object -l
```

5. After the protocol-related definitions and configuration details are added to the cluster definition file, deploy the specified protocols.

```
./spectrumscale deploy
```

You can also use the **mmnetverify** command to identify any network problems before the deployment. For more information, see *mmnetverify command* in *IBM Storage Scale: Command and Programming Reference Guide*.

- a) Use the following command to deploy protocols:

```
./spectrumscale deploy
```

Note:

The system prompts for the Secret Encryption Key that you provided during object configuration or authentication. You can disable prompting to skip this step.

This command does the following tasks:

- Completes pre-deploy checks.
- Deploys protocols as specified in the cluster definition file.
- Completes post-deploy checks.

You can explicitly specify the **--precheck (-pr)** option to do a dry run of pre-deploy checks without starting the deployment. This check is not required because **./spectrumscale deploy** with no argument also runs these checks. Alternatively, you can specify the **--postcheck (-po)** option to do a dry run of post-deploy checks without starting the deployment. These options are mutually exclusive.

- b) After a successful deployment, verify the cluster and CES configuration.

```
/usr/lpp/mmfs/bin/mmlscluster --ces
```

What to do next

You can rerun the **./spectrumscale deploy** command to do the following tasks:

- Add protocol nodes
- Enable other protocols

Installation of performance monitoring tool using the installation toolkit

By default, the performance monitoring tool gets installed when IBM Storage Scale is installed on the system using the **./spectrumscale install** command.

For more information on the performance monitoring tool, see *Using the performance monitoring tool* in *IBM Storage Scale: Problem Determination Guide*.

During the installation, deployment, and upgrade phases, the installation toolkit performs following performance monitoring related configuration changes.

Installation

- Sets nodes to performance monitoring nodes (**mmchnode --perfmon -N Node1,Node2,Node3**)
- Installs **pmsensor** on any nodes that do not have the sensor.
- Installs **pmcollector** on two nodes. Preference is the GUI node first, followed by NSD nodes.
- Starts all sensors or collectors and activates all the default file system or node sensors.
- Resets **GPFSDiskCap** and **GPFSFilesetQuota** sensors to default values: 86400 and 3600 respectively and restricts them to the 1st collector node.

Deployment

- Sets nodes to performance monitoring nodes (**mmchnode --perfmon -N Node1,Node2,Node3**)
- Installs **pmsensor** on any nodes that do not have the sensor.
- Installs **pmcollector** on two nodes. Preference is the GUI node first, followed by NSD nodes.
- Starts all sensors or collectors and activates all the default file system or node sensors.
- Resets **GPFSDiskCap** and **GPFSFilesetQuota** sensors to default values: 86400 and 3600 respectively and restricts them to the 1st collector node.

Upgrade

- Upgrades **pmsensor** and **pmcollector**.
- Resets **GPFSDiskCap** and **GPFSFilesetQuota** sensors to default values: 86400 and 3600 respectively and restricts them to the 1st collector node.

Note: The installation toolkit does not set any NFSIO parameter related setting.

During the installation toolkit prechecks if it is found that the collector nodes require reconfiguring then the **./spectrumscale config perfmon -r on** command needs to be run to enable the installation toolkit to perform this reconfiguration.

Note: A reconfiguration results in the removal of an existing collector and it might move the collector to different nodes and reset sensor data. Custom sensors and data might be erased. To disable reconfiguration, run the **./spectrumscale config perfmon -r off** command.

If reconfiguration is disabled then the installation toolkit performs all installation tasks except for those related to performance monitoring and the management GUI, if GUI servers are specified.

After the installation or the deployment completes and sensors are added, the latest state of events of type TIPS might not get immediately refreshed. As a result, the output of the **mmhealth node show ces** command might contain messages such as the following for the CES component:

```
nfs_sensors_not_configured(NFSIO)
smb_sensors_not_configured(SMBGlobalStats, SMBStats)
```

To resolve this issue, run the **mmhealth node show --refresh** command.

For information on how to install the performance monitoring tool manually, see [“Manually installing the performance monitoring tool”](#) on page 495.

For information on performance monitoring related configuration in a cluster containing ESS, see [“Upgrading protocol nodes in an IBM Storage Scale System cluster”](#) on page 555.

For information on configuring the performance monitoring tool, see *Configuring the Performance Monitoring tool* in *IBM Storage Scale: Problem Determination Guide*.

Logging and debugging for installation toolkit

The installation toolkit reports progress to the console, and certain options also log output to a file.

Console output

The installation toolkit reports progress to the console. By default, only information level messages are displayed. If you need more verbose output, pass the **-v** flag to the **spectrumscale** command.

Note: This flag must be the first argument passed to the **spectrumscale** command. For example:

```
./spectrumscale -v install --precheck
```

Log files

In addition to console output, the `install` and `deploy` functions of the installation toolkit log verbose output to the `logs` subdirectory of the installation toolkit: `/usr/lpp/mmfs/x.x.x.x/ansible-toolkit/logs`

The full log path is also displayed on the console when logging begins. The log file contains verbose output even if the `-v` flag is not specified on the command line.

Installation and deployment checks

The `install` and `deploy` functions have a set of checks performed before and after their main action. These checks can be triggered separately by using either the `--precheck` or `--postcheck` flag.

It is safe to run the pre- and post-checks multiple times, because during these checks cluster data is obtained, but no changes are done.

Note: As with the full `install` and `deploy` functions, the pre- and post-checks create a detailed log file that can be used to provide further details about errors encountered.

Gathering support information

If an unresolvable error is encountered when using the installation toolkit, a support package can be gathered by using the `installer.snap.py` script that is available in the same directory as the installation toolkit.

Note: The `installer.snap.py` script must be run from the installation toolkit directory.

The `installer.snap.py` script gathers information from each of the nodes defined in the configuration and packages it such that it is ready for submission to the IBM Support Center. It creates a tar file with the prefix `installer_snap` in the `/tmp` directory.

The script gathers the following information:

- Cluster configuration file
- installation toolkit logs
- Information from the local node of the protocol services
- A GPFS snap file from each node defined in the configuration
- Basic environment information from the local node including:
 - Current OS and kernel version
 - System message log (`dmesg`)
 - File system usage
 - Network configuration
 - Current processes

Collection of core dump data

For information about changing configuration to enable collection of core dump data from protocol nodes, see *Configuration changes required on protocol nodes to collect core dump data* in *IBM Storage Scale: Problem Determination Guide*.

Enabling and configuring file audit logging using the installation toolkit

You can use the installation toolkit to enable and configure the file audit logging function in the cluster definition file. After enabling this function at the cluster level, you must enable it on file systems. The file audit logging package (`gpfs.librdkafka`) is installed on all supported nodes in the cluster specified to the installation toolkit during the installation, even if file audit logging is not enabled in the cluster configuration. In a cluster containing an ESS system wherein the setup type is ESS or `ess` in the cluster definition file, the file audit logging packages are installed on protocol nodes and client nodes. They are

not installed on ESS EMS and I/O server nodes. Based on the file audit logging configuration options specified in the cluster definition file using the installation toolkit, the function is enabled and configured in the cluster accordingly during the deployment.

For information on required packages for file audit logging, see [“Requirements, limitations, and support for file audit logging”](#) on page 647 and [“Installation prerequisites”](#) on page 470.

Note: A file system must be specified in the cluster definition file before you can enable file audit logging.

You can configure the file audit logging related options in the cluster definition file by using the installation toolkit as follows.

By default, file audit logging is disabled in the cluster definition file.

- To enable file audit logging in the cluster definition file, issue the following command before doing installation or deployment with the installation toolkit:

```
./spectrumscale fileauditlogging enable
```

- To disable file audit logging in the cluster definition file, issue the following command:

```
./spectrumscale fileauditlogging disable
```

- To list the file audit logging configuration in the cluster definition file, issue the following command:

```
./spectrumscale fileauditlogging list
```

You can verify whether file audit logging is enabled in the cluster definition file by viewing the output of the **./spectrumscale node list** command:

```
[ INFO ] List of nodes in current configuration:
[ INFO ] [Installer Node]
[ INFO ] 198.51.100.15
[ INFO ] Setup Type: SpectrumScale
[ INFO ]
[ INFO ] [Cluster Name]
[ INFO ] ESDev1
[ INFO ]
[ INFO ] [Protocols]
[ INFO ] SMB : Enabled
[ INFO ] NFS : Enabled
[ INFO ]
[ INFO ] File Audit logging : Disabled
[ INFO ]
[ INFO ] GPFS Node          Admin  Quorum  Manager  NSD Server  Protocol  GUI Server
OS Arch
[ INFO ] ESDev1-GPFS1      X      X      X              X
rhe17 x86_64
[ INFO ] ESDev1-GPFS2              X              X
rhe17 x86_64
[ INFO ] ESDev1-GPFS3              X              X
rhe17 x86_64
[ INFO ] ESDev1-GPFS4      X      X      X      X
rhe17 x86_64
[ INFO ] ESDev1-GPFS5      X      X      X      X
rhe17 x86_64
```

After enabling the file audit logging function in the cluster definition file, you must enable it on file systems on which you want to enable file audit logging.

- To enable file audit logging on a file system in the cluster definition file, issue the following command:

```
./spectrumscale filesystem modify --fileauditloggingenable FileSystemName
```

You can also specify the retention period and log fileset name with this command. For example, to specify a retention period of 180 days and to specify the log fileset name `testlog`, issue the following command:

```
./spectrumscale filesystem modify --fileauditloggingenable --retention 180 --logfileset testlog FileSystemName
```

- To disable file audit logging on a file system in the cluster definition file, issue the following command:

```
./spectrumscale filesystem modify --fileauditloggingdisable FileSystemName
```

Note: These file audit logging configuration-related changes become effective after the deployment procedure, initiated with **./spectrumscale deploy**, is completed.

Enabling and configuring call home using the installation toolkit

You can use the installation toolkit to enable and configure the call home function in the cluster definition file. Based on the call home configuration options specified in the cluster definition file using the installation toolkit, the call home function is enabled and configured in the cluster accordingly during the installation. By default, the call home function is enabled on all supported nodes in the cluster. If you plan to use the call home function, you must configure call home with the mandatory parameters by using the installation toolkit. For more information, see [this command example](#).

Note: If you want to only configure call home for a subset of the supported cluster nodes, or if you want to manually assign specific nodes to a specific call home group, see *Configuring call home to enable manual and automated data upload* in *IBM Storage Scale: Problem Determination Guide*.

The installation toolkit does not support all operating systems and platforms supported by call home. For example, the installation toolkit supports SLES and Ubuntu only on x86_64 and s390x, while call home also supports SLES and Ubuntu on Power.

The call home node needs to have network connectivity to the IBM Support Center because the call home node needs to upload data to IBM Support Center for all the members of its call home group. The call home node must be able to reach `https://esupport.ibm.com` either directly or using the proxy, which is specified in the call home configuration.

For more information about the call home function, see *Monitoring the IBM Storage Scale system by using call home* in *IBM Storage Scale: Problem Determination Guide*.

You can specify the call home node using the installation toolkit by issuing the following command:

```
./spectrumscale node add -c CallHomeNode
```

Note: This step is not mandatory. If a call home node is not specified, the installation toolkit assigns one of the nodes in the cluster as the call home node.

You can specify more than one node with this command. It is recommended to have at least one call home node for every 128 call home group members to prevent potential performance issues.

To configure call home options in the cluster definition file by using the installation toolkit, use the **./spectrumscale callhome** command as follows.

- To enable the call home in the cluster definition file by using the installation toolkit, issue the following command.

```
./spectrumscale callhome enable
```

A sample output is as follows.

```
[ INFO ] Enabling the callhome.
[ INFO ] Configuration updated.
```

Note: By default, call home is listed as enabled in the cluster definition file.

- To disable the call home in the cluster definition file by using the installation toolkit, issue the following command.

```
./spectrumscale callhome disable
```

A sample output is as follows.


```
[ INFO ] Disabling the callhome.  
[ INFO ] Configuration updated.
```

The call home function is enabled by default in the cluster definition file. If you disable it in the cluster definition file, when you issue the **./spectrumscale install** command, then no call home configuration is performed.

- To specify the call home configuration settings in the cluster definition file by using the installation toolkit, issue the **./spectrumscale callhome config** command. The following command example shows configuring the mandatory call home parameters:

```
./spectrumscale callhome config -n username -i 456123 -e username@example.com -cn US
```

A sample output is as follows.

```
[ INFO ] Setting Customer name to username  
[ INFO ] Setting Customer id to 456123  
[ INFO ] Setting Customer email to username@example.com  
[ INFO ] Setting Customer country to US
```

You are then prompted to accept or decline the support information collection message.

```
By accepting this request, you agree to allow IBM and its subsidiaries to store and use  
your contact information and your support information anywhere they do business worldwide.  
For more  
information, please refer to the Program license agreement and documentation. If you agree,  
please  
respond with 'accept' for acceptance, else with 'not accepted' to decline: accept
```

If you accept, the following output is displayed.

```
[ INFO ] License is accepted so the call home will be configured, if it is enabled.  
[ INFO ] Configuration is updated.
```

Note: You can specify the **-a** parameter with **./spectrumscale callhome config** to accept the support information collection message in advance.

- To clear the call home configuration specified in the cluster definition file by using the installation toolkit, issue the **./spectrumscale callhome clear** command. For example:

```
./spectrumscale callhome clear -i -e -cn
```

A sample output is as follows.

```
[ INFO ] Clearing Customer id  
[ INFO ] Clearing Customer email  
[ INFO ] Clearing Customer country
```

- To schedule the call home data collection in the cluster definition file by using the installation toolkit, issue the **./spectrumscale callhome schedule** command. For example:

```
./spectrumscale callhome schedule -d
```

A sample output is as follows.

```
[ INFO ] Setting scheduling to daily.  
[ INFO ] Configuration is updated.
```

If call home data collection is scheduled daily, data uploads are by default executed at 02:xx AM each day. If call home data collection is scheduled weekly, data uploads are by default executed at 03:xx AM each Sunday. In both cases, xx is a random number from 00 to 59.

You can use the **./spectrumscale callhome schedule -c** command to clear the call home data collection schedule.

- To list the call home configuration settings in the cluster definition file by using the installation toolkit, issue the following command.

```
./spectrumscale callhome list
```

A sample output is as follows.

```
[ INFO ] Current settings are as follows:
[ INFO ] Callhome is enabled.
[ INFO ] Schedule is daily.
[ INFO ] Setting Customer name to username
[ INFO ] Setting Customer id to 456123
[ INFO ] Setting Customer email to username@example.com
[ INFO ] Setting Customer country to US
[ WARN ] No value for Proxy ip in clusterdefinition file.
[ WARN ] No value for Proxy port in clusterdefinition file.
[ WARN ] No value for Proxy user in clusterdefinition file.
[ WARN ] No value for Proxy password in clusterdefinition file.
```

Note: These call home configuration-related changes become effective after the installation procedure, initiated with **./spectrumscale install**, is completed.

For detailed information on the **./spectrumscale callhome** command options, see **spectrumscale command** in *IBM Storage Scale: Command and Programming Reference Guide*.

Installing IBM Storage Scale management GUI by using the installation toolkit

Use the following information for installing the IBM Storage Scale management GUI using the installation toolkit.

The IBM Storage Scale is only installed when nodes have been specified as GUI servers in the cluster definition with the **-g** option:

```
./spectrumscale node add gpfsnode3 -g
```

The IBM Storage Scale GUI requires passwordless ssh to all other nodes in the cluster. Therefore, the GUI server node must be added as an admin node using the **-a** flag:

```
./spectrumscale node add gpfsnode3 -a
```

If no nodes have been specified as GUI servers, then the GUI will not be installed. It is recommended to have 2 GUI interface servers and a maximum of 3 for redundancy.

The GUI will be installed on specified GUI servers when you run the **./spectrumscale install** command and on protocol nodes also specified as GUI servers during the **./spectrumscale deploy** command.

At the end of a successful GPFS installation or protocol deployment, you can access the GUI through a web browser with the following node address: <https://<GUI server IP or hostname>>.

Note: After installing the system and GUI package, you need to create the first GUI user to log in to the GUI. This user can create other GUI administrative users to perform system management and monitoring tasks. When you launch the GUI for the first time after the installation, the GUI welcome page provides options to create the first GUI user from the command line prompt by using the **/usr/lpp/mmfs/gui/cli/mkuser <user_name> -g SecurityAdmin** command.

Installing IBM Storage Scale bridge for Grafana with the Installation Toolkit

Use the installation toolkit to enable and configure the Grafana bridge function in the cluster definition file.

After you enable this function at the cluster level, the installation toolkit installs the Grafana Bridge package (**gpfs.grafana-bridge**) on the designated pmcollector node.

Before running installation or deployment, use the following commands:

- To enable the Grafana bridge feature in the cluster definition file:

```
./spectrumscale grafanabridge enable
```

A sample output is as follows:

```
[ INFO ] Enabling grafana bridge in the cluster configuration file.
[ INFO ] Tip: If this feature is enabled on an existing cluster or a fresh installation, you
must run ./spectrumscale install to install grafana bridge package on pmcollector node.
[ INFO ] Tip: After the Grafana Bridge packages are installed on the pmcollector node, please
refer to the IBM Storage Scale KC documentation to manually configure the Grafana Bridge.
```

- To list Grafana bridge configuration options in the cluster definition file:

```
./spectrumscale grafanabridge -help
```

You can install Grafana bridge with the toolkit command:

```
./spectrumscale install
```

After running this command, the system:

- Installs all GPFS RPMs for all components
- Creates the cluster
- Starts daemons and services
- Installs and configures the GUI
- Enables sensors and collectors
- Installs and configures the Grafana bridge

Note: You must configure the Grafana bridge manually. For setup instructions, see [Setup the IBM Spectrum Scale Performance Monitoring Bridge for classic IBM Spectrum Scale devices](#)

Other available commands for Grafana bridge:

- To verify package installation:

```
dpkg -l | grep gpfs
```

After successful installation, the Grafana bridge connects to the pmcollector node and makes metrics available for visualization in Grafana dashboards.

- To disable the Grafana bridge service:

```
./spectrumscale grafanabridge disable
```

- To list the available commands for grafana bridge service:

```
./spectrumscale grafanabridge list
```

Populating cluster definition file with current cluster state using the installation toolkit

You can use the installation toolkit to populate the cluster definition file with the current cluster state.

Before you begin

Before using this functionality, review its limitations. For more information, see [“Limitations of config populate option of the installation toolkit ” on page 552.](#)

Note: If you are using the object protocol or proxy enabled call home, the config populate functionality does not extract any password from the current cluster configuration. You must enter the password when prompted.

About this task

In the following scenarios, you might need to update the cluster definition file with the current cluster state:

- A manually created cluster in which you want to use the installation toolkit
- A cluster created using the installation toolkit in which manual changes were done

Use the installation toolkit to populate the cluster definition file as follows:

Procedure

1. Define the installer node by issuing the **./spectrumscale setup -s *InstallNode*** command.
2. Repopulate the cluster definition file with the current cluster state by issuing the **./spectrumscale config populate --node *Node*** command.

Note: This command creates the backup of the existing cluster definition file and it creates a new cluster definition file.

- You can pass **--skip nsd** flag with config populate command to skip the nsd and filesystem details.

```
./spectrumscale config populate -N node --skip nsd
```

- In a cluster containing ESS, you must specify the EMS node with the config populate command. For example:

```
./spectrumscale config populate --node EMSNode
```

- You can also use this functionality to populate the cluster definition file with file audit logging and call home related information.

What to do next

If you are performing an upgrade, proceed with the rest of the upgrade steps. For more information, see [“Upgrading IBM Storage Scale components with the installation toolkit” on page 723](#).

Limitations of config populate option of the installation toolkit

You can use the **config populate** option of the installation toolkit to populate the cluster definition file with the current cluster state. This functionality has the following limitations.

Table 46. Limitations of the config populate functionality	
Function	Description
CES Groups	The config populate functionality does not detect CES groups. This might cause issues with CES stretch clusters
Clusters larger than 64 nodes	The config populate functionality does not support more than 64 nodes. For more information, see “Limitations of the installation toolkit” on page 518 .
Configuration parameters such as pagepool	The config populate functionality does not detect configuration parameters such as pagepool.
Custom GPFS profiles	The config populate functionality does not detect custom profiles. This might not be an issue in most scenarios except when you plan to use the config populate functionality to back up a cluster's configuration for a future new installation.

Table 46. Limitations of the config populate functionality (continued)

Function	Description
File placement optimizer (FPO)	The config populate functionality does not detect FPO NSD stanza information. However, the existing FPO NSD stanza information is not affected.
Greater than 3 GUI nodes	The config populate functionality adds in GUI nodes properly if the number of GUI nodes is less than or equal to 3. However, if there are more than 3 GUI nodes, the config populate functionality does not add them. This might leave you with GUI nodes that are not at the latest level of code that must be manually upgraded, and other unexpected side effects.
Greater than 3 performance monitoring collector nodes	The config populate functionality does not support more than 3 performance monitoring collector nodes. If GUI nodes are detected, they are also assigned as the performance monitoring collector nodes.
Highly-available write cache (HAWC)	The config populate functionality does not detect HAWC information. However, the existing HAWC information is not affected.
Heterogeneous cluster	The config populate functionality can be used in a cluster that contains nodes with mixed CPU architectures with some restrictions. The installation toolkit populates the information of nodes in the cluster that have the same CPU architecture as the installer node. Information for nodes of a different CPU architecture is not populated.
IBM Storage Protect for Space Management and IBM Storage Archive Enterprise Edition (EE)	The config populate functionality does not detect IBM Storage Protect for Space Management and IBM Storage Archive Enterprise Edition (EE) information. The installation toolkit provides some DMAPI detection and gives recommendations on how to proceed in case issues are encountered. For example: DMAPI holds an FS open for one of these functions, thereby causing upgrades to fail. In this case, the installation toolkit shows a hint that identifies the node and that it might need to be rebooted to proceed with the upgrade.

Table 46. Limitations of the config populate functionality (continued)

Function	Description
IBM Storage Scale release 4.2.0.2	In the IBM Storage Scale release 4.2.0.2, although the config populate functionality can be used to populate the cluster definition file with the current cluster state, not all parameters are properly updated. For example, CES IPs are not properly detected and you might need to add them manually by using the following command: ./spectrumscale config protocols -e ces_ip1,ces_ip2,... Review your cluster configuration after running ./spectrumscale config populate to ensure that all desired values are populated.
Local read-only cache (LROC)	The config populate functionality does not detect LROC information. However, the existing LROC information is not affected.
Non-federated performance monitoring setup	The config populate functionality does not support non-federated performance monitoring setup. The installation toolkit converts all performance monitoring configuration to federated mode unless specified to ignore performance monitoring completely.
NSD SAN attachment	The config populate functionality does not support retrieval of the NSD information in case of an NSD SAN attachment.
Offline mode in upgrade configuration	The config populate functionality cannot be used if one or more nodes in the cluster are designated as offline in the upgrade configuration.
transparent cloud tiering	The config populate functionality does not detect transparent cloud tiering related information.
Windows, AIX	The config populate functionality does not support these operating systems. If any of these operating systems are detected on a node and if that node has the same CPU architecture and endianness as the node from which the installation toolkit is running, the config populate functionality shows a warning, skips the node, and continues populating information from other supported nodes.

Related concepts

[“Limitations of the installation toolkit” on page 518](#)

Before using the installation toolkit to install IBM Storage Scale and deploy protocols, review the following limitations and workarounds, if any.

Upgrading protocol nodes in an IBM Storage Scale System cluster

By using the installation toolkit, you can install IBM Storage Scale and upgrade protocols nodes in a cluster.

Complete the following steps to upgrade protocol nodes by using the installation toolkit in a cluster that contains an IBM Storage Scale System:

1. Specify one of the newly added protocol nodes as an installer node, and specify an IBM Storage Scale System for setup by issuing the following command:

```
./spectrumscale setup -s EMS_NODE -st ess
```

2. Populate the cluster definition file with the current configuration of the management server node, and the protocol and the client nodes in the cluster. In a cluster that contains an IBM Storage Scale System, specify a management server node by issuing the following command:

```
./spectrumscale config populate -N EMS_NODE
```

Note: If the `config populate` option is incompatible with your configuration, add the nodes and the CES configuration to the installation toolkit manually.

```
./spectrumscale node add NodeName -e
```

For more information, see [“Populating cluster definition file with current cluster state using the installation toolkit”](#) on page 551.

3. Check the list of nodes that the IBM installation toolkit upgrades, and remove any non-CES nodes that you can upgrade manually.

```
./spectrumscale node list
```

4. Check the installation toolkit upgrade by issuing the following command:

```
./spectrumscale upgrade precheck
```

5. Upgrade protocol nodes by issuing the following command:

```
./spectrumscale upgrade run
```

6. Optional: Specify other configuration options, installing, and deploying by using the installation toolkit. For more information, see [“Defining the cluster topology for the installation toolkit”](#) on page 534, [“Installing IBM Storage Scale and creating a cluster”](#) on page 539, and [“Deploying protocols”](#) on page 541.

Several validations are done when you use the installation toolkit in an IBM Storage Scale cluster. However, you can ensure that the configuration using the installation toolkit does not have any impact on management server and I/O server. Also, you must be aware of the following considerations when you use the installation toolkit for an IBM Storage Scale cluster:

- You cannot take the following actions by using the installation toolkit.
 - Add management server node and I/O server nodes as protocol nodes.
 - Add GUI nodes in the IBM Storage Scale cluster. It is assumed that the management server node is running a GUI setup and the ability to add more GUI nodes by using the installation toolkit is disabled to not affect the existing configuration. During upgrade, the installation toolkit does not alter an existing configuration that contains one or more GUI setups. The upgrade of any of these components is done outside of the IBM Storage Scale System management server and I/O server nodes, and their configuration is not changed.
 - Add I/O server nodes as NSD nodes.

- Create NSDs and file systems from disks within the IBM Storage Scale System I/O nodes.
- Specify a cluster name. Because the installation toolkit detects and uses an existing cluster name.
- CES shared root file system must be mounted on all protocol nodes.
- Only performance monitoring sensor packages are installed on all nodes in an IBM Storage Scale cluster other than the IBM Storage Scale System management server and I/O server nodes. They are configured to point to the existing collector on the management server node. Apart from package installation, the installation toolkit adds protocol-specific sensors to the overall cluster performance monitoring configuration during the installation and deployment phases. During upgrade, the installation toolkit does not alter an existing configuration that might have one or more collectors or sensors. The upgrade of any of these components is done outside of the IBM Storage Scale System management server and I/O server nodes, and their configuration is not changed.

Related concepts

[“Using the installation toolkit to perform installation tasks: Explanations and examples” on page 532](#)

Use these explanations and examples of installation toolkit options and tasks to perform installation and deployment by using the installation toolkit.

Configuration of an IBM Storage Scale stretch cluster in an export services environment: a sample use case

This page describes a stretch cluster with NFS and SMB protocols that are enabled, installed, and deployed using the installation toolkit.

Note: Clustered watch folder is not supported in a stretch cluster environment.

Overview of the stretch cluster use case

A single GPFS cluster is defined over three geographically separate sites: two production sites and a tiebreaker site. One or more file systems are created, mounted, and accessed concurrently from the two active production sites that are connected over a reliable WAN network.

The data and metadata replication features of GPFS are used to maintain a secondary copy of each file system block, relying on the concept of disk failure groups to control the physical placement of the individual copies:

1. Separate the set of available disk volumes into two failure groups. Define one failure group at each of the active production sites.
2. Create a replicated file system. Specify a replication factor of 2 for both data and metadata.

With two copies of the data in separate locations, if one site has an unrecoverable disaster, you can recover from a single site with no data loss. Data from two separate sites can share namespace and be accessed by either site. CES groups are enabled to control traffic to the local site. For more information, see *Synchronous mirroring with GPFS replication* in *IBM Storage Scale: Administration Guide*.




Figure 46. Synchronous mirroring with GPFS replication

About the tested use case

The stretch cluster in this use case was configured as active-active, meaning that clients can read and write data from either site of the cluster. For more information, see *Synchronous mirroring with GPFS replication* in *IBM Storage Scale: Administration Guide*. You can also configure an active-passive stretch cluster. For more information, see *An active-passive cluster* in *IBM Storage Scale: Administration Guide*. You can replicate data with GPFS in two ways: synchronous replication and asynchronous replication with Active File Management (AFM). Because there are some differences between the two options, you need to understand both options in order to choose the solution that best fits your use case.

Synchronous replication

Synchronous replication ensures that:

- your data is always available
- you can read and write in both locations
- you do not have to perform recovery actions from applications except for changing the IP address/hostname.

Because data is synchronously replicated, the application gets consistent data and no data is lost during failover or failback. Synchronous replication requires a reliable high-speed, low latency network between sites, and it has a performance impact on the application when the failback occurs.

Asynchronous replication with AFM

This use case covers the synchronous replication solution only. The use case shows how in an active-active stretch cluster when one site's storage is offline, you can still read from and write to a replica of the data to the site that is online.

Note: Stretch CES clusters with SMB must have low latencies. High latencies might result in performance degradation.

Asynchronous replication with AFM can work on a high latency network. Because data is asynchronously replicated, all updates might not be replicated when a failure occurs. Therefore, the application needs to be able to tolerate data loss and to run recovery operations to fail back. This setup requires two separate GPFS clusters instead of one cluster at multiple geographical sites.

Limitations of a stretch cluster that uses GPFS synchronous replication

Consider the following limitations of a stretch cluster with NFS or SMB for synchronously replicating data:

1. The IBM Storage Scale installation toolkit cannot deploy protocols if the CES networks across the two sites cannot communicate. For more information, see *Limitations* in *IBM Storage Scale: Big Data and Analytics Guide*.
2. If your implementation requires you to set up IBM Storage Scale for object to one site, you will not have a seamless failover if you lose all of the protocol nodes on that site. You need to change the object ring configuration so that it points back to the CES Group that is available on the other site. For details, see *Configuration of object for isolated node and network groups* in *IBM Storage Scale: Administration Guide*.
3. When you have object enabled on one site and that entire site goes down unexpectedly, you might have to recover your endpoints manually since you can no longer ping them. In this case, refer to the steps provided in the OpenStack documentation: <https://docs.openstack.org/keystone/pike/install/keystone-install-rdo.html>.

Note: A deployment has a high chance of failure if the CES networks at each site cannot communicate with each other. For more information, see “Limitations of the installation toolkit” on page 518. For this use case, the cluster was deployed with protocols with separate VLAN-connected networks on each site; however, those networks are able to communicate with each other.

Using the spectrumscale installation toolkit to install a stretch cluster

When you set up a stretch cluster, it is important to understand the physical setup of the storage and how it maps from each site to each file system and failure group. Figure 2 shows the tested configuration where each site (A and B) has storage that is only seen by the NSD servers in that site.

Local Site CES Group

The 10GbE GPFS network and the 10GbE CES network will use the same adapter, but different ports on the adapter as well as different vlans/subnets, so there will be no network contention here.




Figure 47. Local Site CES Group

For this use case example, the installation toolkit was used to install the IBM Storage Scale software. You can find the installation toolkit by changing directories to where it was extracted (the default 6.0.x.x extraction path follows. This path might vary depending on the code level).

```
cd /usr/lpp/mmfs/6.0.x.x/installer
```

Use these instructions to install and deploy a stretch cluster.

1. Designate a setup node by issuing the following command:

```
./spectrumscale setup -s InstallNodeIP
```

The setup node is used to run all of the toolkit commands and to specify the protocol and NSD nodes.

2. Specify the protocol and NSD nodes by issuing the following commands:

```
./spectrumscale node add protocol1A -a -p -g
./spectrumscale node add protocol2A -a -p -g
./spectrumscale node add protocol3A -p
./spectrumscale node add protocol4A -p
./spectrumscale node add protocol1B -p
./spectrumscale node add protocol2B -p
./spectrumscale node add protocol3B -p
./spectrumscale node add protocol4B -p
./spectrumscale node add nsd1A -n -q
./spectrumscale node add nsd2A -n -q
./spectrumscale node add nsd1B -n -q
./spectrumscale node add nsd2B -n -q
./spectrumscale node add nsd3C -n -q
```

The `-s` argument identifies the IP address that nodes use to retrieve their configuration. This IP address is one associated with a device on the installation node. (The IP address is automatically validated during the setup phase.)

The `-q` argument indicates the quorum nodes that are to be configured in the cluster. To keep the cluster accessible during a failure, configure most of the quorum nodes to have GPFS active. In this use case, there are five quorum nodes, therefore three must be active to keep the cluster accessible. These nodes were chosen specifically because they are the least likely to become inaccessible at the same time. Because nsd1A and nsd2A are at one site, nsd1B and nsd2B are at a second site, and nsd3C is at a third site, the likelihood of more than three going down at a time is minimal.

No manager nodes were specified with the `-m` argument, but by default, if no `-m` argument is specified, the installation toolkit automatically sets the protocol nodes to manager nodes, leaving an even balance across both sites.

The GUI node designations are specified with the `-g` argument to be on protocol nodes that reside on the same site, but you can choose to have a single GUI, two GUIs on one site, or two GUIs on different sites. In this case, two GUIs were tested on a single site.

3. Define NSD mappings to physical disks and assign those NSDs to failure groups and file systems. The following example NSDs are designated as dataAndMetadata; however, if you have the capacity (disk space and disk speed), set up Metadata disks on SSDs for the best performance.

```
./spectrumscale nsd add -p nsd1A -s nsd2A -u dataAndMetadata -fs ces -fg 2 /dev/mapper/lun_8
./spectrumscale nsd add -p nsd1B -s nsd2B -u dataAndMetadata -fs ces -fg 1 /dev/mapper/
lun_1

./spectrumscale nsd add -p nsd1B -s nsd2B -u dataAndMetadata -fs gpfs0 -fg 2 /dev/mapper/
lun_6
./spectrumscale nsd add -p nsd2B -s nsd1B -u dataAndMetadata -fs gpfs0 -fg 2 /dev/mapper/
lun_4
./spectrumscale nsd add -p nsd1B -s nsd2B -u dataAndMetadata -fs gpfs0 -fg 2 /dev/mapper/
lun_10
./spectrumscale nsd add -p nsd2B -s nsd1B -u dataAndMetadata -fs gpfs0 -fg 2 /dev/mapper/
lun_24
./spectrumscale nsd add -p nsd2A -s nsd1A -u dataAndMetadata -fs gpfs0 -fg 1 /dev/mapper/
lun_2
./spectrumscale nsd add -p nsd1A -s nsd2A -u dataAndMetadata -fs gpfs0 -fg 1 /dev/mapper/
lun_3
./spectrumscale nsd add -p nsd2A -s nsd1A -u dataAndMetadata -fs gpfs0 -fg 1 /dev/mapper/
lun_4
./spectrumscale nsd add -p nsd1A -s nsd2A -u dataAndMetadata -fs gpfs0 -fg 1 /dev/mapper/
lun_5

./spectrumscale nsd add -p nsd3C -u descOnly -fs gpfs0 -fg 3 /dev/sda
./spectrumscale nsd add -p nsd3C -u descOnly -fs ces -fg 3 /dev/sdb
```

Each file system, ces or gpfs0, has multiple disks that have primary and secondary servers at each site. This ensures that the file system stays online when an entire site goes down. With multiple primary and secondary servers for each disk and failure group that is local to each site, the GPFS replication keeps the data up to date across both sites. A disk with a primary and secondary server on site A belongs to failure group 1, and a disk with a primary and secondary server on site B belongs

to failure group 2. This enables the two-way replication across the failure groups, meaning that one replica of data is kept at each site. The nsd3C node is known as the tiebreaker node. The physical disks that reside on that node /dev/sda and /dev/sdb are designated as 'descOnly' disks and are local to that node and are their own failure group. The descOnly argument indicates that the disk contains no file data or metadata. It is used solely to keep a copy of the file system descriptor. It is recommended to have that tiebreaker node in a separate geographical location than the other two sites.

4. Set up the file system characteristics for two-way replication on both the ces and gpfs0 file systems by issuing the following command:

```
./spectrumscale filesystem modify -r 2 -mr 2 ces
./spectrumscale filesystem modify -r 2 -mr 2 gpfs0
```

This sets the metadata and data replication to 2.

5. Designate file system paths for protocols by issuing the following commands:

```
./spectrumscale config protocols -f ces -m /ibm/ces
```

6. Set the cluster name by issuing the following command:

```
./spectrumscale config gpfs -c gumby.tuc.stglabs.ibm.com
```

7. Install the stretch cluster by issuing the following command:

```
./spectrumscale install --precheck
./spectrumscale install
```

8. Set up the IP lists by issuing the following command:

```
./spectrumscale config protocols -e
10.18.52.30,10.18.52.31,10.18.52.32,10.18.52.33,10.18.60.30,10.18.60.31,10.18.60.32,10.18.60
.33
./spectrumscale filesystem list
```

9. Enable the protocols by issuing the following commands:

```
./spectrumscale enable nfs
./spectrumscale enable smb
```

10. Configure authentication by issuing the following command:

```
./spectrumscale auth file ad
./spectrumscale node list
```

11. Deploy the stretch cluster by issuing the following commands:

```
./spectrumscale deploy --precheck
./spectrumscale deploy
```

12. After the deployment completes, check the AD setup and status.

For the use case, the same AD server was on both sites, but you can use any authentication type in a stretch cluster that is supported on a single-site IBM Storage Scale cluster. Note that because a stretch cluster is still one cluster, more than one authentication method per site is not supported.

To check the status of the cluster's authentication issue either of these commands **mmuserauth service list** or **mmuserauth service check --server-reachability**.

Issue the **mmuserauth service list** command. The system displays information similar to the following:

```
FILE access configuration : AD
PARAMETERS                VALUES
-----
ENABLE_NFS_KERBEROS       false
SERVERS                   10.18.2.1
USER_NAME                  Administrator
```

```

NETBIOS_NAME          stretch_cluster
IDMAP_ROLE             master
IDMAP_RANGE            10000-1000000
IDMAP_RANGE_SIZE       10000
UNIXMAP_DOMAINS        DOMAIN1(10000000-29999999)
LDAPMAP_DOMAINS        none

OBJECT access configuration : LOCAL
PARAMETERS             VALUES
-----
ENABLE_KS_SSL          false
ENABLE_KS_CASIGNING     false
KS_ADMIN_USER          admin

```

Issue the **mmuserauth service check --server-reachability** command. The system displays information similar to the following:

```

Userauth file check on node: protocol1A
  Checking nsswitch file: OK
AD servers status
  NETLOGON connection: OK
  Domain join status: OK
  Machine password status: OK
Service 'gpfs-winbind' status: OK

Userauth object check on node: protocol1A
  Checking keystone.conf: OK
  Checking wsgi-keystone.conf: OK
  Checking /etc/keystone/ssl/certs/signing_cert.pem: OK
  Checking /etc/keystone/ssl/private/signing_key.pem: OK
  Checking /etc/keystone/ssl/certs/signing_cacert.pem: OK

Service 'httpd' status: OK

```

Possible steps to convert an IBM Storage Scale cluster to a stretch cluster

1. Add the nodes from the second and third sites to the original cluster either manually or by using the spectrumscale toolkit.
2. Create the tiebreaker disks on the third site.
3. If replicating an existing file system, use the **mmchfs** command to set the replicas of data and metadata blocks to 2. If you are creating a new file system, ensure that the replication factor is set to 2 when it is created. For details, see the section Using the spectrumscale installation toolkit to install a stretch cluster.
4. Restripe your file system by issuing the **mmrestripefs <filesystem> -R** command.
5. Enable CES on the protocol nodes that you have added to the configuration.
6. Create CES groups on both sites.

Configuring the stretch cluster

1. Set up some basic tuning parameters.

For the use case, the following tuning parameters were used to improve the performance and reliability of the cluster. Tuning parameters will vary significantly depending on the hardware resources in your environment.

```

mmchconfig readReplicaPolicy=fastest
mmchconfig unmountOnDiskFail=yes -N nsd3
mmchconfig workerThreads=1024 -N cesNode
mmchconfig -ipagepool=43G -N protocol1A
mmchconfig -ipagepool=31G -N protocol2A
mmchconfig pagepool=48G -N protocol3A
mmchconfig pagepool=48G -N protocol4A
mmchconfig pagepool=48G -N protocol1B
mmchconfig pagepool=48G -N protocol2B
mmchconfig pagepool=48G -N protocol3B
mmchconfig pagepool=48G -N protocol4B
mmchconfig pagepool=12G -N nsd1A
mmchconfig pagepool=16G -N nsd1B

```

```
mmchconfig pagepool=12G -N nsd2B
mmchconfig pagepool=12G -N nsd3C
mmchconfig maxFilesToCache=2M
mmchconfig maxMBpS=5000 -N cesNodes
```

For details on each parameter, see *Parameters for performance tuning and optimization* in *IBM Storage Scale: Administration Guide*. The use case was tested with **readReplicaPolicy=fastest** which is the recommended setting. A known limitation with **readReplicaPolicy=fastest** is that with networks that add ~3 ms latency (which are common in such installations) there is no substantial difference between local and remote disks (assuming the disk latency might be in the 40/50ms range). Thus, you might still read data from the remote site. Therefore, it is acceptable to use **readReplicaPolicy=local** to ensure the data is written/read on the local site as long as the local servers are on the same subnet as the clients and the remote servers are not. An NSD server on the same subnet as the client is also considered as local.

The **readReplicaPolicy=fastest** setting will work with either network topology, both sites on the same subnet or each site on its own subnet, as long as there is a measurable difference in the I/O access time.

2. Set up the CES nodes.

CES group are needed when the CES networks on each site cannot communicate with each other. By having each site's local nodes in the same CES group, the administrator is able to control where the CES IPs failover to when there is an issue with a specific protocol node. If CES groups are not set up, a CES IP from Site A might attempt to failover to a node on Site B, and because there is no adapter for that IP to alias to on Site B (assuming different subnets), the failover will not succeed. CES groups make it easy to manage what CES nodes can host what CES IPs.

Set the CES nodes in the cluster to the corresponding groups by issuing the **mmchnode --ces-group** command (CES group names are not case-sensitive). For example:

```
mmchnode --ces-group SiteA -N protocol1A
mmchnode --ces-group SiteA -N protocol2A
mmchnode --ces-group SiteB -N protocol1B
mmchnode --ces-group SiteB -N protocol2B
```

In the example, protocol nodes protocol1A and protocol2A are set to the Site A CES group, protocol nodes protocol1B and protocol2B are set to the Site B CES group.

For detailed instructions, see *Setting up Cluster Export Services groups in an IBM Storage Scale cluster* in *IBM Storage Scale: Administration Guide*.

3. Assign CES IPs to the corresponding CES groups. This ensures that IPs that reside on nodes in Site A do not fail over to nodes that reside in Site B and vice versa. Issue the **mmces address change** command. For example:

```
mmces address change --ces-ip
10.18.52.30,10.18.52.31,10.18.52.32,10.18.52.33 --ces-group SiteA
mmces address change --ces-ip
10.18.60.30,10.18.60.31,10.18.60.32,10.18.60.33 --ces-group SiteB
```

4. To verify the CES groups your nodes belong to, issue the **mmces node list** command. The sample output is as follows:

Node	Name	Node Flags	Node Groups
10	protocol1B	none	siteB
11	protocol2B	none	siteB
12	protocol13B	none	siteB
13	protocol14B	none	siteB
6	protocol1A	none	siteA
7	protocol2A	none	siteA
8	protocol13A	none	siteA
9	protocol14A	none	siteA

5. To verify the CES groups your CES IPs belong to, issue the **mmces address list** command. The sample output is as follows:

Address	Node	Group	Attribute
10.18.52.30	protocol1AsiteA	object_singleton_node	object_database_node
10.18.52.31	protocol2AsiteA	none	
10.18.52.32	protocol3AsiteA	none	
10.18.52.33	protocol4AsiteA	none	
10.18.60.30	protocol1BsiteB	none	
10.18.60.31	protocol2BsiteB	none	
10.18.60.32	protocol3BsiteB	none	
10.18.60.33	protocol4BsiteB	none	

A load balancer is recommended for the protocol stack to cater to a site loss. The load balancer will ensure that you do not encounter issues if using DNS Round Robin when a site goes down and that the host name in the DNS server can resolve all of the IP addresses.

Using NFS and SMB with a stretch cluster

Using NFS and SMB protocols is similar to using an IBM Storage Scale cluster that is in one geographical location. All clients can read and write to either site and to any CES IP that they connect with depending on access. If a single protocol node fails at one site, a normal IP failover will still occur within the site, and the client seamlessly fails over with NFS I/O continuing. SMB clients, however, might need to be reconnected. On failures, clients can reconnect to another cluster node because the IP addresses of failing nodes are transferred to another healthy cluster node. Windows SMB clients automatically open a new connection without additional intervention, but the application that is running I/O may need to be restarted.

Monitoring and administering a stretch cluster

Monitoring your stretch cluster is the same as monitoring a cluster in a single location, except for the disk setup, and knowing when your disks are down and what site is affected. You can see the disk status using the **mmfsdisk** command. The sample output is as follows:

disk name	driver type	sector size	failure group	holds metadata	holds data	status	storage availability	pool
nsd22	nsd	512	2	Yes	Yes	ready	up	system
nsd23	nsd	512	2	Yes	Yes	ready	up	system
nsd24	nsd	512	2	Yes	Yes	ready	up	system
nsd25	nsd	512	2	Yes	Yes	ready	up	system
nsd26	nsd	512	1	Yes	Yes	ready	up	system
nsd27	nsd	512	1	Yes	Yes	ready	up	system
nsd28	nsd	512	1	Yes	Yes	ready	up	system
nsd29	nsd	512	1	Yes	Yes	ready	up	system
nsd30	nsd	512	1	Yes	Yes	ready	up	system
nsd31	nsd	512	3	No	No	ready	up	system

A healthy cluster shows all the disks that have a status of up. You can also verify the replication settings using the **mmfsfs -m -M -r -R** command. The sample output is as follows:

flag	value	description
-m	2	Default number of metadata replicas
-M	3	Maximum number of metadata replicas
-r	2	Default number of data replicas
-R	3	Maximum number of data replicas

If the default number of data and metadata replicas is set to two, this will indicate that you have no disk failures and that your data is being replicated across both failure groups.

Issue the **mm1sdisk ces** command. The sample output is as follows:

disk name	driver type	sector size	failure group	holds metadata	holds data	status	storage availability	pool
nsd20	nsd	512	2	Yes	Yes	ready	up	system
nsd21	nsd	512	1	Yes	Yes	ready	up	system
nsd32	nsd	512	3	No	No	ready	up	system

If you lose access to one site's storage due to maintenance, network issues, or hardware issues, the disks in the cluster are marked as down and the **mmhealth node show** command results shows them as down. This is acceptable because the stretch cluster can keep operating when an entire site goes down. There can be a negative impact on performance while one site is down, but that is expected.

To see the disk cluster status for the use case, issuing the **mm1sdisk gpfs0** command shows the following information:

disk name	driver type	sector size	failure group	holds metadata	holds data	status	storage availability	pool
nsd22	nsd	512	2	Yes	Yes	ready	down	system
nsd23	nsd	512	2	Yes	Yes	ready	down	system
nsd24	nsd	512	2	Yes	Yes	ready	down	system
nsd25	nsd	512	2	Yes	Yes	ready	down	system
nsd26	nsd	512	1	Yes	Yes	ready	up	system
nsd27	nsd	512	1	Yes	Yes	ready	up	system
nsd28	nsd	512	1	Yes	Yes	ready	up	system
nsd29	nsd	512	1	Yes	Yes	ready	up	system
nsd30	nsd	512	1	Yes	Yes	ready	up	system
nsd31	nsd	512	3	No	No	ready	up	system

For the use case, the results of the **mmhealth node show -N nsd2B disk** command show three disks:

Node name: nsd2B			
Node status: FAILED			
Status Change: 17 min. ago			
Component	Status	Status Change	Reasons
GPFS	FAILED	17 min. ago	gpfs_down, quorum_down
NETWORK	HEALTHY	10 days ago	-
FILESYSTEM	DEPEND	17 min. ago	unmounted_fs_check(gpfs1, ces, gpfs0)
DISK	DEPEND	17 min. ago	disk_down(nsd20, nsd22, nsd23)
PERFMON	HEALTHY	10 days ago	-

To see all of the failed disks, issue the **mmhealth node show nsd2B** command (without the -N attribute). For the use case, the system displays the following information:

Node name: nsd2B			
Component	Status	Status Change	Reasons
DISK	DEPEND	18 min. ago	disk_down(nsd20, nsd22, nsd23)
nsd1	DEPEND	18 min. ago	-
nsd10	DEPEND	18 min. ago	-
nsd11	DEPEND	18 min. ago	-
nsd12	DEPEND	18 min. ago	-
nsd13	DEPEND	18 min. ago	-
nsd14	DEPEND	18 min. ago	-
nsd15	DEPEND	18 min. ago	-
nsd16	DEPEND	18 min. ago	-
nsd17	DEPEND	18 min. ago	-
nsd18	DEPEND	18 min. ago	-
nsd19	DEPEND	18 min. ago	-
nsd2	DEPEND	18 min. ago	-
nsd20	DEPEND	18 min. ago	disk_down
nsd22	DEPEND	18 min. ago	disk_down
nsd23	DEPEND	18 min. ago	disk_down
nsd24	DEPEND	18 min. ago	disk_down
nsd25	DEPEND	18 min. ago	disk_down
nsd3	DEPEND	18 min. ago	-

nsd4	DEPEND	18 min. ago	-
nsd5	DEPEND	18 min. ago	-
nsd6	DEPEND	18 min. ago	-
nsd7	DEPEND	18 min. ago	-
nsd8	DEPEND	18 min. ago	-
nsd9	DEPEND	18 min. ago	-

Event	Parameter	Severity	Active Since	Event Message
disk_down	nsd20	WARNING	16 min. ago	Disk nsd20 is reported as not up
disk_down	nsd22	WARNING	16 min. ago	Disk nsd22 is reported as not up
disk_down	nsd23	WARNING	16 min. ago	Disk nsd23 is reported as not up
disk_down	nsd24	WARNING	16 min. ago	Disk nsd24 is reported as not up
disk_down	nsd25	WARNING	16 min. ago	Disk nsd25 is reported as not up

After the issue is resolved, restart the disks and make sure that the data and metadata replicas are intact. First, ensure that GPFS is active on all nodes. Next, issue the **mmchdisk <filesystem> start-a** command. This informs GPFS to try to access the disks that are marked down and, if possible, to move them back into the up state. This is accomplished by first changing the disk availability from down to recovering. The file system metadata is then scanned and any missing updates (replicated data that was changed while the disk was down) are repaired. If this operation is successful, the availability is then changed to up. If the metadata scan fails, availability is changed to unrecovered. This could occur if too many disks are down. The metadata scan can be re-initiated later by issuing the **mmchdisk** command again. If more than one disk in the file system is down, all of the disks that are down must be started at the same time by issuing **mmchdisk <filesystem> start -a**. If you start them separately and metadata is stored on any disk that remains down, the **mmchdisk start** command fails.

```

mmnsddiscover: Attempting to rediscover the disks. This may take a while ...
mmnsddiscover: Finished.
nsd2A: Rediscovered nsd server access to nsd26.
nsd2A: Rediscovered nsd server access to nsd28.
nsd3C: Rediscovered nsd server access to nsd31.
sdn2B: Rediscovered nsd server access to nsd23.
nsd1B: Rediscovered nsd server access to nsd23.
nsd2B: Rediscovered nsd server access to nsd24.
nsd1B: Rediscovered nsd server access to nsd24.
nsd1A: Rediscovered nsd server access to nsd29.
nsd2A: Rediscovered nsd server access to nsd30.
nsd2A: Rediscovered nsd server access to nsd27.
nsd2B: Rediscovered nsd server access to nsd25.
nsd2B: Rediscovered nsd server access to nsd22.
nsd2A: Rediscovered nsd server access to nsd25.
nsd2A: Rediscovered nsd server access to nsd22.
Scanning file system metadata, phase 1 ...
33 % complete on Fri Feb 3 11:46:41 2017
66 % complete on Fri Feb 3 11:56:57 2017
100 % complete on Fri Feb 3 11:58:24 2017 Scan completed successfully.
Scanning file system metadata, phase 2 ...
Scan completed successfully.
Scanning file system metadata, phase 3 ...
8 % complete on Fri Feb 3 11:58:29 2017
16 % complete on Fri Feb 3 11:58:32 2017
23 % complete on Fri Feb 3 11:58:35 2017
...
91 % complete on Fri Feb 3 11:59:18 2017
95 % complete on Fri Feb 3 11:59:22 2017
98 % complete on Fri Feb 3 11:59:25 2017
100 % complete on Fri Feb 3 11:59:26 2017
Scan completed successfully.
Scanning file system metadata, phase 4 ...
Scan completed successfully.
Scanning user file metadata ...
2.37 % complete on Fri Feb 3 11:59:46 2017 ( 2473984 inodes with total 672770 MB data
processed)
3.86 % complete on Fri Feb 3 12:00:07 2017 ( 4734976 inodes with total 1094807 MB data
processed)
4.59 % complete on Fri Feb 3 12:00:27 2017 ( 7880704 inodes with total 1301307 MB data
processed)
5.30 % complete on Fri Feb 3 12:00:47 2017 ( 11003904 inodes with total 1501577 MB data
processed)
6.01 % complete on Fri Feb 3 12:01:07 2017 ( 14077952 inodes with total 1703928 MB data
processed)
6.70 % complete on Fri Feb 3 12:01:27 2017 ( 17154048 inodes with total 1896877 MB data
processed)
7.36 % complete on Fri Feb 3 12:01:47 2017 ( 20135936 inodes with total 2084748 MB data
processed)

```

```

7.97 % complete on Fri Feb 3 12:02:07 2017 ( 22512640 inodes with total 2257626 MB data
processed)
8.21 % complete on Fri Feb 3 12:02:27 2017 ( 23322624 inodes with total 2327269 MB data
processed)
8.39 % complete on Fri Feb 3 12:02:48 2017 ( 24182784 inodes with total 2377108 MB data
processed)
8.52 % complete on Fri Feb 3 12:03:09 2017 ( 25182208 inodes with total 2414040 MB data
processed)
8.64 % complete on Fri Feb 3 12:03:29 2017 ( 26166272 inodes with total 2447380 MB data
processed)
...
96.58 % complete on Fri Feb 3 12:36:40 2017 ( 198458880 inodes with total 27362407 MB data
processed)
96.82 % complete on Fri Feb 3 12:37:00 2017 ( 202438144 inodes with total 27430464 MB data
processed)
97.06 % complete on Fri Feb 3 12:37:20 2017 ( 206526720 inodes with total 27498158 MB data
processed)
97.30 % complete on Fri Feb 3 12:37:40 2017 ( 210588672 inodes with total 27567944 MB data
processed)
97.46 % complete on Fri Feb 3 12:38:00 2017 ( 266730496 inodes with total 27612826 MB data
processed)
97.52 % complete on Fri Feb 3 12:38:20 2017 ( 302344960 inodes with total 27629694 MB data
processed)
97.59 % complete on Fri Feb 3 12:38:40 2017 ( 330066432 inodes with total 27648547 MB data
processed)
100.00 % complete on Fri Feb 3 12:38:52 2017 ( 394185216 inodes with total 27657707 MB data
processed)
Scan completed successfully.

```

The recovery time for this command can vary depending on how much data was written while the disks were down. If the disks were down for a long time (greater than 24 hours) and a lot of data was written in that time, it is expected that the **mmchdisk** command could take quite a while to complete. The time needed to bring up the disks depends on the quantity of data changed during the time the disks were down. This command is run while the file data remains accessible to the applications so I/O clients can continue to operate.

IBM Storage Scale stretch cluster use case conclusion

Each use case for a stretch cluster will vary. The sample use case represents one tested configuration. For more information see the following topics:

- *Synchronous mirroring with GPFS replication in IBM Storage Scale: Administration Guide*
- *An active-passive cluster in IBM Storage Scale: Administration Guide*
- *Limitations in IBM Storage Scale: Big Data and Analytics Guide*
- *Configuration of object for isolated node and network groups in IBM Storage Scale: Administration Guide*
- *Parameters for performance tuning and optimization in IBM Storage Scale: Administration Guide*
- <https://docs.openstack.org/keystone/pike/install/keystone-install-rdo.html>
- *Setting up Cluster Export Services groups in an IBM Storage Scale cluster in IBM Storage Scale: Administration Guide*

Performing additional tasks using the installation toolkit

Use this information for using the installation toolkit for tasks such as deploying protocols on existing clusters and adding nodes to an existing installation.

Deploying protocols on an existing cluster

Use this information to deploy protocols on an existing cluster by using the installation toolkit.

Procedure

1. Designate the installer node by using the following command:

```
./spectrumscale setup -s Installer_Node_IP
```

2. Designate protocol nodes in your environment to use for the deployment by using the following command:

```
./spectrumscale node add FQDN -p
```

3. Designate GUI nodes in your environment to use for the deployment by using the following command:

```
./spectrumscale node add FQDN -g
```

4. Configure protocols to point to a shared root file system by using the following command:

```
./spectrumscale config protocols -f FS_Name -m FS_Mountpoint
```

5. On the administrator node and all the protocol nodes, add a list of the pool of export IP addresses to the /etc/hosts file.

Example where the list of pool of export IP addresses is added to the hosts file

In this example, the *Comma_Separated_List_of_Exportpool_IPs* is
9.12.53.169,9.12.53.170,9.12.53.171,9.11.53.172.

```
/etc/hosts
9.12.53.169 ces1
9.12.53.170 ces2
9.12.53.171 ces3
9.11.53.172 ces
```

6. Configure the pool of export IP addresses by using the following command:

```
./spectrumscale config protocols -e Comma_Separated_List_of_Exportpool_IPs
```

7. Enable NFS on all protocol nodes by using the following command:

```
./spectrumscale enable nfs
```

8. Enable SMB on all protocol nodes by using the following command:

```
./spectrumscale enable smb
```

9. Review the configuration before deployment by using the following command:

```
./spectrumscale config protocols
./spectrumscale node list
```

10. Deploy protocols on your defined environment by using the following command:

```
./spectrumscale deploy --precheck
./spectrumscale deploy
```

Adding nodes, NSDs, or file systems to an existing cluster

Use this information to add nodes, NSDs, or file systems to an existing cluster by using the installation toolkit.

Procedure

- To add nodes to an existing cluster, do the following.
 - a) Add one or more node types by using the following commands.

Note: For adding nodes to an existing cluster, ensure that the IBM Storage Scale version on the nodes that you are adding is the same as the version on the existing nodes in the cluster.

 - Client nodes:

```
./spectrumscale node add FQDN
```

- NSD nodes:

```
./spectrumscale node add FQDN -n
```

- Protocol nodes:

```
./spectrumscale node add FQDN -p
```

- GUI nodes:

```
./spectrumscale node add FQDN -g -a
```

For the list of supported node types that can be added, see *spectrumscale command* in *IBM Storage Scale: Administration Guide*.

Important: The installation toolkit does not alter anything on the existing nodes in the cluster. You can determine the existing nodes in the cluster by using the **mmlscluster** command.

The installation toolkit might change the performance monitoring collector configuration if you are adding the new node as a GUI node or an NSD node, due to collector node prioritization. However, if you do not want to change the collector configuration then you can use the **./spectrumscale config perfmon -r off** command to disable performance monitoring before initiating the installation procedure.

- b) Install IBM Storage Scale on the new nodes by using the following commands.

```
./spectrumscale install --precheck  
./spectrumscale install
```

- c) If protocol nodes are being added, deploy protocols by using the following commands.

```
./spectrumscale deploy --precheck  
./spectrumscale deploy
```

- To add NSDs to an existing cluster, do the following.
 - a) Verify that the NSD server that connects this new disk runs an OS compatible with the installation toolkit and that the NSD server exists within the cluster.
 - b) Add NSDs to the cluster definition by using the following command.

```
./spectrumscale nsd add -p FQDN_of_Primary_NSD_Server Path_to_Disk_Device_File
```

- c) Run the installation by using the following commands.

```
./spectrumscale install --precheck  
./spectrumscale install
```

- To add file systems to an existing installation, do the following.
 - a) Verify that free NSDs exist and that they can be listed by the installation toolkit by using the following commands.

```
mmlnsd  
./spectrumscale nsd list
```

- b) Define the file system by using the following command.

```
./spectrumscale nsd modify NSD -fs File_System_Name
```

- c) Install the file system by using the following commands.

```
./spectrumscale install --precheck  
./spectrumscale install
```

Enabling another protocol on an existing cluster that has protocols enabled

Use this information to enable protocols on an existing cluster that has other protocols that are enabled by using the installation toolkit.

Procedure

1. Use one of the following steps depending on your configuration.

- Enable NFS on all protocol nodes by using the following command.

```
./spectrumscale enable nfs
```

- Enable SMB on all protocol nodes by using the following command.

```
./spectrumscale enable smb
```

2. Enable the new protocol by using the following command.

```
./spectrumscale deploy --precheck  
./spectrumscale deploy
```

Diagnosing errors during installation, deployment, or upgrade

Use this information to diagnose errors during installation, deployment, or upgrade by using the installation toolkit.

General troubleshooting steps

1. Note the screen output that is indicating the error. This output helps in narrowing down the general failure.
When a failure occurs, the screen output also shows the log file that contains the failure message.
2. Open the log file in an editor such as **vi**.
3. Go to the end of the log file and search upwards for the text FATAL.
4. Find the topmost occurrence of FATAL (or first FATAL error that occurred) and look in preceding and following sections of this error message for further indications of the failure.

For more information, see *Finding deployment related error messages more easily and using them for failure analysis* in *IBM Storage Scale: Problem Determination Guide*.

Installation or upgrade failure due to insufficient disk space

Installation or upgrade can fail with a transaction test error if the root (/) filesystem does not have sufficient free space. This issue is not specific to the installation toolkit and can occur with any RPM-based installation.

Example error message:

```
Transaction test error:  
installing package <package-name> needs <X>MB more space on the / filesystem  
  
Error Summary  
-----  
Disk Requirements:  
  At least <X>MB more space needed on the / filesystem.
```

Resolution:

1. Verify the available disk space on the root (/) filesystem:

```
df -h /
```

Note: Ensure that sufficient free space is available to avoid unexpected issues.

2. Clean up the root filesystem by removing unnecessary files or logs.
3. Retry the installation or upgrade.

Preparing a cluster that contains ESS for adding protocols

Use this information to prepare a cluster that contains ESS for adding functions such as NFS, SMB, GUI, and performance monitoring.

Before you begin

Before you add protocols to a cluster that contains ESS, ensure that:

- The cluster that contains ESS is active and online.
- Red Hat Enterprise Linux 7.x is installed on nodes that are going to serve as protocol nodes.
- Red Hat Enterprise Linux 7.x base repository is set up on nodes that are going to serve as protocol nodes.

Procedure

1. Use the ESS GUI or CLI to create a CES shared root file system.

Note: The CES shared root file system must be at least 4 GB and it must not be encrypted.

2. On all nodes that are going to serve as protocol nodes, download and extract the IBM Storage Scale protocols standard or advanced packages.

For information on extracting packages, see [“Extracting the IBM Storage Scale software on Linux nodes” on page 473](#).

3. On all nodes that are going to serve as protocol nodes, install core GPFS RPMs from the `/usr/lpp/mmfs/5.x.x.x/gpfs_rpms`.



Attention: Do not install packages such as `perfmon`, `gui`, `callhome`, `java`, and `protocols` now. These components are installed and configured with the protocols deployment.

The core GPFS RPMs that need to be installed include:

- `gpfs.base`
 - `gpfs.gpl`
 - `gpfs.license.xx`
 - `gpfs.gskit`
 - `gpfs.docs`
 - `gpfs.msg.en_US`
 - `gpfs.compression`
 - `gpfs.adv` (optional)
 - `gpfs.crypto` (optional)
4. Add the nodes that are going to serve as protocol nodes to the cluster by using the `mmaddnode` command.
 5. Enable the Cluster Configuration Repository (CCR) on the cluster if it is not enabled already.



Attention: If GPFS levels between protocol nodes and ESS differ significantly, ensure that the nodes with the newer code level of GPFS are designated as both quorum and manager nodes. For example, old ESS systems with GPFS 4.1.0-8 are incompatible with CES. These ESS systems can be a part of a cluster with protocols only if they are not designated as quorum and manager nodes.

This aspect requires careful planning and administration because:

- By default, ESS nodes are designated as quorum and manager nodes.
- In some cases, there might not be extra nodes outside of ESS to make sure that only nodes with the newer code level are designated as quorum and manager nodes.

6. Verify that passwordless SSH is working between all nodes in the cluster.
7. Verify that firewall ports are set correctly.
8. Verify that the CES shared root file system is mounted and that it is set to auto mount.
9. Set NFSv4 ACLs for file systems that are going to be used for protocol data, if they are not set already.
10. On one of the nodes that are going to serve as protocol nodes, run the installation toolkit to start deploying protocols.

Note:

- Designate only the nodes that you plan to use as protocol nodes. Do not designate existing ESS nodes such as EMS or I/O nodes as protocol nodes.
- Point to existing file systems only. Do not attempt to create new file systems or NSDs for this purpose.

For information about deploying protocols in a cluster, see **spectrumscale command** in *IBM Storage Scale: Command and Programming Reference Guide* and [“Deploying protocols on an existing cluster” on page 567](#).

Related concepts

[“Upgrading protocol nodes in an IBM Storage Scale System cluster” on page 555](#)

By using the installation toolkit, you can install IBM Storage Scale and upgrade protocols nodes in a cluster.

Adding an IBM Storage Archive Enterprise Edition (EE) node in an IBM Storage Scale cluster

Use this information to add an IBM Storage Archive Enterprise Edition (EE) node in an IBM Storage Scale cluster.

Procedure

1. Use the installation toolkit to install IBM Storage Scale and create a cluster.
For more information, see [“Using the installation toolkit to perform installation tasks: Explanations and examples” on page 532](#).
2. If the node on which IBM Storage Archive Enterprise Edition (EE) is to be installed is running on RHEL 6.x, install IBM Storage Scale packages on that node manually by using the RPMs from a protocol node.
3. Add the IBM Storage Archive Enterprise Edition (EE) node to your cluster at any time after the cluster is created.

For more information, see *Adding nodes to a GPFS cluster* in *IBM Storage Scale: Administration Guide*.

4. Before you install IBM Storage Archive Enterprise Edition (EE), enable Data Management API (DMAPI) on the file system to be used with IBM Storage Archive Enterprise Edition (EE).
 - a) Unmount the file system to be used with IBM Storage Archive Enterprise Edition (EE).
Note: If this file system is also used for protocols, you cannot unmount the file system unless you issue the **mmshutdown** command on all protocol nodes.
 - b) Enable DMAPI by using the **mmchfs Device -z yes** command.
In this command example, *Device* is the device name of the file system.
 - c) If you used the **mmshutdown** command in a preceding step, mount the file system again and issue the **mmstartup** on all protocol nodes.
5. Install and configure IBM Storage Archive Enterprise Edition (EE).

For detailed information on using IBM Spectrum Archive and IBM Storage Scale together, see [Active Archive Implementation Guide with Spectrum Scale Object and IBM Spectrum Archive Redpaper](#). Chapters 2 and 3 of this document contain specific details about cluster configuration and how to add the IBM Spectrum Archive node to the IBM Storage Scale cluster.

Protocol node IP further configuration

The **mmces** commands for adding or moving IPs are as follows.

For more information, see **mmces command** in *IBM Storage Scale: Command and Programming Reference Guide*.

Checking that protocol IPs are currently configured

The `mmces address list` command shows all the currently configured protocol IPs.

Note: The term CES IP refers to the Cluster Export Services IP, or a protocol IP used specifically for NFS, S3, SMB, and HDFS protocol access. All CES IPs must have an associated hostname and reverse DNS lookup must be configured for each. For more information, see *Adding export IPs* in [“Deploying protocols”](#) on page 541.

```
mmces address list
```

The system displays output similar to the following:

```
Address Node Group Attribute
-----
10.0.0.101 prt001st001 none none
10.0.0.102 prt002st001 none none
10.0.0.103 prt003st001 none none
10.0.0.104 prt004st001 none none
10.0.0.105 prt005st001 none none
10.0.0.106 prt006st001 none none
```

Additional IPs can be added to a node. In this example, IPs 10.0.0.108 and 10.0.0.109 are added to protocol node 5 (prt005st001).

```
mmces address add --ces-node prt005st001 --ces-ip 10.0.0.108,10.0.0.109
mmces address list
```

The system displays output similar to the following:

```
Address Node Group Attribute
-----
10.0.0.101 prt001st001 none none
10.0.0.102 prt002st001 none none
10.0.0.103 prt003st001 none none
10.0.0.104 prt004st001 none none
10.0.0.105 prt005st001 none none
10.0.0.106 prt006st001 none none
10.0.0.108 prt005st001 none none
10.0.0.109 prt005st001 none none
```

IPs can also be moved among nodes manually. This is helpful if IP allocation becomes unbalanced among the nodes. Continuing from the prior example, prt005st001 now has three protocol IPs whereas all other nodes have a single protocol IP. To balance this out a bit better, 10.0.0.109 will be manually moved to node prt004st001.

```
mmces address move --ces-ip 10.0.0.109 --ces-node prt004st001
mmces address list
```

The system displays output similar to the following:

```
Address Node Group Attribute
-----
10.0.0.101 prt001st001 none none
10.0.0.102 prt002st001 none none
10.0.0.103 prt003st001 none none
10.0.0.104 prt004st001 none none
10.0.0.105 prt005st001 none none
10.0.0.106 prt006st001 none none
10.0.0.108 prt005st001 none none
10.0.0.109 prt004st001 none none
```

List all protocol services enabled

Run the following command from any protocol node to list the running services on all protocol nodes:

```
mmces service list -a
```

The system displays output similar to the following:

```
Enabled services: SMB S3 NFS OBJ HDFS  
prt001st001: SMB is not running, S3 is running, NFS is not running, OBJ is running, HDFS is  
running,  
prt002st001: SMB is running, S3 is running, NFS is running, OBJ is not running, HDFS is running  
prt003st001: SMB is running, S3 is running, NFS is running, OBJ is running, HDFS is running
```

Chapter 5. Installing IBM Storage Scale on AIX nodes

There are three steps for installing GPFS on AIX nodes.

About this task

Before you begin installation, read [“Planning for GPFS” on page 309](#) and the [IBM Storage Scale FAQ in IBM Documentation](#).

Do not attempt to install GPFS if you do not have the prerequisites listed in [“Hardware requirements” on page 309](#) and [“Software requirements” on page 310](#).

Ensure that the PATH environment variable on each node includes `/usr/lpp/mmfs/bin`.

The installation process includes:

1. [“Creating a file to ease the AIX installation process” on page 575](#)
2. [“Verifying the level of prerequisite software” on page 575](#)
3. [“Procedure for installing GPFS on AIX nodes” on page 576](#)

Creating a file to ease the AIX installation process

Creation of a file that contains all of the nodes in your GPFS cluster prior to the installation of GPFS, will be useful during the installation process. Using either host names or IP addresses when constructing the file will allow you to use this information when creating your cluster through the `mmcrcluster` command.

For example, create the file `/tmp/gpfs.allnodes`, listing the nodes one per line:

```
k145n01.dpd.ibm.com
k145n02.dpd.ibm.com
k145n03.dpd.ibm.com
k145n04.dpd.ibm.com
k145n05.dpd.ibm.com
k145n06.dpd.ibm.com
k145n07.dpd.ibm.com
k145n08.dpd.ibm.com
```

Verifying the level of prerequisite software

Before you can install GPFS, verify that your system has the correct software levels installed.

About this task

If your system *does not* have the prerequisite AIX level, refer to the appropriate installation manual before proceeding with your GPFS installation. See the [IBM Storage Scale FAQ in IBM Documentation](#) for the latest software levels.

To verify the software version, run the command:

```
WCOLL=/tmp/gpfs.allnodes dsh "oslevel"
```

The system displays the current AIX level.

Procedure for installing GPFS on AIX nodes

These installation procedures are generalized for all levels of GPFS. Ensure you substitute the correct numeric value for the modification (*m*) and fix (*f*) levels, where applicable. The modification and fix level are dependent upon the current level of program support.

About this task

Follow these steps to install the GPFS software using the `installp` command:

1. [“Accepting the electronic license agreement” on page 576](#)
2. [“Creating the GPFS directory” on page 576](#)
3. [“Creating the GPFS installation table of contents file” on page 577](#)
4. [“Installing the GPFS man pages” on page 577](#)
5. [“Installing GPFS over a network” on page 577](#)
6. [“Verifying the GPFS installation” on page 578](#)

Accepting the electronic license agreement

The IBM Storage Scale software license agreement is shipped and viewable electronically. The electronic license agreement must be accepted before software installation can continue.

About this task

IBM Storage Scale has following editions based on functional levels:

- IBM Storage Scale Standard Edition
- IBM Storage Scale Data Access Edition
- IBM Storage Scale Advanced Edition
- IBM Storage Scale Data Management Edition

For IBM Storage Scale Standard Edition or IBM Storage Scale Data Access Edition and IBM Storage Scale Advanced Edition or IBM Storage Scale Data Management Edition installations, the installation cannot occur unless the appropriate license agreements are accepted.

When using the `installp` command, use the `-Y` flag to accept licenses and the `-E` flag to view license agreement files on the media.

GPFS license agreements are retained on the AIX system after installation completes. These license agreements can be viewed after installation in the following directories:

```
/usr/swlag/GPFS_standard (IBM Storage Scale Standard Edition or IBM Storage Scale Data  
Access Edition)  
/usr/swlag/GPFS_advanced (IBM Storage Scale Advanced Edition or IBM Storage Scale Data  
Management Edition)
```

Creating the GPFS directory

About this task

To create the GPFS directory:

1. On any node create a temporary subdirectory where GPFS installation images will be extracted. For example:

```
mkdir /tmp/gpfs1pp
```

2. Copy the installation images from the CD-ROM to the new directory, by issuing:

```
bffcreate -qvX -t /tmp/gpfs1pp -d /dev/cd0/AIX all
```

This command places these GPFS installation files in the images directory:

```
gpfs.base
gpfs.docs.data
gpfs.msg.en_US
gpfs.license.xx
gpfs.compression
gpfs.gskit
gpfs.crypto (IBM Storage Scale Advanced Edition or IBM Storage Scale Data Management
Edition only)
gpfs.adv (IBM Storage Scale Advanced Edition or IBM Storage Scale Data Management Edition
only)
```

Creating the GPFS installation table of contents file

About this task

To create the GPFS installation table of contents file:

1. Make the new image directory the current directory:

```
cd /tmp/gpfs1pp
```

2. Use the `inutoc` command to create a `.toc` file. The `.toc` file is used by the `installp` command.

```
inutoc .
```

Installing the GPFS man pages

In order to use the GPFS man pages you must install the `gpfs.docs.data` image.

About this task

The GPFS manual pages will be located at `/usr/share/man/`.

Installation consideration: The `gpfs.docs.data` image need not be installed on all nodes if man pages are not desired or local file system space on the node is minimal.

Installing GPFS over a network

About this task

Install GPFS according to these directions, where *localNode* is the name of the node on which you are running:

1. If you are installing on a shared file system network, ensure the directory where the GPFS images can be found is NFS exported to all of the nodes planned for your GPFS cluster (`/tmp/gpfs.allnodes`).
2. Ensure an acceptable directory or mountpoint is available on each target node, such as `/tmp/gpfs1pp`. If there is not, create one:

```
WCOLL=/tmp/gpfs.allnodes mmdsh "mkdir /tmp/gpfs1pp"
```

3. If you are installing on a shared file system network, to place the GPFS images on each node in your network, issue:

```
WCOLL=/tmp/gpfs.allnodes mmdsh "mount localNode:/tmp/gpfs1pp /tmp/gpfs1pp"
```

Otherwise, issue:

```
WCOLL=/tmp/gpfs.allnodes mmdsh "rcp localNode:/tmp/gpfs1pp/gpfs* /tmp/gpfs1pp"
WCOLL=/tmp/gpfs.allnodes mmdsh "rcp localNode:/tmp/gpfs1pp/.toc /tmp/gpfs1pp"
```

4. Install GPFS on each node:

```
WCOLL=/tmp/gpfs.allnodes mmdsh "installp -agXYd /tmp/gpfs1pp gpfs"
```

For information about using **rcp** with IBM Storage Scale, see [“Remote file copy command”](#) on page 324.

Verifying the GPFS installation

About this task

Verify that the installation procedure placed the required GPFS files on each node by running the `ls1pp` command on *each* node:

```
ls1pp -l gpfs\*
```

The system returns output similar to the following:

Fileset	Level	State	Description

Path: /usr/lib/objrepos			
gpfs.adv	6.0.1.x	COMMITTED	GPFS Advanced Features
gpfs.base	6.0.1.x	COMMITTED	GPFS File Manager
gpfs.crypto	6.0.1.x	COMMITTED	GPFS Cryptographic Subsystem
gpfs.gskit	8.0.55.x	COMMITTED	GPFS GSKit Cryptography Runtime
gpfs.compression	6.0.1.x	COMMITTED	GPFS Compression Libraries
gpfs.msg.en_US	6.0.1.x	COMMITTED	GPFS Server Messages - U.S. English
gpfs.license.std	6.0.1.x	COMMITTED	IBM Spectrum Scale Standard Edition
License			
Path: /etc/objrepos			
gpfs.base	6.0.1.x	COMMITTED	GPFS File Manager
Path: /usr/share/lib/objrepos			
gpfs.docs.data	6.0.1.x	COMMITTED	GPFS Server Manpages and Documentation

The output that is returned on your system can vary depending on the edition that you have installed (IBM Storage Scale Standard Edition or IBM Storage Scale Data Access Edition or IBM Storage Scale Advanced Edition or IBM Storage Scale Data Management Edition).

Note: The path returned by `ls1pp -l` shows the location of the package control data used by `installp`. The listed path does not show GPFS file locations. To view GPFS file locations, use the `-f` flag.

Chapter 6. Installing IBM Storage Scale on Windows nodes

There are several steps for installing GPFS on Windows nodes. The information in this topic points you to the detailed steps.

Before you begin installation, read the following:

- [“Planning for GPFS” on page 309](#)
- The [IBM Storage Scale FAQ in IBM Documentation](#)
- [“GPFS for Windows overview” on page 579](#) and all of its subtopics

Do not install GPFS unless you have the prerequisites listed in [“Hardware requirements” on page 309](#) and [“Software requirements” on page 310](#).

The installation process includes:

1. [“Installing GPFS prerequisites” on page 587](#)
2. [“Procedure for installing GPFS on Windows nodes” on page 590](#)
3. [“Configuring a mixed Windows and UNIX \(AIX or Linux\) cluster” on page 592](#)

To install GPFS for Windows, first configure your Windows systems as described in [“Installing GPFS prerequisites” on page 587](#). This includes steps such as joining an Active Directory domain and installing Cygwin from the [Cygwin website \(www.cygwin.com\)](#), which provides a UNIX-like environment on Windows. GPFS installation is simple once the prerequisites are completed. Finally, if your system will be part of a cluster that includes UNIX nodes, follow the steps described in [“Configuring a mixed Windows and UNIX \(AIX or Linux\) cluster” on page 592](#). This includes creating the GPFS Administration service, installing OpenSSH, and other requirements. Complete this process before performing configuration steps common to all GPFS supported platforms.

Note: Throughout this information, UNIX file name conventions are used. For example, the GPFS cluster configuration data is stored in the `/var/mmfs/gen/mmsdrfs` file. On Windows, the UNIX namespace starts under the Cygwin installation directory, which by default is `%SystemDrive%\cygwin64`, so the GPFS cluster configuration data is stored in the `C:\cygwin64\var\mmfs\gen\mmsdrfs` file.

GPFS for Windows overview

GPFS for Windows participates in a new or existing GPFS cluster in conjunction with AIX and Linux systems.

Support includes the following:

- Core GPFS parallel data services
- Windows file system semantics
- A broad complement of advanced GPFS features
- User identity mapping between Windows and UNIX
- Access to GPFS 3.4 and/or later file systems

Identity mapping between Windows and UNIX user accounts is a key feature. System administrators can explicitly match users and groups defined on UNIX with those defined on Windows. Users can maintain file ownership and access rights from either platform. System administrators are not required to define an identity map. GPFS automatically creates a mapping when one is not defined.

GPFS supports the unique semantic requirements posed by Windows. These requirements include case-insensitive names, NTFS-like file attributes, and Windows file locking. GPFS provides a bridge between a Windows and POSIX view of files, while not adversely affecting the functions provided on AIX and Linux.

GPFS for Windows provides the same core services to parallel and serial applications as are available on AIX and Linux. GPFS gives parallel applications simultaneous access to files from any node that has GPFS mounted, while managing a high level of control over all file system operations. System administrators and users have a consistent command interface on AIX, Linux, and Windows. With few exceptions, the commands supported on Windows are identical to those on other GPFS platforms. See [“GPFS limitations on Windows”](#) on page 581 for a list of commands that Windows clients do not support.

Related concepts

Installing GPFS prerequisites

This topic provides details on configuring Windows systems prior to installing GPFS.

Procedure for installing GPFS on Windows nodes

IBM provides GPFS as Windows Installer packages (MSI), which allow both interactive and unattended installations. Perform the GPFS installation steps as the Administrator or some other member of the Administrators group.

Running GPFS commands

Once GPFS is installed, the account you use for GPFS administrative operations (such as creating a cluster and file systems) will depend on your cluster type.

Configuring a mixed Windows and UNIX (AIX or Linux) cluster

Configuring the Windows HPC server

In order for Windows HPC Server to function with GPFS, disable dynamic hosts file and re-enable dynamic DNS updates by doing the following:

GPFS support for Windows

The [IBM Storage Scale FAQ in IBM Documentation](#) lists the levels of the Windows operating system that are supported by GPFS.

Limited GPFS support for the Windows platform first appeared in GPFS 3.2.1. Subsequent GPFS releases have expanded the set of available functions and features so that now most GPFS capabilities supported on the UNIX platforms are also supported on Windows.

Related concepts

GPFS limitations on Windows

GPFS for Windows supports most of the GPFS features that are available on AIX and Linux, but some limitations apply.

File system name considerations

File name considerations

File names created on UNIX-based GPFS nodes that are not valid on Windows are transformed into valid short names. A name may not be valid either because it is a reserved name like NUL or COM2, or because it has a disallowed character like colon (:) or question mark (?).

Case sensitivity

Native GPFS is case-sensitive; however, Windows applications can choose to use case-sensitive or case-insensitive names.

Antivirus software

If more than one GPFS Windows node is running antivirus software that scans directories and files, shared files only need to be scanned by one GPFS node. It is not necessary to scan shared files more than once.

Differences between GPFS and NTFS

GPFS differs from the Microsoft Windows NT File System (NTFS) in its degree of integration into the Windows administrative environment, Windows Explorer, and the desktop.

Access control on GPFS file systems

GPFS provides support for the Windows access control model for file system objects.

GPFS limitations on Windows

GPFS for Windows supports most of the GPFS features that are available on AIX and Linux, but some limitations apply.

The following limitations apply to configuring and operating a file system on a Windows node:

- File systems must have been created with GPFS 3.1.2.5 or later.
- You cannot upgrade an existing file system that was created with GPFS 3.1 or earlier.

The remaining GPFS for Windows limitations apply only to the Windows nodes in a cluster:

- The following GPFS commands are not supported on Windows:
 - **mmafm*** commands
 - **mmapplypolicy**
 - **mmaudit*** commands
 - **mmbackup**
 - **mmbackupconfig, mmrestoreconfig**
 - **mmcall*** commands
 - **mmcdp*** commands
 - **mmces*** commands
 - **mmcheckquota, mmdefedquota, mmedquota, mmlsquota, mmrepquota, mmsetquota**
 - **mmclone**
 - **mmcloud*** commands
 - **mmdelacl, mmeditACL**
 - **mmhadoop*** commands
 - **mmhealth**
 - **mmimgbackup, mmimgrestore**
 - **mmkeyserv**
 - **mmnetverify**
 - **mmnfs*** commands
 - **mmperfmon, mmpmon**
 - **mmrestorefs**
 - **mmsed**
 - **mm smb*** commands
 - **mm snmp*** commands
 - **mm sysmon*** commands
 - **mmuserauth**
 - **spectrumscale**
 - All **GNR** commands
- The GPFS Application Programming Interfaces (APIs) are not supported on Windows.
- The native Windows backup utility is not supported.
- Symbolic links that are created on UNIX-based nodes are specially handled by GPFS Windows nodes; they appear as regular files with a size of 0 and their contents cannot be accessed or modified.
- GPFS on Windows nodes attempts to preserve data integrity between memory-mapped I/O and other forms of I/O on the same computation node. However, if the same file is memory mapped on more than one Windows node, data coherency is not guaranteed between the memory-mapped sections on

these multiple nodes. In other words, GPFS on Windows does not provide distributed shared memory semantics. Therefore, applications that require data coherency between memory-mapped files on more than one node might not function as expected.

- GPFS Windows compute nodes are not supported in AFM clusters.

Related concepts

[GPFS support for Windows](#)

[File system name considerations](#)

[File name considerations](#)

File names created on UNIX-based GPFS nodes that are not valid on Windows are transformed into valid short names. A name may not be valid either because it is a reserved name like NUL or COM2, or because it has a disallowed character like colon (:) or question mark (?).

[Case sensitivity](#)

Native GPFS is case-sensitive; however, Windows applications can choose to use case-sensitive or case-insensitive names.

[Antivirus software](#)

If more than one GPFS Windows node is running antivirus software that scans directories and files, shared files only need to be scanned by one GPFS node. It is not necessary to scan shared files more than once.

[Differences between GPFS and NTFS](#)

GPFS differs from the Microsoft Windows NT File System (NTFS) in its degree of integration into the Windows administrative environment, Windows Explorer, and the desktop.

[Access control on GPFS file systems](#)

GPFS provides support for the Windows access control model for file system objects.

File system name considerations

GPFS file system names should be no longer than 32 characters if the file system will be mounted on a Windows node. GPFS file system names are used as Windows file system labels, and Windows limits the length of these labels to 32 characters. Any attempt to mount a file system with a long name will fail.

You can rename a file system using a command like the following:

```
mmchfs gpfs_file_system_name_that_is_too_long -W gpfs_name_1
```

Related concepts

[GPFS support for Windows](#)

[GPFS limitations on Windows](#)

GPFS for Windows supports most of the GPFS features that are available on AIX and Linux, but some limitations apply.

[File name considerations](#)

File names created on UNIX-based GPFS nodes that are not valid on Windows are transformed into valid short names. A name may not be valid either because it is a reserved name like NUL or COM2, or because it has a disallowed character like colon (:) or question mark (?).

[Case sensitivity](#)

Native GPFS is case-sensitive; however, Windows applications can choose to use case-sensitive or case-insensitive names.

[Antivirus software](#)

If more than one GPFS Windows node is running antivirus software that scans directories and files, shared files only need to be scanned by one GPFS node. It is not necessary to scan shared files more than once.

[Differences between GPFS and NTFS](#)

GPFS differs from the Microsoft Windows NT File System (NTFS) in its degree of integration into the Windows administrative environment, Windows Explorer, and the desktop.

[Access control on GPFS file systems](#)

GPFS provides support for the Windows access control model for file system objects.

File name considerations

File names created on UNIX-based GPFS nodes that are not valid on Windows are transformed into valid short names. A name may not be valid either because it is a reserved name like NUL or COM2, or because it has a disallowed character like colon (:) or question mark (?).

For more information, see [Naming Files, Paths, and Namespaces](#) in Microsoft Windows documentation.

Windows applications can use short names to access GPFS files with Windows file names that are not valid. GPFS generates unique short names using an internal algorithm. You can view these short names using `dir /x` in a DOS command prompt.

Table 47 on page 583 shows an example:

Table 47. Generating short names for Windows	
UNIX	Windows
foo+bar.foobar	FO~23Q_Z.foo
foo bar.-bar	FO~TD}C5._ba
f bar.-bary	F_AMJ5!._ba

Related concepts

[GPFS support for Windows](#)

[GPFS limitations on Windows](#)

GPFS for Windows supports most of the GPFS features that are available on AIX and Linux, but some limitations apply.

[File system name considerations](#)

[Case sensitivity](#)

Native GPFS is case-sensitive; however, Windows applications can choose to use case-sensitive or case-insensitive names.

[Antivirus software](#)

If more than one GPFS Windows node is running antivirus software that scans directories and files, shared files only need to be scanned by one GPFS node. It is not necessary to scan shared files more than once.

[Differences between GPFS and NTFS](#)

GPFS differs from the Microsoft Windows NT File System (NTFS) in its degree of integration into the Windows administrative environment, Windows Explorer, and the desktop.

[Access control on GPFS file systems](#)

GPFS provides support for the Windows access control model for file system objects.

Case sensitivity

Native GPFS is case-sensitive; however, Windows applications can choose to use case-sensitive or case-insensitive names.

This means that case-sensitive applications, such as those using Windows support for POSIX interfaces, behave as expected. Native Win32 applications (such as Windows Explorer) have only case-aware semantics.

The case specified when a file is created is preserved, but in general, file names are case insensitive. For example, Windows Explorer allows you to create a file named **Hello.c**, but an attempt to create **hello.c** in the same folder will fail because the file already exists. If a Windows node accesses a folder that contains two files that are created on a UNIX node with names that differ only in case, Windows inability to distinguish between the two files might lead to unpredictable results.

Related concepts

[GPFS support for Windows](#)

[GPFS limitations on Windows](#)

GPFS for Windows supports most of the GPFS features that are available on AIX and Linux, but some limitations apply.

[File system name considerations](#)

[File name considerations](#)

File names created on UNIX-based GPFS nodes that are not valid on Windows are transformed into valid short names. A name may not be valid either because it is a reserved name like NUL or COM2, or because it has a disallowed character like colon (:) or question mark (?).

[Antivirus software](#)

If more than one GPFS Windows node is running antivirus software that scans directories and files, shared files only need to be scanned by one GPFS node. It is not necessary to scan shared files more than once.

[Differences between GPFS and NTFS](#)

GPFS differs from the Microsoft Windows NT File System (NTFS) in its degree of integration into the Windows administrative environment, Windows Explorer, and the desktop.

[Access control on GPFS file systems](#)

GPFS provides support for the Windows access control model for file system objects.

Antivirus software

If more than one GPFS Windows node is running antivirus software that scans directories and files, shared files only need to be scanned by one GPFS node. It is not necessary to scan shared files more than once.

When you run antivirus scans from more than one node, schedule the scans to run at different times to allow better performance of each scan, as well as to avoid any conflicts that might arise because of concurrent exclusive access attempts by the antivirus software from multiple nodes. Note that enabling real-time antivirus protection for GPFS volumes could significantly degrade GPFS performance and cause excessive resource consumption.

Tip: Consider using a single, designated Windows node to perform all virus scans.

Latest versions of Windows such as Windows 10 come with a built-in antivirus component known as Windows Defender. While performing real-time scanning of files, Windows Defender might memory-map these files even when they are not in use by any user application. This memory-mapping of files on GPFS file systems by Windows Defender in the background, can result in performance degradation. Therefore, it is recommended that GPFS drives or volumes be excluded from Windows Defender scans all together. Another side effect of this background memory-mapping of GPFS files by antivirus programs is that an attempt to compress or uncompress GPFS files from these Windows nodes might fail. For more information, see *File compression and memory mapping* in *IBM Storage Scale: Administration Guide*.

Related concepts

[GPFS support for Windows](#)

[GPFS limitations on Windows](#)

GPFS for Windows supports most of the GPFS features that are available on AIX and Linux, but some limitations apply.

[File system name considerations](#)

[File name considerations](#)

File names created on UNIX-based GPFS nodes that are not valid on Windows are transformed into valid short names. A name may not be valid either because it is a reserved name like NUL or COM2, or because it has a disallowed character like colon (:) or question mark (?).

[Case sensitivity](#)

Native GPFS is case-sensitive; however, Windows applications can choose to use case-sensitive or case-insensitive names.

[Differences between GPFS and NTFS](#)

GPFS differs from the Microsoft Windows NT File System (NTFS) in its degree of integration into the Windows administrative environment, Windows Explorer, and the desktop.

[Access control on GPFS file systems](#)

GPFS provides support for the Windows access control model for file system objects.

Differences between GPFS and NTFS

GPFS differs from the Microsoft Windows NT File System (NTFS) in its degree of integration into the Windows administrative environment, Windows Explorer, and the desktop.

The differences are as follows:

- Manual refreshes are required to see any updates to the GPFS namespace.
- You cannot use the recycle bin.
- You cannot use distributed link tracking. This is a technique through which shell shortcuts and OLE links continue to work after the target file is renamed or moved. Distributed link tracking can help you locate the link sources in case the link source is renamed or moved to another folder on the same or different volume on the same computer, or moved to a folder on any computer in the same domain.
- You cannot use NTFS change journaling. This also means that you cannot use the Microsoft Indexing Service or Windows Search Service to index and search files and folders on GPFS file systems.

GPFS does not support the following NTFS features:

- File compression (on individual files or on all files within a folder)
- Encrypted directories
- Encrypted files (GPFS file encryption is administered through GPFS-specific commands. For more information, see *Encryption* in *IBM Storage Scale: Administration Guide*.)
- Quota management (GPFS quotas are administered through GPFS-specific commands)
- Reparse points
- Defragmentation and error-checking tools
- Alternate data streams
- Directory Change Notification
- The assignment of an access control list (ACL) for the entire drive
- Any Access Control Entry (ACE) type other than ALLOW, DENY, AUDIT and ALARM.

Upon an attempt to set any unsupported ACE type (such as OBJECT or CALLBACK/Conditional type), GPFS will filter and skip the unsupported ACE type from the ACL, resulting in a lossy ACL getting applied to the GPFS file or directory. Since conditional ACEs get filtered, GPFS does not support Dynamic Access Control (DAC)

- Generation of AUDIT and ALARM events specified in a System Access Control List (SACL). GPFS is capable of storing SACL content, but will not interpret it.
- The scanning of all files or directories that a particular SID owns (**FSCTL_FIND_FILES_BY_SID**)
- Windows sparse files API
- A change journal for file activity
- Transactional NTFS (also known as TxF)

Related concepts

[GPFS support for Windows](#)

[GPFS limitations on Windows](#)

GPFS for Windows supports most of the GPFS features that are available on AIX and Linux, but some limitations apply.

[File system name considerations](#)

[File name considerations](#)

File names created on UNIX-based GPFS nodes that are not valid on Windows are transformed into valid short names. A name may not be valid either because it is a reserved name like NUL or COM2, or because it has a disallowed character like colon (:) or question mark (?).

Case sensitivity

Native GPFS is case-sensitive; however, Windows applications can choose to use case-sensitive or case-insensitive names.

Antivirus software

If more than one GPFS Windows node is running antivirus software that scans directories and files, shared files only need to be scanned by one GPFS node. It is not necessary to scan shared files more than once.

Access control on GPFS file systems

GPFS provides support for the Windows access control model for file system objects.

Access control on GPFS file systems

GPFS provides support for the Windows access control model for file system objects.

Each GPFS file or directory has a Security Descriptor (SD) object associated with it and you can use the standard Windows interfaces for viewing and changing access permissions and object ownership (for example, Windows Explorer Security dialog panel). Internally, a Windows SD is converted to an NFSv4 access control list (ACL) object, which ensures that access control is performed consistently on other supported operating systems. GPFS supports all discretionary access control list (DACL) operations, including inheritance. GPFS is capable of storing system access control list (SACL) objects, but generation of AUDIT and ALARM events specified in SACL contents is not supported.

An important distinction between GPFS and Microsoft Windows NT File Systems (NTFS) is the default set of permissions for the root (top-level) directory on the file system. On a typical NTFS volume, the DACL for the top-level folder has several inheritable entries that grant full access to certain special accounts, as well as some level of access to nonprivileged users. For example, on a typical NTFS volume, the members of the local group **Users** would be able to create folders and files in the top-level folder. This approach differs substantially from the traditional UNIX convention where the root directory on any file system is only writable by the local **root** superuser by default. GPFS adheres to the latter convention; the root directory on a new file system is only writable by the UNIX user **root**, and does not have an extended ACL when the file system is created. This is to avoid impacting performance in UNIX-only environments, where the use of extended ACLs is not common.

When a new GPFS file system is accessed from a Windows client for the first time, an immutable security descriptor object is created for the root directory automatically. This immutable security descriptor for the root directory contains a non-inheritable DACL that grants full access to the local **Administrators** group and read-only access to the **Everyone** group. This allows only privileged Windows users (members of the local **Administrators** group) to create new files and folders immediately under the root directory. Because the root directory DACL has no inheritable entries, new top-level objects under the root directory are created with a default non-inheritable DACL that only grants local **Administrators** and **SYSTEM** accounts full access. Therefore, as a best practice, privileged Windows users must create top-level objects under the root directory and then explicitly set inheritable DACLs on these top-level directories as appropriate (granting the necessary level of access to non-privileged users).

Note: Some applications expect to find NTFS-style permissions on all file systems and they might not function properly when that is not the case. Running such an application in a GPFS folder where permissions have been set similar to NTFS defaults might correct this.

Related concepts

GPFS support for Windows

GPFS limitations on Windows

GPFS for Windows supports most of the GPFS features that are available on AIX and Linux, but some limitations apply.

File system name considerations

File name considerations

File names created on UNIX-based GPFS nodes that are not valid on Windows are transformed into valid short names. A name may not be valid either because it is a reserved name like NUL or COM2, or because it has a disallowed character like colon (:) or question mark (?).

Case sensitivity

Native GPFS is case-sensitive; however, Windows applications can choose to use case-sensitive or case-insensitive names.

Antivirus software

If more than one GPFS Windows node is running antivirus software that scans directories and files, shared files only need to be scanned by one GPFS node. It is not necessary to scan shared files more than once.

Differences between GPFS and NTFS

GPFS differs from the Microsoft Windows NT File System (NTFS) in its degree of integration into the Windows administrative environment, Windows Explorer, and the desktop.

Installing GPFS prerequisites

This topic provides details on configuring Windows systems prior to installing GPFS.

Perform the following steps:

1. [“Configuring Windows” on page 588](#)
 - a. [“Assigning a static IP address” on page 588](#)
 - b. [“Joining an Active Directory domain” on page 588](#)
 - c. [“Disabling the Windows firewall” on page 588](#)
 - d. [“Installing the Tracefmt and Tracelog programs \(optional\)” on page 588](#)
2. From here on, all GPFS installation steps must be executed as a special user with Administrator privileges. Further, all GPFS administrative operations (such as creating a cluster and file systems) must also be issued as the same user. This special user will depend on your cluster type. For Windows homogeneous clusters, run GPFS commands as a member of the Domain Admins group or as a domain account that is a member of the local Administrators group. For clusters with both Windows and UNIX nodes, run GPFS commands as root, a special domain user account described in [“Creating the GPFS administrative account” on page 593](#).
3. [“Installing Cygwin” on page 588](#)

See the [IBM Storage Scale FAQ in IBM Documentation](#) for the latest:

- Software recommendations
- Configuration information

For additional requirements when setting up clusters that contain both Windows nodes and AIX or Linux nodes, see [“Configuring a mixed Windows and UNIX \(AIX or Linux\) cluster” on page 592](#).

Related concepts

GPFS for Windows overview

GPFS for Windows participates in a new or existing GPFS cluster in conjunction with AIX and Linux systems.

Procedure for installing GPFS on Windows nodes

IBM provides GPFS as Windows Installer packages (MSI), which allow both interactive and unattended installations. Perform the GPFS installation steps as the Administrator or some other member of the Administrators group.

Running GPFS commands

Once GPFS is installed, the account you use for GPFS administrative operations (such as creating a cluster and file systems) will depend on your cluster type.

Configuring a mixed Windows and UNIX (AIX or Linux) cluster

Configuring the Windows HPC server

In order for Windows HPC Server to function with GPFS, disable dynamic hosts file and re-enable dynamic DNS updates by doing the following:

Configuring Windows

This topic provides some details on installing and configuring Windows on systems that will be added to a GPFS cluster.

The [IBM Storage Scale FAQ in IBM Documentation](#) lists the versions of the Windows operating system that are supported by GPFS.

Assigning a static IP address

GPFS communication requires invariant static IP addresses for each GPFS node.

Joining an Active Directory domain

All Windows systems in the same GPFS cluster should be members of the same Active Directory domain. Join these systems to the Windows domain before adding them to a GPFS cluster.

GPFS expects that all Windows nodes in a cluster are members of the same domain. This gives domain users a consistent identity and consistent file access rights independent of the system they are using. The domain controllers, which run the Active Directory Domain Services, are not required to be members of the GPFS cluster.

Refer to your Windows Server documentation for information on how to install and administer Active Directory Domain Services.

Disabling the Windows firewall

GPFS requires that you modify the default Windows Firewall settings.

The simplest change that will allow GPFS to operate properly is to disable the firewall. Open **Windows Firewall** in the Control Panel and click **Turn Windows firewall on or off**, and select **Off** under the General tab. For related information, see *GPFS port usage* in *IBM Storage Scale: Administration Guide*.

Installing the Tracefmt and Tracelog programs (optional)

GPFS diagnostic tracing (mmtracectl) on Windows uses the Microsoft programs called tracefmt.exe and tracelog.exe. These programs are not included with Windows but can be downloaded from Microsoft. The tracefmt.exe and tracelog.exe programs are only for tracing support and are not required for normal GPFS operations.

The tracefmt.exe and tracelog.exe programs are included with the Windows Driver Kit (WDK) as well as the Windows SDK. You can download either of these packages from the [Microsoft Download Center \(www.microsoft.com/download\)](http://www.microsoft.com/download).

To allow GPFS diagnostic tracing on Windows using the WDK, follow these steps:

1. Download the latest version of the Windows Driver Kit (WDK) from Microsoft.
2. Install the Tools feature of the WDK on some system to obtain a copy of tracefmt.exe and tracelog.exe.
3. Locate and copy tracefmt.exe and tracelog.exe to both the /usr/lpp/mmfs/bin and /usr/lpp/mmfs/win directories to ensure mmtracectl properly locates these programs.

For additional information about GPFS diagnostic tracing, see *The GPFS trace facility* in *IBM Storage Scale: Problem Determination Guide*.

Installing Cygwin

Cygwin is a POSIX environment available for Windows and can be downloaded from the [Cygwin website \(www.cygwin.com\)](http://www.cygwin.com). GPFS uses this component to support its many administrative scripts. System

administrators have the option of using either the GPFS Admin Command Prompt or GPFS Admin Korn Shell to run GPFS commands.

Cygwin must be installed before installing GPFS. It is a software package that provides a Unix-like environment on Windows and provides runtime support for POSIX applications and includes programs such as `grep`, `ksh`, `ls`, and `ps`.

When running Cygwin setup, only the standard packages are installed by default. GPFS requires installation of additional packages, which are listed in [“Installing the 64-bit version of Cygwin”](#) on page 589.

Note: Starting with GPFS 4.1.1, the 32-bit version of Cygwin is no longer supported for Windows nodes running GPFS. Users that are running GPFS 4.1 with the 32-bit version of Cygwin installed must upgrade to the 64-bit version of Cygwin before installing GPFS 4.1.1. For more information, see [“Upgrading from the 32-bit version of Cygwin to the 64-bit version of Cygwin”](#) on page 589. For users on GPFS releases prior to 4.1 (SUA based), see [“Offline upgrade with complete cluster shutdown”](#) on page 769.

Upgrading from the 32-bit version of Cygwin to the 64-bit version of Cygwin

Follow these instructions to upgrade from the 32-bit version of Cygwin to the 64-bit version of Cygwin:

1. Uninstall GPFS 4.1 and reboot.
2. Uninstall IBM GSKit for GPFS and reboot.
3. Uninstall the GPFS 4.1 license.
4. Stop and delete any Cygwin 32-bit services, such as OpenSSH, that might have been configured.
5. Do *not* uninstall the 32-bit version of Cygwin yet, or you may lose GPFS configuration information.
6. Install the 64-bit version of Cygwin using the instructions in [“Installing the 64-bit version of Cygwin”](#) on page 589.
7. Install the GPFS 4.1.1 license for the appropriate edition; for example, `gpfs.ext-4.1.1-Windows-license.msi`.
8. Install the appropriate GPFS 4.1.1 edition; for example, `gpfs.ext-4.1.1.x-Windows.msi`.
9. Install IBM GSKit for GPFS.
10. Uninstall the 32-bit version of Cygwin completely.
11. Follow the procedures in [“Installing and configuring OpenSSH on Windows nodes”](#) on page 594.

Installing the 64-bit version of Cygwin

To install the 64-bit version of Cygwin for Windows, follow these steps:

1. Logon to the Windows node as the account you will use for GPFS administrative operations.
2. Go to the [Cygwin website \(www.cygwin.com\)](http://www.cygwin.com), and click the **Install Cygwin** link in the upper-left pane.
3. Download and start the installation of the `setup-x86_64.exe` file.
4. Follow the prompts to continue with the install. The default settings are recommended until the Select Packages dialog is displayed. Then, select the following packages (use the Search field to quickly find these packages):
 - `diffutils`: A GNU collection of diff utilities
 - `flip`: Convert text file line endings between Unix and DOS formats
 - `m4`: GNU implementation of the traditional Unix macro processor
 - `mksh`: MirBSD Korn Shell
 - `perl`: Perl programming language interpreter
 - `procps-ng`: System and process monitoring utilities
 - `openssh`: The OpenSSH server and client programs (only required if you plan on mixing Linux, AIX and Windows nodes in the same cluster)

5. Click the **Next** button, and continue to follow the prompts to complete the installation.
6. If you are using a mixed cluster environment (Windows and Linux/AIX), follow the steps in [“Installing and configuring OpenSSH on Windows nodes”](#) on page 594.

Procedure for installing GPFS on Windows nodes

IBM provides GPFS as Windows Installer packages (MSI), which allow both interactive and unattended installations. Perform the GPFS installation steps as the Administrator or some other member of the Administrators group.

Before installing GPFS on Windows nodes, verify that all the installation prerequisites have been met. For more information, see [“Installing GPFS prerequisites”](#) on page 587.

To install GPFS on a Windows node, follow these steps:

1. Run the following license installation package from the product media and accept the license:

`gpfs.base-6.0.x.x-Windows-license.msi` (IBM Storage Scale Standard Edition or IBM Storage Scale Data Access Edition)

Rebooting is normally not required.

2. Download and install the latest service level of GPFS from the [IBM Support Portal: Downloads for General Parallel File System \(www.ibm.com/support/entry/portal/Downloads/Software/Cluster_software/General_Parallel_File_System\)](#). The GPFS package name includes the GPFS version (for example, `gpfs.base-6.0.x.x-Windows.msi`). For the latest information on the supported Windows operating system versions, see the [IBM Storage Scale FAQ](#) in IBM Documentation.

Under some circumstances, the installation process will prompt you to reboot the systems when it is required. You do not need to install the GPFS 6.0.x.x package included on the media before installing the latest update.

3. Download and install the latest level of IBM GSKit for GPFS. GSKit (`gpfs.gskit-x.x.x.x.msi`) comes as a single package that works with the IBM Storage Scale Standard Edition or IBM Storage Scale Data Access Edition.

Note: For this release, the IBM GSKit version must be at least 8.0.55.x or higher.

To upgrade GPFS, do the following:

1. Uninstall your current GPFS packages (for example, `gpfs.base-5.x.x.0-Windows.msi`, `gpfs.ext-5.x.x.0-Windows.msi`, and `gpfs.gskit-8.0.x.x.msi`).
2. Reboot.
3. Install the latest PTF packages (for example, `gpfs.base-6.0.x.x-Windows.msi`).

When upgrading, you do not need to uninstall and reinstall the license package unless you are explicitly instructed to do so by IBM. In addition, `gpfs.gskit` only needs to be upgraded if the update (zip file) contains a newer version.

For more information, refer to [Chapter 3, “Establishing and starting your IBM Storage Scale cluster,”](#) on page 467.

The GPFS installation package provides some options that you can select by installing GPFS from a Command Prompt. Property values control the options. The command line syntax is:

```
msiexec.exe /package <Package> [Optional Parameters] [Property=Value]
```

The following property values are available:

AgreeToLicense=yes

This property applies to the license package `gpfs.base-6.0.x.x-Windows-license.msi`. It allows GPFS to support an installation procedure that does not require user input. IBM recommends that you perform at least one interactive installation on a typical system before attempting an unattended installation. This will help identify any system configuration issues that could cause the installation to fail.

The following command installs GPFS without prompting you for input or generating any error dialog boxes:

```
msiexec /package gpfs.base-6.0.x.x-Windows-license.msi /passive AgreeToLicense=yes
```

The `msiexec.exe` executable file supports display options other than `/passive`. See `msiexec` documentation for details.

RemoteShell=no

This property applies to the product package (for example, `gpfs.base-6.0.x.x-Windows-license.msi`). It is equivalent to running `mmwinervctl set -remote-shell no`, but performs this configuration change before `mmwinerv` initially starts. This option is available to satisfy security policies that restrict network communication protocols.

You can verify that GPFS is installed correctly by opening the **Programs and Features** control panel.

IBM Storage Scale should be included in the list of installed programs. The program's version should match the version of the update package.

The GPFS software license agreement is shipped and is viewable electronically. The license agreement will remain available in the `%SystemDrive%\Program Files\IBM\GPFS\6.0.1.x\license` directory for future access.

Related concepts

[GPFS for Windows overview](#)

GPFS for Windows participates in a new or existing GPFS cluster in conjunction with AIX and Linux systems.

[Installing GPFS prerequisites](#)

This topic provides details on configuring Windows systems prior to installing GPFS.

[Running GPFS commands](#)

Once GPFS is installed, the account you use for GPFS administrative operations (such as creating a cluster and file systems) will depend on your cluster type.

[Configuring a mixed Windows and UNIX \(AIX or Linux\) cluster](#)

[Configuring the Windows HPC server](#)

In order for Windows HPC Server to function with GPFS, disable dynamic hosts file and re-enable dynamic DNS updates by doing the following:

Running GPFS commands

Once GPFS is installed, the account you use for GPFS administrative operations (such as creating a cluster and file systems) will depend on your cluster type.

For Windows clusters, run GPFS commands as a member of the **Domain Admins** group or as a domain account that is part of the local **Administrators** group. For clusters with both Windows and UNIX nodes, run GPFS commands as **root**, a special domain user account described in [“Creating the GPFS administrative account”](#) on page 593.

You must run all GPFS commands from either a **GPFS Admin Command Prompt** or a **GPFS Admin Korn Shell**. The GPFS administrative commands may not work properly if issued from the **Cygwin Terminal**. Open a new **GPFS Admin Command Prompt** or **GPFS Admin Korn Shell** after installing GPFS so that it uses the updated **PATH** environment variable required to execute GPFS commands.

GPFS Admin Command Prompt and **GPFS Admin Korn Shell** can be accessed using their desktop shortcuts or from under the **IBM GPFS Start** menu group. Each of these program links starts its respective shell using the **Run as Administrator** option.

Related concepts

[GPFS for Windows overview](#)

GPFS for Windows participates in a new or existing GPFS cluster in conjunction with AIX and Linux systems.

[Installing GPFS prerequisites](#)

This topic provides details on configuring Windows systems prior to installing GPFS.

Procedure for installing GPFS on Windows nodes

IBM provides GPFS as Windows Installer packages (MSI), which allow both interactive and unattended installations. Perform the GPFS installation steps as the Administrator or some other member of the Administrators group.

Configuring a mixed Windows and UNIX (AIX or Linux) cluster

Configuring the Windows HPC server

In order for Windows HPC Server to function with GPFS, disable dynamic hosts file and re-enable dynamic DNS updates by doing the following:

Configuring a mixed Windows and UNIX (AIX or Linux) cluster

For GPFS clusters that include both Windows and UNIX (AIX or Linux) nodes, this topic describes the additional configuration steps needed beyond those described in [“Installing GPFS prerequisites” on page 587](#).

For mixed clusters, perform the following steps:

1. Optionally, install and configure identity mapping on your Active Directory domain controller (see [“Identity Mapping for Unix \(IDMU\) / RFC 2307 Attributes” on page 592](#)).
2. Create the root administrative account (see [“Creating the GPFS administrative account” on page 593](#)).
3. Edit the Domain Group Policy to give root the right to log on as a service (see [“Allowing the GPFS administrative account to run as a service” on page 593](#)).
4. Configure the GPFS Administration service (mmwinser) to run as root (see [“Configuring the GPFS Administration service” on page 593](#)).
5. Install and configure OpenSSH (see [“Installing and configuring OpenSSH on Windows nodes” on page 594](#)).

Complete this process before performing configuration steps common to all GPFS supported platforms.

Identity Mapping for Unix (IDMU) / RFC 2307 Attributes

GPFS can exploit a Windows Server feature called RFC 2307 attributes to provide consistent identities among all nodes in a cluster.

GPFS expects that all Windows nodes in a cluster are members of the same Active Directory domain. This gives domain users a consistent identity and consistent file access rights independent of the system they are using.

GPFS can exploit the RFC 2307 attributes for mapping users and groups between Windows and UNIX. These attributes can be administered by using Identity Mapping for Unix (IDMU) from Microsoft in Windows Server versions up to and including Windows Server 2012 R2. Beginning with Windows Server 2016 (since the IDMU MMC-snap-in has been removed) you can specify these RFC 2307 attributes by using the **Active Directory Users and Computers MMC Snap-in** or the **Active Directory Administrative Center**. For more information, see *Configuring ID mappings in Active Directory Users and Computers for Active Directory Users and Computers (ADUC)* for instructions on editing RFC 2307 attributes in *IBM Storage Scale: Administration Guide*.

The only way to achieve Windows-Unix user-mapping in GPFS is with RFC 2307 attributes. These attributes can be administered by using Identity Mapping for Unix (IMU) from Microsoft in Windows Server versions up to and including Windows Server 2012 R2. Beginning Windows Server 2016, these RFC 2307 attributes can be specified by using the Active Directory Users and Computers (ADUC) MMC Snap-in.

For more information, see *Identity management on Windows / RFC 2307 Attributes* instructions on editing IDMU or RFC 2307 attributes in *IBM Storage Scale: Administration Guide*.

GPFS uses RFC 2307 attributes, specifically `uidNumber`, to map users and `gidNumber` to map groups between Windows and Unix. From the perspective of GPFS, IDMU is synonymous with RFC 2307

attributes. This means that any references to IDMU in the documentation should be interpreted as RFC 2307 attributes.

For IDMU or ADUC installation and configuration information, see *Configuring ID mappings in Active Directory Users and Computers for Active Directory Users and Computers (ADUC) in IBM Storage Scale: Administration Guide*.

Creating the GPFS administrative account

GPFS uses an administrative account in the Active Directory domain named **root** in order to interoperate with UNIX nodes in the cluster. Create this administrative account as follows:

1. Create a domain user with the logon name **root**.
2. Add user **root** to the **Domain Admins** group or to the local Administrators group on each Windows node.
3. In **root** Properties/Profile/Home/LocalPath, define a HOME directory such as `C:\Users\root\home` that does not include spaces in the path name and is not the same as the profile path.
4. Give **root** the right to log on as a service as described in [“Allowing the GPFS administrative account to run as a service”](#) on page 593.

Step “3” on page 593 is required for the Cygwin environment (described in [“Installing Cygwin”](#) on page 588) to operate correctly. Avoid using a path that contains a space character in any of the path names. Also avoid using **root**'s profile path (for example, `C:\User\root`). OpenSSH requires specific permissions on this directory, which can interfere with some Windows applications.

You may need to create the HOME directory on each node in your GPFS cluster. Make sure that **root** owns this directory.

Allowing the GPFS administrative account to run as a service

Clusters that depend on a **root** account to interoperate with UNIX nodes in a cluster will need to configure the GPFS Administrative Service (**mmwinser**) to run as the **root** account. For this, **root** needs to be assigned the right to log on as a service. See [“Configuring the GPFS Administration service”](#) on page 593 for details.

The right to log on as a service is controlled by the Local Security Policy of each Windows node. You must set the Local Security Policy on all Windows nodes in a GPFS cluster.

For many Windows servers, this setting is enabled. The following procedure allows checking and assigning for the *log on as a service* right to a local account for a system that is running Windows Server 2022:

1. Open **Open Local Security Policy** (`secpol.msc`).
2. In the console tree, expand **Local Policies/User Rights Assignment**.
3. Double-click **Log on as a service**.
4. Use **Add User or Group** to include an **Administrator** and **NT SERVICE\ALL SERVICES**, and then click **OK**.

Refer to your *Windows Server* documentation for a full explanation of Local Security Policy.

Configuring the GPFS Administration service

GPFS for Windows includes a service called **mmwinser**. In the Windows Services management console, this service has the name GPFS Administration. **mmwinser** supports GPFS operations such as **autoload** and remote command execution in Windows GPFS clusters. The Linux and AIX versions of GPFS do not have a similar component. The **mmwinser** service is used on all Windows nodes starting with GPFS 3.3.

The GPFS installation package configures **mmwinser** to run using the default **LocalSystem** account. This account supports Windows GPFS clusters. For clusters that include both Windows and UNIX nodes, you must configure **mmwinser** to run as **root**, the GPFS administrative account. Unlike **LocalSystem**,

root can access the IDMU service and can access other GPFS nodes as required by some cluster configurations.

For IDMU installation and configuration information, see *Identity management on Windows* in *IBM Storage Scale: Administration Guide*. For information on supporting administrative access to GPFS nodes, see the *Requirements for administering a GPFS file system* topic in the *IBM Storage Scale: Administration Guide*.

Before configuring `mmwinssd` to run as **root**, you must first grant **root** the right to run as a service. For details, see “Allowing the GPFS administrative account to run as a service” on page 593.

Use the GPFS command `mmwinssdctl` to set and maintain the GPFS Administration service's log on account. `mmwinssdctl` must be run on a Windows node. You can run `mmwinssdctl` to set the service account before adding Windows nodes to a cluster. You can also use this command to change or update the account on nodes that are already in a cluster. GPFS can be running or stopped when executing `mmwinssdctl`, however, refrain from running other GPFS administrative commands at the same time.

In this example, `mmwinssdctl` configures three nodes before they are added to a GPFS cluster containing both Windows and UNIX:

```
mmwinssdctl set -N node1,node2,node3 --account mydomain/root --password mypwd --remote-shell no
```

Whenever `root`'s password changes, the `mmwinssd` logon information needs to be updated to use the new password. The following command updates on all Windows nodes in a cluster with a new password:

```
mmwinssdctl set -N all --password mynewpwd
```

As long as `mmwinssd` is running, the service will not be affected by an expired or changed password and GPFS will continue to function normally. However, GPFS will not start after a system reboot when `mmwinssd` is configured with an invalid password. If for any reason the Windows domain or `root` password changes, then `mmwinssdctl` should be used to update the domain and password. The domain and password can also be updated on a per node basis by choosing Administrative Tools > Computer Management > Services and Applications > Services, and selecting GPFS Administration. Choose File > Properties > Logon and update the `<domain>\username` and the password.

For more information, see **`mmwinssdctl` command** in *IBM Storage Scale: Command and Programming Reference Guide*.

Installing and configuring OpenSSH on Windows nodes

If using a mixed cluster, OpenSSH must be configured on the Windows nodes. Refer to the [Cygwin FAQ \(www.cygwin.com/faq.html\)](http://www.cygwin.com/faq.html) and documentation on how to setup `sshd`. Replace the usage of the account `cyg_server` in the Cygwin documentation with `root` when setting up a privileged account for `sshd`.

The following are some guidelines in addition to the Cygwin instructions on setting up `sshd`:

1. Verify that all nodes can be pinged among themselves by host name, Fully Qualified Domain Name (FQDN) and IP address.
2. If not using IPv6, disable it. For more information, see [How to disable IPv6 or its components in Windows \(support.microsoft.com/kb/929852\)](http://support.microsoft.com/kb/929852).
3. Check that `passwd` contains the privileged user that you plan to use for GPFS operations, as well as its correct home path:

```
cat /etc/passwd | grep "root"
```

```
root:unused:11103:10513:U-WINGPFS\root,S-1-5-21-3330551852-1995197583-3793546845-1103:/cygdrive/c/home/  
root:/bin/bash
```

If the user is not listed, rebuild your `passwd`:

```
mkpasswd -l -d wingpfs > /etc/passwd
```


4. From the Cygwin shell, run `/usr/bin/ssh-host-config` and respond yes to the prompts. When prompted to enter the value of CYGWIN for the daemon, enter `ntsec`. Specify `root` in response to the query for the new user name. You may receive the following warning:

```
***Warning: The specified account 'root' does not have the
***Warning: required permissions or group memberships. This may
***Warning: cause problems if not corrected; continuing...
```

As long as the account (in this case, `root`) is in the local Administrators group, you can ignore this warning.

5. When the installation is complete, enter the following:

```
net start sshd

The CYGWIN sshd service is starting.
The CYGWIN sshd service was started successfully.
```

Note: The OpenSSH READMEs are available at `/usr/share/doc/openssh`. Also see the [IBM Storage Scale FAQ](#) in IBM Documentation.

Once OpenSSH is installed, the GPFS administrative account `root` needs to be configured so that it can issue `ssh` and `scp` commands without requiring a password and without producing any extraneous messages. This kind of passwordless access is required from any node used for GPFS administration to all other nodes in the cluster.

For additional information, see *Requirements for administering a GPFS file system* in *IBM Storage Scale: Administration Guide* and *Troubleshooting Windows problems* in *IBM Storage Scale: Problem Determination Guide*.

Related concepts

[GPFS for Windows overview](#)

GPFS for Windows participates in a new or existing GPFS cluster in conjunction with AIX and Linux systems.

[Installing GPFS prerequisites](#)

This topic provides details on configuring Windows systems prior to installing GPFS.

[Procedure for installing GPFS on Windows nodes](#)

IBM provides GPFS as Windows Installer packages (MSI), which allow both interactive and unattended installations. Perform the GPFS installation steps as the Administrator or some other member of the Administrators group.

[Running GPFS commands](#)

Once GPFS is installed, the account you use for GPFS administrative operations (such as creating a cluster and file systems) will depend on your cluster type.

[Configuring the Windows HPC server](#)

In order for Windows HPC Server to function with GPFS, disable dynamic hosts file and re-enable dynamic DNS updates by doing the following:

Configuring the Windows HPC server

In order for Windows HPC Server to function with GPFS, disable dynamic hosts file and re-enable dynamic DNS updates by doing the following:

1. On all nodes in the Windows HPC cluster, open the hosts file `%systemroot%\system32\drivers\etc\hosts` and change `ManageFile = true` to `ManageFile = false`.
2. On the HPC head node, execute the following from HPC PowerShell, assuming all nodes are part of the head node's Active Directory domain.

Go to Start > All Programs > Microsoft HPC Pack > HPC PowerShell and execute:

```
Set-HpcNetwork -EnterpriseDnsRegistrationType WithConnectionDnsSuffix
```

Related concepts

[GPFS for Windows overview](#)

GPFS for Windows participates in a new or existing GPFS cluster in conjunction with AIX and Linux systems.

[Installing GPFS prerequisites](#)

This topic provides details on configuring Windows systems prior to installing GPFS.

[Procedure for installing GPFS on Windows nodes](#)

IBM provides GPFS as Windows Installer packages (MSI), which allow both interactive and unattended installations. Perform the GPFS installation steps as the Administrator or some other member of the Administrators group.

[Running GPFS commands](#)

Once GPFS is installed, the account you use for GPFS administrative operations (such as creating a cluster and file systems) will depend on your cluster type.

[Configuring a mixed Windows and UNIX \(AIX or Linux\) cluster](#)

Chapter 7. Installing IBM Storage Scale on public cloud by using cloudkit

Use this information for installing IBM Storage Scale on public cloud with cloudkit.

Note: Currently, Amazon Web Services (AWS), Google Cloud Platform (GCP), and Microsoft Azure are supported for deployment on public clouds.

You can install IBM Storage Scale on public cloud by using the cloudkit. You can also use the cloudkit for tasks such as cloud resource provisioning, IBM Storage Scale install and configuration.

Overview of the cloudkit

The cloudkit is a command line interface that provides automation of deployment, upgrade, and configuration of IBM Storage Scale on public clouds.

For a list of features available in the cloudkit, see [Table 1](#). Cloudkit uses the underlying configured cloud account, and dynamically fetches the information that is required for provisioning IBM Storage Scale resources on the public cloud. Thereafter, the cloudkit installs, configures, and deploys the specified configuration.

The cloudkit enables the following tasks:

- Dynamically provision required resources, install, configure, and deploy IBM Storage Scale into either a newly created Virtual Private Cloud environment or an existing one.
- Allow the creation of virtual machine images, network, DNS, and jump host that can be used to rapidly deploy IBM Storage Scale on public cloud.
- Ability to create storage only, compute only and combined clusters (compute and storage cluster).
- Perform a remote mount of a filesystem from one IBM Storage Scale storage cluster to one or multiple compute clusters.
- Performs storage cluster prerequisites needed by IBM Storage Scale Container Native Access.
- Revoke a previously remote mounted filesystem.
- Enabling IBM Storage Scale active file management (AFM) to cloud object storage on AWS, GCP, and Microsoft Azure.
- Destroy a previously deployed IBM Storage Scale cluster.
- Ability to scale-out IBM Storage Scale cluster.
- Ability to upgrade IBM Storage Scale cluster on AWS, GCP, or Microsoft Azure.

Installation and configuration are driven through commands that can be run either interactively or by using command line options.

Cloudkit offers a flexible, composable deployment architecture:

- Allows administrators to pick and choose the resource creation through cloudkit; or to use already created infrastructure resources.
- Offers flexibility in terms of components like network, DNS, network access methods, disk types, instance types.
- Administrators can use cloudkit also to create isolated resources required for IBM Storage Scale workload end-to-end, instead of reusing resources from existing infrastructure.

The cloudkit operation consists of three main layers, the command layer, infrastructure provisioning layer, and IBM Storage Scale Ansible configuration layer:

- The command layer is the front-end interface for cloudkit users. This layer provides a Command Line Interface (CLI) that allows users to deploy and manage applications with ease. The CLI provides a

simple interactive query and filters syntax for user input. The command layer is designed to be user friendly and intuitive, making it easy for users to get started with cloudkit.

- The infrastructure provisioning layer is responsible for creating and managing the underlying cloud infrastructure needed to support the IBM Storage Scale need. This layer includes terraform templates for creating and managing virtual machines, storage resources, security rules, and networking components.
- The Ansible configuration layer is responsible for configuring IBM Storage Scale by installing the required RPMS and building the cluster on the cloud infrastructure.

Overview of the cloudkit installation options

This topic covers comprehensive command options available for deploying and managing IBM Storage Scale cluster on public cloud.

The cloudkit provides an interactive experience guiding the user through its prompts, the list of commands outlined below are the starting points. Use these commands to start the interaction with the cloudkit.

Preparation

The cloudkit needs to be installed on a Linux-based host before it can be used for an IBM Storage Scale deployment on public cloud. Such Linux-based host is referred to as *installer node*. For information about setting up an installer node, see [“Preparing the installer node” on page 421](#). After the **cloudkit** setup is complete, log in to the installer node.

The **cloudkit** binary is found at the `/usr/lpp/mmfs/<release_version>/cloudkit` directory. In this directory, the IBM Storage Scale cloudkit can be invoked through the **cloudkit** command. Optionally, this directory can be added to the path.

Before attempting to create an IBM Storage Scale cluster on a public cloud, the cloudkit must be configured as described in the next sections.

Initialization

1. Use the `cloudkit init` command to install the prerequisites needed for the utility.

To configure, run the `cloudkit init` command:

```
./cloudkit init
I: Logging at /root/scale-cloudkit/logs/cloudkit-25-11-2024_0-11-59.log
? Passphrase file path for encrypting DB contents: /root/secrets/cloudkit_config.ini
```

The `passPhrase` file need to pass during the **init** command run. For more information, see [Preparing the cloudkit environment file](#).

Note: When a new version of IBM Storage Scale data bundle is downloaded from [IBM Fix Central](#) and extracted to a node, it is mandatory to rerun the **cloudkit init** command even if the command was previously run for a different version of IBM Storage Scale.

2. Use the `cloudkit configure` command to configure local machine to use your cloud account. For more information, see *Configuring the cloudkit* in the *IBM Storage Scale: Administration Guide*.

Deployment

Before deploying IBM Storage Scale on a public cloud, make sure to complete the procedures described in [“Initialization” on page 598](#).

To understand the deployment option provided by the cloudkit, you need to know the way **cloudkit** deploys IBM Storage Scale on a cloud and the stages it goes through:

1. Cloudkit uploads require a GPFS binary to cloud repository.

- Use the `cloudkit create repository` command to optionally create a package repository on the cloud object store.
- 2. Cloudkit prepares the cloud operating system image based on a cloud repository.
 - Use the `cloudkit create image` command to optionally create a virtual machine image containing all IBM Storage Scale packages preinstalled.
- 3. Cloudkit creates a virtual private network that is later used for the deployment of an IBM Storage Scale storage cluster.
 - Use the `cloudkit create network` command to optionally create a virtual private network.
- 4. Cloudkit creates or associates a domain name system (DNS) to facilitate hostname resolution.
 - Use the `cloudkit create dns` command to optionally create a DNS domain.
- 5. Cloudkit creates a jump host or bastion host by using the previously created virtual private network.
 - Use the `cloudkit create jumphost` command to optionally create a jump host.
- 6. Cloudkit deploys an IBM Storage Scale cluster using the previously created operating system image.
 - Use the `cloudkit create cluster` command to create an IBM Storage Scale cluster. This command can be used to create an IBM Storage Scale storage, compute or combined cluster.

To help you plan your requirement deployment architecture, refer to [“Planning the virtual private cloud \(VPC\) architecture for AWS” on page 424](#) and [“Planning the virtual private cloud \(VPC\) architecture for GCP” on page 430](#).

Administering

The cloudkit can be used to manage a previously deployed cloudkit cluster using the following options.

1. Use the `cloudkit grant filesystem` command to remote mount a filesystem from a storage cluster to a compute cluster previously created by the same instance of cloudkit.
2. Use the `cloudkit grant repository` command to provide access to a package repository located on the cloud object store to a specific Virtual Private Cloud.
3. Use the `cloudkit port-forward` command to provide access to the IBM Storage Scale GUI through a jump host.
4. Use the `cloudkit revoke filesystem` command to remove a previous remote mount configuration.
5. Use the `cloudkit revoke repository` command to remove the access from a virtual private cloud to a repository.
6. Use the `cloudkit edit cluster` command to scale out cluster resources.
7. Use the `cloudkit caching setup` command to set up an AFM relationship from a local scale cluster to a remote cluster or cloud object storage.

For more information, see *Administering cloudkit* in the *IBM Storage Scale: Administration Guide*.

To see an end-to-end process of using interactive command, see the [demo](#).

Upgrade

The cloudkit can be used to upgrade existing package repository and an IBM Storage Scale cluster using the following options:

1. Use `cloudkit upgrade repository` command to upgrade the existing repository to specified cloudkit version.
2. Use `cloudkit upgrade cluster` command to upgrade the existing cluster to specified cloudkit version.

Note: Upgrade of IBM Storage Scale cluster is only supported on AWS and GCP.

For more information, see [“Upgrading IBM Storage Scale on cloud” on page 688](#).

Cleanup

The cloudkit can be used to delete the resources which we provisioned:

1. Use the `cloudkit delete cluster` command to delete the cluster.
2. Use the `cloudkit delete jumphost` command to delete the jump host.
3. Use the `cloudkit delete dns` command to delete the DNS domain.
4. Use the `cloudkit delete network` command to delete the virtual private cloud or virtual network.
5. Use the `cloudkit delete image` command to delete the image.
6. Use the `cloudkit delete repo` command to delete the repository.
7. Use the `cloudkit delete caching-target` command to delete the caching target.

Note: Cloudkit keeps track of resources created using it. When the 'cluster with a new vpc' is created by cloudkit, make sure this VPC does not contain any active resources before proceeding with deletion of cluster. As this cluster stack contains VPC resources and if there are other resources created beyond cloudkit using this VPC resources could block the cluster deletion.

In scenarios of cluster with jumphost created via cloudkit, it will be deleted as part of cluster deletion operation. If this jumphost is being used by other clusters, their access might be impacted. Hence it is advised to verify the usage of jumphost before proceeding with deletion.

The following table lists the command options to perform cloud resource provisioning, IBM Storage Scale install and configuration.

Table 48. cloudkit command options	
cloudkit command option	Purpose
configure	Configure local machine to use your cloud account
create	Create a resource from stdin
delete	Delete a specific resource
describe	Show details of a specific resource
edit	Edit a specific resource
grant	Grant access to a specific resource
help	Help about any command
init	Installs prerequisite(s) required for the utility
list	List a resource from stdin
port-forward	Redirects the IBM Storage Scale GUI access through a jump host
revoke	Revoke filesystem mount access
upgrade	Upgrade a resource from stdin
validate	Validate resources
setup	Set up features related to IBM Storage Scale
version	Prints the version number of the tool

Other Considerations

Compute cluster with bastion:

```
-1      icmp      Allow ICMP traffic from bastion to compute instances
```

22	TCP	"Allow SSH traffic from bastion to compute instances"
-1	icmp	"Allow ICMP traffic within compute instances"
22	TCP	"Allow SSH traffic within compute instances"
1191	TCP	"Allow GPFS intra cluster traffic within compute instances"
60000-61000	TCP	"Allow GPFS ephemeral port range within compute instances"
47080	TCP	"Allow management GUI (http/localhost) TCP traffic within compute instances"
47443	UDP	"Allow management GUI (https/localhost) TCP traffic within compute instances"
4444	TCP	"Allow management GUI (https/localhost) TCP traffic within compute instances"
4739	TCP	"Allow management GUI (localhost) TCP traffic within compute instances"
4739	"UDP"	"Allow management GUI (localhost) UDP traffic within compute instances"
9080	TCP	"Allow performance monitoring collector traffic within compute instances"
9081	TCP	"Allow performance monitoring collector traffic within compute instances"
80	TCP	"Allow http traffic within compute instances"
443	TCP	"Allow https traffic within compute instances"
443	TCP	"Allow GUI traffic from bastion/jumphost"

Note: "Allow ICMP traffic from bastion to compute instances" and "Allow SSH traffic from bastion to compute instances" are not added if direct connect is used.

Storage cluster with bastion:

-1	icmp	Allow ICMP traffic from bastion to storage instances
22	TCP	"Allow SSH traffic from bastion to storage instances"
-1	icmp	"Allow ICMP traffic within storage instances"
22	TCP	"Allow SSH traffic within storage instances"
1191	TCP	"Allow GPFS intra cluster traffic within storage instances"
60000-61000	TCP	"Allow GPFS ephemeral port range within storage instances"
47080	TCP	"Allow management GUI (http/localhost) TCP traffic within storage instances"
47443	UDP	"Allow management GUI (https/localhost) TCP traffic within storage instances"
4444	TCP	"Allow management GUI (https/localhost) TCP traffic within storage instances"
4739	TCP	"Allow management GUI (localhost) TCP traffic within storage instances"
4739	UDP	"Allow management GUI (localhost) UDP traffic within storage instances"
9080	TCP	"Allow performance monitoring collector traffic within storage instances"
9081	TCP	"Allow performance monitoring collector traffic within storage instances"
80	TCP	"Allow http traffic within storage instances"
443	TCP	"Allow https traffic within storage instances"
443	TCP	"Allow GUI traffic from bastion/jumphost"

Note: "Allow ICMP traffic from bastion to storage instances" and "Allow SSH traffic from bastion to storage instances" are not added if direct connect is used.

Compute cluster with remote mount:

```
-1, ICMP, Allow ICMP traffic from spectrum scale cluster
1191, TCP, Allow GPFS intra cluster traffic from spectrum scale cluster
443, TCP, Allow management GUI (http/localhost) TCP traffic from spectrum scale cluster
60000-61000, TCP, Allow spectrum scale ephemeral port range
```

Deploying cloudkit

To ensure the proper functioning of **cloudkit** for IBM Storage Scale, consider the capabilities and characteristics of the deployment options of this command-line interface.

Create repository (**cloudkit create repository**)

Creates the repository that holds the IBM Storage Scale rpms and debs that are needed to install IBM Storage Scale packages.

- Uploads made through **cloudkit** require a GPFS binary to cloud repository.
- Offers the option to create a package repository on the cloud object store.

Create network (**cloudkit create network**)

Creates the network and subnets that are necessary for the cluster creation.

- Creates the network resources that are suited for IBM Storage Scale deployment.
- Creates the following network resources:
 - Private subnet for storage
 - Private subnet for compute
 - Public subnet for jump host

- Private subnet for protocol nodes
- The subnet range (CIDR) is calculated based on the number of IPs that are requested.

Create images(`cloudkit create image`)

Prepares an IBM Storage Scale image with the required IBM Storage Scale `xpms` that can be later used for deployment, which significantly reduces the deployment time.

- Prepares the image of the cloud operating system based on a cloud repository.
- Offers the option to create a virtual machine image that contains all IBM Storage Scale packages preinstalled.

Create domain name system(`cloudkit create dns`)

Creates or associates the privately hosted domain name system (DNS) to the targeted virtual private network.

- Creates forward DNS zone
- Creates reverse DNS zone

Create jump host(`cloudkit create jumphost`)

Creates a jump host or bastion to use it to access the cluster nodes.

- Facilitates access to the private mode of an on-premises IBM Storage Scale cluster.
- Creates a jump host that can be used for cluster access or administration.
- This configuration is optional when you are using a cloud virtual machine or a virtual machine is within the cloud itself.

Create cluster(`cloudkit create cluster`)

Creates the cluster with specified input values.

- Cloudkit deploys an IBM Storage Scale cluster using the previously created operating system image.
- Creates an IBM Storage Scale cluster. The cluster that gets created can be a storage cluster, a compute cluster, or a combined cluster.

Limitations of cloudkit

To ensure the proper functioning of cloudkit, consider its limitations when using this command line interface for IBM Storage Scale.

The limitations of cloudkit are:

- Deleting a cluster may lead to failure if there are other resources that reference the resources created by cloudkit, or if there are additional resources sharing the network resources created by cloudkit that are not recognized by cloudkit.
- To ensure the secure usage of the bastion private key, cloudkit fails to accept if the permissions of bastion private key is other than 400 or 600.
- Cloudkit stores its states, temporary files, and logs under `$HOME/scale-cloudkit`, it is user's responsibility to replicate and keep it secure. If this metadata is lost, it cannot be reconstructed. However, loss of this metadata does not delete the resources on the cloud.
- As part of a deployment made by using cloudkit, one namespace (file system) is configured per `Storage-Only` or `Combined-compute-storage`.
- If any of the ebs volumes is detached and reattached after the initial cluster creation, these ebs volumes must be manually deleted after the cluster is destroyed.
- Run `grant` from the same cloudkit version which is used to deploy the other resources like compute and storage cluster.
- The cloudkit is currently limited to be executed as root user. Attempting to execute cloudkit as non root user, even with `sudo` privileges will result in failure.
- While running an upgrade by using the **`cloudkit upgrade`** command, the interactive mode displays the possible upgrade versions for the targeted cluster.

- Ensure to have the file system mounted after you perform edit operation. Use **mmmount all -a** command.
- Edit operation on the cluster (that is, adding new nodes to an existing cluster) lacks automation for quorum node increase as per the node count change. Administrator must manually assign quorum roles to nodes based on the requirement.
- Addition of new storage nodes creates new NSDs and adds them to the file system. This characteristic results in uneven distribution of data. The administrator needs to manually run the **mmrestripefs -b** command to distribute data to new disks
- In the installer method based on AWS cloud VM, the installer VM must have its own security group, which must be different from the default security group that is created as part of VPC.
- Cloud accounts come with a default network (VPC) that has a preconfigured subnet and security rules. Using VPC when performing an IBM Storage Scale cluster installation is not recommended. If you are using a new account, it is recommended to use the **cloudkit create network** command to create a network that works properly with IBM Storage Scale.
- For IBM Storage Scale Developer Edition, the size for file system creation is limited to 12 TB per cluster.
- Before using an instance type, you must ensure that it is available in the chosen availability zone across the respective cloud providers (AWS, GCP, Microsoft Azure).
- The following Microsoft Azure resources cannot be filtered or viewable for reuse if they are created in resource groups that are different to the resource group that is operated by the cloudkit: DNS, key vault, encryption set controllers, images, network, and storage accounts.
- If the file system version of a cluster was updated by using the **mmchfs -V** command, editing that upgraded cluster may fail.
- IBM Storage Scale 5.2.2.0 introduces a new metadata field, which is called Deployment Purpose and offers two values to select from: Production and Non-Production. For clusters that are upgraded from previous releases to 5.2.2.0, this metadata field is automatically set as Production.
- For GCP deployments, if the selected disk type is **hyperdisk-throughput**, the minimum disk size must be 2 TB.

Supported features of cloudkit

Know the different public clouds, operating system versions, IBM Storage Scale features, and cloudkit commands that are supported.

Prerequisites

- Access to Red Hat base repository, a repository created by using the local ISO image.
- Ability to read/write into the user home directory.

Table 49. Supported public clouds			
Public cloud	Support	Limitations	Description
Amazon Web Services (AWS)	Yes	Red Hat Enterprise Linux only	Limited stock options of RHEL versions are available through the cloudkit. Custom images that contain any RHEL version that is supported by base IBM Storage Scale can be pointed through the cloudkit.

Table 49. Supported public clouds (continued)

Public cloud	Support	Limitations	Description
Google Cloud Platform (GCP)	Yes	Red Hat Enterprise Linux only	Limited stock options of RHEL versions are available through the cloudkit. Custom images that contain any RHEL version that is supported by base IBM Storage Scale can be pointed through the cloudkit.
IBM Cloud	Yes	Red Hat Enterprise Linux only	Limited stock options of RHEL versions are available through the cloudkit. Custom images that contain any RHEL version that is supported by base IBM Storage Scale can be pointed through the cloudkit.
Microsoft Azure	Yes	Red Hat Enterprise Linux only	Limited stock options of RHEL versions are available through the cloudkit. Custom images that contain any RHEL version that is supported by base IBM Storage Scale can be pointed through the cloudkit.

Note: IBM's statements about its plans, directions, and intent are subject to change or withdrawal without notice at IBM's sole discretion. The provided information is intended to outline IBM's general product direction and it should not be relied on in making a purchasing decision.

Table 50. Operating systems that the cloudkit installer nodes support

Supported operating system	Version	Limitations	Additional information
Red Hat Enterprise Linux (RHEL)	8.8, 9.0, 9.1, 9.2, 9.3, 9.4, 9.5 and 9.6	N/A	On-premises physical host or VM only.
Red Hat Enterprise Linux (RHEL)	9.6	N/A	Cloud VM only.

Table 51. IBM Storage Scale features

IBM Storage Scale function	Automatic configuration through cloudkit	Manual operation after cloudkit usage
Storage cluster creation option, compute cluster creation option, combined storage and compute cluster creation option	Yes	-
Remote mount storage cluster to client cluster	Yes	-
Remote mount from CNSA to cloudkit created cluster	Prerequisites configured automatically	Yes, through IBM Storage Scale Container Native Storage installation
Remote mount from Fusion to cloudkit created storage cluster	Prerequisites configured automatically	Yes, through Fusion installation
Data Management Edition license application	Yes	-
Data Access Edition license application	Yes	-
Developer Edition license application	Yes	-
Erasur Code Edition license application	No (cloudkit not available for ECE)	Not supported
Call home	Not supported	Yes, through installation toolkit or mm command CLI, if wanted
Performance Data collection	Yes	-
Storage cluster encryption	Yes (Encryption at Rest)	Yes, through mm command CLI and manually created key server that IBM Storage Scale supports
Quotas	-	Yes, through mm command CLI
ACLs	-	Yes, through mm command CLI
ILM	No (single pool is configured)	Not supported as only a single pool is configured
File clones	-	Yes, through mm command CLI
Snapshots	-	Yes, through mm command CLI
CSI volume snapshot	-	Yes, through CSI
GUI	Yes	-
REST-API server	Yes (as part of the GUI)	-
CSI	-	Yes, with manual setup through CSI documentation, prerequisites of open source Kubernetes (AKS, EKS, GKE are not supported), prerequisites of IBM Storage Scale CSI, and understanding of CSI networking requirements.

Table 51. IBM Storage Scale features (continued)

IBM Storage Scale function	Automatic configuration through cloudkit	Manual operation after cloudkit usage
File Audit Logging	-	Yes, through mm command CLI
Clustered watch folder	-	Yes, through mm command CLI
CES protocols	Yes	Yes, on AWS, GCP, and Microsoft Azure
AFM	Yes	Yes, with manual setup as AFM cache through the cloud site
AFM-COS	Yes	Yes, on AWS, GCP, and Microsoft Azure
AFM-DR	Yes	Yes, through mm command CLI
AWS protocol	Yes	-
Microsoft Azure protocol	Yes	-
Compression	-	Not supported
GPU Direct storage	-	Not supported
Scale up (nodes)	No	Not supported
Scale up CNSA compute cluster	-	Supported through IBM Storage Scale Container Native Storage independent of storage cluster and cloudkit
Scale up (file system capacity)	No	Not supported
Scale down (nodes)	No	Not supported
Scale down (file system capacity)	No	Not supported
Scale out (adding new nodes)	Yes	-

Table 52. Feature matrix for cloudkit

Usage	Feature	Support	Limitations	Additional information
cloudkit create repository	IBM Storage Scale or ZIMon RPMs static web-hosting, creates one of the following items: • AWS S3 bucket • GCP artifact registry • Azure blob • IBM Cloud protocols bucket	Yes	Repository scope is limited to the target region.	

Table 52. Feature matrix for cloudkit (continued)

Usage	Feature	Support	Limitations	Additional information
cloudkit create image	Creates an IBM Storage Scale image by using a repository	Yes		As the base image, you can use a stock or custom image of the AMI or GCP operating system, or a Microsoft Azure image or an IBM Cloud image.
cloudkit create network	Creates a virtual private cloud or a virtual network	Yes		
cloudkit create dns	Creates a DNS domain	Yes		
cloudkit create cluster	Creates an IBM Storage Scale cluster	Yes		Can create storage only cluster, compute only cluster, and combined storage and compute cluster.
cloudkit grant filesystem	Remote mount file system from storage to compute cluster	Yes		One remote mount relation per compute.
cloudkit edit cluster	Scales out an IBM Storage Scale cloud cluster	Yes		
cloudkit upgrade cluster	Updates an IBM Storage Scale cluster	Yes		Only AWS, GCP, and Microsoft Azure are supported.
cloudkit setup caching	Configures the AFM relationship from a local IBM Storage Scale cluster to a remote cluster or cloud object storage	Yes		Only AWS, GCP, and Microsoft Azure are supported.
cloudkit setup protocol	Enables NFS, SMB, and S3 protocols on an IBM Storage Scale cloud cluster.	Yes		Only AWS, GCP, and Microsoft Azure are supported.

Accessing the cluster on the cloud

Accessing nodes within the cluster

To access a node through a bastion host using SSH, you can use the following command:

```
ssh -i <path_to_key> <user>@<destination_host> -o "proxycommand ssh -W %h:%p -i <path_to_key> <bastion-user>@<Bastion_host>"
```

Replace `user` with your username, `bastion-user` with bastion host username, `bastion-host` with the hostname or IP address of your bastion host, and `destination-host` with the hostname or IP address of the node you want to access.

This command uses the `-J` flag to specify the bastion host. It sets up a SSH jump connection through the bastion host to the destination host, allowing you to securely access the destination node even if it is behind a firewall or other security measures.

Accessing the IBM Storage Scale management GUI

After the output for a `cloudkit create cluster` command run, the following GUI access details are displayed.

```
I: IBM Storage Scale GUI:
    https://XX.XXX.XXX.XXX:3124/

I: To use this cluster via REST interface, you can run (an example
API):
    curl -u '<gui_username>:<gui_password>' -X GET -k https://XX.XXX.XXX.XXX:3124/
    scalemgmt/v2/info

I: IBM Storage Scale cluster 'scale-cls' has been created.
```

Deploying IBM Storage Scale on AWS

Follow these steps to deploy IBM Storage Scale on AWS:

1. Prepare your AWS account:
 - a. Create an AWS account at [AWS](#), if you do not have one.
 - b. Use the region selector in the navigation bar to choose the AWS region where you want to deploy IBM Storage Scale on AWS.
 - c. Create a key pair in your preferred region. For more information, see [Amazon EC2 Key Pairs](#)
 - d. Verify the desired AWS Region and Availability Zone supports the EC2 instance types that are chosen to be provisioned for the IBM Storage Scale storage and compute nodes. You can verify the AWS EC2 instance type from **AWS console > EC2 Service Page > Instance Type**.
 - e. Verify the AWS EC2 quota limits, and ensure that enough quotas exist for the EC2 instance types that you intend to deploy. You can verify the AWS EC2 quota limits from **AWS console > EC2 Service Page > Limits**.

If necessary, [request a service limit increase](#) for the EC2 instance types that you intend to deploy. To request a service limit increase, in the AWS Support Center, choose **Create Case > Service Limit Increase > EC2 instances**, and then complete the fields in the **Limit Increase** form.



Attention: The IBM Storage Scale deployment on AWS uses CloudFormation templates. These CloudFormation templates rely on the creation and use of IAM Roles to facilitate seamless access to required AWS resources. Therefore, it is important to ensure that while launching the CloudFormation template, the user has sufficient permissions to create and manage IAM Roles.

Note: If you intend to provision a cluster with many storage and compute nodes, it is recommended to increase the AWS Systems Manager's Parameter Store throughput limit. You can change

the throughput limit from **AWS Console > Systems Manager > Parameter store > Setting > Parameter Store throughput**.

2. Launch the IBM Storage Scale AWS stack:

- a. Choose one of the following options to launch the AWS CloudFormation template into your AWS account.

- [Option 1: Deploying IBM Spectrum Scale on a new Amazon VPC with a single availability zone](#)
- [Option 2: Deploying IBM Spectrum Scale on a new Amazon VPC with multiple availability zones](#)
- [Option 3: Deploying IBM Spectrum Scale on an existing Amazon VPC.](#)

For a single availability zone deployment configuration, IBM Storage Scale does not replicate any of the file system data. It is recommended that this option is used only for cases in which a higher probability of data loss is acceptable to potentially improve performance, and save on data movement costs. For example, in a scratch file system scenario.

To better tolerate EBS failures, it is recommended to use the multiple availability zone option, in which IBM Storage Scale replicates the data so that there are two total copies of each data element. In this case, the replication is intended to deal with a loss of a volume. However, the probability of data loss depends on how the failure rates impact the total number of volumes that might fail at any point. It is recommended that users carefully read Amazon's statement about the durability of EBS volumes, particularly the annual failure rates, to better assess the possibility of potential data loss that results from EBS volume failures. For more information, see [Amazon EBS features](#).

Note: The region in which the IBM Storage Scale CloudFormation template is launched is where the network infrastructure for the IBM Storage Scale cluster is built. The template is launched in the US East (Ohio) region by default. You can change the region.

- b. On the **Select Template** page, keep the default setting for the **Template URL**, and then choose **Next**.
- c. On the **Specify Details** page, change the stack name if needed. Review the parameters for the template.

Note: Provide values for the parameters that require input. For all other parameters, review the default settings and customize them as necessary.

- d. In the **Options** page, you can specify tags or key-value pairs for resources in your stack, and set **Advanced options**.

Note: For more information on how to specify tags, see [AWS CloudFormation Resource Tags Type](#). For setting stack options in the **Advanced options** section, see [Setting AWS CloudFormation Stack Options](#).

- e. In the **Review** page, review and confirm the template settings. Under **Capabilities**, select the checkbox to acknowledge that the template creates IAM resources.
- f. Click **Create** to deploy the stack.
- g. Monitor the status of the stack. When the status displays CREATE_COMPLETE, the IBM Storage Scale environment is ready.
- h. View the resources that were created for the stack in the URLs displayed in the **Outputs** tab.

Note: This IBM Storage Scale AWS stack deployment is automated by the nested AWS CloudFormation templates. The main template builds the network-related resources first, using the VPC template, and then launches separate stacks for the bastion host and IBM Storage Scale cluster. Deleting the stack that is created by the main template deletes the entire IBM Storage Scale deployment stack. However, you still need to delete CloudWatch alerts manually.

3. Connect to the IBM Storage Scale cluster.

When the AWS CloudFormation template successfully creates the stack, all instances of compute and NSD servers launched by the AWS setup is up and running with the IBM Storage Scale file system that is mounted on it.

Follow these steps to connect to the IBM Storage Scale cluster:

- Open the Amazon EC2 console at <https://console.aws.amazon.com/ec2/>.
- In the navigation pane, under **Instances**, check for the public DNS (IPv4) value for the instance named LinuxBastion.
- Use your AWS private key to connect to the bastion host by using SSH.

Note: This is the key that you specify in the **Key Name** parameter of the AWS CloudFormation template during deployment.

- From the bastion host, use the SSH agent to log in to any of the compute or storage nodes that were launched by the AWS CloudFormation templates.

For more information about using an SSH agent to forward your private key on connection, see [Using SSH agent forwarding](#).

Important: Do not copy your private key to the bastion host instance.

Note: The IBM Storage Scale controller node is not part of the IBM Storage Scale cluster, and is used for cluster lifecycle management.




Figure 48. High-level IBM Storage Scale cluster architecture to connect from host

- Test the deployment by using IBM Storage Scale commands.

After you log in to a compute or storage nodes, you can administer IBM Storage Scale. For more information on administering IBM Storage Scale, see the *IBM Storage Scale: Administration Guide*.

Note: You must be a root user to run IBM Storage Scale administration commands.

The controller is not part of the IBM Storage Scale cluster, and therefore cannot run any of the following administrative commands. The controller node can be used only for Lifecycle management commands.

The **mm1scluster** command displays the details of the IBM Storage Scale cluster. The command gives the following output:

```
[ec2-user@ip-10-0-1-110]/usr/lpp/mmfs/bin/mm1scluster
GPFS cluster information
=====
```

```

GPFS cluster name:      ip-10-0-1-110.ap-south-1.compute.internal
GPFS cluster id:       12901386493707864068
GPFS UID domain:       ip-10-0-1-110.ap-south-1.compute.internal
Remote shell command:  /usr/bin/ssh
Remote file copy command: /usr/bin/scp
Repository type:       CCR

```

Node Designation	Daemon node name	IP address	Admin node name
1	ip-10-0-1-110.ap-south-1.compute.internal	10.0.1.110	ip-10-0-1-110.ap-south-1.compute.internal
2	ip-10-0-3-42.ap-south-1.compute.internal	10.0.3.42	ip-10-0-3-42.ap-south-1.compute.internal
3	ip-10-0-1-72.ap-south-1.compute.internal	10.0.1.72	ip-10-0-1-72.ap-south-1.compute.internal
4	ip-10-0-3-82.ap-south-1.compute.internal	10.0.3.82	ip-10-0-3-82.ap-south-1.compute.internal

For more information on the **mmclscluster** command, see the *IBM Storage Scale: Command and Programming Reference* guide in the [IBM Storage Scale documentation](#).

The **mmclsnsd** command displays the NSD server information. The command gives the following output:

```

[ec2-user@ip-10-0-1-110] /usr/lpp/mmfs/bin/mmclsnsd

```

File system	Disk name	NSD servers
fs1	0a9503b2cf264bf08	ip-10-0-3-42.ap-south-1.compute.internal
fs1	0d7e3d725140a0cdb	ip-10-0-1-110.ap-south-1.compute.internal
fs1	0fb10f7d9a981f39e	ip-10-0-1-72.ap-south-1.compute.internal

For more information on the **mmclsnsd** command, see the *IBM Storage Scale: Command and Programming Reference* guide in the [IBM Storage Scale documentation](#).

The **mmclsdisk** command displays disk details. The command gives the following output:

```

[ec2-user@ip-10-0-1-110] mmclsdisk fs1 -L

```

disk name	driver type	sector size	failure group	holds metadata	holds data	status	availability	disk id	storage pool
0a9503b2cf264bf08	nsd	512	1	Yes	Yes	ready	up		1
0d7e3d725140a0cdb	nsd	512	2	Yes	Yes	ready	up		2
0fb10f7d9a981f39e	nsd	512	3	No	No	ready	up		3

Number of quorum disks: 3
 Read quorum value: 2
 Write quorum value: 2

For more information on the **mmclsdisk** command, see the *IBM Storage Scale: Command and Programming Reference* guide in the [IBM Storage Scale documentation](#).

The **mmclsdf** command displays the fs1 filesystem information. The command gives an output similar to the following:

```

[ec2-user@ip-10-0-1-110] mmclsdf fs1

```

disk name	disk size in KB	failure group	holds metadata	holds data	free in KB in full blocks	free in KB in fragments
Disks in storage pool: system (Maximum disk size allowed is 3.97 TB)						
0a9503b2cf264bf08	524288000	1	Yes	Yes	520880128 (99%)	8792 (0%)
0d7e3d725140a0cdb	524288000	2	Yes	Yes	520880128 (99%)	8792 (0%)
0fb10f7d9a981f39e	5242880	3	No	No	0 (0%)	0 (0%)
(pool total)	1053818880				1041760256 (99%)	17584 (0%)
(total)	1053818880				1041760256 (99%)	17584 (0%)

Inode Information
 Number of used inodes: 4038
 Number of free inodes: 497722
 Number of allocated inodes: 501760
 Maximum number of inodes: 1025024

For more information on the **mmdf** command, see the *IBM Storage Scale: Command and Programming Reference* guide in the [IBM Storage Scale documentation](#).

The **mmgetstate** command shows the state of the instances in the cluster. The command gives an output similar to the following:

```
[ec2-user@ip-10-0-1-110]sudo /usr/lpp/mmfs/bin/mmgetstate -a
```

Node number	Node name	GPFS state
1	ip-10-0-1-110	active
2	ip-10-0-3-42	active
3	ip-10-0-1-72	active
4	ip-10-0-3-82	active

For more information on the **mmgetstate** command, see the *IBM Storage Scale: Command and Programming Reference* guide in the [IBM Storage Scale documentation](#).

Option 1: Deploying IBM Storage Scale on a new Amazon VPC with a single Availability Zone

The following section details the parameters for a new VPC deployment with a single Availability Zone.

Table 53. File system configurations		
Parameter label (name)	Default	Description
Block Size (BlockSize)	1 M	The file system block size. You can choose a value from 256 KiB - 16 MiB.
GPFS Mount Point (GpfsMountPoint)	/gpfs/fs1	The mount point for the IBM Storage Scale file system. Note: fs1 is the file system name that is created by default, and is mounted on location /gpfs/fs1. Changing the mount path to /gpfs/fs2 creates a file system with name fs2, and the file system is mounted on /gpfs/fs2.

Table 54. NSD Configurations		
Parameter label (name)	Default	Description
EBS Type (EBSType)	gp2	The EBS volume type for IBM Storage Scale storage that is attached to each NSD server node. The following options are available: - General Purpose SSD (gp2) - Provisioned IOPS SSD (io1) - Cold HDD (sc1) - Throughput Optimized HDD (st1) - EBS Magnetic (standard) For more information about EBS volume type choices, see Amazon EBS . To maximize the benefit of io1 EBS volume type, it is recommended to use the appropriate EBS optimized instances. For more information, see the AWS documentation.
Disk Per Node (DiskPerNode)	1	The number of NSD volumes to attach to each NSD server node. You can choose 1 - 15 disks.

Table 54. NSD Configurations (continued)		
Disk Size (DiskSize)	500	The disk size of NSD volumes that are attached to each NSD server node, in GiBs. Supported disk sizes are 500 - 16,384 GiB.

Table 55. Server Node Configurations		
Parameter label (name)	Default	Description
Storage Node Count (StorageNodeCount)	2	The number of EC2 instances used for the NSD server on the GPFS cluster. You can select 2 - 64 instances.
Storage Instance Type (StorageInstanceType)	t2.medium	The instance type to use for the storage node. A minimum t2.medium instance type is needed. However, other types are available, and a value can be selected based on your requirements. Ensure that the chosen instance type is supported in the desired region.
StorageAMIID		AMI-ID to be used for the storage node. Only Red Hat Enterprise Linux Operating system versions 7.7, 7.8, 8.1, 8.2, and 9.4 are supported. For more information, see Amazon Machine Images (AMI). Customers must specify an appropriate AMI running a supported Operating Systems. Red Hat Enterprise Linux AMIs are available from the AWS Marketplace and the AWS Community. Additionally, customers can choose to create their own custom AMI. For more information, see <i>Creating custom AMI in IBM Storage Scale: Administration Guide</i>. Important: The customer is responsible to ensure that they have the necessary support agreements in place for the Red Hat Enterprise Linux Operating Systems AMIs they specify. The support agreement is needed to facilitate resolutions with Red Hat support for any issues.

Note: By default, all server nodes are assigned with an IBM Storage Scale server license.

Table 56. Compute Node Configurations		
Parameter label (name)	Default	Description
Compute Node Count (ComputeNodeCount)	2	The number of IBM Storage Scale compute node instances. You can select 1 - 64 instances.
Compute Instance Type (ComputeInstanceType)	t2.medium	The instance type to use for the compute node. A minimum t2.medium instance type is needed. However, other types are available, and a value can be selected based on your requirements. Ensure that the chosen instance type is supported in the desired region.

Table 56. Compute Node Configurations (continued)

ComputeAMIID		AMI-ID to be used for the compute node. Only Red Hat Enterprise Linux Operating system versions 7.7, 7.8, 8.1, 8.2, and 9.4 are supported. For more information, see Amazon Machine Images (AMI). Customers must specify an appropriate AMI running a supported Operating Systems. Red Hat Enterprise Linux AMIs are available from the AWS Marketplace and the AWS Community. Additionally, customers can choose to create their own custom AMI. For more information on how to build custom AMI, see <i>Creating custom AMI</i> in <i>IBM Storage Scale: Administration Guide</i>. Important: The customer is responsible to ensure that they have the necessary support agreements in place for the Red Hat Enterprise Linux Operating Systems AMIs they specify. The support agreement is needed to facilitate resolutions with Red Hat support for any issues.

Table 57. Network Configuration

Parameter label (name)	Default	Description
Availability Zone (AvailabilityZones)	Requires input	Availability Zone to use for the subnet in the VPC.
VPC CIDR (VPCCIDR)	10.0.0.0 /16	The CIDR block for the VPC.
Private Subnet CIDR (PrivateSubnet1CIDR)	10.0.1.0 /24	The CIDR block for the private subnet located in Availability Zone 1.
Public Subnet CIDR (PublicSubnet1CIDR)	10.0.0.0 /24	The CIDR block for the public subnet located in Availability Zone 1.
Allowed External Access CIDR (RemoteAccessCIDR)	Requires input	The CIDR block that is allowed external SSH access to the bastion hosts. For example, x.x.x.x/16 - 28. It is recommended that you set this value to a trusted CIDR block. For example, you might want to restrict access to your corporate network.

Table 58. Amazon EC2 Configuration

Parameter label (name)	Default	Description
Key Pair Name (KeyPairName)	Requires input	A public or private key pair, which enables you to connect securely to your instance after it is launched. When you created an AWS account, this is the key pair that you created in your preferred region.
Bastion AMI OS (BastionAMIOS)	Amazon-Linux2023	The Linux distribution for the AMI to be used for the bastion host instances. If you choose CentOS, make sure that you have a subscription to the CentOS AMI in AWS Marketplace..
Bastion Instance Type (BastionInstanceType)	t2.micro	The EC2 instance type for the bastion host instances.

Table 58. Amazon EC2 Configuration (continued)		
Controller Instance Type	t2.micro	The instance type to use for the controller node. The drop-down lists all the available instance types. The instance type of the controller node affects the deployment time of the IBM Storage Scale cluster. The larger the chosen cluster, the larger the controller instance type must be for optimal deployment time. The default value is usually preferred only for proof-of-concept deployments. For IBM Storage Scale clusters with larger than 30 nodes, a value with a large instance type must be specified for the controller instance type.

Table 59. Personal Configuration		
Parameter label (name)	Default	Description
Spectrum S3 Bucket (SpectrumS3Bucket)	Requires input	S3 bucket name to be created and used for storing state files. Ensure that the S3 bucket does not exist. You can review the rules for naming a bucket at Bucket Restriction .
Operator Email (OperatorEmail)	Requires input	The email address to which the notifications about scaling operations are sent.

Table 60. License Information		
Parameter label (name)	Default	Description
IBM Customer Number (IBMCustomerNumber)	Requires input	Needed to validate the customer entitlement for BYOL offering.
License Agreement Terms (LicenseAgreementTerms)	Requires input	Review the licensing terms at Licence Information , and if you agree to the terms, choose Accept .

Option 2: Deploying IBM Storage Scale on a new Amazon VPC with multiple Availability Zones

The following section details the parameters for a new VPC deployment with multiple Availability Zones.

Table 61. File System Configurations		
Parameter label (name)	Default	Description
Block Size (BlockSize)	1 M	The file system block size. You can choose a value from 256 KiB - 16 MiB.
GPFS Mount Point (GpfsMountPoint)	/gpfs/fs1	The mount point for the IBM Storage Scale volume. Note: fs1 is the file system name that is created by default, and is mounted on location /gpfs/fs1. Changing the mount path to /gpfs/fs2 creates a file system with name fs2, and the file system is mounted on /gpfs/fs2.

Table 62. NSD Configurations		
Parameter label (name)	Default	Description

Table 62. NSD Configurations (continued)		
EBS Type (EBSType)	gp2	The EBS volume type for IBM Storage Scale storage that is attached to each NSD server node. The following options are available: • General Purpose SSD (gp2) • Provisioned IOPS SSD (io1) • Cold HDD (sc1) • Throughput Optimized HDD (st1) • EBS Magnetic (standard) For more information about EBS volume type choices, see Amazon EBS. To maximize the benefit of io1 EBS volume type, it is recommended to use the appropriate EBS optimized instances. For more information, see AWS documentation.
Disk Per Node (DiskPerNode)	1	The number of NSD volumes to attach to each NSD server node. You can choose from 1 - 15 disks.
Disk Size (DiskSize)	500	The disk size of NSD volumes that are attached to each NSD server node, in GiBs. Supported disk sizes are 10 - 16,384 GiB.

Table 63. Server Node Configurations		
Parameter label (name)	Default	Description
Storage Node Count (StorageNodeCount)	2	The number of EC2 instances to launch for the NSD server on the GPFS cluster. You can select from 2 - 64 instances.
Storage Instance Type (StorageInstanceType)	t2.medium	The instance type to use for the storage node. A minimum t2.medium instance type is needed. However, other types are available, and a value can be selected based on your requirements. Ensure that the chosen instance type is supported in the desired region.
StorageAMIID		AMI-ID to be used for the storage node. Only Red Hat Enterprise Linux Operating system versions 7.7, 7.8, 8.1, 8.2, and 9.4 are supported. For more information, see Amazon Machine Images (AMI). Customers must specify an appropriate AMI running a supported Operating Systems. Red Hat Enterprise Linux AMIs are available from the AWS Marketplace and the AWS Community. Additionally, customers can choose to create their own custom AMI. For more information, see the <i>Creating custom AMI</i> section in the <i>IBM Storage Scale: Administration Guide</i>. Important: The customer is responsible to ensure that they have the necessary support agreements in place for the Red Hat Enterprise Linux Operating Systems AMIs they specify. The support agreement is needed to facilitate resolutions with Red Hat support for any issues.

Note: By default, all server nodes are assigned with an IBM Storage Scale server license.

Table 64. Compute Node Configurations		
Parameter label (name)	Default	Description
Compute Node Count (ComputeNodeCount)	2	The number of IBM Storage Scale compute node instances. You can select 1 - 64 instances.
Compute Instance Type (ComputeInstanceType)	t2.medium	The instance type to use for the compute node. A minimum t2.medium instance type is needed. However, other types are available, and a value can be selected based on your requirements. Ensure that the chosen instance type is supported in the desired region.
ComputeAMIID		AMI-ID to be used for the compute node. Only Red Hat Enterprise Linux Operating system versions 7.7, 7.8, 8.1, 8.2, and 9.4 are supported. For more information, see Amazon Machine Images (AMI). Customers must specify an appropriate AMI running a supported Operating Systems. Red Hat Enterprise Linux AMIs are available from the AWS Marketplace and the AWS Community. Additionally, customers can choose to create their own custom AMI. For more information on how to build custom AMI, see the <i>Creating custom AMI</i> section in the <i>IBM Storage Scale: Administration Guide</i>. Important: The customer is responsible to ensure that they have the necessary support agreements in place for the Red Hat Enterprise Linux Operating Systems AMIs they specify. The support agreement is needed to facilitate resolutions with Red Hat support for any issues.

Table 65. Network Configuration		
Parameter label (name)	Default	Description
Availability Zone (AvailabilityZones)	Requires input	The list of Availability Zones to use for the subnets in the VPC. Only two Availability Zones are used for this deployment, and the logical order of your selections is preserved.
VPC CIDR(VPCCIDR)	10.0.0.0 /16	The CIDR block for the VPC.
Private Subnet 1 CIDR (PrivateSubnet1CIDR)	10.0.1.0 /24	The CIDR block for the private subnet located in Availability Zone 1.
Private Subnet 2 CIDR (PrivateSubnet2CIDR)	10.0.3.0 /24	The CIDR block for the private subnet located in Availability Zone 2.
Private Subnet 3 CIDR (PrivateSubnet3CIDR)	10.0.5.0 /24	The CIDR block for the private subnet located in Availability Zone 3.
Public Subnet 1 CIDR (PublicSubnet1CIDR)	10.0.0.0 /24	The CIDR block for the public subnet located in Availability Zone 1.
Public Subnet 2 CIDR (PublicSubnet2CIDR)	10.0.2.0 /24	The CIDR block for the public subnet located in Availability Zone 2.
Public Subnet 3 CIDR (PublicSubnet3CIDR)	10.0.4.0 /24	The CIDR block for the public subnet located in Availability Zone 3.

Table 65. Network Configuration (continued)		
Allowed External Access CIDR (RemoteAccessCIDR)	Requires input	The CIDR block that is allowed external SSH access to the bastion hosts. For example, <i>x.x.x.x/16-28</i> . It is recommended that you set this value to a trusted CIDR block. For example, you might want to restrict access to your corporate network.

Table 66. Amazon EC2 Configuration		
Parameter label (name)	Default	Description
Key Pair Name (KeyPairName)	Requires input	A public or private key pair, which allows you to connect securely to your instance after it launches. When you created an AWS account, this is the key pair that you created in your preferred region.
Bastion AMI OS (BastionAMIOS)	Amazon-Linux2023	The Linux distribution for the AMI to be used for the bastion host instances. If you choose CentOS, make sure that you have a subscription to the CentOS AMI in AWS Marketplace .
Tie Breaker Instance Type	t2.medium	Instance type to use for the Tie breaker node instance.
Bastion Instance Type (BastionInstanceType)	t2.micro	The EC2 instance type for the bastion host instances.
Controller Instance Type	t2.micro	The instance type to use for the controller node instance. The list displays all the available instance types. The instance type of the controller node affects the deployment time of the IBM Storage Scale cluster. The larger the chosen cluster, the larger the controller instance type must be for optimal deployment time. The default value is usually preferred only for proof-of-concept deployments. For IBM Storage Scale clusters with larger than 30 nodes, a value with a large instance type must be specified for the controller instance type.

Table 67. Personal Configuration		
Parameter label (name)	Default	Description
Spectrum S3 Bucket (SpectrumS3Bucket)	Requires input	S3 bucket name to be created and used for storing state files. Ensure that the S3 bucket does not exist. You can review the rules for naming a bucket at Bucket Restriction .
Operator Email (OperatorEmail)	Requires input	The email address to which the notifications about scaling operations are sent.

Table 68. License Information		
Parameter label (name)	Default	Description
IBM Customer Number (IBMCustomerNumber)	Requires input	Needed to validate the customer entitlement for BYOL offering.
License Agreement Terms (LicenseAgreementTerms)	Requires input	Review the licensing terms at Licence Information , and if you agree to the terms, choose Accept .

Option 3: Deploying IBM Storage Scale on an existing Amazon VPC

The following requirements must be met for IBM Storage Scale to be deployed on an existing VPC.

- One private and one public network for single Availability Zone. Three private and three public networks for three Availability Zones.

Note: It is recommended for the user to have an existing VPC in three Availability Zones.

- The following are the private subnet requirements:
 1. Internet gateway (IGW) is not configured in the route table.
 2. Auto-assign public IPv4 address is disabled. The IBM Storage Scale nodes with public IP are not supported due to security issues.
 3. Configure the Network Address Translation (NAT) for this subnet. The IBM Storage Scale stack creation requires access to AWS S3 and RHUI servers.
- Public subnets are needed for Bastion nodes. It is recommended to use Amazon templates to create Bastion nodes. The IBM Storage Scale stack creation does not need public subnets.

Note: It is recommended to use the following setup:

1. Dedicated private subnet for the IBM Storage Scale cluster.
 2. Subnet masks that are large enough to support maximum number of nodes.
 3. Same subnet mask for all private subnet.
 4. Subnets that are not defined as default subnet.
- Create an S3 endpoint. An S3 endpoint is created by the CF template for VPC that is being used to create VPC.
 - VPC endpoints must be created and added to the route tab.

Note: The VPC created by the template has the value for DNS hostname set to **Yes**. The DNS hostname value is set to **No** when a new VPC is created. For more information on DNS hostnames, see <https://docs.aws.amazon.com/AmazonVPC/latest/UserGuide/vpc-dns.html#vpc-dns-hostnames>.

Table 69. File System Configurations		
Parameter label (name)	Default	Description
Block Size (BlockSize)	1 M	The file system block size. You can choose a value from 256 KiB - 16 MiB.
GPFS Mount Point(GpfsMountPoint)	/gpfs/fs1	The mount point for the IBM Storage Scale volume. Note: fs1 is the file system name that is created by default, and is mounted on location /gpfs/fs1. Changing the mount path to /gpfs/fs2 creates a file system with name fs2, and the file system is mounted on /gpfs/fs2.

Table 70. NSD Configurations		
Parameter label (name)	Default	Description

Table 70. NSD Configurations (continued)		
EBS Type (EBSType)	gp2	The EBS volume type for IBM Spectrum Scale storage that is attached to each NSD server node. The following options are available: • General Purpose SSD (gp2) • Provisioned IOPS SSD (io1) • Cold HDD (sc1) • Throughput Optimized HDD (st1) • EBS Magnetic (standard) For more information about EBS volume type choices, see Amazon EBS. To maximize the benefit of io1 EBS volume type, it is recommended to use the appropriate EBS optimized instances. For more information, refer to the AWS documentation.
Disk Per Node (DiskPerNode)	1	The number of NSD volumes to attach to each NSD server node. You can choose 1 - 15 disks.
Disk Size (DiskSize)	500	The disk size of NSD volumes that are attached to each NSD server node, in GiBs. Supported disk sizes are 500 - 16,384 GiB.

Table 71. Server Node Configurations		
Parameter label (name)	Default	Description
Storage Node Count (StorageNodeCount)	2	The number of EC2 instances to launch for the NSD server on the GPFS cluster. You can select 2 - 64 instances.
Storage Instance Type (StorageInstanceType)	t2.medium	The instance type to use for the storage node. A minimum t2.medium instance type is needed. However, other types are available, and a value can be selected based on your requirements. Ensure that the chosen instance type is supported in the desired region.
StorageAMIID		AMI-ID to be used for the storage node. Only Red Hat Enterprise Linux Operating system versions 7.7, 7.8, 8.1, 8.2, and 9.4 are supported. For more information, see Amazon Machine Images (AMI). Customers must specify an appropriate AMI running a supported Operating Systems. Red Hat Enterprise Linux AMIs are available from the AWS Marketplace and the AWS Community. Additionally, customers can choose to create their own custom AMI. For more information, see <i>Creating custom AMI</i> in <i>IBM Storage Scale: Administration Guide</i>. Important: The customer is responsible to ensure that they have the necessary support agreements in place for the Red Hat Enterprise Linux Operating Systems AMIs they specify. The support agreement is needed to facilitate resolutions with Red Hat support for any issues.

Table 72. Compute Node Configurations		
Parameter label (name)	Default	Description
Compute Node Count (ComputeNodeCount)	2	The number of IBM Storage Scale compute node instances. You can select 1 - 64 instances.
Compute Instance Type (ComputeInstanceType)	t2.medium	The instance type to use for the compute node. A minimum t2.medium instance type is needed. However, other types are available, and a value can be selected based on your requirements. Ensure that the chosen instance type is supported in the desired region.
ComputeAMIID		AMI-ID to be used for the compute node. Only Red Hat Enterprise Linux Operating system version 7.7, 7.8, 8.1, 8.2, and 9.4 are supported. For more information, see Amazon Machine Images (AMI). Customers must specify an appropriate AMI running a supported Operating Systems. Red Hat Enterprise Linux AMIs are available from the AWS Marketplace and the AWS Community. Additionally, customers can choose to create their own custom AMI. For more information, see <i>Creating custom AMI in IBM Storage Scale: Administration Guide</i>. Important: The customer is responsible to ensure that they have the necessary support agreements in place for the Red Hat Enterprise Linux Operating Systems AMIs they specify. The support agreement is needed to facilitate resolutions with Red Hat support for any issues.

Table 73. Network Configuration		
Parameter label (name)	Default	Description
VPC ID		ID of your existing VPC. For example, vpc-0343606e. Value can be select from the list.
Private Subnet 1 ID (PrivateSubnet1ID)		Private subnet ID corresponding to Availability Zone 1 in your VPC. For example, subnet-a0246dcd.  CAUTION: Ensure that the values for Private Subnet ID 1, Private Subnet ID 2, and Private Subnet ID 3 are: • Either the same value that belongs to a single Availability Zone. • Or unique values that belong to a three distinct Availability Zones. The same value in each Private Subnet ID field signifies that only one Availability Zone is used.
Private Subnet 2 ID (PrivateSubnet2ID)		Private subnet ID corresponding to Availability Zone 2. You can choose private subnet ID corresponding to Availability Zone 1 again when only one private subnet exists in your VPC. For example, subnet-a0246dcd.
Private Subnet 3 ID (PrivateSubnet3ID)		Private subnet ID corresponding to Availability Zone 3. You can choose private subnet ID corresponding to Availability Zone 1 again when only one private subnet exists in your VPC. For example, subnet-a0246dcd.

Table 73. Network Configuration (continued)		
Allowed External Access CIDR (RemoteAccessCIDR)	Requires input	The CIDR block that is allowed external SSH access to the Bastion hosts. For example, <i>x.x.x.x/16</i> - 28. It is recommended that you set this value to a trusted CIDR block. For example, you might want to restrict access to your corporate network.

Note: For a single-AZ deployment, add the same PrivateSubnetID in all three places. For multi-AZ deployment, select a unique ID in each place.

Table 74. Amazon EC2 Configuration		
Parameter label (name)	Default	Description
Key Pair Name (KeyPairName)	Requires input	A public or private key pair, which is used to connect securely to your instance after it launches. When you created an AWS account, this is the key pair that you created in your preferred region.
Tie Breaker Instance Type	t2.medium	Instance type to use for the Tie breaker node instance.
Bastion Security Group ID		Bastion host security group ID to enable SSH connections. For example, sg-5f16e910.
Controller Instance Type	t2.micro	The instance type to use for the controller node instance. The list displays all the available instance types. The instance type of the controller node affects the deployment time of the IBM Storage Scale cluster. The larger the chosen cluster, the larger the controller instance type must be for optimal deployment time. The default value is usually preferred only for proof-of-concept deployments. For IBM Storage Scale clusters with larger than 30 nodes, a value with a large instance type must be specified for the controller instance type.

Table 75. Personal Configuration		
Parameter label (name)	Default	Description
Spectrum S3 Bucket (SpectrumS3Bucket)	Requires input	S3 bucket name to be created and used for storing state files. Ensure that the S3 bucket does not exist. You can review the rules for naming a bucket at Bucket Restriction .
Operator Email (OperatorEmail)	Requires input	The email address to which the notifications about scaling operations are sent.

Table 76. License Information		
Parameter label (name)	Default	Description
License Agreement Terms (LicenseAgreementTerms)	Requires input	Review the licensing terms at Licence Information , and if you agree to the terms, choose Accept .
IBM Customer Number		IBM Customer number to validate entitlement for BYOL offering.

Deploying AFM on AWS

IBM Storage Scale Active File Management (AFM) is used for data movement and caching between the on-premises and public cloud environments. Public cloud end points are only supported as cache sites.




Figure 49. Example of an on-premises to AWS hybrid cloud networking architecture with AFM

Follow these steps to deploy AFM:

1. Set up the on-premises cluster.

The on-premises IBM Storage Scale cluster acts as the AFM home cluster. An AFM home cluster is defined as an IBM Storage Scale cluster that can make the NFS v3 exports available to other IBM Storage Scale clusters. Follow these steps to set up the on-premises cluster:

- Run the **mmnfs** command to define the NFS exports and provide access to the AFM gateway nodes:

```
mmnfs export add /ibm/gpfs0/onprem-aws -client "172.16.1.237  
(Access_Type=RW,Squash=root_squash)"
```

The command gives the following output:

```
mmnfs: The NFS export was created successfully
```

For more information, see the *IBM Storage Scale: Command and Programming Reference Guide* guide in the [IBM Storage Scale Documentation](#).

- Run the **mmafmconfig** command to enable support of extended attributes or sparse files on the AFM home exports:

```
mmafmconfig enable /ibm/gpfs0/onprem-aws
```

For more information, see the *IBM Storage Scale: Command and Programming Reference Guide* guide in the [IBM Storage Scale Documentation](#).

2. Set up the AWS cluster.

The IBM Storage Scale cluster on AWS acts as the AFM cache cluster. The AFM cache cluster hosts the filesets that are configured to participate in a relationship with the NFS exports in the AFM home cluster. Follow these steps to set up the AWS cluster:

- a. Run the following command to verify that the NFS client is installed on the nodes that are going to server as the AFM gateway:

```
rpm -qi nfs-utils
```

If the NFS client package is not installed, run the following command to install the NFS client on the nodes that are going to serve as the AFM gateway.

```
yum install nfs-utils
```

Note: A gateway node can communicate with the home cluster to transfer data. Refer the “[Configuration best practices](#)” on page 625 section before you select a node to act as the AFM gateway node.

- b. Run the **mmchnode** command to define the AFM gateway nodes on the nodes selected in [Step 2a](#).

```
mmchnode --gateway -N ip-172-16-1-237.eu-central-1.compute.internal
```

The command gives the following output:

```
mmchnode: Processing node ip-172-16-1-237.eu-central-1.compute.internal
mmchnode: Propagating the cluster configuration data to all
affected nodes. This is an asynchronous process.
```

For more information, see the *IBM Storage Scale: Command and Programming Reference Guide* guide in the [IBM Storage Scale documentation](#).

- c. Run the **mmcrfileset** command to create an AFM fileset with a relationship to the NFS exports server from the on-premises or home cluster.

Note: AFM cache mode must be specified based on the needs of the customer use case. For more information, see .

```
mmcrfileset fs1 onprem-aws-cache -p afmTarget=10.0.240.43:/ibm/gpfs0/onprem-aws -p
afmmode=read-only --inode-limit 999999 --inode-space new
```

The command gives the following output:

```
Fileset onprem-aws-cache created with id 21 root inode 5767171.
```

For more information, see the *IBM Storage Scale: Command and Programming Reference Guide* guide in the [IBM Storage Scale documentation](#).

- d. Run the **mmlinkfileset** command to link the fileset to the junction folder:

Note: A junction folder is a special directory entry, like a POSIX hard link, that connects a name in the directory of the parent fileset to the root directory of a child fileset.

```
mmlinkfileset fs1 onprem-aws-cache -J /gpfs/fs1/onprem-aws-cache
```

The command gives following output:

```
Fileset onprem-aws-cache linked at /gpfs/fs1/onprem-aws-cache
```

For more information, see the *IBM Storage Scale: Command and Programming Reference Guide* guide in the [IBM Storage Scale documentation](#).

You can prefetch the file metadata and data from the AFM home cluster before an application requests the contents. Prefetching files before an application starts can reduce the network delay when an application requests a file. Prefetch can be used to pro-actively manage WAN traffic patterns by moving files over the WAN during a period of low WAN usage. For more information, see the *Prefetch* section in the [IBM Storage Scale documentation](#).

For more information, see the *Setting up the cache cluster* and *Creating an AFM relationship by using the NFS protocol* sections in the [IBM Storage Scale documentation](#).

Configuration best practices

The following section describes the best practices for configuring AFM on AWS.

- Ensure that the network link between the home and cache clusters has enough bandwidth to support the volume of data that is being copied between them.
- Ensure that the site-to-site VPN servers on the on-premise cluster and the AWS cloud cluster have sufficiently large resources. This is because all the data between the home and cache clusters must traverse through the VPN servers. If the VPN servers do not have the requisite amount of CPU and memory availability, the VPN servers might introduce latencies that result in AFM disconnects.
- Set the following configuration parameters on the AWS cluster to a sufficiently high value to compensate for the possible large latencies that are found in the network between the on-premise and AWS clusters:
 - `afmAsyncOpWaitTimeout`
 - `afmRevalOpWaitTimeout`
 - `afmSyncOpWaitTimeout`

Since all the data between the home and cache clusters traverses through the VPN servers, it is inevitable that larger latencies are introduced. These parameters must be set to sufficiently large values to account for larger latency and to ensure that AFM does not disconnect. They are typically set to 1800. These configuration parameters can be modified using the **mmchconfig** command. For more information on the **mmchconfig** command, see *IBM Storage Scale: Command and Programming Reference Guide* in the [IBM Storage Scale documentation](#).

- Ensure that the AFM gateway nodes have enough resources for optimal performance. Therefore, it is important that the AWS instances with the requisite amount of CPU and memory availability are chosen during the initial cluster provisioning.
- Ensure that enough AFM gateways are provisioned on the cache cluster to provide the desired HA redundancy and performance characteristics.
- Ensure that a node is not assigned both the `NSD server` node role and the `AFM gateway` node role simultaneously to avoid resource contention.
- AFM gateways must ideally not be assigned to a node that is running any customer workload.

Limitations of AFM on AWS

The following section describes the limitations of AFM on AWS.

- The AWS IBM Storage Scale cluster can only act as an AFM cache.
- No support for AFM over GPFS protocol. Only AFM over NFS v3 is supported.
- No support for AFM DR.

Chapter 8. Installing transparent cloud tiering (discontinued)

The cloud services installation is not included in the IBM Storage Scale installation toolkit workflow. It needs to be installed separately on top of the IBM Storage Scale cluster. However, the cloud services RPMs are available with the IBM Storage Scale package.

Before you begin, ensure that you have created a node class. For more information, see the *Creating a node class* topic in the *IBM Storage Scale: Administration Guide*. Also, ensure that your server meets the prerequisites. For more information, see [“Software requirements for cloud services”](#) on page 407.

Creating a user-defined node class for transparent cloud tiering or cloud data sharing

This topic provides step-by-step instructions for creating and configuring the required node class to specify which IBM Storage Scale nodes are used for cloud services.

Note: You need to run this command only once on any node on the IBM Storage Scale cluster.

You can install cloud services on a maximum of four nodes per node group (any combination of NSD or CES nodes) on the IBM Storage Scale cluster (which can have a maximum of four node groups). Before you install cloud services, you need to create a node class that specifies the IP addresses (or fully qualified host names) of the nodes where the server packages are going to be installed. You can use any host name that the `-N` option accepts. For details, see `mmcrnodeclass` command in the *IBM Storage Scale: Command and Programming Reference Guide*.

You can enable and manage independent groups of cloud services nodes in different node classes for use with different network configurations per node class. Nodes are mutually exclusive to each node class and cannot be shared with another node class that has cloud services nodes. For example, the node class `TCTNodeClass1` has `node1` and `node2` and the node class `TCTNodeClass2` has `node3` and `node4`. All nodes are cloud services enabled nodes. In this case, the node class `TCTNodeClass2` cannot enable `node1` as a cloud services node under the `TCTNodeClass2` node class because it is already marked for cloud services usage under the `TCTNodeClass1` node class. Additionally, each group of cloud services enabled nodes under a node class cannot share file system with cloud services enabled nodes from a different node class. And each cloud services enabled node under a node class is limited to use only one common file system with each node in that node class. Therefore, you must manage each group of cloud services nodes within a node class as a pool of nodes that require a single common file system.

Note: It is recommended that you set up at least two nodes so that you have good availability for the service if a node were to go down. Ensure to use the GPFS cluster IP address that gets displayed when you run the `mmfscluster` command.

To create a node class, issue a command according to this syntax:

```
mmcrnodeclass ClassName -N {Node[,Node...]} [NodeFile | NodeClass]
```

For example, to create a node class called `TCTNodeClass` by using three nodes `10.10.10.10`, `11.11.11.11`, and `12.12.12.12`, issue this command:

```
mmcrnodeclass TCTNodeClass -N 10.10.10.10,11.11.11.11,12.12.12.12
```

To verify that the node class is created, issue this command:

```
mmnlsnodeclass
```

The system displays output similar to this output:

Node	Class	Name	Members

TCTNodeClass		node10,node11,node12	

Next step: [“Installation steps” on page 628](#)

Installation steps

This topic describes the procedure for installing cloud services on IBM Storage Scale nodes.

Before you begin

cloud services creates the following folders after installation:

- /var/MCStore
- <filesystem>/ .mcstore
- /opt/ibm/MCStore

Do not make any changes to these folders in order to avoid any undesired results.

About this task

The cloud services RPMs are available with IBM Storage Scale Advanced Edition, IBM Storage Scale Data Management Edition, IBM Storage Scale Developer Edition, or IBM Storage Scale Erasure Code Edition. You can verify this by issuing this command: **rpm -qa | grep gpfs.adv**. Depending on the target Linux distribution, the RPMs are available in the following paths where *_rpms or *_debs is the component specific directory:

- Red Hat Enterprise Linux 7.x: /usr/lpp/mmfs/5.0.x.x/*_rpms/rhel7
- Ubuntu 16.04 LTS: /usr/lpp/mmfs/5.0.x.x/*_debs/ubuntu16
- SLES 12: /usr/lpp/mmfs/5.0.x.x/*_rpms/sles12

With each build of IBM Storage Scale, two RPMs are available for cloud services:

- gpfs.tct.server-x.x.x.x86_64.rpm. This is the server package that is installed on the cloud services nodes. These nodes must have outbound WAN connectivity to be able to support data migration between IBM Storage Scale cluster and the cloud object storage. When installed on non-cloud gateway nodes, files are recalled locally without routing the requests through the gateway nodes. To enable the client-assisted recall, you must run **mmcloudgateway clientassist enable**.
- gpfs.tct.client-x.x.x.x86_64.rpm. This is the client package that contains the cloud services binary files (specifically the data manager commands). This package must be installed on all remaining cluster nodes to activate ILM-based file migration or file recall that is initiated from that node (every node in the cluster is safe). The client RPM is not required on the nodes where the server package is installed.

Note: It is important to install the client package on every node so that files that are deleted from the IBM Storage Scale cluster are deleted efficiently from the cloud storage tier. If this package is not installed, files are deleted less aggressively and only as a result of running the reconcile service.

Note: For policy-based migrations, it is recommended that the file system manager node (rather all nodes with manager roles) is installed at least with the client RPM. This ensures that if any migration request goes to the file system manager node that is not a Transparent cloud tiering node, the request is forwarded to one of the server nodes in the node class, and the migration seamlessly takes place. If the RPM is not there on the file system manager, then migration fails. Transparent cloud tiering is needed on all clients (not just for migration requests) or transparent recall does not work properly and delete processing is not handled efficiently for files transferred to the cloud.

Procedure

1. Copy the server RPMs to the desired nodes.

2. Install cloud services package on each node by issuing the following command depending on your architecture:

rpm -ivh gpfs.tct.server-x.x.x.x86_64.rpm

Note: By default, the cloud services is installed under the /opt/ibm/MCStore directory.

3. Install the client package on the desired nodes by issuing the following command depending on your OS and architecture:

- For SLES or Red Hat Enterprise Linux: **rpm -ivh gpfs.tct.client-x.x.x.<arch>.rpm**

- For Ubuntu Linux: **dpkg -i gpfs.tct.client-x.x.x.<arch>.deb**

Setting up a cloud services cluster

This topic provides step-by-step instructions for quickly setting up a cloud services node class that can support transparent cloud tiering or cloud data sharing on the IBM Storage Scale cluster.

About this task

In the multi-node setup, a maximum of 4 of any combination of CES or NSD nodes are supported as a cloud services node group -- a minimum of two are recommended for high availability. Data migration or recall workload is shared across all the nodes. During a transparent recall or migration, the ILM policy gathers a list of files to be migrated, and the list is divided equally among all the nodes.

Note: Once a multi-node setup is up and running, the transparent cloud tiering GUI components are available on the IBM Storage Scale GUI.

The `mtime` attribute on the IBM Storage Scale file system indicates when a file is updated. For a cloud services managed file system, ensure that the `mtime` parameter is set to `yes`, so that when cloud services internally invokes `stat()` or `fstat()` calls during migration, file system reports exact `mtime` value for a file. If `mtime` parameter is set to "no", there is a potential data loss situation, when a modified / updated co-resident file does not get migrated as a new version to the cloud. You can check the `mtime` status by using the **mmfsfs <file system> -E** command.

Perform the following steps for setting up a cloud services cluster:

Procedure

1. Install the cloud services RPMs on the required IBM Storage Scale nodes. For more information, see [Chapter 8, "Installing transparent cloud tiering \(discontinued\)," on page 627](#).
2. Designate the nodes where you installed the RPMs as cloud services nodes. For more information, see the *Designating the transparent cloud tiering nodes* topic in the *IBM Storage Scale: Administration Guide*.
3. Start cloud services on all nodes. For more information, see the *Starting the cloud services* topic in the *IBM Storage Scale: Administration Guide*.
4. Create a cloud account. For simpler configurations, use the built-in key manager that comes with cloud services. Use KLM when KLM is already running on IBM Storage Scale (for other reasons) or if specific security requirements that drive its usage. For more information, see the *Creating a cloud storage account* topic in the *IBM Storage Scale: Administration Guide*.
5. Define cloud storage access points (CSAP). For more information, see the *Defining cloud storage access points (CSAP)* topic in the *IBM Storage Scale: Administration Guide*.
6. Create a cloud services. For more information, see the *Creating a cloud services* topic in the *IBM Storage Scale: Administration Guide*.
7. Configure a key manager. For more information, see the *Creating a key manager* topic in the *IBM Storage Scale: Administration Guide*.
8. Create a container pair set and test your configuration. For more information, see the *Creating a container pair set* topic in the *IBM Storage Scale: Administration Guide*.

9. Make sure you back up your configuration and your key manager. For more information, see the *Backing up the Cloud services configuration* topic in the *IBM Storage Scale: Administration Guide*.

What to do next

If you are doing transparent cloud tiering, do the following steps to verify that the cluster is ready for transparent recall/migration based on an ILM policy:

- Create a sample policy, *migrate.policy*, and apply this policy by using the cloud services nodes by with the command: **mmapplypolicy gpfs -P migrate.policy -N FirstCloud**

Note: If you administer the defaultHelperNodes settings, you do not need to specify the -N parameter to distribute workload here.

You can confirm that data migration is happening across multiple nodes by looking at the following events:

- Each transparent cloud tiering node has policy files under /tmp.
- Each transparent cloud tiering node starts reflecting the migration metrics for the amount of data that is passed through it.

cloud services run maintenance automatically in maintenance windows in the background. Do not forget to review the maintenance windows and adjust them appropriately. For more information, see the *Setting up maintenance tasks* topic in the *IBM Storage Scale: Administration Guide*.

Note: If you change the IP address or network settings of any of the cloud services nodes once it is operational, do the following steps:

1. Stop the cloud services (by using the **mmcloudgateway service stop** command), the way it is mandated to stop IBM Storage Scale services while you change the node IP address.
2. To change the IP address, it is recommended to follow the steps that are provided in the *Changing IP addresses and host names* topic in the *IBM Storage Scale: Administration Guide*.

For more information, see *Configuring and tuning your system for Cloud services* topic in the *IBM Storage Scale: Administration Guide*.

Adding a cloud services node to an existing cloud services cluster

This topic describes the procedure for adding a cloud services client or a server to an existing cloud services cluster.

About this task

Procedure

1. Add the new node to the existing node class by using the **mmchnodeclass add** command. For example,

```
mmchnodeclass TCTNodeClass1 add -N node5
```

2. Install the cloud services RPMs (either client or server) on the node that is added. For more information, see [“Installation steps”](#) on page 628.
3. Designate the node as cloud services node.

```
mmchnode --cloud-gateway-enable -N node5 --cloud-gateway-nodeclass TCTNodeClass1
```

For more information, see the *Designating the Cloud services nodes* topic in the *IBM Storage Scale: Administration Guide*.

4. Start the cloud services software on the new node by using the **mmcloudgateway service start** command.

```
mmcloudgateway service start -N node5
```

Chapter 9. Installing and configuring IBM Storage Scale management API

The IBM Storage Scale management API implementation is based on the IBM Storage Scale GUI stack. Installation and configuration of GUI sets up the API infrastructure. That is, you do not need to perform any API specific deployment procedures.

Ensure that the following tasks are completed to start using the IBM Storage Scale management API to monitor and manage the IBM Storage Scale system:

1. Install the GUI server. Use any of the following options to install the GUI:
 - a. *Manually installing IBM Storage Scale management GUI in IBM Storage Scale: Concepts, Planning, and Installation Guide*
 - b. *Installing IBM Storage Scale management GUI by using the spectrumscale installation toolkit in IBM Storage Scale: Concepts, Planning, and Installation Guide*
2. After installing the system and GUI package, you need to create the first GUI user to log in to the GUI. This user can create other GUI administrative users to perform system management and monitoring tasks. When you launch the GUI for the first time after the installation, the GUI welcome page provides options to create the first GUI user from the command line prompt by using the `/usr/lpp/mmfs/gui/cli/mkuser <user_name> -g SecurityAdmin` command.
3. Create new users and provide the required permissions. You can either use the **Access > GUI Users** GUI page or the `usr/lpp/mmfs/gui/cli/mkuser` command to create new users and assign roles. For more information on how to create GUI users and assigning user roles, see *Managing GUI users* section in the *IBM Storage Scale: Administration Guide*.

Every API call to the IBM Storage Scale needs a valid GUI user and password. For example, `curl -k -u admin:admin001 -XGET -H content-type:application/json "https://<IP address>:443/scalemgmt/v2/info"`. The user roles assigned to the user determine actions that can be performed.

Chapter 10. Installing IBM Storage Scale native REST API

Use this information for installing IBM Storage Scale native REST API.

Manually installing IBM Storage Scale native REST API with the RPM

You can install the IBM Storage Scale native REST API feature by using a uniquely named set of RPM packages.

This feature is included in the following IBM Storage Scale self-extracting packages:

- IBM Storage Scale Standard Edition
- IBM Storage Scale Advanced Edition
- IBM Storage Scale Data Management Edition
- IBM Storage Scale Data Access Edition
- IBM Storage Scale Developer Edition

The installation packages include a dedicated directory named **scaleapi_rpms**, which contain all the RPMs that is needed to install this feature.

Before you begin, ensure that you are familiar with the manual installation steps described in the [“Manually installing the IBM Storage Scale software packages on Linux nodes” on page 477](#) section of the IBM Storage Scale product documentation.

After you install the traditional RPM packages based on the appropriate edition and license, install the additional `gpfs.scaleapi` RPM or DEB package from the `scaleapi_rpms` or `scaleapi_debs` directory. This package contains all components that are required for the native REST API feature. You must install the `gpfs.scaleapi` package by using the appropriate package manager **after** all other IBM Storage Scale packages are installed. After the native REST API feature specific rpm or deb package is installed, the administration daemon (called **scaleadmd**) is started with **systemd** through a post-installation scriptlet. To verify that the daemon is running, enter the following command:

```
systemctl status scaleadmd
```

Note: The administration daemon starts in an unconfigured state until the cluster is created by using the **scalectl** CLI.

After you install the traditional IBM Storage Scale RPM packages and IBM Storage Scale native REST API feature packages, you must build the GPL layer. After the GPL layer is built, the cluster can be created by using the `scalectl` command.

Installing IBM Storage Scale native REST API with the installation toolkit

You can use the installation toolkit to install the IBM Storage Scale native REST API feature.

Prerequisites

- Know how to use the installation toolkit for the installation of IBM Storage Scale. For more information, see [“Performing additional tasks using the installation toolkit” on page 567](#).
- Root and passwordless SSH access are necessary to deploy IBM Storage Scale native REST API cluster.

Procedure

1. Download and extract the IBM Storage Scale self-extracting package on the installer node.

```
/home/Storage_Scale_Advanced-5.2.3.0-x86_64-Linux-install --silent
```

2. Go to the IBM Storage Scale installation toolkit directory.

```
cd /usr/lpp/mmfs/5.2.3.0/ansible-toolkit/
```

3. Set up the installer node.

```
./spectrumscale setup -s <installerNodeIP>
```

4. Enable installation toolkit to install the administration daemon package.

```
./spectrumscale scaleadmd enable
```

5. Add and define the first node on the install toolkit configuration.

```
./spectrumscale node list
```

```
./spectrumscale node add node-1 -a
```

6. For multi-node cluster, add the remaining nodes.

```
./spectrumscale node list
```

```
./spectrumscale node add node-2
```

```
./spectrumscale node add node-3
```

```
./spectrumscale node add node-4
```

```
./spectrumscale node add node-5
```

```
./spectrumscale node add node-6
```

7. If required, create NSDs.

```
./spectrumscale node add node-1 -n
```

```
./spectrumscale node add node-2 -n
```

```
./spectrumscale nsd add /dev/sdd /dev/sde /dev/sdf -p node-1 -s node-2
```

8. Import the node identity certificate. For more information about generating a certificate, see *Create a node certificate* section in the *IBM Storage Scale: Administration Guide*.

```
./spectrumscale nodeid define --cert <path to certificate> --key <path to private key> --chain <path to ca chain> -N all
```

For example,

```
./spectrumscale nodeid define --cert /root/tls/server.pem --key /root/tls/server.key --chain /root/tls/ca.crt -N all
```

9. Install the IBM Storage Scale native REST API cluster:

```
./spectrumscale install --precheck
```

```
./spectrumscale install
```


After successful installation, deploy other required protocols and coexistence component by using the existing IBM Storage Scale installation toolkit. For more information, see [“Deploying protocols” on page 541](#).

Chapter 11. Installing GPUDirect Storage for IBM Storage Scale

IBM Storage Scale support for GPUDirect Storage (GDS) enables a direct path between GPU memory and storage.

Before you begin

For information about the prerequisites, see [“Planning for GPUDirect Storage”](#) on page 343.

About this task

Perform the following steps to install GDS:

Procedure

1. Install IBM Storage Scale with NSD servers for your file system. For more information, see [Chapter 4, “Installing IBM Storage Scale on Linux nodes and deploying protocols,”](#) on page 469.
2. Set up the InfiniBand fabric (or RoCE fabric by analogy). For details about the installation instructions, see [Mellanox OFED installation](#). Complete the following steps to configure InfiniBand fabric:
 - a) Use the latest MOFED driver. For details about the supported driver versions, see [Components required for GDS in IBM Storage Scale FAQ in IBM Documentation](#).
 - b) Set up IP over InfiniBand. For more information, see [Configuring InfiniBand and RDMA networks](#).
 - c) If your NSD servers are part of ESS, then reinstall the MOFED driver with the **--upstream-libs** flag. Use the MOFED version that ships with the ESS.
3. Install CUDA on the NSD clients running the GPUs. For more information about the supported version of CUDA, see [Components required for GDS in IBM Storage Scale FAQ in IBM Documentation](#).
For CUDA installation details, see [NVIDIA GPUDirect Storage Installation and Troubleshooting Guide](#).
4. Configure IBM Storage Scale. For more information, see [Configuring the GPFS cluster in IBM Storage Scale: Administration Guide](#).
5. Configure the NVIDIA components. For more information, see [Configuring GDS in IBM Storage Scale: Administration Guide](#).
6. Start IBM Storage Scale by using the **mmstartup -a** command.

What to do next

Checking the installation

Perform the following steps to check the installation of GDS:

1. Run the NVIDIA GDS utility **gdscheck -p** before you run the GDS workloads to verify the environment. You need Python3 installed on the node to run this utility.
2. Verify the status of PCIe Access Control Services (ACS) and PCIe Input/Output Memory Management Unit (IOMMU), as these components affect GDS function and performance. The output of the **gdscheck -p** utility must display the following status for IOMMU and ACS components:

```
IOMMU disabled
ACS disabled
```

3. Check for IBM Storage Scale support in the output of **gdscheck -p** as shown in the following example:

```
# gdscheck -p | grep "Spectrum Scale"
IBM Spectrum Scale : Supported
```

4. If the GPFS log (mmfs log) contains the following information after you start IBM Storage Scale, it indicates that the GDS support is successfully enabled:

```
[I] VERBS DC API loaded.
```

and

```
[I] VERBS DC API initialized.
```

Related concepts

[“GPUDirect Storage support for IBM Storage Scale” on page 28](#)

IBM Storage Scale's support for NVIDIA's GPUDirect Storage (GDS) enables a direct path between GPU memory and storage. This solution addresses the need for higher throughput and lower latencies. File system storage is directly connected to the GPU buffers to reduce latency and load on the CPU. For IBM Storage Scale, this means that data can be read or written directly from / to an NSD server's pagepool and it is sent to the GPU buffer of the IBM Storage Scale clients by using RDMA. IBM Storage Scale with GDS requires an InfiniBand or RoCE fabric. In IBM Storage Scale, the **mmdiag** command is enhanced to print diagnostic information for GPUDirect Storage.

[“Planning for GPUDirect Storage” on page 343](#)

IBM Storage Scale support for GPUDirect Storage (GDS) enables a direct path between GPU memory and storage. You need to ensure that certain conditions are met before you start installing the feature.

Related tasks

[“Upgrading GPUDirect Storage” on page 692](#)

The recommended version for use of GDS with IBM Storage Scale is 5.1.6 as it supports accelerated writes.

Chapter 12. Installation of Active File Management (AFM)

Consider the following while you install Active File Management (AFM).

The home and cache are two separate operational clusters that are separated from each other through LAN or WAN. The cluster setup has no considerations for using AFM functions. AFM functions are available on all editions of IBM Storage Scale.

The home cluster can be a legacy NAS or a cluster that is running on an IBM Storage Scale cluster. User-extended attributes, ACLs, and sparse files are not supported on a home target that is running a legacy NFS export.

The NFS version 3 and 4 or GPFS protocol can be used for communication. The native NSD protocol can be used when the home cluster is also an IBM Storage Scale cluster and the cache can remotely mount the file system on the home cluster.

Nodes that can act as gateway nodes must be identified on the cache cluster. Gateway nodes must preferably be configured in the cache cluster before you create filesets and start applications.

Nodes that can act as NFS servers must be identified on the home cluster. If the GPFS protocol is planned to be used, the home file system must be remote mounted on all the gateway nodes in the cache cluster. Gateway nodes and NFS servers must preferably be configured in the cache and home clusters before you create filesets and start applications.

AIX AFM gateway nodes in the AFM cache cluster must be running the Linux operating system.. AIX cannot be configured either as a home or AFM-DR secondary.

Chapter 13. Installing AFM Disaster Recovery

The primary and the secondary are two separate operational clusters that are separated from each other through LAN or WAN. AFM-DR functions are available in Advanced Edition, Data Management Edition, Developer Edition, or IBM Storage Scale Erasure Code Edition. All nodes in both the primary and secondary clusters must be running the Advanced Edition, Data Management Edition, Developer Edition, or Erasure Code Edition.

The NFS version 3 and 4 or native NSD protocol can be used for communication. Similar to AFM filesets, gateway nodes and NFS servers must be planned on the primary and secondary clusters. Gateway nodes and NFS servers must preferably be configured in the primary and secondary clusters before you create filesets and start applications. If you are planning to use native NSD protocol, the primary file system must be remote mounted on all the gateway nodes in the secondary cluster.

Note: While you upgrade AFM-DR to IBM Storage Scale 4.2.2 or later, the primary cluster must be upgraded before you upgrade the secondary cluster.

Chapter 14. Installing call home

The call home feature collects files, logs, traces, and details of certain system health events from different nodes and services. These details are shared with the IBM support center for monitoring and problem determination.

Prerequisites for installing call home

The following criteria must be met to use the call home function on your GPFS cluster:

- Red Hat Enterprise Linux 7.x, Red Hat Enterprise Linux 8.x, Ubuntu 20.04 LTS, and SLES 15 nodes are supported.
- Intel x86_64, PowerPC LE/BE and IBM Z are supported.
- GPFS CCR (Clustered Configuration Repository) must be enabled.
- The call home node must be able to access the following IP addresses and ports. In the configurations with a proxy, these IP addresses and ports must be accessible over the proxy connection.
 - Host name: `esupport.ibm.com`
 - IP address: 129.42.56.189, 129.42.60.189, 129.42.54.189, 129.42.21.70, 129.42.18.70, 129.42.19.70
 - Port number: 443

The recommendation is to open 129.42.0.0/18.

`esupport.ibm.com` works for the Call Home uploads as a proxy by forwarding data to ECuRep. Therefore, allowing any direct connections to ECuRep IPs is not required.

- The call home node needs to have at least 1 GB, preferably 3 GB of free space in the `dataStructureDump` directory. The current `dataStructureDump` directory can be determined by using the following command:

```
mmdiag --config | grep "dataStructureDump "
```

The current `dataStructureDump` directory is `tmp/mmfs` by default. You can change the default value of the `dataStructureDump` directory by using the **mmchconfig** command.

- There must not be multiple nodes with the same short host name within the GPFS cluster.
- In the case of the customer provided remote shell or remote file-copy commands, the customer must make sure that these commands support the following options:

```
'-p -o BatchMode=yes -o ConnectTimeout=3'
```

Note: `ssh` and `scp` support these options.

The configured commands can be seen by running the following command:

```
mmclscluster | grep Remote
Remote shell command: /usr/bin/ssh
Remote file copy command: /usr/bin/scp
```

- When you are using the `sudo` wrappers and a custom `dataStructureDump` directory, you must ensure that this directory is recursively executable for the `sudo` user, as specified in the IBM Storage Scale settings.

Note: In this case, the `dataStructureDump` directory is not set to the default value of `/tmp/mmfs`.

Steps for installing call home

The call home functionality is included in the `gpfs.base` package. If the IBM Storage Scale packages are installed and the cluster is created, no other steps are required for installing call home.

The installation toolkit can be used for removing call home RPMs from the old IBM Storage Scale versions, and for enabling and configuring the call home feature. For more information, see [“Enabling and configuring call home using the installation toolkit”](#) on page 548.

For more information on configuring call home manually, see the *Monitoring the IBM Storage Scale system by using call home* section in the *IBM Storage Scale: Problem Determination Guide*.

Chapter 15. Installing file audit logging

Before file audit logging can be installed, consider the requirements, limitations, and steps associated with it.

For more information about enabling file audit logging, see *Configuring file audit logging* in the *IBM Storage Scale: Administration Guide*.

Requirements, limitations, and support for file audit logging

Use this information to understand requirements, limitations, and support for installing file audit logging.

OS requirements

File audit logging is supported on all Linux OS versions supported by IBM Storage Scale. For more information, refer to [IBM Storage Scale FAQ in IBM Documentation](#).

Security requirements and limitations

- Root authority is required to run **mmaudit**.

Restrictions imposed by mixed environments and protocols

- Events generated on non-Linux nodes are not audited.
- IBM Storage Scale file audit logging has full support for the following protocols (support for all other protocols must be considered limited):
 - NFS ganesha
 - SMB
 - Native UNIX file access
- Events are not generated at or below the `cesSharedRoot` path.

File audit logging attributes availability and limitations

- For more information about the availability and limitations of the file audit logging attributes, see the [“JSON attributes in file audit logging” on page 252](#) topic.
- File audit logging in AFM supports only the following scenarios:
 - File audit logging is enabled on home and cache clusters but each cluster uses a file audit logging fileset that is local to the cluster and is not an AFM fileset. That is, the file audit logging fileset of the home cluster is a local fileset on the home cluster and the file audit logging fileset of the cache cluster is a local fileset on the cache cluster. In this scenario, the AFM fileset events on the home or cache cluster appear in the home file audit logging fileset and the cache file audit logging fileset. Also, when the **mmdiag --eventproducer** command is issued, the AFM fileset events appear on both clusters.
 - The audit logging fileset of the home cluster is AFM enabled and exported to the cache cluster in the read-only mode. In this scenario, file audit events on the home cluster within the same AFM enabled file audit logging fileset appear on both the home and the cache cluster. When the **mmdiag --eventproducer** command is issued, the file audit events appear only on the home cluster.
- Specifying an AFM fileset as a file audit logging fileset for the cache cluster is not supported.

GPFS file system requirements and limitations

- File audit logging can be enabled only for file systems that are created or upgraded to IBM Storage Scale 5.0.0 or later.
- Space provisioning must be considered to store the generated events in the `.audit_log` fileset.
- The `.audit_log` fileset is protected from tampering. It cannot be easily deleted to free up space in the file system. This is done by creating the fileset in the IAM noncompliant mode (default) or

compliant mode, which allows expiration dates to be set on the files containing the audit records within the fileset. If the fileset is created in IAM mode noncompliant, then the root user can change the expiration date to the current date so that audit files can be removed to free up disk space. If the fileset is created in IAM mode compliant (because of the use of the **--compliant** option), not even the root user can change the expiration dates of the audit logging files and they cannot be removed until the expiration date.

- Events are not generated for file system activity within the file audit logging fileset itself.
- There is a limit of 20 filesets for file system prior to 5.1.3 (27.0 file system version).

GPFS and spectrumscale functional limitations

- Conversion of a file audit logging fileset to AFM DR is not supported.
- When file audit logging or clustered watch folder is enabled on a file system, changing the file system name or deleting the file system is not allowed. To change file system name or delete the file system, file audit logging and clustered watch folder must first be disabled.

Miscellaneous requirements, limitations, recommendations, and support statements

- File audit logging is available in IBM Storage Scale Advanced Edition, IBM Storage Scale Data Management Edition, IBM Storage Scale Developer Edition, or IBM Storage Scale Erasure Code Edition.
- File audit logging is supported in SELinux enforcing, permissive, and disabled modes. When file audit logging is run in enforcing mode, there is an extra event generated that is related to attributes due to the SELinux labeling of files.
- Event generation occurs in the I/O path. To avoid delay in I/O operations during the authentication and the configuration, I/O operations are not be audited. On a per node basis, the authentication and the configuration occurs on the first I/O after enabling audit, starting mmfsd, and every time the file system is mounted.
- File audit logging uses buffered IO. This means that when a file system operation generates an event, the event is then appended to a buffer rather than writing directly to the disk. Therefore, if a file system operation returns successfully, it is not guaranteed that the operation would be logged in the audit log. For example, if the node goes down after an operation completes, but before the buffered IO is written to disk, then that operation will never be written to the audit log. Another example is when the buffers we use in memory to store events before writing to disk become full. Rather than throttle the filesystem waiting for buffer space to free up, we will try for a short period of time to find available buffer space before dropping the event.
- File audit logging is not supported by Kafka message queue in IBM Storage Scale 5.1.1 or later.
- The **mmrestorefs** command does not restore files that are located in the `.audit_log` fileset or configuration fileset. The current configuration would not be overwritten and audit records would not be removed or restored.

Requirements for using file audit logging with remotely mounted file systems

Use this information to determine the requirements for using file audit logging with remotely mounted file systems.

All clusters that own or access file systems that are being audited by file audit logging through remote mounts must meet the following requirements:

- All clusters must be at IBM Storage Scale 5.0.2 or later.
- For any cluster that is upgraded to IBM Storage Scale 5.0.2 or later, the **mmchconfig release=LATEST** command must be run.
- Any file system that is audited must be at IBM Storage Scale 5.0.2 or later. If upgraded, the **mmchfs -V full** command must be run after the cluster is installed at or upgraded to IBM Storage Scale 5.0.2 or later.

- The **mmaudit** command can be run only on the cluster that owns the file system with file audit logging enabled.
- Clusters that are remotely mounting file systems with file audit logging enabled must be upgraded before the owning cluster.

Chapter 16. Installing clustered watch folder

Before clustered watch folder can be installed, consider the requirements, limitations, and steps associated with it.

For more information, see [“Clustered watch folder” on page 256](#).

Requirements, limitations, and support for clustered watch folder

Use this information to understand the requirements, limitations, and support for installing clustered watch folder.

RPM and package requirements

In order for watch folder events to be generated from a node, the node must have the following packages installed:

- For RHEL and SLES, the `librdkafka` package requires the `openssl-devel` and `cyrus-sasl-devel` packages.
- For Ubuntu, the `librdkafka` package requires the `libssl-dev` and `libsasl2-dev` packages.
- `librdkafka` (`gpfs.librdkafka rpm/package`).

OS and hardware requirements

Note: Not all of these requirements apply to the support of remotely mounted file systems.

Watch folder is supported on all Linux® OS versions supported by IBM Storage Scale. For more information, refer to [IBM Storage Scale FAQ in IBM® Documentation](#).

Security requirements and limitations

- Root authority is required to run `mmwatch`.

Restrictions imposed by mixed environments and protocols

- Events will not be generated on non-Linux nodes.
- IBM Storage Scale clustered watch folder has full support for the following protocols (support for all other protocols must be considered limited):
 - NFS-Ganesha
 - SMB
 - Native UNIX file access

Clustered watch folder attributes

For more information about the availability and limitations of the clustered watch folder attributes, see [“JSON attributes in clustered watch folder” on page 258](#).

GPFS file system requirements and limitations

- Clustered watch folder can be enabled only for file systems that are created or upgraded to IBM Storage Scale 5.0.3 or later.
- When clustered watch folder is enabled on a file system, changing the file system name is not allowed. Clustered watch folder must first be disabled to change the file system name.

Miscellaneous requirements, limitations, recommendations, and support statements

- Clustered watch folder is available in IBM Storage Scale Advanced Edition, IBM Storage Scale Data Management Edition, IBM Storage Scale Developer Edition, or IBM Storage Scale Erasure Code Edition.
- Only 25 cluster watches are supported per cluster.

Note:

- For more information about setting up a basic external sink to push events to, see <https://kafka.apache.org/quickstart>.
- If you want to use a specific network interface for the brokers in the external Kafka sink, specify the **advertised.listeners** configuration parameter for the brokers with the IP address on the specified network interface card.
- Clustered watch folder is supported in SELinux enforcing, permissive, and disabled modes. When clustered watch folder is run in enforcing mode, there is an extra event generated that is related to attributes due to the SELinux labeling of files.
- Event generation occurs in the I/O path. So, to avoid delaying I/O operations while authentication and configuration occurs, the I/O operations that occur during that time span are not audited. On a per node basis, authentication and configuration occurs on the first I/O path after clustered watch is enabled, during initialization of the mmfsd server, and every time the file system is mounted.
- Events of clustered watch are sent in batches to the Kafka brokers. This means that when a file system operation generates an event, the event is then appended to a batch rather than sent directly to Kafka. Therefore, if a file system operation returns successfully, it is not guaranteed that the operation is sent to the external Kafka queue. For example, if the node goes down after an operation completes but before the event batch is sent, that operation is never be received by Kafka. Rather than throttle the file system while waiting for Kafka to receive events, the file system attempts (for a short period of time) to send the events before dropping them.
- Do not deploy Kafka on I/O nodes, as it might affect system performance. It is recommended to deploy Kafka on separate infrastructure.

Requirements for using clustered watch folder with remotely mounted file systems

These requirements must be followed to use clustered watch folder with remotely mounted file systems.

All clusters that own or access file systems that are using clustered watch folder through remote mounts must meet the following requirements:

- All clusters must be at IBM Storage Scale 5.0.3 or later.
- For any cluster that is upgraded to IBM Storage Scale 5.0.3 or later, the **mmchconfig release=LATEST** command must be run.
- Any file system that is watched must be at IBM Storage Scale 5.0.3 or later. If upgraded, the **mmchfs -V full** command must be run after the cluster is installed at or upgraded to IBM Storage Scale 5.0.3 or later.
- The **mmwatch** command can be run on the cluster that owns the file system that is using clustered watch folder only.
- Clusters that are remotely mounting file systems that are using clustered watch folder must be upgraded before the owning cluster.

Manually installing clustered watch folder

Use this information to manually install clustered watch folder on your system.

Important: Before you can install clustered watch folder, you must have the required packages for IBM Storage Scale installed on the node.

Ensure that the following software package is installed on every node in the cluster that acts as producer.

- For Red Hat Enterprise Linux:
 - `gpfs.librdkafka-5.2*.rpm`

Note: This package is available in `/usr/lpp/mmfs/5.2.x.x/gpfs_rpms/rhelX`.

- For Ubuntu Linux:

– `gpfs.librdkafka-5.2*.deb`

Note: This package is available in `/usr/lpp/mmfs/5.2.x.x/gpfs_debs/ubuntu`.

Note: For more information, see [“Requirements, limitations, and support for clustered watch folder”](#) on page 651.

1. If you are installing the package for clustered watch folder on existing IBM Storage Scale nodes and this is not part of a cluster installation or upgrade, you must shut down and then restart the daemon in order for the producer to be able to load the required library.
2. After the IBM Storage Scale cluster is available, and the file systems are mounted, use the **mmwatch** command for each file system that you want to enable for clustered watch folder.

Note: For more information about enabling clustered watch folder, see the *mmwatch* command in the *IBM Storage Scale: Command and Programming Reference Guide*.

Chapter 17. Installing the signed kernel modules for UEFI secure boot on x86_64

Starting with 5.1.9.0, IBM Storage Scale supports the UEFI secure boot feature on x86_64 systems. Secure boot is a UEFI verification mechanism that ensures the code launched by the system firmware is trusted. IBM Storage Scale uses this feature to cryptographically sign its kernel modules so that their integrity can be verified during system startup.

Before you begin

For more information, see [“Signed kernel modules for UEFI secure boot on x86_64 and secure boot Linux on Z” on page 8](#).

Ensure that the prerequisites for installing signed kernel modules are complete. For more information, see [“Planning for UEFI secure boot on x86_64 and secure boot on Linux on Z” on page 462](#).

About this task

To install the signed kernel modules for IBM Storage Scale on x86_64 systems, complete the following steps:

Procedure

1. Install the IBM Storage Scale packages, except for the `gpfs.gplbin` RPM. You can install the packages manually (RPMs) or by using the installation toolkit. For more information, see [Chapter 4, “Installing IBM Storage Scale on Linux nodes and deploying protocols,” on page 469](#).
2. If the `gpfs.gplbin` RPM for your kernel version is not yet available, obtain it before proceeding. For more information, see [“Signed kernel modules for UEFI secure boot on x86_64 and secure boot Linux on Z” on page 8](#)

Note: The `gpfs.gplbin` RPM is available in the extracted IBM Storage Scale installation package or from [IBM Fix Central](#). If the required version is not available, contact IBM Support to request a compatible build.

3. Enable module signature validation by issuing the following command:

```
mokutil --enable-validation
```

When the **input password:** prompt appears, enter a temporary one-time password. You will use this password to confirm and enable module signature validation in the machine owner key (MOK) manager during the next system reboot.

4. Convert the IBM provided public key to DER format.

```
openssl x509 -in <Storage_Scale_public_key.pem.cer> -inform PEM -out <scale.cer> -outform DER
```

5. Import the converted public key.

```
mokutil --import <scale.cer>
```

When the **input password:** prompt appears, enter a temporary one-time password. You will use this password to authorize the public key enrollment in the MOK manager during the next boot. For more information, see [Enrolling public key on target system by adding the public key to the MOK list in the Red Hat documentation](#).

6. Reboot the IBM Storage Scale node.

```
sudo reboot
```

7. During reboot, when the **Press any key to perform MOK Management** message appears, press any key to launch the MOK manager UI.
8. In the MOK Manager UI:
 - a) Select **Enroll MOK** to enroll the public key in UEFI key management.
 - b) When the **password:** prompt appears, enter the one-time password that was set during the **mokutil --import <scale.cer>** to authorize the key.
 - c) On the **Enroll the key?** screen, select **Yes** and press Enter.
 - d) On the **Enable Secure Boot validation?** screen, select **Yes** and press Enter.
 - e) When the **password:** prompt appears, enter the one-time password that was set during the **mokutil --enable-validation** to enable the secure boot validation.
 - f) When the **The system will reboot now** message appears, press **Enter**.
9. After the system reboots, verify that the public key is successfully enrolled.

```
mokutil --list-enrolled
```

10. Install the `gpfs.gplbin` package.

```
rpm -ivh gpfs.gplbin-5.14.0-284.25.1.el9_2.x86_64-6.0.1.0.x86_64.rpm
```

Where,

- 5.14.0-284.25.1 is the kernel for which the modules in the RPM are built.
- `el9_2.x86_64` stands for RHEL 9.2 on `x86_64`.
- 6.0.1.0 refers to the IBM Storage Scale release.

11. Start IBM Storage Scale services.

```
mmstartup
```

A sample output is as follows:

```
Thu Aug 12 13:22:40 EDT 2004: 6027-1642 mmstartup: Starting GPFS ...
```

Chapter 18. Installing the signed kernel modules for secure boot on Linux on Z

About this task

Starting with 5.2.1.0, IBM Storage Scale supports the secure boot feature on Linux on Z. For IBM Storage Scale, using the secure boot means that the kernel modules are cryptographically signed by IBM so that their integrity can be verified when the system starts.

Follow the next steps to install the signed kernel modules for IBM Storage Scale on Linux on Z.

Procedure

1. Verify that the virtual machine (VM) or logical partition (LPAR) of the Linux on Z system supports secure boot.

Use the following command:

```
cat /sys/firmware/ipl/has_secure
1
```

2. Download from [Fix Central](#) the RPM package that holds the signed kernel modules. The public key is either part of the self-extracting package or it is included as part of the signed kernel stand-alone package.
3. Import the public key by using the Linux on Z HMC.
4. Assign the public key to the appropriate VM or LPAR.

For more information, see [Managing a secure boot certificate for an LPAR in Linux on IBM Systems documentation](#).

5. Restart the VM or LPAR and verify that the system is in secure boot mode by issuing the next command:

```
cat /sys/firmware/ipl/secure
1
```

For information about how to restart the VM or LPAR in secure boot mode, see [Initiating secure boot in Linux on IBM Systems documentation](#).

6. Check that the key import was successful by issuing the command:

```
keyctl list %:.platform
```

7. Install the `gpfs.gplbin` RPM.

Chapter 19. Uninstalling IBM Storage Scale

IBM Storage Scale maintains a number of files that contain configuration and file system related data. Since these files are critical for the proper functioning of IBM Storage Scale and must be preserved across releases, they are not automatically removed when you uninstall GPFS.

About this task

Follow these steps if you do not intend to use IBM Storage Scale on any of the nodes in your cluster and you want to remove all traces of IBM Storage Scale:

Attention: After following these steps and manually removing the configuration and file system related information, you will permanently lose access to all of your current IBM Storage Scale data.

1. Unmount all GPFS file systems on all nodes by issuing the `mmumount all -a` command.
2. Issue the `mmde1fs` command for each file system in the cluster to remove GPFS file systems.
3. Issue the `mmde1nsd` command for each NSD in the cluster to remove the NSD volume ID from the device.

If the NSD volume ID is not removed and the disk is again used with GPFS at a later time, you receive an error message when you specify the `mmcrnsd` command. For more information, see *NSD creation fails with a message referring to an existing NSD* in *IBM Storage Scale: Problem Determination Guide*.

4. Issue the `mmshutdown -a` command to shutdown GPFS on all nodes.
5. Uninstall IBM Storage Scale from each node:
 - For your Linux nodes, run the uninstallation program to remove GPFS for the correct version of the packages for your hardware platform and Linux distribution. For example, on SLES and Red Hat Enterprise Linux nodes:

```
rpm -e gpfs.crypto (IBM Storage Scale Advanced Edition, IBM Storage Scale Data
Management Edition, or IBM Storage Scale Developer Edition only)
rpm -e gpfs.gpl
rpm -e gpfs.license.xx
rpm -e gpfs.msg.en_US
rpm -e gpfs.compression
rpm -e gpfs.librdkafka
rpm -e gpfs.java
rpm -e gpfs.afm.cos
rpm -e gpfs.docs
rpm -e gpfs.gskit
rpm -e gpfs.adv (IBM Storage Scale Advanced Edition, IBM Storage Scale Data Management
Edition, or IBM Storage Scale Developer Edition only)
rpm -e gpfs.base
```

On Ubuntu Linux nodes:

```
dpkg -P gpfs.crypto (IBM Storage Scale Advanced Edition and IBM Storage Scale Data
Management Edition only)
dpkg -P gpfs.gpl
dpkg -P gpfs.license.xx
dpkg -P gpfs.msg.en-us
dpkg -P gpfs.compression
dpkg -P gpfs.librdkafka
dpkg -P gpfs.java
dpkg -P gpfs.afm.cos
```

```
dpkg -P gpfs.docs
dpkg -P gpfs.gskit
dpkg -P gpfs.adv (IBM Storage Scale Advanced Edition and IBM Storage Scale Data
Management Edition only)
dpkg -P gpfs.base
```

- For your AIX nodes:

```
installp -u gpfs
```

- For your Windows nodes, follow these steps:

- a. Open Programs and Features in the Control Panel.
- b. Uninstall IBM General Parallel File System.
- c. Reboot the system.
- d. From a Command Prompt, run the following command:

```
sc.exe delete mmwinserv
```

6. Remove the /var/mmfs and /usr/lpp/mmfs directories.
7. Remove all files that start with mm from the /var/adm/ras directory.
8. Remove /tmp/mmfs directory and its content, if present.

Cleanup procedures required if reinstalling with the installation toolkit

Before you can reinstall with the installation toolkit, you must perform some cleanup procedures first.

Note:

CES Swift Object protocol feature is not supported on IBM Storage Scale 5.2.0 and later. If the Swift Object protocol is configured on a cluster, refer the [this cleanup procedure](#) in [IBM Storage Scale 5.2.0 documentation](#).

You can use the following procedures to clean up various stages of the previous GPFS installation and protocols deployment.

Starting the cleanup process

1. Clean up the installer directory by issuing the following command:

```
mmdsh -N NodeList rm -rf /usr/lpp/mmfs/6.0.1.x
```

Note: The installer directory in this command depends upon the release version.

2. Check /etc/resolv.conf for any changes that pointed the authentication directory to DNS, and remove them.

Uninstalling Ansible

The installation toolkit is built on the Ansible automation platform. For more information, see [“Supported Ansible version”](#) on page 528.

You might need to uninstall Ansible as a part of the cleanup process or if the supported version is not installed. Check the Ansible version that is installed and, if needed, uninstall Ansible after verifying that it is not being used for any other purpose.

1. Check the version of Ansible that is installed by using one of the following commands.

```
ansible --version
```


or

```
rpm -qa | grep ansible
```

2. Uninstall Ansible by using one of the following commands depending on your environment.

- Uninstall by using the **rpm** command on RHEL and SLES nodes.

```
rpm -qa | grep ansible | xargs rpm -e
```

- Uninstall by using the **yum** command on RHEL nodes.

```
yum remove ansible
```

- Uninstall by using the **zypper remove** command on SLES nodes.

```
zypper remove ansible
```

- Uninstall by using the **pip3** command on RHEL, SLES, and Ubuntu nodes.

```
pip3 uninstall ansible
```

- Uninstall by using the **apt-get** command on Ubuntu nodes.

```
apt-get purge ansible
```

Cleaning up the authentication configurations

Use these commands to completely remove the current authentication configuration for all protocols, NFS/SMB only. Be aware that doing so may cause loss of access to previously written data.

1. Remove the configurations. To do so, issue the following commands:

```
mmuserauth service remove --data-access-method file  
mmuserauth service remove --data-access-method object
```

2. Delete the ID-mapping information when authentication was configured with AD. To do so, use the following command.

Note: Once this ID-mapping information is deleted, you might not be able to access existing files (or you might experience other unforeseen access issues), so do this only if you do not need access to existing data.

```
mmuserauth service remove --data-access-method file --idmapdelete
```

Disabling CES

Disable CES on every node *nodeNameX*. To do so, issue the following command:

```
mmchnode -N nodeNameX --ces-disable
```

Protocol cleanup steps

1. Clear the *cesSharedRoot* configuration by issuing the following command:

```
mmchconfig cesSharedRoot=DEFAULT
```

2. Check your recent package installs to determine which packages need to be removed by issuing the following commands:

RHEL and SLES

```
rpm -qa --last|more
```

Ubuntu

```
dpkg -l
```

After you have determined which packages need to be removed, do so by issuing the following commands, as needed:

RHEL

```
mmdsh -N NodeList yum erase gpfs.smb nfs-ganesha-mount \
nfs-ganesha-vfs PyQt4 nfs-ganesha-utils phonon-backend-gstreamer phonon sip qt-x11 kde-filessystem \
qt qt-settings libmng nfs-ganesha-proxy nfs-ganesha-nullfs nfs-ganesha-gpfs -y

mmdsh -N NodeList yum erase python-webob MySQL-python mariadb-server perl-DBD-MySQL mariadb \
python-futures python-warlock python-jsonpatch python-jsonschema python-jsonpointer python-cmd2 \
python-cliff pyparsing memcached python-oslo-utils python-oslo-serialization python-oslo-i18n \
python-babel babel python-keyring python-stevedore python-pytableparser python-pbr python-oslo-config \
python-netaddr python-iso8601 python-dnspython python-urllib3 python-six python-requests \
python-paste-deploy python-tempita python-paste python-netifaces python-simplejson \
python-greenlet python-eventlet -y

mmdsh -N NodeList yum erase python2-keystoneauth1 python2-keystoneclient python2-keystonemiddleware \
python2-openstackclient python2-oslo-cache python2-oslo-concurrency python2-oslo-config \
python2-oslo-context python2-oslo-db python2-oslo-i18n python2-oslo-middleware python2-oslo-policy \

python3-keystonemiddleware python3-openstackclient python3-openstacksdk python3-oslo-cache \
python3-oslo-concurrency python3-oslo-config python3-oslo-context python3-oslo-db python3-oslo-i18n \
python3-oslo-log python3-oslo-messaging python3-oslo-middleware python3-oslo-policy \
python3-oslo-serialization python3-oslo-service python3-oslo-upgradecheck python3-oslo-utils \

python-oslo-cache-lang python-oslo-concurrency-lang python-oslo-db-lang python-oslo-i18n-lang \
python-oslo-log python-oslo-log-lang python-oslo-messaging python-oslo-middleware-lang \

mmdsh -N NodeList yum erase gpfs.smb-debugsource -y

mmdsh -N NodeList yum erase sssd-tools krb5-workstation

mmdsh -N NodeList yum erase python-routes python-repoze-lru python-oauthlib python-crypto \
python-qpid-common python-qpid python-dogpile-cache python-fixtures python-dogpile-core \
python-testtools python-extras python-oslo-context python-kombu python-anyjson python-amqp \
python-mimeparse python-markupsafe PyYAML python-passlib libyaml python-ibm-db-sa python-ibm-db \
python-sqlparse python-memcached python-posix_ipc python-sqlalchemy -y

mmdsh -N NodeList yum erase gpfs.gss.pmsensors gpfs.gss.pmcollector gpfs.pm-ganesha boost-regex -y

mmdsh -N NodeList yum erase openstack-utils -y

mmdsh -N NodeList yum erase python-ordereddict -y

mmdsh -N NodeList yum erase mariadb-libs -y

python-keystone python-cinderclient python-glanceclient python-keystoneclient python-neutronclient \
python-novaclient python-keystoneauth1 python-openstackclient -y
```

SLES

```
mmdsh -N NodeList zypper remove gpfs.smb nfs-ganesha-mount \
nfs-ganesha-vfs PyQt4 nfs-ganesha-utils phonon-backend-gstreamer phonon sip qt-x11 kde-filessystem \
qt qt-settings libmng nfs-ganesha-proxy nfs-ganesha-nullfs nfs-ganesha-gpfs -y

mmdsh -N NodeList zypper remove python-webob MySQL-python mariadb-server perl-DBD-MySQL mariadb \
python-futures python-warlock python-jsonpatch python-jsonschema python-jsonpointer python-cmd2 \
python-cliff pyparsing memcached python-oslo-utils python-oslo-serialization python-oslo-i18n \
python-babel babel python-keyring python-stevedore python-pytableparser python-pbr python-oslo-config \
python-netaddr python-iso8601 python-dnspython python-urllib3 python-six python-requests \
python-paste-deploy python-tempita python-paste python-netifaces python-simplejson \
python-greenlet python-eventlet -y

mmdsh -N NodeList zypper remove python2-keystoneauth1 python2-keystoneclient python2-keystonemiddleware \
python2-openstackclient python2-oslo-cache python2-oslo-concurrency python2-oslo-config \
python2-oslo-context python2-oslo-db python2-oslo-i18n python2-oslo-middleware python2-oslo-policy \

python3-keystonemiddleware python3-openstackclient python3-openstacksdk python3-oslo-cache \
python3-oslo-concurrency python3-oslo-config python3-oslo-context python3-oslo-db python3-oslo-i18n \
python3-oslo-log python3-oslo-messaging python3-oslo-middleware python3-oslo-policy \
python3-oslo-serialization python3-oslo-service python3-oslo-upgradecheck python3-oslo-utils \

python-oslo-cache-lang python-oslo-concurrency-lang python-oslo-db-lang python-oslo-i18n-lang \
python-oslo-log python-oslo-log-lang python-oslo-messaging python-oslo-middleware-lang \

mmdsh -N NodeList zypper remove gpfs.smb-debugsource -y
```

```

mmdsh -N NodeList zypper remove sssd-tools krb5-workstation

mmdsh -N NodeList zypper remove python-routes python-repoze-lru python-oauthlib python-crypto \
python-qpidd-common python-qpidd python-dogpile-cache python-fixtures python-dogpile-core \
python-testtools python-extras python-oslo-context python-kombu python-anyjson python-amqp \
python-mimeparse python-markupsafe PyYAML python-passlib libyaml python-ibm-db-sa python-ibm-db \
python-sqlparse python-memcached python-posix_ipc python-sqlalchemy -y

mmdsh -N NodeList zypper remove gpfs.gss.pmsensors gpfs.gss.pmcollector gpfs.pm-ganesha boost-regex -y

mmdsh -N NodeList zypper remove openstack-utils -y

mmdsh -N NodeList zypper remove python-ordereddict -y

mmdsh -N NodeList zypper remove mariadb-libs -y

python-keystone python-cinderclient python-glanceclient python-keystoneclient python-neutronclient \
python-novaclient python-keystoneauth1 python-openstackclient -y

```

Ubuntu

```

mmdsh -N NodeList dpkg -P gpfs.smb nfs-ganesha-mount nfs-ganesha-vfs \
PyQt4 nfs-ganesha-utils phonon-backend-gstreamer phonon sip qt-x11 kde-filesystem qt qt-settings libmng \
nfs-ganesha-proxy nfs-ganesha-nullfs gpfs.nfs-ganesha-gpfs gpfs.nfs-ganesha -y

mmdsh -N NodeList dpkg -P python-webob MySQL-python mariadb-server perl-DBD-MySQL mariadb \
python-futures python-warlock python-jsonpatch python-jsonschema python-jsonpointer python-cmd2 \
python-cliff pyparsing memcached python-oslo-utils python-oslo-serialization python-oslo-i18n \
python-babel babel python-keyring python-stevedore python-pytablewriter python-pbr python-oslo-config \
python-netaddr python-iso8601 python-dnspython python-urllib3 python-six python-requests \
python-paste-deploy python-tempita python-paste python-netifaces python-simplejson \
python-greenlet python-eventlet -y

mmdsh -N NodeList dpkg -P python-routes python-repoze-lru python-oauthlib python-crypto \
python-qpidd-common python-qpidd python-dogpile-cache python-fixtures python-dogpile-core \
python-testtools python-extras python-oslo-context python-kombu python-anyjson python-amqp \
python-mimeparse python-markupsafe PyYAML python-passlib libyaml python-ibm-db-sa python-ibm-db \
python-sqlparse python-memcached python-posix_ipc python-sqlalchemy -y

mmdsh -N NodeList dpkg -P gpfs.gss.pmsensors gpfs.gss.pmcollector gpfs.pm-ganesha boost-regex -y

mmdsh -N NodeList dpkg -P openstack-utils gpfs.gpfs.nfs-ganesha-gpfs -y

mmdsh -N NodeList dpkg -P python-ordereddict -y

mmdsh -N NodeList dpkg -P mariadb-libs -y

```

File protocols authentication packages cleanup

RHEL

```

mmdsh -N NodeList yum erase bind-utils krb5-workstation openldap-clients sssd-tools ypbind yp-tools -y

```

SLES

```

mmdsh -N NodeList zypper remove bind-utils krb5-client openldap2-client sssd-tools ypbind yp-tools -y

```

Ubuntu

```

mmdsh -N NodeList dpkg -P dnsutils krb5-user ldap-utils sssd-tools nis libslp1 -y

```

Note: Plan to erase the SSSD packages only if you do not plan to use the SSSD service on the nodes for any other purpose.

SSSD packages cleanup commands are as follows:

RHEL

```

mmdsh -N NodeList yum erase sssd*

```

SLES

```

mmdsh -N NodeList zypper remove sssd*

```

Ubuntu

```
mmdsh -N NodeList dpkg -P sssd*
```

3. Issue **rpm -qa --last|more** or **dpkg -l** on all nodes again, as the preceding list does not contain everything in the new builds.

You might also need to remove `gpfs.smb`.

4. Clean up performance monitoring tool files. To do so, issue the following commands:

```
mmdsh -N NodeList rm -rf /opt/IBM/zimon*
mmdsh -N NodeList rm -rf /usr/IBM/zimon*
mmdsh -N NodeList rm -rf /var/log/cnlog/zimon*
mmdsh -N NodeList rm -rf /var/lib/yum/repose/x86_64/7Server/*zimon*
mmdsh -N NodeList rm -rf /var/lib/yum/repose/ppc64/7Server/*zimon*
```

5. Clean up CTDB. To do so, issue the following commands:

```
mmdsh -N NodeList rm -rf /var/lib/ctdb
```

6. Clean up NooBaa.

- a. Check the NooBaa version.

```
rpm -qa | grep noobaa
```

- b. Remove RPM from all the nodes.

```
mmdsh -N NodeList rpm -e <noobaa_rpm_name>
```

7. Clean up Swift files. To do so, issue the following commands:

```
mmdsh -N NodeList rm -rf /etc/swift
mmdsh -N NodeList rm -rf /var/lib/mysql
mmdsh -N NodeList rm -rf /etc/keystone
mmdsh -N NodeList rm -rf /var/lib/keystone
mmdsh -N NodeList rm -rf /root/openrc
mmdsh -N NodeList rm -rf /var/cache/yum/x86_64/7Server/*
mmdsh -N NodeList rm -rf /var/cache/yum/ppc64/7Server/*
mmdsh -N NodeList rm -rf /var/log/maria*
mmdsh -N NodeList rm -rf /usr/share/pixmaps/comps/maria*
mmdsh -N NodeList rm -rf /var/log/keystone
mmdsh -N NodeList rm -rf /var/spool/cron/keystone
mmdsh -N NodeList rm -rf /usr/lib/python2.7/site-packages/sos/plugins/openstack_keystone*
mmdsh -N NodeList rm -rf /tmp/keystone-signing-swift
mmdsh -N NodeList rm -rf /usr/lib/python2.7/site-packages/sos/plugins/*
mmdsh -N NodeList rm -rf /var/lib/yum/repos/x86_64/7Server/icm_openstack
mmdsh -N NodeList rm -f /usr/lib/python2.7/site-packages/swiftonfile-2.5.0-py2.7.egg*
mmdsh -N NodeList rm -f /usr/bin/*objectizer*.pyc
mmdsh -N NodeList rm -f /usr/bin/generate_dbmap.pyc
mmdsh -N NodeList systemctl disable openstack-keystone.service
mmdsh -N NodeList systemctl disable openstack-swift-container-updater.service
mmdsh -N NodeList systemctl disable openstack-swift-container-update
mmdsh -N NodeList systemctl disable openstack-swift-object-updater.service
mmdsh -N NodeList systemctl disable openstack-swift-container-auditor.service
mmdsh -N NodeList systemctl disable openstack-swift-container-replicator.service
mmdsh -N NodeList systemctl disable openstack-swift-container.service
mmdsh -N NodeList systemctl disable openstack-swift-object-replicator.service
mmdsh -N NodeList systemctl disable openstack-swift-object.service
mmdsh -N NodeList systemctl disable openstack-keystone.service
mmdsh -N NodeList systemctl disable openstack-swift-account-reaper.service
mmdsh -N NodeList systemctl disable openstack-swift-account-auditor.service
mmdsh -N NodeList systemctl disable openstack-swift-account-replicator.service
mmdsh -N NodeList systemctl disable openstack-swift-account.service
mmdsh -N NodeList systemctl disable openstack-swift-proxy.service
mmdsh -N NodeList systemctl disable openstack-swift-object-auditor.service
mmdsh -N NodeList systemctl disable openstack-swift-object-expirer.service
mmdsh -N NodeList systemctl disable openstack-swift-container-reconciler.service
mmdsh -N NodeList rm -rf
/var/lib/yum/yumdb/o/0b01eb65826df92befd8c161798cb842fa3c941e-openstack-
utils-2015.1-201502031913.ibm.e17.7-noarch
```

8. Reboot the nodes. To do so, issue the following command:

```
mmdsh -N NodeList shutdown -r now
```

If some nodes do not fully come up, do the following:

- a. Power off that node.
- b. Wait for a minute, then power on the node back.

Note: Some nodes might need to be powered off and on more than once.

9. Clean up yum on all nodes. To do so, issue the following commands:

```
mmdsh -N NodeList yum clean all
mmdsh -N NodeList rm -rf /etc/yum.repos.d/gpfs.repo /etc/yum.repos.d/icm* /etc/yum.repos.d/ces.repo
mmdsh -N NodeList rm -rf /etc/yum.repos.d/epel* /etc/yum.repos.d/ido*
mmdsh -N NodeList rm -rf /var/lib/yum/repos/x86_64/7Server/*
mmdsh -N NodeList rm -rf /var/lib/yum/repos/ppc64/7Server/*
```

10. Remove GPFS.

Notes:

- This step is not required if you know that GPFS has not changed between old and new builds.
- This step is not required if you prefer to perform an upgrade of GPFS.
- This step is not required if you would like to keep the base GPFS installation intact and merely rerun the protocols deployment.
- (To permanently remove GPFS, see [Chapter 19, “Uninstalling IBM Storage Scale,”](#) on page 659.)

- a. Check which GPFS packages are on each node by issuing the following command:

RHEL and SLES

```
mmdsh -N NodeList rpm -qa|grep gpfs
```

Ubuntu

```
mmdsh -N NodeList dpkg -l|grep gpfs
```

The system displays output similar to the following:

```
gpfs.msg.en_US-6.0.1-x.noarch
gpfs.gskit-8.0.55.x.x86_64
gpfs.license.xx-6.0.1-x.x86_64
gpfs.crypto-6.0.1-x.x86_64
gpfs.docs-6.0.1-x.noarch
gpfs.gpl-6.0.1-x.noarch
gpfs.librdkafka-6.0.1-x.x86_64
gpfs.java-6.0.1-x.x86_64
gpfs.afm.cos-6.0.1-x.x86_64
gpfs.adv-6.0.1-x.x86_64
gpfs.base-6.0.1-x.x86_64
```

- b. Before removing anything, make sure that GPFS is shut down on all nodes. To do so, issue the following command:

```
mmshutdown -a
```

- c. Remove the packages by issuing the following commands *in the order shown*.

Note: When you remove the `gpfs . base` package, you lose the `mmdsh` access.

Therefore, be sure to remove `gpfs . base` last, as shown here.

RHEL and SLES

```
mmdsh -N NodeList rpm -e gpfs.base-debuginfo-6.0.1-x.x86_64
mmdsh -N NodeList rpm -e gpfs.crypto-6.0.1-x.x86_64
mmdsh -N NodeList rpm -e gpfs.msg.en_US-6.0.1-x.noarch
mmdsh -N NodeList rpm -e gpfs.gskit-8.0.55.x.x86_64
mmdsh -N NodeList rpm -e gpfs.compression-6.0.1-x.x86_64
mmdsh -N NodeList rpm -e gpfs.license.xx-6.0.1-x.x86_64
mmdsh -N NodeList rpm -e gpfs.docs-6.0.1-x.noarch
mmdsh -N NodeList rpm -e gpfs.gpl-6.0.1-x.noarch
mmdsh -N NodeList rpm -e gpfs.librdkafka-6.0.1-x.x86_64
```

```
mmdsh -N NodeList rpm -e gpfs.java-6.0.1-x.x86_64
mmdsh -N NodeList rpm -e gpfs.afm.cos-6.0.1-x.x86_64
mmdsh -N NodeList rpm -e gpfs.adv-6.0.1-x.x86_64
mmdsh -N NodeList rpm -e gpfs.base-6.0.1-x.x86_64
```

Ubuntu

```
mmdsh -N NodeList dpkg -P gpfs.base-debuginfo_6.0.1-x.x86_64
mmdsh -N NodeList dpkg -P gpfs.crypto_6.0.1-x.x86_64
mmdsh -N NodeList dpkg -P gpfs.msg.en-us_6.0.1-x.all
mmdsh -N NodeList dpkg -P gpfs.gskit_8.0.55.x.x86_64
mmdsh -N NodeList dpkg -P gpfs.compression_6.0.1-x.x86_64
mmdsh -N NodeList dpkg -P gpfs.license.xx_6.0.1-x.x86_64
mmdsh -N NodeList dpkg -P gpfs.docs_6.0.1-x.noarch
mmdsh -N NodeList dpkg -P gpfs.gpl_6.0.1-x.noarch
mmdsh -N NodeList dpkg -P gpfs.librdkafka-6.0.1-x.x86_64
mmdsh -N NodeList dpkg -P gpfs.java-6.0.1-x.x86_64
mmdsh -N NodeList dpkg -P gpfs.afm.cos-6.0.1-x.x86_64
mmdsh -N NodeList dpkg -P gpfs.adv_6.0.1-x.x86_64
mmdsh -N NodeList dpkg -P gpfs.base_6.0.1-x.x86_64
```

Note: In the preceding commands, the package name `gpfs.license.xx` needs to be changed depending on the IBM Storage Scale product edition.

11. Reinstall GPFS.

12. Proceed to [“Installation prerequisites”](#) on page 470 and [“Using the installation toolkit to perform installation tasks: Explanations and examples”](#) on page 532.

Note: If you want to remove all cluster configurations, you can also issue the `mmde1node -f` command to each node. However, if you choose to do so, you will also have to recreate cluster, BSDs, and file systems.

Uninstalling the performance monitoring tool

You can uninstall the performance monitoring tool by running the following commands on all nodes that have monitoring enabled on it.

RHEL or SLES nodes:

- To uninstall the sensor on the node, use the `rpm -evh gpfs.gss.pmsensors` command.
- To uninstall the collector on the node, use the `rpm -evh gpfs.gss.pmcollector` command.
- To uninstall the NFS proxy on the node, use the `rpm -evh gpfs.pm-ganesha` command.

Ubuntu nodes:

- To uninstall the sensor on the node, use the `dpkg -r gpfs.gss.pmsensors` command.
- To uninstall the collector on the node, use the `dpkg -r gpfs.gss.pmcollector` command.
- To uninstall the NFS proxy on the node, use the `dpkg -r gpfs.pm-ganesha` command.

For reinstalling sensors and the collector, see [“Manually installing the performance monitoring tool”](#) on page 495.

Uninstalling the IBM Storage Scale management GUI

Do the following to uninstall management GUI and remove the performance monitoring components that are installed for the GUI:

1. Issue the `systemctl stop` command as shown in the following example:

```
systemctl stop gpfsgui
```

2. If the GUI runs on an IBM Storage Scale cluster where sudo wrapper is enabled, export the name of the user that was configured as the file system administrator as an environment variable by issuing the following command:

```
export SUDO_USER=gpfsadmin
```

In this example, the name of the sudo user is *gpfsadmin*. Exporting the environment variable is necessary so that the uninstall process get to know which user name must be used to run the administrative file system commands that are necessary while uninstalling the GUI. For more information on how to configure the IBM Storage Scale GUI to use sudo wrapper, see *Configuring IBM Storage Scale GUI to use sudo wrapper* in *IBM Storage Scale: Administration Guide*.

3. Issue the following command to clean up the GUI database:

```
psql postgres postgres -c "drop schema fscc cascade"
```

4. Issue one of the following commands depending on the platform to remove the GUI package.

- Red Hat Enterprise Linux and SLES

```
rpm -e gpfs.gui-6.0.1-x.noarch
```

- Ubuntu

```
dpkg -r gpfs.gui_6.0.1-x_all.deb
```

Removing nodes from management GUI-related node class

If you are reinstalling IBM Storage Scale, in some scenarios you might need to remove some nodes from the node classes that are used by the management GUI.

About this task

For information about these node classes, see [“Node classes used for the management GUI” on page 503](#). If you want to remove the GUI designation of a node, you must remove it from the GUI_MGT_SERVERS node class. After you remove a node from this node class, it is no longer designated as a GUI node and GUI services are not started on this node after reinstalling IBM Storage Scale. To remove a node from the GUI_MGT_SERVERS node class, use the **mmchnodeclass** command only when the GUI services are not running.

Procedure

On the node that you want to remove from the GUI_MGT_SERVERS node class, run the following commands:

```
systemctl stop gpfsdgui
mmchnodeclass GUI_MGMT_SERVERS delete -N guinode
```

These commands stop the GUI services and remove the GUI node from the GUI_MGT_SERVERS node class.

Note: If you use **mmchnodeclass** to change the GUI_MGT_SERVERS node class while the GUI services are running, the management GUI adds the removed node to the GUI_MGT_SERVERS node class again.

Permanently uninstall cloud services and clean up the environment

Before setting up a fresh multi-node, you must clean up the environment.

About this task

Perform the following steps to do a cleanup of the environment and be able to create a fresh multi-node setup.

Procedure

1. Recall any data that is migrated to the cloud storage by Transparent cloud tiering. For more information, see the *Recalling files from the cloud storage tier* topic in the *IBM Storage Scale: Administration Guide*.
2. Delete the container pair set by issuing a command similar to this:

```
mmcloudgateway containerpairset delete --cloud-nodeclass cloud1 --container-pair-set-name vm206
```

For more information, see *Creating a container pair set* topic in the *IBM Storage Scale: Administration Guide*.

3. Delete the cloud services by issuing a command similar to this:

```
mmcloudgateway cloudservice delete --cloud-nodeclass cloud --cloud-service-name newServ
```

For more information, see *Creating a cloud services* topic in the *IBM Storage Scale: Administration Guide*.

4. Delete the CSAP by issuing a command similar to this:

```
mmcloudgateway cloudStorageAccessPoint delete --cloud-nodeclass cloud --cloud-storage-access-point-name csap1
```

For more information, see *Defining cloud storage access points* topic in the *IBM Storage Scale: Administration Guide*.

5. Delete the cloud storage account by issuing a command similar to this:

```
mmcloudgateway account delete --cloud-nodeclass cloud --account-name new
```

For more information, see the *Managing a cloud storage account* topic in the *IBM Storage Scale: Administration Guide*.

6. Stop the cloud services on all nodes by issuing a command similar to this:

```
mmcloudgateway service stop -N cloud1
```

For more information, see the *Stopping the Transparent cloud tiering* topic in the *IBM Storage Scale: Administration Guide*.

7. Disable cloud services nodes on the node group by issuing a command similar to this:

```
mmchnode --cloud-gateway-disable -N vm1,vm2 --cloud-gateway-nodeclass cloud1
```

For more information, see the *Designating the Transparent cloud tiering nodes* topic in the *IBM Storage Scale: Administration Guide*.

8. Delete the node class if required. For more information, see the **mmdelnodeclass** command in the *IBM Storage Scale: Command and Programming Reference Guide*
9. Uninstall the cloud services RPMs (both server and client) from all the nodes by issuing the following command:

- `rpm -e gpfs.tct.server-x.x.x.x86_64`
- `rpm -e gpfs.tct.client-x.x.x.x86_64.rpm`

Chapter 20. Upgrading

To upgrade to a newer version of IBM Storage Scale, first consider the version that you are upgrading from and then consider coexistence and compatibility issues.

IBM Storage Scale supports a limited form of backward compatibility between two adjacent releases. Limited backward compatibility allows you to temporarily operate with a mixture of IBM Storage Scale nodes running on the newer version and nodes running an earlier version:

- Within a cluster, limited backward compatibility enables you to perform an online upgrade to the new IBM Storage Scale version of the code if an online upgrade from your current version to the newer version is supported.
- In a multicluster environment, limited backward compatibility allows the individual clusters to be upgraded on their own schedules. Access to the file system data can be preserved even though some of the clusters might still be running at an earlier version.
- In multicluster environments, it is recommended to upgrade the home cluster before the cache cluster especially if file audit logging, watch folder, clustered watch, and AFM functions are being used.

Tip: Ensure that there is approximately 5 GB of free space on the nodes to download, extract, and install the packages.

Note: Cluster Export Services (CES) nodes that have the Samba service enabled, do not support backward compatibility. For more information about upgrading CES nodes, see [“Upgrade process flow”](#) on page 725.

Online upgrades allow you to install new IBM Storage Scale code on one node at a time while maintaining access to the file systems without shutting down IBM Storage Scale on other nodes. However, you must upgrade all nodes within a short time. The time dependency exists because some IBM Storage Scale features in the newer version become available on each node as soon as the node is upgraded, while other features are not available until you upgrade all participating nodes.

You can also perform offline upgrades, if you can shut down the entire cluster. An offline upgrade is similar to the online upgrade procedure. However, due to the entire cluster being offline, it is possible to jump straight to the latest code level instead of hopping to an intermediate level, as might be required during an online upgrade. For information about an offline upgrade procedure, see [“Offline upgrade with complete cluster shutdown”](#) on page 769.

Important: Online upgrade from an unsupported version of IBM Storage Scale or GPFS is not supported.

For information about the upgrade process flow used by the installation toolkit, see [“Upgrade process flow”](#) on page 725. You can use this process flow as a reference for doing manual upgrades on Linux nodes.

After all nodes are upgraded to the new code, you must finalize the upgrade by running the commands **mmchconfig release=LATEST** and **mmchfs -V full** (or **mmchfs -V compat**). Also, certain new features might require you to run the **mmmigratefs** command to enable them. For more information, see [“Completing the upgrade to a new level of IBM Storage Scale”](#) on page 753.

For the latest information on upgrade, coexistence, and compatibility, see the *Supported OS and software versions* section of [IBM Storage Scale FAQ](#) in IBM Documentation.

IBM Storage Scale upgrade consists of these topics:

- [“IBM Storage Scale supported upgrade paths”](#) on page 672
- [“Online upgrade support for protocols and performance monitoring”](#) on page 673
- [“Upgrade of IBM Storage Scale nodes”](#) on page 674
- [“Upgrading IBM Storage Scale non-protocol nodes that run on Linux”](#) on page 677
- [“Upgrading IBM Storage Scale protocol nodes”](#) on page 682
- [“Upgrading AFM and AFM DR”](#) on page 692

- [“Upgrading call home” on page 702](#)
- [“Upgrading SMB packages” on page 697](#)
- [“Upgrading NFS packages” on page 699](#)
- [Upgrading S3 packages](#)
- [“Manually upgrading pmswift” on page 706](#)
- [“Manually upgrading the performance monitoring tool” on page 704](#)
- [“Manually upgrading the IBM Storage Scale management GUI” on page 708](#)
- [“Upgrading transparent cloud tiering \(discontinued\)” on page 710](#)
- [“Upgrading to IBM Cloud Object Storage software level 3.7.2 and above” on page 722](#)
- [“Upgrading IBM Storage Scale components with the installation toolkit” on page 723](#)
- [“Changing the IBM Storage Scale product edition” on page 747](#)
- [“Completing the upgrade to a new level of IBM Storage Scale” on page 753](#)
- [“Reverting to the previous level of IBM Storage Scale” on page 758](#)
- [“Coexistence considerations” on page 761](#)
- [“Compatibility considerations” on page 761](#)
- [“Considerations for IBM Storage Protect for Space Management” on page 761](#)
- [“Applying maintenance to your IBM Storage Scale system” on page 761](#)
- [“Guidance for upgrading the operating system on IBM Storage Scale nodes” on page 762](#)
- [“Servicing IBM Storage Scale protocol nodes” on page 768](#)
- [“Offline upgrade with complete cluster shutdown” on page 769](#)

IBM Storage Scale supported upgrade paths

Use this information to understand the online upgrade paths that are supported for IBM Storage Scale.

The following table lists the supported online upgrade paths for IBM Storage Scale.

<i>Table 77. IBM Storage Scale supported online upgrade paths</i>							
To → From ↓	6.0.1.x	6.0.0.x	5.2.3.x	5.2.2.x	5.2.1.x	5.2.0.x	5.1.9.x
6.0.1.x	☑	--	--	--	--	--	--
6.0.0.x	☑	☑	--	--	--	--	--
5.2.3.x	☑	☑	☑	--	--	--	--
5.2.2.x	☑	☑	☑	☑	--	--	--
5.2.1.x	☑	☑	☑	☑	☑	--	--
5.2.0.x	☑	☑	☑	☑	☑	☑	--
5.1.9.x	⊗	⊗	☑	☑	☑	☑	☑

- --: Not applicable
- ☑: Supported
- ⊗: Not supported

Note: After an IBM Storage Scale version is past its end of service (EOS) date, no further testing or proof of upgrade path to later supported releases is done. It is strongly recommended to upgrade to a newer version of IBM Storage Scale that is supported before your currently installed version is past its EOS date. For more information, see [Product lifecycle](#) in IBM Support documentation.

Important:


- To upgrade to IBM Storage Scale 6.0.0.0 from 5.1.9.x or an earlier version, you must perform an offline upgrade because the online upgrade is not supported for said upgrade paths.
- The versions that are not shown in the upgrade paths table are past their EOS date. You can access upgrade procedures for these versions in the IBM Storage Scale documentation of earlier releases. For example, if you have IBM Storage Scale 4.2.3.x version, refer to IBM Storage Scale 5.0.x documentation. However, online upgrade from an unsupported version is not supported. You must use an offline upgrade for such upgrade paths.

Remember:

- Online upgrade support for protocol nodes depends on which services are enabled and deployed. Different protocol services have different levels of online upgrade support. For more information, see [“Online upgrade support for protocols and performance monitoring”](#) on page 673.
- IBM Storage Scale installation toolkit can be used for upgrades from version 4.1.1.x onwards.

Coexistence of different releases

Different releases of IBM Storage Scale can coexist, that is, be active in the same cluster and simultaneously access the same file system. For release coexistence, IBM Storage Scale follows the N-1 rule. According to this rule, a particular IBM Storage Scale release (N) can coexist with the prior release of IBM Storage Scale (N-1). This allows IBM Storage Scale to support an online (rolling) upgrade, that is a node by node upgrade.



For more information, see [2.8 Can different IBM Storage Scale maintenance levels coexist?](#) in the [IBM Storage Scale Frequently Asked Questions and Answers](#).

Related information

[IBM Spectrum Scale Software Version Recommendation Preventive Service Planning](#)
[IBM Spectrum Scale RAID FAQ in IBM Documentation](#)

Online upgrade support for protocols and performance monitoring

Use this information to understand the online upgrade support for protocols and performance monitoring in IBM Storage Scale.

Online upgrade support for protocol nodes depends on which services are enabled and deployed. Different protocol services have different levels of online upgrade support.

- NFS: If there are different NFS versions on protocol nodes in a cluster, then performing an online upgrade on one node at a time can be tolerated.
- SMB: All protocol nodes running the SMB service must have the same version installed at any time. In addition, to a brief outage of the SMB service, this also requires a brief outage of the NFS service in case Active Directory (AD) authentication is being used. SMB upgrade in an IBM Storage Scale cluster is done in two phases. The protocol nodes in the cluster are divided in two halves and SMB is upgraded on each half one by one.

- S3: If there are different S3 versions on protocol nodes in a cluster, then performing an online upgrade on one node at a time can be tolerated.
- Performance monitoring: Having two different versions on collector nodes will result in a loss of communication and could potentially result in missing performance metrics until collector versions match. This limitation is not applicable if the source and the target versions are later than 5.0.1.

Upgrade of IBM Storage Scale nodes

IBM Storage Scale supports upgrade to the latest release from a previous supported IBM Storage Scale release. For more information, see [“IBM Storage Scale supported upgrade paths”](#) on page 672 and [“Online upgrade support for protocols and performance monitoring”](#) on page 673.

To upgrade an IBM Storage Scale cluster, do the following steps for each node in the cluster.

For upgrading nodes that run on Linux, use the following procedures.

- [“Upgrading IBM Storage Scale non-protocol nodes that run on Linux”](#) on page 677.
- [“Upgrading IBM Storage Scale protocol nodes”](#) on page 682.

For upgrading Hadoop, use the following procedures.

CDP Private Cloud Base distribution

- To upgrade the CDP Private Cloud Base distribution, see *Upgrading CDP* in *IBM Storage Scale: Big Data and Analytics Guide*.

HDP Ambari distribution

- To upgrade the IBM Storage Scale Hadoop Connector through Ambari, see *Upgrading HDFS Transparency* in *IBM Storage Scale: Big Data and Analytics Guide*.
- To upgrade the IBM Storage Scale Hadoop Connector manually, see *Upgrading the HDFS Transparency cluster* in *IBM Storage Scale: Big Data and Analytics Guide*.

For upgrading nodes that run on AIX or Windows , use the following procedures.

- [“Upgrading IBM Storage Scale nodes that run on AIX”](#) on page 674
- [“Upgrading IBM Storage Scale nodes that run on Windows”](#) on page 676

When all the nodes in the cluster are successfully upgraded to the new level, the next step is to complete the upgrade. For more information, see [“Completing the upgrade to a new level of IBM Storage Scale”](#) on page 753.

Upgrading IBM Storage Scale nodes that run on AIX

Use these steps to upgrade IBM Storage Scale nodes that run on AIX.

Before you begin

Before you start, make sure to comply with [“Software requirements”](#) on page 310.

Procedure

1. On the node that you are upgrading, stop all user and applications activities in the IBM Storage Scale file systems.
2. Back up your file systems.
Important: Creating a backup of your file systems is crucial to help ensure that your data is protected.
3. If file systems are mounted on the node, issue the **mmumount** command to unmount them.



Warning: Do not use the **-f** option of the **mmumount** command.

4. Shut down IBM Storage Scale on the node.

- a) Issue the following command to shut down IBM Storage Scale.

```
mmshtutdown -a
```

Note: Postinstallation scripts of an IBM Storage Scale upgrade handle **mmccrmonitor** and other mm processes, so you do not need to close them.

- b) Corroborate that IBM Storage Scale is turned off.

```
mmgetstate -aL
```

In the output, the GPFS status column must report down.

5. If you do not plan to restart the nodes after the upgrade, verify that the GPFS daemon is stopped and that the kernel extensions are unloaded.

```
mmfsenv -u
```

Note:

- The **mmfsenv -u** command is normally used with assistance from support or for cases where its usage is documented such as in this scenario.
- The **mmfsenv -u** command checks only the local node from which it is run, not all the nodes in the cluster.

If the kernel extensions are unloaded, you get an output similar to the following example:

```
mmfsenv -u
GPFS: 6027-501 /usr/lpp/mmfs/bin/mmfskxunload: module /usr/lpp/mmfs/bin/mmfs unloaded
```

If the kernel extensions are unloaded, you can continue with the upgrade and no restart of the node is necessary to be able to use GPFS.

If the kernel extensions are still loaded, you get a Device busy error, similar to the following example.

```
mmfsenv -u
sysconfig(SYS_CFGKMOD-term): Device busy
```

If the kernel extensions are still loaded, you can continue with the upgrade. However, to be able to use GPFS, you must wait until the node is restarted or the error is cleared.

Troubleshooting: The Device busy error occurs when a process or an open file is keeping GPFS busy, thus preventing the kernel driver from unloading. The solution is to wait until after the installation is complete and then restart the node by using **shutdown** or **reboot** commands. For more information about these commands, see [shutdown command](#) and [reboot or fastboot command](#) in [IBM AIX documentation](#). Alternatively, if you can identify the process that is keeping the kernel extensions busy, stop or end that process and then retry **mmfsenv -u**.

6. Download the installation package for the target version of IBM Storage Scale licensed programs package from [IBM Fix Central](#).

The package name is `Storage_Scale_<edition>-<target_version>-<architecture>-AIX-update.tgz`. For example: `Storage_Scale_Advanced-5.2.3.7-ppc64-AIX-update.tgz`.

7. Prepare the files for the upgrade.

- a) Copy the installation images to a temporary location.

```
# cd /tmp/scale_<target_version>
# ls -ltr
total 90232
-rw-r--r-- 1 root system 46197337 May 30 07:42 Storage_Scale_<edition>-<target_version>-<architecture>-AIX-update.tgz
```

b) Decompress the file.

```
# gunzip Storage_Scale_<edition>-<target_version>-<architecture>-AIX-update.tgz
# ls -ltr
total 209128
-rw-r--r-- 1 root      system    107069440 May 30 07:42 Storage_Scale_<edition>-<target_version>-<architecture>-AIX-update.tar
```

c) Extract the TAR file.

```
# tar xvf Storage_Scale_<edition>-<target_version>-<architecture>-AIX-update.tar
```

8. Upgrade the node to a new version of IBM Storage Scale.

a) Create the .toc file that is necessary to run the **installp** command.

```
# pwd
/tmp/scale_<target_version>
# inutoc .
# ls -ltr
```

b) Install the packages of the target version of IBM Storage Scale.

```
# installp -acgXd /tmp/scale_<target_version> all
```

c) Verify that the target version is installed.

```
# lslpp -l | grep gpfs
```

9. One by one, update all the nodes by following steps 6 through 8 on them.

10. Start IBM Storage Scale on the node by issuing the following command:

```
# /usr/lpp/mmfs/bin/mmstartup -a
```

11. Check the status of the cluster.

```
# /usr/lpp/mmfs/bin/mmgetstate -aL
```

In the output, the GPFS status column must report active.

12. Corroborate that the target version is the one that is running on the cluster.

```
# /usr/lpp/mmfs/bin/mmdia -v
```

In the output, check the value for Current GPFS build.

13. Mount any file systems that are not mounted automatically when the IBM Storage Scale daemon starts.

What to do next

When all the nodes in the cluster are successfully upgraded to the new level, the next step is to complete the upgrade. For more information, see [“Completing the upgrade to a new level of IBM Storage Scale” on page 753](#).

Upgrading IBM Storage Scale nodes that run on Windows

Use these steps to upgrade IBM Storage Scale nodes that run on Windows.

Before you begin

Before you start, make sure to comply with [“Software requirements” on page 310](#).

Procedure

1. On the node that you are upgrading, stop all user and applications activities in the IBM Storage Scale file systems.

2. If the call home service is enabled on the cluster, issue the following command to disable it:

```
mmcallhome capability disable
```

3. Back up your file systems.

Important: Creating a backup of your file systems is crucial to help ensure that your data is protected.

4. If file systems are mounted on the node, issue the **mmumount** command to unmount them.



Warning: Do not use the **-f** option of the **mmumount** command.

```
mmumount all
```

5. Shut down IBM Storage Scale on the node by issuing the following command:

```
mmshutdown
```

Note: Postinstallation scripts of an IBM Storage Scale upgrade handle **mmccrmonitor** and other mm processes, so you do not need to close them.

6. If you do not plan to restart the nodes after the upgrade, verify that the GPFS daemon is stopped and that the kernel extensions are unloaded.
7. Upgrade the node to a new version of IBM Storage Scale.
 - a) Open the **Programs and Features** window and remove **IBM Storage Scale Standard Edition 5.0.y**.
 - b) Uninstall the version of IBM Storage Scale from which you are upgrading and restart the system.
 - c) Uninstall IBM GPFS GSKit applicable for the IBM Storage Scale version from which you are upgrading and restart the system.
 - d) Uninstall the IBM Storage Scale license.
 - e) Copy the installation images and install the IBM Storage Scale licensed program as described in [Installing IBM Storage Scale on Windows nodes](#).

Important: For Windows upgrades, you must issue all the IBM Storage Scale administration commands from the 6.0.x node.

8. Start IBM Storage Scale on the node by issuing the following command:

```
mmstartup
```

9. Mount any file systems that are not mounted automatically when the IBM Storage Scale daemon starts.

What to do next

When all the nodes in the cluster are successfully upgraded to the new level, the next step is to complete the upgrade. For more information, see [“Completing the upgrade to a new level of IBM Storage Scale” on page 753](#).

Upgrading IBM Storage Scale non-protocol nodes that run on Linux

Use these steps to upgrade non-protocol nodes that are running on Linux distributions in an IBM Storage Scale cluster.

Before you begin

Before you start, make sure to comply with [“Software requirements” on page 310](#).

About this task

Do the following steps for each non-protocol node in the cluster.

Procedure

1. On the node that you are upgrading, stop all user and applications activities in the IBM Storage Scale file systems.
2. If the call home service is enabled on the cluster, run the following command to disable it:

```
mmcallhome capability disable
```

Important: Back up your file systems to ensure that your data is protected.

3. If file systems are mounted on the node, issue the **mmumount** command to unmount them.



Warning: Do not use the **-f** option of the **mmumount** command.

```
mmumount all
```

Note: Before you shut down the current cluster manager node, run the **mmchmgr** command to move the cluster manager role to another node to avoid unexpected I/O interruption.

4. Shut down IBM Storage Scale on the node by running the following command.

```
mmshutdown
```

Note: You do not need to close the **mmccrmonitor** command and other mm processes because they are handled by IBM Storage Scale upgrade post-install scripts.

5. If you do not plan to restart the nodes after the upgrade, verify that the GPFS daemon is stopped and that the kernel extensions are unloaded.

```
mmfsenv -u
```

If the **mmfsenv -u** command reports that it cannot unload the kernel extensions because they are busy, then the installation can proceed, but the node must be rebooted after the installation. Kernel extensions are busy means that a process has a current directory in some GPFS file system directory or it has an open file descriptor. The Linux utility **lsof** can identify the process and that process can then be killed. Retry **mmfsenv -u** after killing the process and if the command succeeds, then a reboot of the node can be avoided.

Note:

- The **mmfsenv -u** command is normally used with assistance from support or for cases where its usage is specifically documented such as in this scenario.
 - The **mmfsenv -u** command only checks the local node from which it is run, not all the nodes in the cluster.
6. Extract the IBM Storage Scale software packages as described in the topic [“Extracting the IBM Storage Scale software on Linux nodes”](#) on page 473. For information about the location of extracted packages, see [“Location of extracted packages”](#) on page 480.
 7. Change to the directory that contains the IBM Storage Scale package files.

- If Red Hat Enterprise Linux or SUSE Linux Enterprise Server (SLES) is installed on the node, run the following command:

```
cd /usr/lpp/mmfs/<x.x.x>/gpfs_rpms
```

- If Ubuntu Linux is installed on the node, run the following command:

```
cd /usr/lpp/mmfs/<x.x.x>/gpfs_debs
```

Where **<x.x.x>** is the version level of IBM Storage Scale that you are upgrading to.

8. Upgrade the node by using one of the following set of commands. The command choice depends on the operating system and the product edition that is installed on the node.

Note: The following commands are specific to the current release. If you are upgrading to an earlier release, see the documentation for that release.

Tip: To determine the currently installed edition, run the following command:

```
mmlslicense -L
```

Note:

- Starting with IBM Storage Scale 5.1.0 or later, there are no separate packages for call home, since call home is integrated into the `gpfs.base` package.
- If you are upgrading to IBM Storage Scale 5.1.0 or later, you do not need to upgrade the `gpfs.kafka` package.
- Red Hat Enterprise Linux 7.x, 8.x, and 9.x

Note: For rpm packages specific to operating systems such as rhel7, rhel8, and rhel9, you need to install them from specific directories such as rhel7/, rhel8/, and rhel9/.

- If IBM Storage Scale Standard Edition or IBM Storage Scale Data Access Edition is installed on the node, run the following command. The entire command must be on one line.

```
rpm -Fvh ./gpfs.gskit*.rpm ./gpfs.gpl*.rpm ./gpfs.docs*.rpm ./gpfs.compression*.rpm  
./gpfs.base*.rpm ./gpfs.msg.en_US*.rpm  
./gpfs.java*.rpm ./rhel/gpfs.librdkafka*.rpm
```

- If IBM Storage Scale Advanced Edition or IBM Storage Scale Data Management Edition or IBM Storage Scale Developer Edition is installed on the node, run the following command. The entire command must be on one line:

```
rpm -Fvh ./gpfs.msg.en_US*.rpm ./gpfs.java*.rpm ./gpfs.gskit*.rpm  
./gpfs.gpl*.rpm ./gpfs.docs*.rpm ./gpfs.crypto*.rpm ./gpfs.compression*.rpm  
./gpfs.base*.rpm ./gpfs.adv*.rpm ./rhel/gpfs.librdkafka*.rpm
```

- SUSE Linux Enterprise Server (SLES)

- If IBM Storage Scale Standard Edition or IBM Storage Scale Data Access Edition is installed on the node, run the following command. The entire command must be on one line:

```
rpm -Fvh ./gpfs.gskit*.rpm ./gpfs.gpl*.rpm ./gpfs.docs*.rpm ./gpfs.compression*.rpm  
./gpfs.base*.rpm ./gpfs.msg.en_US*.rpm  
./gpfs.java*.rpm
```

- If IBM Storage Scale Advanced Edition or IBM Storage Scale Data Management Edition is installed on the node, run the following command. The entire command must be on one line:

```
rpm -Fvh ./gpfs.msg.en_US*.rpm ./gpfs.java*.rpm ./gpfs.gskit*.rpm  
./gpfs.gpl*.rpm ./gpfs.docs*.rpm ./gpfs.crypto*.rpm ./gpfs.compression*.rpm  
./gpfs.base*.rpm ./gpfs.adv*.rpm
```

- Ubuntu Linux

- If IBM Storage Scale Standard Edition or IBM Storage Scale Data Access Edition is installed on the node, run the following command. The entire command must be on one line:

```
apt-get install --only-upgrade ./gpfs.gskit*.deb ./gpfs.gpl*.deb  
./gpfs.docs.deb ./gpfs.compression.deb  
./gpfs.base*.deb ./gpfs.msg.en-us*.deb  
./gpfs.java*.deb ./ubuntu/gpfs.librdkafka*.deb
```

- If IBM Storage Scale Advanced Edition or IBM Storage Scale Data Management Edition is installed on the node, run the following command. The entire command must be on one line:

```
apt-get install --only-upgrade ./gpfs.gskit*.deb ./gpfs.gpl*.deb  
./gpfs.docs*.deb ./gpfs.crypto*.deb ./gpfs.compression*.deb  
./gpfs.base*.deb ./gpfs.adv*.deb
```

```
./gpfs.msg.en-us*.deb ./gpfs.java*.deb
./ubuntu/gpfs.librdkafka*.deb
```

9. Install the following packages, if needed.

a. For all editions, if you are upgrading from IBM Storage Scale 4.2.3.x or earlier, install the `gpfs.compression` package:

- If Red Hat Enterprise Linux or SUSE Linux Enterprise Server Linux is installed on the node, run the following command from the `gpfs_rpms` directory:

```
rpm -ivh gpfs.compression*.rpm
```

- If Ubuntu Linux is installed on the node, run the following command from the `gpfs_debs` directory:

```
apt-get install gpfs.compression*.deb
```

b. For all editions, for the `gpfs.ext` package:

- If you are upgrading to IBM Storage Scale 5.0.1 or earlier, ensure that you upgrade the `gpfs.ext` package.
- In versions 5.0.2 and later, you do not need to upgrade the `gpfs.ext` package. This package is removed and its functions are merged into other packages.



Attention: On Ubuntu Linux nodes, in version 5.0.2 and later, you might see an older version of the `gpfs.ext` package. You can safely remove this package by running the following command:

```
dpkg -P gpfs.ext
```

c. For IBM Storage Scale Advanced Edition and IBM Storage Scale Data Management Edition, if you plan to use the AFM-based Asynchronous Disaster Recovery (AFM DR) feature, the file encryption feature, or the transparent cloud tiering feature, you must install the following packages:

- `gpfs.adv`
- `gpfs.crypto`

10. If needed, upgrade the IBM Storage Scale native REST API packages as follows:

a. Change to the directory that contains native REST API packages:

- If Red Hat Enterprise Linux or SUSE Linux Enterprise Server Linux is installed on the node, issue the following command:

```
cd /usr/lpp/mmfs/<x.x.x.x>/scaleapi_rpms
```

- If Ubuntu Linux is installed on the node, issue the following command:

```
cd /usr/lpp/mmfs/<x.x.x.x>/scaleapi_debs
```

Where `<x.x.x.x>` is the version level of IBM Storage Scale that you are upgrading to. The version must be 5.2.3.1 or later for nodes installed with the native REST API.

b. Upgrade the native REST API package:

- If Red Hat Enterprise Linux or SUSE Linux Enterprise Server Linux is installed on the node, issue the following command:

```
rpm -Fvh ./gpfs.scaleapi*.rpm
```

- If Ubuntu Linux is installed on the node, issue the following command:

```
apt-get install --only-upgrade ./gpfs.scaleapi*.deb
```

After the native REST API feature specific rpm or deb package is installed successfully, the administration daemon (**scaleadmd**) is started with **systemd** through a post-installation scriptlet.

- c. Verify that the administration daemon is active.

```
systemctl status scaleadmd
```

A sample output is as follows:

```
scaleadmd.service - Storage Scale Admin Daemon Service
Loaded: loaded (/usr/lib/systemd/system/scaleadmd.service; enabled; vendor preset: disabled)
Active: active (running) since Thu 2025-04-17 15:20:38 CEST; 6 days ago
Process: 498376 ExecStartPre=/bin/sh -c targetLink=$(/bin/readlink -f /var/mmfs/run) ;
Main PID: 498382 (scaleadmd)
Tasks: 13
Memory: 102.2M
CGroup: /system.slice/scaleadmd.service
498382 /usr/lib/mmfs/bin/scaleadmd
[-n "targetLink" -a "${targetLink##*/}" = "mmfs" ] &6
{ /bin/m>
Apr 17 15:20:38 mos-dev-1.novalocal systemd[1]: Starting Storage Scale Admin Daemon Service...
Apr 17 15:20:38 mos-dev-1.novalocal systemd[1]: Started Storage Scale Admin Daemon Service. Apr 17 15:20:38 mos-dev-1.novalocal scaleadmd[498382]: REST server stopped with error: unable to start admin daemon rest server: node identity not configured
Apr 17 15:20:38 mos-dev-1.novalocal scaleadmd[498382]: TCP gRPC server stopped with error: unable to start admin daemon internal server: node identity not configur>
Apr 17 15:20:38 mos-dev-1.novalocal scaleadmd[498382]: Starting GRPC server on unix:///var/mmfs/run/scaleadmd.grpc.sock, boot mode 1
Apr 17 15:40:32 mos-dev-1.novalocal scaleadmd[498382]: REST server stopped with error: unable to start admin daemon rest server: node identity not configured
Apr 17 15:40:44 mos-dev-1.novalocal scaleadmd[498382]: Starting GRPC server on https://[::]:50052, boot mode 2
Apr 17 15:40:44 mos-dev-1.novalocal scaleadmd[498382]: Starting REST API server on https://0.0.0.0:46443, boot mode
```

Note: If endpoint operations fail unexpectedly after upgrading the native REST API, restart the administration daemon on each upgraded node. These failures might be caused by timing-related interactions between the administration daemon and the cluster configuration repository (CCR) during configuration access.

11. After a successful upgrade, rebuild the GPFS portability layer (GPL). For more information, see [“Building the GPFS portability layer on Linux nodes”](#) on page 483.
12. Upgrade performance monitoring packages. For more information, see [“Manually upgrading the performance monitoring tool”](#) on page 704 and [“Manually upgrading pmswift”](#) on page 706.

Note: You might see the error message `pmcollector service critical error` on GUI nodes. If this problem occurs, restart the `pmcollector` service by running the following command on all GUI nodes:

```
systemctl restart pmcollector
```

13. Upgrade management GUI packages. For more information, see [“Manually upgrading the IBM Storage Scale management GUI”](#) on page 708.
14. Start IBM Storage Scale on the node by running the following command:

```
mmstartup
```

15. Mount any file systems that are not mounted automatically when the IBM Storage Scale daemon starts.
16. If you disable the call home service in a preceding step, run the following command to enable it:

```
mmcallhome capability enable accept
```

Repeat the preceding steps on all non-protocol nodes in the cluster. After these steps are done for all non-protocol nodes in the cluster, proceed with the following steps.

17. Upgrade the license package corresponding to the installed IBM Storage Scale product edition on all nodes in the cluster.

- For Red Hat Enterprise Linux and SLES:

```
mmdsh -N NodeList rpm -Fvh /usr/lpp/mmfs/N.N.N.N/gpfs_rpms/gpfs.license-xxx*.rpm
```

- For Ubuntu:

```
mmdsh -N NodeList apt-get install --only-upgrade /usr/lpp/mmfs/N.N.N.N/gpfs_debs/  
gpfs.license-xxx*.deb
```

- *NodeList* is the list of all nodes in the cluster.
- *N.N.N.N* is the product version. For example, 5.0.2.0.
- *xxx* is the indicator of the product edition:
 - da: Data Access Edition
 - dm: Data Management Edition
 - dev: Developer Edition
 - std: Standard Edition
 - adv: Advanced Edition

18. If the product edition is Advanced, Data Management, or Developer and file audit logging or clustered watch folder is enabled in the cluster, see [“Upgrade paths and commands for file audit logging and clustered watch folder” on page 722](#).

- If you are upgrading to IBM Storage Scale 5.1.2 or later and you have previously run the **mmmsgqueue config --remove-msgqueue** command, no further action is required.
- If you are upgrading to IBM Storage Scale 5.1.2 or later and you have previously not run the **mmmsgqueue config --remove-msgqueue** command, do one of the following steps.
 - If you have never enabled the message queue, manually remove the `gpfs.kafka` package.
 - In all other cases, run the **mmmsgqueue config --remove-msgqueue** command. This command removes the `gpfs.kafka` package at the end of the cleanup.

What to do next

When all the nodes in the cluster are successfully upgraded to the new level, the next step is to complete the upgrade. For more information, see [“Completing the upgrade to a new level of IBM Storage Scale” on page 753](#).

Upgrading IBM Storage Scale protocol nodes

Use these steps to upgrade protocol nodes in an IBM Storage Scale cluster.

Before you begin

Before you start, make sure to comply with [“Software requirements” on page 310](#).

About this task



Attention:

- S3 users must run I/O workloads with an increased value for the variables `AWS_MAX_ATTEMPTS` and `AWS_MAX_RETRY` to sustain the high availability. Because CES IP movement might take time based on the active I/O from multiple protocols, if enabled in the cluster. Refer AWS CLI documentation for how to set these.
- Before the cluster is ready for upgrade, back up the S3 configuration (if the S3 protocol is enabled in the cluster). For more information, see *Manual backup of the S3 configuration data in the Backing up the S3 configuration data* section in *IBM Storage Scale: Administration Guide*.

Perform the following steps for each protocol node in the cluster:

Procedure

1. On the node that you are upgrading, stop all user and applications activities in the IBM Storage Scale file systems.
2. Stop any protocol services that are running on the node:
 - a. Issue the following command to suspend Cluster Export Services (CES) and stop the protocol services on the node:

```
mmces node suspend --stop
```

If you are upgrading from IBM Storage Scale version 5.0.2.0 or earlier, issue the following commands to suspend the protocol node and stop the protocol services:

```
mmces node suspend  
mmces service stop Protocol
```

- b. If you are upgrading from IBM Storage Scale version 5.0.1.2 or earlier to version 5.0.2 or later, issue the following command to ensure that the **mmunmount** of the CES file system completes successfully:

```
mmcesmonitor stop
```

3. If the call home service is enabled on the cluster, issue the following command to disable it:

```
mmcallhome capability disable
```

Important: Back up your file systems to ensure that your data is protected.

4. If file systems are mounted on the node, issue the **mmumount** command to unmount them.



Warning: Do not use the **-f** option of the **mmumount** command.

```
mmumount all
```

5. Shut down IBM Storage Scale on the node by issuing the following command:

```
mmshutdown
```

Note: You do not need to close **mmccrmonitor** and other mm processes because they are handled by IBM Storage Scale upgrade post-install scripts.

6. If you do not plan to restart the nodes after the upgrade, verify that the GPFS daemon is stopped and that the kernel extensions are unloaded.

```
mmfsenv -u
```

If the **mmfsenv -u** command reports that it cannot unload the kernel extensions because they are busy, then the installation can proceed, but the node must be rebooted after the installation. Kernel extensions are busy means that a process has a current directory in some GPFS file system directory or it has an open file descriptor. The Linux utility **lsof** can identify the process and that process can then be killed. Retry **mmfsenv -u** after killing the process and if the command succeeds, then a reboot of the node can be avoided.

Note:

- The **mmfsenv -u** command is normally used with assistance from support or for cases where its usage is specifically documented such as in this scenario.
- The **mmfsenv -u** command only checks the local node from which it is run, not all the nodes in the cluster.

7. Extract the IBM Storage Scale software packages as described in the topic [“Extracting the IBM Storage Scale software on Linux nodes”](#) on page 473. For more information about the location of extracted packages, see [“Location of extracted packages”](#) on page 480.

8. Change to the directory that contains the IBM Storage Scale package files.

- If Red Hat Enterprise Linux or SUSE Linux Enterprise Server (SLES) is installed on the node, issue the following command:

```
cd /usr/lpp/mmfs/<x.x.x.x>/gpfs_rpms
```

- If Ubuntu Linux is installed on the node, issue the following command:

```
cd /usr/lpp/mmfs/<x.x.x.x>/gpfs_debs
```

9. Upgrade the protocol node by using one of the following set of steps for depending on the operating system and the product edition that is installed on the node.

Note: The following commands are specific to the current release. If you are upgrading to an earlier release, see the documentation for that release.

Tip: To determine the currently installed edition, issue the following command:

```
mmfslicense -L
```

Note:

- Starting with IBM Storage Scale 5.1.0 or later, there are no separate packages for call home, since call home is integrated into the `gpfs.base` package.
- If you are upgrading to IBM Storage Scale 5.1.0 or later, you do not need to upgrade the `gpfs.kafka` package.
- Red Hat Enterprise Linux 7.x, 8.x, and 9.x

Note: For rpm packages specific to operating systems such as rhel7, rhel8, and rhel9, you need to install them from specific directories such as rhel7/, rhel8/, and rhel9/.

- If IBM Storage Scale Standard Edition or IBM Storage Scale Data Access Edition is installed on the node, run the following command. The entire command must be on one line.

```
rpm -Fvh ./gpfs.gskit*.rpm ./gpfs.gpl*.rpm ./gpfs.docs*.rpm ./gpfs.compression*.rpm  
./gpfs.base*.rpm ./gpfs.msg.en_US*.rpm  
./gpfs.java*.rpm ./rhel/gpfs.librdkafka*.rpm
```

- If IBM Storage Scale Advanced Edition or IBM Storage Scale Data Management Edition or IBM Storage Scale Developer Edition is installed on the node, run the following command. The entire command must be on one line:

```
rpm -Fvh ./gpfs.msg.en_US*.rpm ./gpfs.java*.rpm ./gpfs.gskit*.rpm  
./gpfs.gpl*.rpm ./gpfs.docs*.rpm ./gpfs.crypto*.rpm ./gpfs.compression*.rpm  
./gpfs.base*.rpm ./gpfs.adv*.rpm ./rhel/gpfs.librdkafka*.rpm
```

- SUSE Linux Enterprise Server (SLES)

- If IBM Storage Scale Standard Edition or IBM Storage Scale Data Access Edition is installed on the node, run the following command. The entire command must be on one line:

```
rpm -Fvh ./gpfs.gskit*.rpm ./gpfs.gpl*.rpm ./gpfs.docs*.rpm ./gpfs.compression*.rpm  
./gpfs.base*.rpm ./gpfs.msg.en_US*.rpm  
./gpfs.java*.rpm
```

- If IBM Storage Scale Advanced Edition or IBM Storage Scale Data Management Edition is installed on the node, run the following command. The entire command must be on one line:

```
rpm -Fvh ./gpfs.msg.en_US*.rpm ./gpfs.java*.rpm ./gpfs.gskit*.rpm
```



```
./gpfs.gpl*.rpm ./gpfs.docs*.rpm ./gpfs.crypto*.rpm ./gpfs.compression*.rpm  
./gpfs.base*.rpm ./gpfs.adv*.rpm
```

- Ubuntu Linux

- If IBM Storage Scale Standard Edition or IBM Storage Scale Data Access Edition is installed on the node, run the following command. The entire command must be on one line:

```
apt-get install --only-upgrade ./gpfs.gskit*.deb ./gpfs.gpl*.deb  
./gpfs.docs*.deb ./gpfs.compression*.deb  
./gpfs.base*.deb ./gpfs.msg.en-us*.deb  
./gpfs.java*.deb ./ubuntu/gpfs.librdkafka*.deb
```

- If IBM Storage Scale Advanced Edition or IBM Storage Scale Data Management Edition is installed on the node, run the following command. The entire command must be on one line:

```
apt-get install --only-upgrade ./gpfs.gskit*.deb ./gpfs.gpl*.deb  
./gpfs.docs*.deb ./gpfs.crypto*.deb ./gpfs.compression*.deb  
./gpfs.base*.deb ./gpfs.adv*.deb  
./gpfs.msg.en-us*.deb ./gpfs.java*.deb  
./ubuntu/gpfs.librdkafka*.deb
```

10. If needed, install or upgrade the following packages.

- a. For all editions, if you are upgrading from IBM Storage Scale 4.2.3.x or earlier, install the `gpfs.compression` package:

- If Red Hat Enterprise Linux or SUSE Linux Enterprise Server Linux is installed on the node, issue the following command from the `gpfs_rpms` directory:

```
rpm -ivh gpfs.compression*.rpm
```

- If Ubuntu Linux is installed on the node, issue the following command from the `gpfs_debs` directory:

```
apt-get install gpfs.compression*.deb
```

- b. For all editions, for the `gpfs.ext` package:

- If you are upgrading to IBM Storage Scale 5.0.1 or earlier, ensure that you upgrade the `gpfs.ext` package.
- In versions 5.0.2 and later, you do not need to upgrade the `gpfs.ext` package. The package is not present because its functionality is merged into other packages.



Attention: On Ubuntu Linux nodes, in version 5.0.2 and later, you might see an older version of the `gpfs.ext` package. You can safely remove this package by issuing the following command:

```
dpkg -P gpfs.ext
```

- c. For IBM Storage Scale Advanced Edition and IBM Storage Scale Data Management Edition, if you plan to use the AFM-based Asynchronous Disaster Recovery feature, the file encryption feature, or the transparent cloud tiering feature, you must install the following packages:

- `gpfs.adv`
- `gpfs.crypto`

11. If needed, upgrade the IBM Storage Scale native REST API packages as follows:

- a. Change to the directory that contains native REST API packages:

- If Red Hat Enterprise Linux or SUSE Linux Enterprise Server Linux is installed on the node, issue the following command:

```
cd /usr/lpp/mmfs/<x.x.x.x>/scaleapi_rpms
```

- If Ubuntu Linux is installed on the node, issue the following command:

```
cd /usr/lpp/mmfs/<x.x.x.x>/scaleapi_debs
```

Where `<x.x.x.x>` is the version level of IBM Storage Scale that you are upgrading to. The version must be 5.2.3.1 or later for nodes installed with the native REST API.

- Upgrade the native REST API package:

- If Red Hat Enterprise Linux or SUSE Linux Enterprise Server Linux is installed on the node, issue the following command:

```
rpm -Fvh ./gpfs.scaleapi*.rpm
```

- If Ubuntu Linux is installed on the node, issue the following command:

```
apt-get install --only-upgrade ./gpfs.scaleapi*.deb
```

After the native REST API feature specific rpm or deb package is installed successfully, the administration daemon (**scaleadmd**) is started with **systemd** through a post-installation scriptlet.

- Verify that the administration daemon is active.

```
systemctl status scaleadmd
```

A sample output is as follows:

```
scaleadmd.service - Storage Scale Admin Daemon Service
Loaded: loaded (/usr/lib/systemd/system/scaleadmd.service; enabled; vendor preset:
disabled)
Active: active (running) since Thu 2025-04-17 15:20:38 CEST; 6 days ago
Process: 498376 ExecStartPre=/bin/sh -c targetLink=$(/bin/readlink -f /var/mmfs/run) ;
Main PID: 498382 (scaleadmd)
Tasks: 13
Memory: 102.2M
CGroup: /system.slice/scaleadmd.service
498382 /usr/lpp/mmfs/bin/scaleadmd
[-n "targetLink" -a "${targetLink#*/}" = "mmfs" ] &6
{ /bin/m>
Apr 17 15:20:38 mos-dev-1.novalocal systemd[1]: Starting Storage Scale Admin Daemon
Service...
Apr 17 15:20:38 mos-dev-1.novalocal systemd[1]: Started Storage Scale Admin Daemon
Service. Apr 17 15:20:38 mos-dev-1.novalocal scaleadmd[498382]: REST server stopped with
error: unable to start admin daemon rest server: node identity not configured
Apr 17 15:20:38 mos-dev-1.novalocal scaleadmd[498382]: TCP gRPC server stopped with
error: unable to start admin daemon internal server: node identity not configur>
Apr 17
15:20:38 mos-dev-1.novalocal scaleadmd[498382]: Starting GRPC server on unix:///var/
mmfs/run/scaleadmd.grpc.sock, boot mode 1
Apr 17
15:40:32 mos-dev-1.novalocal scaleadmd[498382]: REST server stopped with error: unable
to start admin daemon rest server: node identity not configured
Apr 17 15:40:44 mos-dev-1.novalocal scaleadmd[498382]: Starting GRPC server on https://
[:]:50052, boot mode 2
Apr 17 15:40:44 mos-dev-1.novalocal scaleadmd[498382]: Starting REST API server on
https://0.0.0.0:46443, boot mode
```

Note: If endpoint operations fail unexpectedly after upgrading the native REST API, restart the administration daemon on each upgraded node. These failures might be caused by timing-related interactions between the administration daemon and the cluster configuration repository (CCR) during configuration access.

- After a successful upgrade, rebuild the GPFS portability layer (GPL). For more information, see [“Building the GPFS portability layer on Linux nodes”](#) on page 483.
- Upgrade management GUI packages. For more information, see [“Manually upgrading the IBM Storage Scale management GUI”](#) on page 708.
- Start IBM Storage Scale on the node by issuing the following command:

```
mmstartup
```

15. Mount any file systems that are not mounted automatically when the IBM Storage Scale daemon starts.
16. Upgrade NFS packages on the protocol node. For more information, see [“Upgrading NFS packages” on page 699](#).
17. Upgrade S3 packages on the protocol node. For more information, see [“Upgrading S3 packages” on page 700](#).
18. If you disabled the call home service in a preceding step, issue the following command to enable it:

```
mmcallhome capability enable accept
```

Repeat the preceding steps on all protocol nodes in the cluster. After these steps are done for all protocol nodes in the cluster, proceed with the following steps.

19. If Swift Object is enabled on cluster then do object version sync and start Swift object on all protocol nodes. For more information, see [“Upgrading object packages” on page 695](#).
20. Upgrade SMB packages on all protocol nodes. For more information, see [“Upgrading SMB packages” on page 697](#).
21. If you are upgrading from IBM Storage Scale 5.1.0.x or earlier to IBM Storage Scale 5.1.1.x or later, remove the `gpfs.protocols-support` package, if it is present.

- For Red Hat Enterprise Linux and SLES:

```
mmdsh -N ProtocolNodeList rpm -ev --nodeps gpfs.protocols-support
```

- For Ubuntu (The entire command must be on one line):

```
mmdsh -N ProtocolNodeList apt-get remove gpfs.protocols-support
```

ProtocolNodeList is the list of all protocol nodes in the cluster.

22. Do the protocol authentication configuration depending on your setup. For more information, see [“Protocol authentication configuration changes during upgrade” on page 744](#).
23. Upgrade the license package corresponding to the installed IBM Storage Scale product edition on all nodes in the cluster.

- For Red Hat Enterprise Linux and SLES:

```
mmdsh -N NodeList rpm -Fvh /usr/lpp/mmfs/N.N.N.N/gpfs_rpms/gpfs.license-xxx*.rpm
```

- For Ubuntu:

```
mmdsh -N NodeList apt-get install --only-upgrade /usr/lpp/mmfs/N.N.N.N/gpfs_debs/  
gpfs.license-xxx*.deb
```

- *NodeList* is the list of all nodes in the cluster.
- *N.N.N.N* is the product version. For example, 5.0.2.0.
- *xxx* is the indicator of the product edition:
 - da: Data Access Edition
 - dm: Data Management Edition
 - dev: Developer Edition
 - std: Standard Edition
 - adv: Advanced Edition

24. If the product edition is Advanced, Data Management, or Developer and file audit logging or clustered watch folder is enabled in the cluster, see [“Upgrade paths and commands for file audit logging and clustered watch folder” on page 722](#).
 - If you are upgrading to IBM Storage Scale 5.1.2 or later and you have previously run the **mmmsgqueue config --remove-msgqueue** command, no further action is required.

- If you are upgrading to IBM Storage Scale 5.1.2 or later and you have previously not run the **mmmsgqueue config --remove-msgqueue** command, do one of the following steps.
 - If you have never enabled the message queue, manually remove the `gpfs.kafka` package.
 - In all other cases, run the **mmmsgqueue config --remove-msgqueue** command. This command removes the `gpfs.kafka` package at the end of the cleanup.

What to do next

When all the nodes in the cluster are successfully upgraded to the new level, the next step is to complete the upgrade. For more information, see [“Completing the upgrade to a new level of IBM Storage Scale” on page 753](#).

Upgrading IBM Storage Scale on cloud

Cloudkit offers an easy way to upgrade the IBM Storage Scale version of a cluster created on cloud.

For support matrix and prerequisites information, refer to the following tables.

Note: This upgrade process does not include any upgrade at operating system level. This upgrading procedure applies only at the software level of IBM Storage Scale.

Table 78. Upgrade matrix	
Source version	Target version
6.0.0.x	6.0.1
5.2.3.x	
5.2.2.x	
5.2.1.x	
5.2.0.x	

Table 79. Supported cloud platforms	
Feature	Support
Amazon Web Services (AWS)	Yes
Google Cloud Platform (GCP)	Yes
Azure Cloud Services	Yes

Cluster upgrade

A **cloudkit** cluster upgrade performs the online upgrade of an IBM Storage Scale cluster on cloud and makes it possible to install a new version of IBM Storage Scale on one node at a time, while access to the file systems is maintained.

Prerequisites

Before you start the upgrade process, ensure that you meet the following prerequisites:

- An IBM Storage Scale cluster that is already deployed, with a version that is previous to the target upgrade version.
- A repository that has the target upgrade version is available. This repository is later used to perform the cluster level upgrade.

You can upgrade an existing repository to the target version by using the next command.

```
./cloudkit create repository
```

Or use the **cloudkit upgrade repository** command to upgrade the existing repository to the target version.

- Consult the [Upgrade matrix](#) to make sure that the upgrade path is supported.

Upgrading a cluster from 6.0.0.x to 6.0.1.0

To upgrade an IBM Storage Scale cluster from 6.0.0.x to 6.0.1.0, follow the next steps.

Important: To safely migrate the database, contact IBM Support before you attempt to use cloudkit to upgrade a cluster from 6.0.0.x to the 6.0.1.0 version.

1. Update the database to clean up the **template_file** data source.
2. Run the **cloudkit init** command.
3. Run the **cloudkit upgrade cluster** command.

Upgrading a cluster from 5.2.3.x to 6.0.1.0

To upgrade an IBM Storage Scale cluster from 5.2.3.x to 6.0.1.0, follow the next steps.

1. Run the **cloudkit init** command.
2. Run the **cloudkit upgrade cluster** command.

Upgrading a cluster from 5.2.2.x to 6.0.1.0

To upgrade an IBM Storage Scale cluster from 5.2.2.x to 6.0.1.0, follow the next steps.

Important: To safely migrate the database, contact IBM Support before you attempt to use cloudkit to upgrade a cluster from 5.2.2.x to the 6.0.1.0 version.

1. Migrate database from 5.2.2.x to 6.0.1.0.
2. Create a DNS by using the **cloudkit create dns** command.
3. Run the **cloudkit upgrade cluster** command.
4. After a successful upgrade, update the node names with the new host names by using the fully qualified domain name (FQDN).
5. Log in to the cluster and run the **hostnamectl set-hostname** command on each node of the cluster.

```
hostnamectl set-hostname --static <hostname>
echo 'preserve_hostname: True' > /etc/cloud/cloud.cfg.d/<hostname>.cfg
```

6. Create a temporary specification file named "specfile" and use the following format to add the next content to it.

Format:

```
<current node host name> --daemon-interface=<new dns hostname> --admin-interface==<new dns hostname>
```

Example:

```
ip-10-0-3-150.ec2.internal --daemon-interface=ibm-storage-scale.strgscale.com --admin-
interface=ibm-storage-scale.strgscale.com
ip-10-0-3-210.ec2.internal --daemon-interface=ibm-storage-scale.strgscale.com --admin-
interface=ibm-storage-scale.strgscale.com
ip-10-0-3-222.ec2.internal --daemon-interface=ibm-storage-scale.strgscale.com --admin-
interface=ibm-storage-scale.strgscale.com
ip-10-0-3-230.ec2.internal --daemon-interface=ibm-storage-scale.strgscale.com --admin-
interface=ibm-storage-scale.strgscale.com
```

Remember: In this temporary specification file, make the next value replacements:

- Replace the values like `ip-10-0-3-150.ec2.internal` with the corresponding host name for each cluster node.
- Replace `--daemon-interface` and `--admin-interface` values with the new DNS host name.

Upgrading a cluster from 5.2.1.x to 6.0.1.0

To upgrade a cluster from version 5.2.1.x to 6.0.1.0, you must first upgrade the cluster to 5.2.2.0. Then, follow [“Upgrading a cluster from 5.2.2.x to 6.0.1.0” on page 689](#) to upgrade the cluster to 6.0.1.0.

To upgrade a cluster from 5.2.1.x to 6.0.1.0, you must follow these steps before you start the upgrade process.

Important: To safely migrate the database, contact IBM Support before you attempt to use `cloudkit` to upgrade a cluster from 5.2.1.x to the 6.0.1.0 version.

1. Migrate database from 5.2.1.x to 5.2.2.0.
2. Run the **`cloudkit init`** command.
3. Create a new repository in the target version (5.2.2.0) by issuing the **`cloudkit create repository`** command.
4. After you complete the previous steps, upgrade the IBM Storage Scale cluster by using the **`cloudkit upgrade cluster`** command.
5. Follow [“Upgrading a cluster from 5.2.2.x to 6.0.1.0” on page 689](#) to upgrade the cluster to 6.0.1.0.

Verifying the cluster version after an upgrade

1. After a successful cluster upgrade, log in to a cluster node and run the **`mmdiag --version`** command to verify that the cluster is upgraded to the correct target version:
2. Then, run the **`./cloudkit describe cluster`** command to check that the IBM Storage Scale version is updated with the target version.

Upgrading IBM Storage Scale

Follow these steps to upgrade your version of IBM Storage Scale.

1. Download the latest IBM Storage Scale Data Management Bundle from [IBM Fix Central](#) and upload it to the IBM Storage Scale EC2 controller node.

Note: All these steps must be run on the controller node.

2. Run the following script to extract IBM Storage Scale Data Management Bundle:

```
< Path_of_bundle> --silent
```

3. Go to the IBM Storage Scale installation toolkit directory.

Note: The IBM Storage Scale installation toolkit is used to upgrade the cluster to the wanted IBM Storage Scale version. For more information, see *Upgrading IBM Spectrum Scale components with the installation toolkit* section in the *IBM Storage Scale: Concepts, Planning, and Installation Guide*.

4. Specify the newly created modified `id_rsa` key during the following installation toolkit setup command:

```
cd /usr/lpp/mmfs/{{ scale_version }}/ansible_toolkit; ./spectrumscale setup -s
<controller_node_ip>
```

5. Run the **`config populate`** command:

```
./spectrumscale config populate -N <node>
```

Note: The node is the IP address or hostname of any storage node.

6. Run the **setup** command again:

```
./spectrumscale setup -s <controller_node_ip>
```

Note: The **setup** command is run as the **config populate** command overrides the `id_rsa` key configuration if the `-i` option is not specified.

7. Run the following command to export the flag for SNC:

```
export SS_SNC_ROLLING_UPGRADE=true
```

8. Run the following upgrade commands:

```
./spectrumscale upgrade precheck  
./spectrumscale upgrade run
```

9. Run the following command to create a new directory on the controller node that can be used to store all IBM Storage Scale packages for further use after IBM Storage Scale is upgraded:

```
mkdir -p /opt/IBM/<scale version>/gpfs_cloud_packages_<scale version>
```

10. Go to the `gpfs_rpms` package directory of the upgraded IBM Storage Scale, and copy the following packages to the directory created in step 9.

```
cd /usr/lpp/mmfs/{{ scale_version }}/gpfs_rpms  
cp gpfs.adv* /opt/IBM/<scale version>/gpfs_cloud_packages_<scale version>/  
cp gpfs.base* /opt/IBM/<scale version>/gpfs_cloud_packages_<scale version>/  
cp gpfs.crypto* /opt/IBM/<scale version>/gpfs_cloud_packages_<scale version>/  
cp gpfs.docs* /opt/IBM/<scale version>/gpfs_cloud_packages_<scale version>/  
cp gpfs.gpl* /opt/IBM/<scale version>/gpfs_cloud_packages_<scale version>/  
cp gpfs.gskit* /opt/IBM/<scale version>/gpfs_cloud_packages_<scale version>/  
cp gpfs.gui* /opt/IBM/<scale version>/gpfs_cloud_packages_<scale version>/  
cp gpfs.crypto* /opt/IBM/<scale version>/gpfs_cloud_packages_<scale version>/  
cp gpfs.java* /opt/IBM/<scale version>/gpfs_cloud_packages_<scale version>/  
cp gpfs.license.dm* /opt/IBM/<scale version>/gpfs_cloud_packages_<scale version>/  
cp gpfs.msg.en_US* /opt/IBM/<scale version>/gpfs_cloud_packages_<scale version>/
```

11. Go to the `zimon_rpms` package directory of the upgraded IBM Storage Scale, and copy the following packages to the directory created in step 9:

```
cd /usr/lpp/mmfs/{{ scale_version }}/zimon_rpms/rhel7/  
cp gpfs.pmcollector* /opt/IBM/<scale version>/gpfs_cloud_packages_<scale version>/  
cp gpfs.pmsensors* /opt/IBM/<scale version>/gpfs_cloud_packages_<scale version>/  
  
cd /usr/lpp/mmfs/{{ scale_version }}/zimon_rpms/rhel8/  
cp gpfs.pmcollector* /opt/IBM/<scale version>/gpfs_cloud_packages_<scale version>/  
cp gpfs.pmsensors* /opt/IBM/<scale version>/gpfs_cloud_packages_<scale version>/
```

Note: After you upgrade the IBM Storage Scale cluster, any **mmcloudworkflows** commands must be run only after the following environment variables are set:

```
export SPECTRUM_SCALE_PACKAGE_LOC= <Location of the IBM Spectrum Scale packages created in  
step 9>  
For example: export SPECTRUM_SCALE_PACKAGE_LOC=/opt/IBM/<scale version>/  
gpfs_cloud_packages_<scale version>/
```

12. Go to the `/usr/lpp/mmfs/<scale version>/Public_Keys` package directory of the upgraded IBM Storage Scale, and copy the following packages to the directory created in step 9:

```
cd /usr/lpp/mmfs/<scale version>/Public_Keys  
cp SpectrumScalepublic_key.pgp /opt/IBM/<scale version>/gpfs_cloud_packages_<scale version>/
```

For the instructions to use the installation toolkit to do an upgrade of almost all IBM Storage Scale components from the current version of IBM Storage Scale to the wanted version, see [“Upgrading IBM Storage Scale components with the installation toolkit” on page 723](#).

Upgrading GPUDirect Storage

The recommended version for use of GDS with IBM Storage Scale is 5.1.6 as it supports accelerated writes.

About this task

There are two types of upgrade paths:

- Upgrading clusters that are configured with IBM Storage Scale 5.1.1 technical preview
- Upgrading clusters with IBM Storage Scale versions that do support GDS, that is, 5.1.2 or later

Upgrading from 5.1.1 technical preview

The GDS technical preview for IBM Storage Scale 5.1.1 and all the following IBM Storage Scale versions starting with 5.1.2 are not concurrently upgradeable. You need to perform the following steps when you need to use GDS on IBM Storage Scale 5.1.2:

1. Stop all GDS I/O on IBM Storage Scale 5.1.1.x.
2. Upgrade IBM Storage Scale to 5.1.2 or later. For more information, see [Chapter 20, “Upgrading,” on page 671](#).
3. Update MOFED to supported level. For more information, see [Mellanox OFED installation](#).
4. Update CUDA to the supported version. For more information, see [Installing GDS](#).
5. Update the NVIDIA GDS driver. For more information, see [GDS package installation](#).

Upgrading clusters with IBM Storage Scale versions that do support GDS, that is, 5.1.2 or later

The various versions of GDS in IBM Storage Scale are compatible backwards and thus the general IBM Storage Scale upgrade procedure can be followed. For more information, see [Chapter 20, “Upgrading,” on page 671](#).

Related concepts

[“GPUDirect Storage support for IBM Storage Scale” on page 28](#)

IBM Storage Scale's support for NVIDIA's GPUDirect Storage (GDS) enables a direct path between GPU memory and storage. This solution addresses the need for higher throughput and lower latencies. File system storage is directly connected to the GPU buffers to reduce latency and load on the CPU. For IBM Storage Scale, this means that data can be read or written directly from / to an NSD server's pagepool and it is sent to the GPU buffer of the IBM Storage Scale clients by using RDMA. IBM Storage Scale with GDS requires an InfiniBand or RoCE fabric. In IBM Storage Scale, the **mmdiag** command is enhanced to print diagnostic information for GPUDirect Storage.

[“Planning for GPUDirect Storage” on page 343](#)

IBM Storage Scale support for GPUDirect Storage (GDS) enables a direct path between GPU memory and storage. You need to ensure that certain conditions are met before you start installing the feature.

Related tasks

[“Installing GPUDirect Storage for IBM Storage Scale” on page 639](#)

IBM Storage Scale support for GPUDirect Storage (GDS) enables a direct path between GPU memory and storage.

Upgrading AFM and AFM DR

Consider the following while you upgrade Active File Management (AFM) or Active File Management - DR (AFM DR).

Before you upgrade to a newer version of IBM Storage Scale, consider the version from which you are upgrading. IBM Storage Scale supports a limited form of compatibility with an earlier version between two adjacent releases and hence, coexistence and compatibility measures are required. For more information, see [“IBM Storage Scale supported upgrade paths” on page 672](#). The limited compatibility with an earlier version allows temporarily operating with some IBM Storage Scale nodes that are running on the newer version, and some nodes that are running an earlier version. Within a cluster, because of this you can

perform an online upgrade to the new IBM Storage Scale version, if upgrade from your current version to the newer version is supported.

In AFM and multi-cluster environment, individual clusters can be upgraded at different schedules. The access to the file system data can be preserved even though some of the clusters might still be running on an earlier version. Home or the cache cluster must be upgraded independent of the other.

During an offline upgrade, the IBM Storage Scale service is interrupted. For an offline upgrade, you must shut down the cluster and suspend the application workload of the cluster. During an online upgrade, IBM Storage Scale service is not interrupted. In an online upgrade, the system is upgraded node-by-node or failure group-by-failure group. During the upgrade, IBM Storage Scale runs on a subset of nodes. You can also perform offline upgrades, if you can shut down the entire cluster. An offline upgrade is similar to the online upgrade procedure. As the entire cluster is offline, it is possible to upgrade to the latest code level instead of upgrading to an intermediate level, as might be needed during an online upgrade.

Before you start an AFM upgrade, the home cluster or primary cluster for AFM-DR must be upgraded first before the cache or secondary cluster.

Before you consider an online upgrade of the home or cache, ensure that:

- The cluster is healthy and operational.
- IBM Storage Scale is running on all nodes that are defined in the cluster.
- All protocols that are defined on the protocol node are running.
- Ensure that the storage at the cache cluster is adequate by using **mmrepquota** command during the home upgrade. Specifically, check the storage when the cache eviction feature is enabled for the storage management.

Cache cluster - In multiple gateway environment, gateway nodes can be upgraded one by one. In these cases, filesets, which are associated with the gateway node to upgrade, are transferred to another gateway node, and any write-class operation triggers the recovery feature. This recovery feature builds the queue on the associated gateway node to continue processing the operations to home. Thus, cache to home are not disconnected. However, some performance degradation can be seen due to another gateway node that is working for the connect for those filesets that are hosted on the upgrade node previously. In heavy load systems, transferring the filesets to another gateway node might have a performance impact. It is advised to choose a time for such upgrades where the load on the system or the number of data transfers is minimal.

Parallel data transfers enabled with GW mappings - in multiple-gateways environment, where parallel data transfer is enabled with multiple gateways, upgrading any of these mapped nodes results in a normal data transfer path.

Note: Fileset mapping might remain intact after an upgrade, depending on the `afmHashing` version that is being used.

When gateway nodes manage many filesets, it is recommended to perform the rolling upgrade by using start stop replication. For more information about the start and stop replication, see [“Restarting replication to upgrade AFM and AFM DR”](#) on page 694.

If you want the replication to continue as-is during the upgrade, you can tune the number of parallel recoveries that are run on a gateway node by using the **afmMaxParallelRecoveries** parameter. The number of filesets that are specified as the **afmMaxParallelRecoveries** value are accessed for recovery. After recoveries are complete on these specified filesets, other filesets are accessed for recovery. By default, the **afmMaxParallelRecoveries** parameter is set to 0, and the recovery process is run on all filesets. You can restrict the number of recoveries by specifying the **afmMaxParallelRecoveries** value. The restriction on recoveries conserves hardware resources.

Home cluster - Cluster Export Services (CES) provides highly available file and object services to an IBM Storage Scale cluster by using Network File System (NFS), or Server Message Block (SMB) protocols. With the CES environment, the exports at home can be seen from cache by using the CES IP addresses. These IP addresses can align to protocol nodes when the CES node that already holds the CES IP address is shut down for an upgrade. The IP addresses alignment is according to the CES IP address distribution policies.

Cache might see a short disruption at the time of CES failover at home but cache filesets continue to operate.

In a non-CES or NFS Server at Home environments, home to cache disconnects until the upgrade of NFS home server is complete. In a disconnected mode, the cache builds up the queue for application operations. After Home is available, these operations are pushed to Home.

Note: Shut down the NFS server before you issue the **mmshutdown** command on home nodes. When the IBM Storage Scale system upgrade is complete and the node mounts the file system, you can start the NFS server.

Note: The **mmclone** command is not supported on cache for AFM and primary for AFM DR. Clones that are created at home or secondary are replicated as different files. While you upgrade to IBM Storage Scale 4.2.2 or later, the cache cluster must be upgraded before you consider upgrading the home cluster.

The cache cluster and the home cluster can be upgraded by using the same method. If any of these clusters has protocol nodes, upgrade these nodes. For more information, see [“Online upgrade support for protocols and performance monitoring”](#) on page 673.

After you completed upgrade, see [“Completing the upgrade to a new level of IBM Storage Scale”](#) on page 753. To know upgrade support for protocols and performance monitoring, see [“Online upgrade support for protocols and performance monitoring”](#) on page 673.

Upgrading an AFM-DR system

Complete the following steps to upgrade an AFM-DR system:

Procedure

1. Stop replication from the AFM-DR primary site to the AFM-DR secondary site.
2. Upgrade the AFM-DR secondary site. If IBM Storage Scale does not automatically start, restart IBM Storage Scale.
3. Restart replication from the AFM-DR primary site to the AFM-DR secondary site.
4. When the primary and secondary sites are synchronized, fail over to the secondary site.
5. Upgrade the AFM-DR primary site. If IBM Storage Scale does not automatically start, restart IBM Storage Scale.
6. When the two sites are again synchronized, fail over back to the original AFM-DR primary site.
7. Run the **mmchconfig release=LATEST** and **mmchfs -V full** commands on both clusters.

Restarting replication to upgrade AFM and AFM DR

You can use replication stop and start method to assist an upgrade. Replication must be stopped on all filesets of a gateway node you want to upgrade.

When the replication is stopped, you can shut down the gateway node for an upgrade. Local data availability is not affected as the filesets can be accessed by other nodes. However, replication activity stops during upgrade. You must ensure that the replication is restarted after the gateway node is upgraded.

To restart replication, do the following steps:

1. Plan downtime to upgrade IBM Storage Scale. AFM is stopped on the gateway node and the GPFS file system is unmounted and shutdown.
2. Issue the following command to identify filesets that belong to the gateway node you want to upgrade:

```
mmafctl fsname getstate |grep GatewayNode
```

3. Before upgrade, check if any messages are left in the queue. Wait until all pending messages are completed.

4. Issue the following command to stop replication of the fileset:

```
mmafmctl fsname stop -j filesetname
```

5. Before you upgrade the gateway node, issue the following command to verify that the replication stopped:

```
mmafmctl fsname getstate
```

6. Unmount the file system and shut down GPFS on the gateway node.
7. Upgrade the gateway node.
8. Issue **mmstartup** to start the gateway node daemon.
9. Check the node state by issuing **mmgetstate**.
10. After the node is active, mount the file systems and issue the following command:

```
mmafmctl fsname start -j filesetname
```

Recovery is initiated and replication starts.

Note: Synchronization of data between home and cache stops during the upgrade process. Recovery is run after you start replication to ensure that data at home and cache is synchronized.

Upgrading object packages

About this task

Important:

- CES Swift Object protocol feature is not supported from IBM Storage Scale 5.2.0 onwards.
- IBM Storage Scale 5.1.8 is the last release that has CES Swift Object protocol.
- IBM Storage Scale 5.2.0 will tolerate the update of a CES node from IBM Storage Scale 5.1.8.
 - *Tolerate* means:
 - The CES node will be updated to 5.2.0.
 - Swift Object support will not be updated as part of the 5.2.0 update.
 - You may continue to use the version of Swift Object protocol that was provided in IBM Storage Scale 5.1.8 on the CES 5.2.0 node.
 - IBM will provide usage and known defect support for the version of Swift Object that was provided in IBM Storage Scale 5.1.8 until you migrate to a supported object solution that IBM Storage Scale provides.
 - CES Swift Object is replaced with IBM Storage Scale S3. For more details, refer the *S3 support overview* section the *IBM Storage Scale: Concepts, Planning, and Installation Guide*.
 - For more information about Swift Object in IBM Storage Scale, refer to the [IBM Storage Scale 5.2.0 documentation](#).
- Contact IBM for further details and migration planning.



Attention:

- All protocol nodes must be upgraded to the newer code level along with the base GPFS.
- The object storage packages on all protocol nodes must be upgraded to the same level at the same time.
- Online upgrade of the object protocol is not supported because of changes that make it incompatible with servers that are running on the previous level. The upgrade must be scheduled at a time when downtime is appropriate.

- If you are upgrading from an IBM Storage Scale version earlier than 5.1.x and the node operating system version is earlier than Red Hat Enterprise Linux 8.x, the nodes need to be upgraded to Red Hat Enterprise Linux 8.x during the upgrade. For more information, see [“Instructions for removing object protocol packages when upgrading protocol nodes to Red Hat Enterprise Linux 8.x” on page 767.](#)
- The object protocol is not supported in IBM Storage Scale 5.1.0.0. If you want to deploy object, upgrade to the IBM Storage Scale 5.1.0.1 or a later release.

Procedure

1. From one of the protocol nodes, run the following command to stop all object services on all protocol nodes:

```
mmces service stop obj -a
```

2. On all protocol nodes, update the `spectrum-scale-object` package to the newer version as follows:

- Use the **yum** command to update `spectrum-scale-object` on Red Hat Enterprise Linux.
 - a. If needed, set up a yum repository for object packages by creating a new file in `/etc/yum.repos.d`, such as `/etc/yum.repos.d/gpfs_object.repo`, with the following content:

```
[gpfs_object]
name=gpfs_object
baseurl=file:///usr/lpp/mmfs/VERSION/object_rpms/rhel8
enabled=1
gpgcheck=0
```

Make sure that the path in the `baseurl` property corresponds to the path to the newer object packages on the system.

- b. The Object protocol in IBM Storage Scale 5.1.2.0 and earlier 5.1.x releases requires OpenStack 16 repositories to be available on all protocol nodes to satisfy the necessary dependencies.
- c. After the repository is set up, run the following command to update the `spectrum-scale-object` package:

```
yum upgrade -y spectrum-scale-object
```

- d. Repeat these steps on each protocol node to update the `spectrum-scale-object` package on all protocol nodes to the newer version.
- e. If a repository file in `/etc/yum.repos.d` is created in a preceding step, remove the file after the upgrade completes.

Migrate the existing object protocol configuration to be consistent with the level of installed packages after you upgrade all protocol nodes to the newer version of the `spectrum-scale-object` package.

Note: Use the following command to find the list of configured protocol nodes:

```
mmllsnode -N CESNodes
```

3. From one of the protocol nodes, run the following command:

```
mmobj config manage --version-sync
```

This command needs to be issued only one time in the cluster to update the configuration on all protocol nodes.

4. Run the following command to start object protocol-related services on all protocol nodes:

```
mmces service start obj -a
```

Related tasks

[“Upgrading Swift Object packages to IBM Storage Scale 5.1.8 or earlier by using the installation toolkit” on page 741](#)

Upgrading SMB packages

Use these steps to upgrade IBM Storage Scale SMB packages.

About this task



Attention: All protocols must also be upgraded to the newer code level along with base GPFS.

SMB package update in an IBM Storage Scale cluster is done in two phases to minimize the downtime of the SMB service. The protocol nodes in the cluster are divided in two halves and SMB is updated on each half one by one. In the following steps, *NodeList1* is the comma-separated list of nodes in the first half of the cluster and *NodeList2* is the comma-separated list of nodes in the second half of the cluster.

Note: All protocol nodes that are running the SMB service must have the same version of `gpfs.smb` installed at any time. A brief outage of the SMB service is required to upgrade `gpfs.smb` to the newer version across all protocol nodes. The procedure that is outlined here is intended to reduce the outage to a minimum.

Procedure

1. On the first half of the protocol nodes, do the following steps.

- a) Issue the following command to suspend the nodes and stop all the protocol services that are running on these nodes.

```
mmces node suspend -N NodeList1 --stop
```

Note: Suspending nodes triggers IP address reassignment and client failover. For more information, see [“SMB fail-over scenarios and upgrade” on page 394](#).

- b) On each node in the first half of protocol nodes, issue one of the following commands, depending on the operating system, to upgrade the SMB package.

- Red Hat Enterprise Linux or SLES

```
rpm -Uvh gpfs.smb*.rpm
```

The name of the IBM Storage Scale SMB package for SLES is in this format:
`gpfs.smb_version_string.sles15.architecture.rpm`.

Note: After you upgrade the operating system on the protocol nodes from SLES 12, you might need to uninstall any obsolete packages.

- Ubuntu

```
apt install ./gpfs.smb*.deb
```

The name of the IBM Storage Scale SMB package for Ubuntu 20.04 LTS is in this format:
`gpfs.smb_version_string~focal_architecture.deb`.

Note: After you upgrade the operating system on the protocol nodes from Ubuntu 16.04 or Ubuntu 18.04, uninstall any obsolete packages by issuing the following command on each of these nodes.

```
apt-get autoremove
```

2. On the second half of the protocol nodes, do the following steps.

- a) Issue the following command to suspend the nodes and stop all the protocol services that are running on these nodes.

```
mmces node suspend -N NodeList2 --stop
```

Note: This command suspends the remaining protocol nodes and results in a brief outage of the SMB service till the upgraded nodes are started again in the following step. Suspending nodes triggers IP address reassignment and client failover. For more information, see [“SMB fail-over scenarios and upgrade”](#) on page 394.

3. On the first half of the protocol nodes, issue the following command to resume the nodes and start all the protocol services that are enabled on these nodes.

```
mmces node resume -N NodeList1 --start
```

For more information on viewing the health status of a node, and detailed information about the health status and potential corrective actions, see *System health monitoring use cases and Events in IBM Storage Scale: Problem Determination Guide*.

4. On the second half of the protocol nodes, do the following steps.

- a) On each node in the second half of protocol nodes, issue one of the following commands, depending on the operating system, to upgrade the SMB package.

- Red Hat Enterprise Linux or SLES

```
rpm -Uvh gpfs.smb*.rpm
```

The name of the IBM Storage Scale SMB package for SLES is in this format:
`gpfs.smb_version_string.sles15.architecture.rpm`.

Note: After you upgrade the operating system on the protocol nodes from SLES 12, you might need to uninstall any obsolete packages.

- Ubuntu

```
apt install ./gpfs.smb*.deb
```

The name of the IBM Storage Scale SMB package for Ubuntu 20.04 LTS is in this format:
`gpfs.smb_version_string~focal_architecture.deb`.

Note: After you upgrade the operating system on the protocol nodes from Ubuntu 16.04 or Ubuntu 18.04, uninstall any obsolete packages by issuing the following command on each of these nodes.

```
apt-get autoremove
```

- b) Issue the following command to resume the nodes and start all the protocol services that are enabled on these nodes.

```
mmces node resume -N NodeList2 --start
```

For more information on viewing the health status of a node, and detailed information about the health status and potential corrective actions, see *System health monitoring use cases and Events in IBM Storage Scale: Problem Determination Guide*.

Related concepts

[“Protocol authentication configuration changes during upgrade”](#) on page 744

During IBM Storage Scale protocol nodes upgrade, do protocol authentication-related configuration depending on your authentication setup.

Upgrading the SMB package after upgrading OS

If you have upgraded the operating system on your protocol nodes, ensure that the SMB package is upgraded to the version required for the new OS.

About this task

After upgrading the operating system and rebuilding the GPFS portability layer (GPL), upgrade to the SMB package that is specific to the new OS by using [this upgrade procedure](#).

For SMB package names, see [“Manually installing the IBM Storage Scale software packages on Linux nodes”](#) on page 477.

Upgrading NFS packages

Upgrade the IBM Storage Scale NFS packages on each protocol node one by one. During the upgrade, a node is suspended, NFS packages are updated based on the operating system, and the node is resumed.

About this task



Attention:

- Use only NFS packages that are included in the IBM Storage Scale or provided as e-fix.
- Do not change configuration by using the **mmnfs config** command until all nodes are upgraded and the **mmnfs config upgrade** is run.

Procedure

1. Suspend the node.

```
mmces node suspend --stop
```

Note: When nodes are suspended, IP addresses are reassigned and the client failover occurs.

2. Create a temporary folder and copy all NFS packages for operating system and architecture to the folder.

Note: Enable the OS package repository to install any additional operating system packages that might be required.

- For Red Hat Enterprise Linux:

```
dnf update gpfs.nfs-ganesha*
```

If the Ganesha debuginfo packages are not installed, the system skips these packages. Install these packages by running the following command:

```
dnf install gpfs.nfs-ganesha*debuginfo*
```

- For SLES:

```
zypper update gpfs.nfs-ganesha*
```

If the Ganesha debuginfo packages are not installed, the system skips these packages. Install these packages by running the following command:

```
zypper install gpfs.nfs-ganesha*debuginfo*
```

- For Ubuntu:

```
apt-get update gpfs.nfs-ganesha*
```

If the Ganesha debuginfo packages are not installed, the system skips these packages. Install these packages by running the following command:

```
apt-get install gpfs.nfs-ganesha*debuginfo*
```

3. Resume the suspended node.

```
mmces node resume --start
```

For more information about viewing the health status of a node, and detailed information about the health status and potential corrective actions, see *System health monitoring use cases and Events* in *IBM Storage Scale: Problem Determination Guide*.

4. After the NFS packages are upgraded on all nodes, issue the following command once from any node to apply NFS configuration changes for Ganesha V7:

Note: Do not make any configuration changes by using the **mmnfs** during the upgrade.

```
mmnfs config upgrade
```

Note:

- If the NFS protocol is not upgraded on some CES nodes, upgrade the affected nodes when those nodes become available.
- After the NFS is upgraded on all nodes, issue the following command from one of the nodes:

```
mmnfs config upgrade
```

For more information about migrating from CNFS to CES NFS, see *Migration of CNFS clusters to CES clusters* in *IBM Storage Scale: Administration Guide*.

Related concepts

[“Protocol authentication configuration changes during upgrade” on page 744](#)

During IBM Storage Scale protocol nodes upgrade, do protocol authentication-related configuration depending on your authentication setup.

Upgrading S3 packages

Use these steps on each protocol node one by one to upgrade IBM Storage Scale S3 packages.

About this task



Attention:

- If one or more CES nodes cannot be upgraded, then issue the command **mmces node suspend -N Node[,Node...] --stop** to suspend the nodes and stop all the protocol services that are running on these nodes. This suspension ensures that the S3 service does not automatically start when the nodes are available again.
- All protocol nodes must also be upgraded to the newer code level along with base GPFS.
- S3 users must run I/O workloads with an increased value for the variables **AWS_MAX_ATTEMPTS** and **AWS_MAX_RETRY** to sustain the high availability. Because CES IP movement might take time based on the active I/O from multiple protocols, if enabled in the cluster. Refer AWS CLI documentation for how to set these.
- Back up S3 protocol configuration before an offline or online upgrade. For more information, see *Manual backup of the S3 configuration data* in the *Backing up the S3 configuration data* section in *IBM Storage Scale: Administration Guide*.

Complete the following steps on all nodes:

Procedure

1. Suspend the node by issuing the following command.

```
mmces node suspend --stop
```

Note: Suspending nodes triggers IP address reassignment and client failover.

2. From IBM Storage Scale 5.2.2 onward, the S3 protocol stack logging mechanism has some improvements. If you are using IBM Storage Scale 5.2.1.x, copy the `/etc/rsyslog.d/noobaa_syslog.conf` to `/etc/rsyslog.d/noobaa_syslog_521.conf`.
3. Upgrade S3 packages on the node.
 - For Red Hat Enterprise Linux, issue the following command.

```
xpm -Uvh gpfs.mms3 noobaa-core
```

4. When you upgrade from IBM Storage Scale 5.2.1.x, the following files are listed in the `/etc/logrotate.d` directory:

```
-rw-r--r-- 1 root root 971 Oct 22 23:21 noobaa-logrotate
-rw-r--r-- 1 root root 791 Oct 22 23:56 logrotate_noobaa.conf
```

where, the `noobaa-logrotate` is the new file created, update this file with any modifications performed in the `env` previously (that is, in `logrotate_noobaa.conf`).

Rename the `logrotate_noobaa.conf` as `logrotate_noobaa_521.conf`.

5. Copy the `noobaa_syslog.conf.rpmnew` file, which NooBaa creates after the upgrade, to the `noobaa_syslog.conf` file in the `/etc/rsyslog.d` directory.
6. Resume the node by issuing the following command.

```
mmces node resume --start
```



Attention: If a cluster is upgrading from an earlier release of IBM Storage Scale 5.2.3.0, issue the following command to disable the NooBaa service:

```
systemctl disable noobaa
```

The NooBaa service does not start automatically after a reboot.

7. After upgrading S3 on all nodes, issue the following command only one time from any one node to finish the S3 protocol upgrade.

```
mms3 upgrade start [--upgrade_version]
```



Attention: After issuing the **mms3 upgrade start** command, if the **mms3** command is used on nodes where S3 service is not upgraded, then it might lead to unexpected behavior.

Note:

If the S3 service on some CES nodes was not upgraded, whenever these nodes are available for upgrade complete:

- a. For each non-upgraded CES node, complete steps 1 through 6 in this procedure.
- b. After upgrading S3 on non-upgraded nodes, run the following command from one of the nodes:

```
mms3 upgrade start [--upgrade_version]
```

Related concepts

[“Protocol authentication configuration changes during upgrade” on page 744](#)

During IBM Storage Scale protocol nodes upgrade, do protocol authentication-related configuration depending on your authentication setup.

Upgrading call home

If you are upgrading from IBM Storage Scale releases 4.2.0 or 4.2.1 to release 5.0.x, call home related configuration changes are needed.

The call home function is upgraded with IBM Storage Scale upgrade. For more information, see [“Upgrade of IBM Storage Scale nodes” on page 674](#).

Call home configuration changes to be made while upgrading to IBM Storage Scale 5.0.x from IBM Storage Scale 4.2.1

When you upgrade from IBM Storage Scale 4.2.1 to IBM Storage Scale 5.0.0 or later, the existing group configuration will be incompatible with the new format.

Follow one of the given methods to deploy the existing group configuration after an upgrade:

- If the existing call home groups do not need to be saved, you can use the `mmcallhome group auto --force` command to automatically delete all the current groups, and create new groups with the necessary configuration.
- If you want to save the existing call home group configuration, follow the given steps:
 1. Find the current call home group configuration by using the `mmcallhome group list` command.
 2. Delete all listed call home groups by using the `mmcallhome group delete <groupName>` command.
 3. Add all the call home groups that were listed in step [“1” on page 702](#), by using the following command:

```
mmcallhome group add GroupName server [--node {all | ChildNode[,ChildNode...]}]
```

4. Set the new groups to track the global settings.
 - a. Run the `mmcallhome capability list`, `mmcallhome proxy list`, or `mmcallhome schedule list` commands to view the current global settings and default settings for the new groups.

The system displays output similar to this:

```
mmcallhome capability list
group      callHomeNode      status
-----
global     ---                     enabled
g1         g5001-21.localnet.com   disabled
g2         g5001-22.localnet.com   disabled
```

- b. For each group, execute the `mmcallhome capability`, `mmcallhome proxy`, the `mmcallhome schedule` commands to configure the system with the existing settings available in the global field.

After each change for each group, the corresponding group entry will disappear from the list command output. Therefore, the new group uses the global value for their setting. The system displays output similar to this:

```
mmcallhome capability list
group      callHomeNode      status
-----
global     ---                     enabled
g1         g5001-21.localnet.com   disabled
g2         g5001-22.localnet.com   disabled
```

Change the capability to enabled for the group g1. The g1 group is set to use the global value, and is not listed in the output of the **mmcallhome capability list** command anymore, as shown in the sample output:

```
mmcallhome capability enable accept
Call home enabled has been set to true

Additional messages:
License acceptance specified on command line. Callhome enabled.
mmcallhome capability list
group      callHomeNode      status
-----
global     ---                      enabled
g2         g5001-22.localnet.com     disabled
```

After this procedure is successfully completed, the **mmcallhome capability list**, **mmcallhome proxy list**, and **mmcallhome schedule list** commands must list only one global set of values.

Call home configuration changes to be made while upgrading to IBM Storage Scale 4.2.1 from IBM Storage Scale 4.2.0

The existing call home configurations are lost while upgrading from IBM Storage Scale 4.2.0 to IBM Storage Scale 4.2.1 or later versions.

If you want to deploy the existing configuration after an upgrade, follow the given steps:

1. Query and save your existing configuration before the upgrade by using the following commands:

```
mmcallhome group list
mmcallhome capability list
mmcallhome info list
mmcallhome proxy list
mmcallhome schedule list
# cat /usr/lpp/mmfs/data/callhome/etc/ENmain.properties | grep 'CustomerId\|Email\|
CountryCode'
# cat /usr/lpp/mmfs/data/callhome/etc/ENlocation.properties | grep CustomerName
```

2. Run the upgrade.
3. Reconfigure call home with the configuration that are saved in step 1, using the **mmcallhome** command. For more information about how to configure call home, see the *Configuring call home to enable manual and automated data upload* section in the *IBM Storage Scale: Problem Determination Guide*.

Removing residual configuration files while upgrading to IBM Storage Scale 5.1.x

The `gpfs.callhome-ecc-client` package is not needed when upgrading to IBM Storage Scale version 5.1.x from an earlier release. This is because the whole Call Home feature is now contained within the base IBM Storage Scale package, `gpfs.base`.

After the upgrade, some residual configuration files that are not required anymore might still be present on Ubuntu. For example:

```
root@shortwing-vm1:~# apt list gpfs.callhome-ecc-client
Listing... Done
gpfs.callhome-ecc-client/now 5.0.2-3 all [residual-config]
```

Use the following command to remove these residual configuration files:

```
apt --purge remove gpfs.callhome-ecc-client
```

Upgrading the performance monitoring tool

Upgrading performance monitoring does not upgrade the existing perfmon configuration. For example, new sensors that are introduced in newer releases are not enabled automatically when you upgrade the performance monitoring tool to the later release. Enable the new sensors manually afterward.

To list the new sensors/releases and the steps to enable these sensors, see `mmperfmon config add` command in *mmperfmon command* in the *IBM Storage Scale: Command and Programming Reference Guide*.

The update path does not touch existing perfmon config and new sensors needs to be enabled manually after you upgrade the performance monitoring tool. When you enable the sensor that has an impact on the pmcollector memory consumption as well.

Note: There are two new sensors, GPFSPool and GPFSFileset for the pmsensor service. If an older version of the IBM Storage Scale performance monitoring system is upgraded, these sensors are not automatically enabled. This is because automatically enabling the sensors might cause the collectors to consume more main memory than what was set aside for monitoring. Changing the memory footprint of the collector database might cause issues for the users if the collectors are tightly configured.

Manually upgrading the performance monitoring tool

Use these steps to upgrade performance monitoring packages.

Prerequisites:

- Before uninstalling, make sure that you stop the sensors and the collector. For more information see, *Starting and stopping the performance monitoring tool* in *IBM Storage Scale: Problem Determination Guide*.
- For more information about uninstalling, see [“Uninstalling the performance monitoring tool” on page 666](#).
- For more information about manual installation, see [“Manually installing the performance monitoring tool” on page 495](#).
- If federation is used for performance monitoring, the collectors participating in a federation must have the same version number.

To upgrade the performance monitoring tool, you must upgrade the `gpfs.gss.pmsensors` and `gpfs.gss.pmcollector` packages:

1. Upgrade to the new version of the performance monitoring tool by using the operating system native package upgrade mechanism as shown:

```
rpm -Uvh gpfs.gss.pmsensors-<version>.<os>.<arch>.rpm.
```

Note: Make sure that `pmsensors` and `pmcollector` are updated using the appropriate RPM packages for the platform.

2. Restart the sensors on the protocol nodes.
3. Restart the collector on the node that previously ran the collector.

Note: If multiple collectors are used, make sure that they all have the same version.

4. Verify that the collector and the sensor are running on the node by issuing the `systemctl status pmcollector` or `systemctl status pmsensors` command.

Note: If you use additional sensor packages such as `pmswift` or `gpfs.pm-ganesha`, upgrade these packages as well.

Reinstalling performance monitoring collector on Red Hat Enterprise Linux 8.x nodes

If performance collection functionality is enabled on an IBM Storage Scale node and you run a leapp upgrade from Red Hat Enterprise Linux 7.x to 8.x, the `el7_gpfs.gss.pmcollector` package is removed and the `el8` package needs to be reinstalled after the leapp upgrade.

Reinstall `gpfs.gss.pmcollector` on all Red Hat Enterprise Linux 8.x collector nodes as follows.

Note: For this procedure, it is assumed that when the `gpfs.gss.pmcollector` package was removed when leapp upgrade was invoked, the performance monitoring collector configuration was saved in `/opt/IBM/zimon/ZIMonCollector.cfg.rpmsave`.

1. Install the applicable `gpfs.gss.pmcollector` package.

The performance monitoring `el8` package is located in the `/usr/lpp/mmfs/5.x.x.x/zimon_rpms/rhel8` directory.

2. Back up the current performance monitoring collector configuration.

```
mv /opt/IBM/zimon/ZIMonCollector.cfg /opt/IBM/zimon/ZIMonCollector.cfg.orig
```

3. Restore the earlier saved version of the performance monitoring collector configuration.

```
cp /opt/IBM/zimon/ZIMonCollector.cfg.rpmsave /opt/IBM/zimon/ZIMonCollector.cfg
```

4. Restart the collector.

```
systemctl restart pmcollector
```

For more information about the performance monitoring tool, see *Configuring the Performance Monitoring tool* in *IBM Storage Scale: Problem Determination Guide*.

Upgrading the performance monitoring packages after upgrading Ubuntu OS

If you have upgraded the Ubuntu operating system on your IBM Storage Scale nodes, more steps are required to ensure that performance monitoring packages are upgraded to the version required for the latest version of Ubuntu.

About this task

After upgrading the operating system and rebuilding the GPFS portability layer (GPL), upgrade to the performance monitoring packages that Ubuntu version by using [this upgrade procedure](#).

For more information about performance monitoring package names for Ubuntu, see [“Manually installing the IBM Storage Scale software packages on Linux nodes”](#) on page 477.

Upgrading signed kernel modules for UEFI secure boot on x86_64 and Linux on Z

When the kernel of the node where IBM Storage Scale runs or the kernel of IBM Storage Scale itself is to be upgraded, make sure that the available signed kernel modules are an exact match.



Warning: If secure boot is enabled, do not upgrade the kernel or IBM Storage Scale unless you get a `gpfs.gplbin` package with matching IBM Storage Scale kernel modules. If the kernel versions do not match, IBM Storage Scale does not start.

To install the new kernel modules, follow the process that is described in [Chapter 17, “Installing the signed kernel modules for UEFI secure boot on x86_64,”](#) on page 655 or [Chapter 18, “Installing the signed kernel modules for secure boot on Linux on Z,”](#) on page 657.

Manually upgrading pmswift

To upgrade pmswift you can either uninstall and reinstall the pmswift rpm or use the native rpm upgrade command: `rpm -Uvh pmswift-version-release.noarch.rpm`. After upgrading, restart the service by using the following command: `systemctl restart pmswiftd.service`.

Uninstall pmswift-version-release.noarch.rpm

1. Stop the pmsensors.service by using the following command:

```
systemctl stop pmsensors.service
```

2. If uninstalling pmswift-4.1.1-4 or later, stop the pmswiftd.service by using the following command:

```
systemctl stop pmswiftd.service
```

If uninstalling pmswift-4.1.1-3, stop the pmprovider.service by using the following command:

```
systemctl stop pmprovider.service
```

3. Uninstall the pmswift rpm by using the following command:

```
rpm -evh --nodeps pmswift
```

If you are uninstalling pmswift-4.1.1-3, it should edit the Object configuration files for all Object servers and remove the entries created at the time of installation. The Object configuration files in `/etc/swift/` directory are:

- account - *.conf
- container - *.conf
- object - *.conf
- proxy - *.conf

This should also edit the sensors configuration file, `/opt/IBM/zimon/ZIMonSensors.cfg`, to remove the Object related sensors entries created at the time of installation. If you are uninstalling pmswift-4.1.1-4 or later these files will be left alone.

4. Ensure that following directories/files are removed. If they are not removed, you can, remove them manually.
 - a. `/usr/local/swiftmon` directory or `/usr/local/pmswift` directory
 - b. `/var/log/swiftmon` directory or `/var/log/pmswift` directory
 - c. `/var/run/swiftmon` directory or `/var/run/pmswift.pid` file
 - d. For pmswift-4.1.1-4 and later remove `/etc/rc.d/init.d/pmswift` file and for pmswift-4.1.1-3 remove `/etc/rc.d/init.d/pmprovider` file
 - e. For pmswift-4.1.1-3 `SwiftAccount.cfg`, `SwiftContainer.cfg`, `SwiftObject.cfg` and `SwiftProxy.cfg` files from within the Performance Monitoring tool's installation directory, `/opt/IBM/zimon/`.
5. Ensure that for pmswift-4.1.1-3 the pmprovider.service and for pmswift-4.1.1-4 and later the pmswiftd.service is not available anymore by running the following command:

```
systemctl daemon-reload
```

Install pmswift-version-release.noarch.rpm

1. Install the pmswift rpm by using the following command:

```
rpm -ivh pmswift-version-release.noarch.rpm
```

2. Ensure that following directories/files are created:

- a. /usr/local/pmswift directory
 - b. /var/log/pmswift directory
 - c. /etc/logrotate.d/pmswift file
 - d. /etc/rc.d/init.d/pmswiftd file
 - e. SwiftAccount.cfg, SwiftContainer.cfg, SwiftObject.cfg and SwiftProxy.cfg files in the Performance Monitoring tool's installation directory, /opt/IBM/zimon/.
3. Edit the Object configuration files for all Object servers that reside in CCR, using the /usr/local/pmswift/bin/pmswift-config-swift set command. CCR will then propagate modified configuration files to /etc/swift/ directory on all the protocol nodes within the cluster. The modified configuration files are:
- account - *.conf
 - container - *.conf
 - object - *.conf
 - proxy - *.conf
4. Edit the sensors configuration file information stored in the CCR using the /usr/local/pmswift/bin/pmswift-config-zimon set command to add the following Object related sensors entries:

```

{
    # SwiftAccount operational metrics
    name = "SwiftAccount"
    period = 1
    type = "generic"
    restrict= "cesNodes"
},
{
    # SwiftContainer operational metrics
    name = "SwiftContainer"
    period = 1
    type = "generic"
    restrict= "cesNodes"
},
{
    # SwiftObject operational metrics
    name = "SwiftObject"
    period = 1
    type = "generic"
    restrict= "cesNodes"
},
{
    # SwiftProxy operational metrics
    name = "SwiftProxy"
    period = 1
    type = "generic"
    restrict= "cesNodes"
},

```

These entries are then automatically propagated to the ZIMonSensors.cfg file in /opt/IBM/zimon on all the nodes in the cluster.

5. Start the pmswiftd.service by using the following command:

```
systemctl start pmswiftd.service
```

6. Start the pmsensors.service by using the following command:

```
systemctl start pmsensors.service
```

Manually upgrading the IBM Storage Scale management GUI

You can upgrade the IBM Storage Scale management GUI to the latest version to get the latest features. You can upgrade one GUI node at a time without shutting down IBM Storage Scale on other nodes to ensure high availability.

Prerequisite

Ensure that you are aware of the following details before you start the upgrade process:

- The management GUI package must be upgraded after all others required IBM Storage Scale components such as `gpfs.base`, `gpfs.java`, or `gpfs.gss.pmcollector` are successfully updated.
- If you are upgrading from IBM Storage Scale 5.0.0 or earlier versions, and you store your LDAP, AD, or GUI SSL configuration in `/opt/ibm/wlp/usr/servers/gpfsGui/server.xml` then you need to back up your settings. After the upgrade, recreate your configuration with **mkldap** for LDAP or AD and with **sethttpskeystore** for GUI SSL certificates. For more information, see *Configuring external authentication for GUI users* in *IBM Storage Scale: Administration Guide*.
- Local users, user groups, user roles, and snapshot rules are not affected with an upgrade, because they exist in the cluster configure repository (CCR).
- Data in the postgres database is retained.
- All IBM Storage Scale packages must be of the same release on the GUI node. For example, do not mix the 4.2.3 GUI rpm with a 4.2.2 base rpm. However, GUI PTFs and fixes can usually be applied without having to install the corresponding PTF or fix of the base package. This method is helpful if you want to resolve a GUI issue without changing anything on the base layer.
- Release levels can be different in GUI nodes and other nodes of the cluster. However, the minimum release level of the cluster must be 4.2.0.0 or later for the GUI to function.
- The user ID `scalemgmt` must not be already used as the GUI requires this user ID to run IBM Storage Scale GUI Websphere Java process.
- The ports 80, 443, 474080, and 47443 must not be already used by other processes.
- You can directly upgrade the IBM Storage Scale GUI from 4.2.0.0 or later to the latest version.

Perform the following steps to manually upgrade the management GUI:

1. Upgrade the GUI package.

For upgrading the previously installed package, use **rpm -Fvh** or **rpm -Uvh** options on RHEL or SLES. The **rpm -Fvh** is used for upgrading the existing installed package and **rpm -Uvh** is used for installing the package and upgrading the package as well.

For upgrading on Ubuntu, use **dpkg -i**. By issuing the **dpkg -s gpfs.gui** command, you can check more properties about your management GUI installation. If the status in the command output is `install ok installed` then the upgrade is performed successfully.

- On RHEL or SLES, issue the following command:

```
rpm -Fvh gpfs.gui-6.0.1-x.noarch.rpm
```

- On Ubuntu, issue the following command:

```
dpkg -i gpfs.gui_6.0.1-x_all.deb
```

2. After installing the packages, you can configure the performance monitoring tools by using the management GUI, if it is required. For more information, see [“Enabling performance tools in management GUI”](#) on page 501.
3. If the minimum release level set for IBM Storage Scale is not same as the GUI version, change the release level by issuing the **mmchconfig release=LATEST** command. As changing the minimum release level affects the cluster behavior, refer the **mmchconfig** command man page and other related topics before you make this configuration change. For more information, see [“Completing the upgrade to a new level of IBM Storage Scale”](#) on page 753.

4. Issue the **systemctl status gpfsGUI** command to verify the GUI service status.
5. Issue the **systemctl status pmcollector** and **systemctl status pmsensors** commands to verify the status of the performance tool.

Upgrading GUI in a FIPS-enabled environment

You can upgrade the GUI in a FIPS-enabled environment.

Perform the following steps on all nodes where the GUI is enabled:

1. Extract the IBM Storage Scale SE package.
2. Navigate to the `gpfs_rpms` directory and upgrade the Java RPM with flags `--nofiledigest` and `--nodigest`:

```
rpm -Uvh --nofiledigest --nodigest gpfs.java-5.2.3-3.x86_64.rpm
```

3. Upgrade the GUI RPM with flags `--nofiledigest` and `--nodigest`:

```
rpm -Uvh --nofiledigest --nodigest gpfs.gui-5.2.3-3.noarch.rpm
```

4. Restart the GUI service:

```
systemctl restart gpfsGUI
```

Note: Once the installation is complete, deploy the cluster using the Ansible Toolkit.

Creating a CSI administrator after new installation or upgrade

The Container Storage Interface (CSI) needs an administrator user in the GUI and REST API user base to manage the CSI Driver for IBM file storage configuration. These users need to hold the CSI Administrator role. You need to assign this user role manually as no user holds this role by default. The steps that are used to create a user with this role differ depending on upgrade or a new installation scenario.

Adding a group that holds the CSI admin role after system upgrade

The CSI administrator group is created by default when you install the IBM Storage Scale management GUI. If you are upgrading the system from 5.0.3 or earlier, this user group does not get created by default. To create the user group with CSI administrator role, perform the following:

1. Go to **Services > GUI > Groups**.
2. Click **Create Group**. The Create User Group dialog appears.
3. Type the name of the user group, for example, *csiAdmin*, in the **User group name** field.
4. Select **CSI Administrator** from the list of user roles.
5. Click **Create** to create the user group with the CSI Administrator user role.

Assigning CSI Administrator user role after new installation or upgrade

For new installations and for upgrades, you must add members to the new group with the CSI Administrator role. If you create a new user, the approach is as follows:

1. Go to **Services > GUI > User**.
2. Click **Create User**. The Create User dialog appears.
3. Type the name of the user in the **Name** field.
4. Select *csiAdmin* from the list of user groups.
5. Create a temporary password for the user.
6. Click **Create** to complete user creation process.

Alternatively, you can add the existing users to a group by using the **Manage Group** option that is available in the **Actions** menu.

Upgrading transparent cloud tiering (discontinued)

To upgrade to a newer version of cloud services, first consider the version that you are migrating from and then consider coexistence and compatibility issues.

You can upgrade cloud services from one version to another. When you do so, the metrics that are associated with cloud services are upgraded as well. For example, in the current version, if there are 10 elements in the metrics, and when you upgrade, there might be 15 elements in the metrics. For more information, see *List of metrics for cloud services* topic in the *IBM Storage Scale: Administration Guide*.

Note: Before performing any migration, ensure that you back up the configuration files. For more information, see the *Backing up the cloud services configuration* topic in the *IBM Storage Scale: Administration Guide*.

cloud services migration consists of these topics:

- [“Upgrading to Transparent cloud tiering 1.1.2 from Transparent cloud tiering 1.1.0 or 1.1.1 ” on page 710](#)
- [“Upgrading to cloud services 1.1.2.1 from 1.1.2 ” on page 711](#)
- [“Upgrading to cloud services 1.1.3 from 1.1.2 ” on page 712](#)
- [“Upgrading to cloud services 1.1.3 from 1.1.0 and 1.1.1 ” on page 712](#)
- [“Upgrading to cloud services 1.1.4 from 1.1.2.x ” on page 713](#)
- [“Upgrading to cloud services 1.1.4 from 1.1.3 ” on page 714](#)

Upgrading to Transparent cloud tiering 1.1.2 from Transparent cloud tiering 1.1.0 or 1.1.1

This topic describes the procedure for upgrading to Transparent cloud tiering 1.1.2 from Transparent cloud tiering 1.1.0 or 1.1.1.

Before you begin

Ensure the following before you start the upgrade the process:

- Transparent cloud tiering service is stopped on all the server nodes by using the **mmcloudgateway service stop** command.
- IBM Storage Scale is at release 4.2.2. Migration might fail if your IBM Storage Scale release is different.
- No data migration or recall is in progress.

About this task

Perform the following steps:

Procedure

1. Copy the `gpfs.tct.server-1.1.2.x86_64.rpm` file to each of the nodes that is specified as Transparent cloud tiering server nodes.
2. Run this command:

```
rpm -Uvh gpfs.tct.server-1.1.2.x86_64.rpm
```

Note: Transparent cloud tiering statistics on the IBM Storage Scale GUI is enabled only when the `ENABLE_MCSTORE` parameter in the `/usr/lpp/mmfs/gui/conf/gpfsgui.properties` is set to "true". During upgrade from IBM Storage Scale 4.2.1.1 to 4.2.2, this parameter is reset to "false", even though its original value is "true", and the Transparent cloud tiering statistics does not show up on the IBM Storage Scale GUI. Therefore, it is recommended to manually set the value of the property to "true" after an upgrade. To activate this change, you must restart the `gpfsgui.service`. All

other configurations will be intact after the upgrade, and you can continue to perform data migration seamlessly as before.

3. To upgrade the client package, issue the following command:

```
rpm -Uvh gpfs.tct.client-1.1.2.1.rpm
```

4. To upgrade the Debian package, issue the following command:

```
sudo dpkg -i gpfs.tct.client*.deb
```

5. To verify the version after upgrade, issue the following command:

```
mmcloudgateway service version
```

Upgrading to cloud services 1.1.2.1 from 1.1.2

This topic describes the procedure for upgrading to cloud services 1.1.2.1 from 1.1.2.

Before you begin

Ensure the following before you start the upgrade the process:

- Transparent cloud tiering service is stopped on all the server nodes by using the **mmcloudgateway service stop** command.
- IBM Storage Scale is at release 4.2.2.1. Migration might fail if your IBM Storage Scale release is different.
- No data migration or recall is in progress.

About this task

Perform the following steps:

Procedure

1. Copy the `gpfs.tct.server-1.1.2.1.x86_64.rpm` file to each of the nodes that is specified as Transparent cloud tiering server nodes.
2. Run this command:

```
rpm -Uvh gpfs.tct.server-1.1.2.1.x86_64.rpm
```

Note: Transparent cloud tiering statistics on the IBM Storage Scale GUI is enabled only when the `ENABLE_MCSTORE` parameter in the `/usr/lpp/mmfs/gui/conf/gpfsGUI.properties` is set to "true". During upgrade from IBM Storage Scale 4.2.2 to 4.2.2.1, this parameter is reset to "false", even though its original value is "true", and the Transparent cloud tiering statistics does not show up on the IBM Storage Scale GUI. Therefore, it is recommended to manually set the value of the property to "true" after an upgrade. To activate this change, you must restart the `gpfsGUI.service`. All other configurations will be intact after the upgrade, and you can continue to perform data migration seamlessly as before.

3. For upgrading the client system, issue this command:

```
rpm -Uvh gpfs.tct.client-1.1.2.1.rpm
```

4. For upgrading the Debian package, issue this command:

```
sudo dpkg -i gpfs.tct.client-1.1.3*.deb
```

Upgrading to cloud services 1.1.3 from 1.1.2

This topic describes the procedure for upgrading to cloud services 1.1.3 from 1.1.2.

Before you begin

Ensure the following before you start the upgrade the process:

- Transparent cloud tiering service is stopped on all the server nodes by using the **mmcloudgateway service stop** command.
- IBM Storage Scale is at release 4.2.3. Migration might fail if your IBM Storage Scale release is different.
- No data migration or recall is in progress.

About this task

Perform the following steps:

Procedure

1. Copy the `gpfs.tct.server-1.1.3.x86_64.rpm` file to each of the nodes that is specified as Transparent cloud tiering server nodes.
2. Run this command:

```
rpm -Uvh gpfs.tct.server-1.1.3.x86_64.rpm
```

3. To upgrade the client nodes, issue this command:

```
rpm -Uvh gpfs.tct.client-1.1.3*.rpm
```

4. To upgrade the Debian package, issue this command:

```
sudo dpkg -i gpfs.tct.client-1.1.3*.deb
```

5. Verify the version by using the **mmcloudgateway service version** command. The system displays output similar to this:

```
Cluster minReleaseLevel: 4.2.2.0

Node  Daemon node name          TCT Type          TCT Version      Equivalent Product
Version
-----
1    c350f1u1b7.pok.stglabs.ibm.com  Server            1.1.3            4.2.3
```

Note: This upgrade takes longer than others because of the database schema change. The larger the database of migrated files, the longer this upgrade will take.

Upgrading to cloud services 1.1.3 from 1.1.0 and 1.1.1

This topic describes the procedure for upgrading to cloud services 1.1.3 from 1.1.0 or 1.1.1.

Before you begin

Ensure the following before you start the upgrade the process:

- Transparent cloud tiering service is stopped on all the server nodes by using the **mmcloudgateway service stop** command.
- IBM Storage Scale is at release 4.2.3. Migration might fail if your IBM Storage Scale release is different.
- No data migration or recall is in progress.

About this task

Perform the following steps:

Procedure

1. Copy the `gpfs.tct.server-1.1.3.x86_64.rpm` file to each of the nodes that is specified as Transparent cloud tiering server nodes.
2. Run this command:

```
rpm -Uvh gpfs.tct.server-1.1.3.x86_64.rpm
```

3. To upgrade the client nodes, issue the following command:

```
rpm -Uvh gpfs.tct.client-1.1.3*.rpm
```

4. To upgrade the Debian package, issue the following command:

```
sudo dpkg -i gpfs.tct.client-1.1.3*.deb
```

5. Verify the version by using the **mmcloudgateway service version** command. The system displays output similar to this:

```
Cluster minReleaseLevel: 4.2.2.0

Node Daemon node name          TCT Type          TCT Version      Equivalent Product
Version
-----
---
    1  c350f1u1b7.pok.stglabs.ibm.com  Server           1.1.3            4.2.3
```

Upgrading to cloud services 1.1.4 from 1.1.2.x

This topic describes the procedure for upgrading to cloud services 1.1.4 from 1.1.2.x

Before you begin

Ensure the following before you start the upgrade the process:

- Uninstalling the old rpm (**rpm -e** command) and then installing the new rpm (**rpm -i**) will not work. The upgrade must be done with the **rpm -Uvh** command only.
- Transparent cloud tiering service is stopped on all the server nodes by using the **mmcloudgateway service stop** command.
- IBM Storage Scale is at release 5.0.0. Upgrade might fail if your IBM Storage Scale release is different.
- File system version is upgraded to the latest (18.00 and beyond) by using the **mmchfs** command, to leverage the default transparent recall policy. If file system version is not upgraded to a given level prior to cloud services rpm upgrade, enabling transparent recalls for a container that is created with earlier cloud services versions will fail. If file system version is upgraded to 18.00 and beyond after upgrading the cloud services rpms, you must restart cloud services by using the **mmcloudgateway service** command to be able to effectuate transparent recalls for containers.
- No data migration or recall is in progress.

About this task

Perform the following steps:

Procedure

1. Copy the `gpfs.tct.server-1.1.4.x86_64.rpm` file to each of the nodes that is specified as Transparent cloud tiering server nodes.
2. Run this command:

```
rpm -Uvh gpfs.tct.server-1.1.4.x86_64.rpm
```

3. To upgrade the client nodes, issue this command:

```
rpm -Uvh gpfs.tct.client-1.1.4*.rpm
```

4. To upgrade the Debian package, issue this command:

```
sudo dpkg -i gpfs.tct.client-1.1.4*.deb
```

5. Verify the version by using the **mmcloudgateway service version** command. The system displays output similar to this:

```
Cluster minReleaseLevel: 5.0.0.0

Node  Daemon node name          TCT Type          TCT Version      Equivalent Product
Version
-----
---
1     c350f1u1b7.pok.stglabs.ibm.com  Server            1.1.4            5.0.0
```

Note: This upgrade takes longer than others because of the database schema change. The larger the database of migrated files, the longer this upgrade will take.

What to do next

Note:

- After the upgrade, the previous sensors are no more valid and cloud services performance monitoring will fail on the GUI. Therefore, you must remove the old sensors and populate new ones. For more information, see [“Upgrading the cloud services sensors ” on page 718](#).
- After the upgrade, do not use any transparent recall policy that was applied in IBM Storage Scale 4.2.x and earlier releases. You might experience hangs during transparent recalls if the old policies are reused. Therefore, you must remove the older policy (which was applied before upgrade using **mmchpolicy <file-system> -P Default**) and enable the transparent recall for the container using the **mmcloudgateway containerpairset update** command, after the upgrade.

Upgrading to cloud services 1.1.4 from 1.1.3

This topic describes the procedure for upgrading to cloud services 1.1.4 from 1.1.3.

Before you begin

Ensure the following before you start the upgrade the process:

- Uninstalling the old rpm (**rpm -e** command) and then installing the new rpm (**rpm -i**) will not work. The upgrade must be done with the **rpm -Uvh** command only.
- Transparent cloud tiering service is stopped on all the server nodes by using the **mmcloudgateway service stop** command.
- IBM Storage Scale is at release 5.0.0. Migration might fail if your IBM Storage Scale release is different.
- File system version is upgraded to the latest (18.00 and beyond) by using the **mmchfs** command, to leverage the default transparent recall policy. If file system version is not upgraded to a given level prior to cloud services rpm upgrade, enabling transparent recalls for a container that is created with earlier cloud services versions will fail. If file system version is upgraded to 18.00 and beyond after upgrading the cloud services rpms, you must restart cloud services by using the **mmcloudgateway service** command to be able to effectuate transparent recalls for containers.
- No data migration or recall is in progress.

About this task

Perform the following steps:

Procedure

1. Copy the `gpfs.tct.server-1.1.4.x86_64.rpm` file to each of the nodes that is specified as Transparent cloud tiering server nodes.
2. Run this command:

```
rpm -Uvh gpfs.tct.server-1.1.4.x86_64.rpm
```

3. To upgrade the client nodes, issue this command:

```
rpm -Uvh gpfs.tct.client-1.1.4*.rpm
```

4. To upgrade the Debian package, issue this command:

```
sudo dpkg -i gpfs.tct.client-1.1.4*.deb
```

5. To upgrade the SLES package, issue this command:

```
rpm -Uvh gpfs.tct.server-1.1.4*.rpm
```

6. Verify the version by using the **mmcloudgateway service version** command. The system displays output similar to this:

```
Cluster minReleaseLevel: 5.0.0.0

Node Daemon node name          TCT Type          TCT Version      Equivalent Product
Version
-----
---
1   c350f1u1b7.pok.stglabs.ibm.com Server             1.1.4            5.0.0
```

Note: This upgrade takes longer than others because of the database schema change. The larger the database of migrated files, the longer this upgrade will take.

What to do next

- After the upgrade, the previous sensors are no more valid and cloud services performance monitoring will fail on the GUI. Therefore, you must remove the old sensors and populate new ones. For more information, see [“Upgrading the cloud services sensors ” on page 718](#).
- After the upgrade, do not use any transparent recall policy that was applied in IBM Storage Scale 4.2.x and earlier releases. You might experience hangs during transparent recalls if the old policies are reused. Therefore, you must remove the older policy (which was applied before upgrade using **mmchpolicy <file-system> -P Default**) and enable the transparent recall for the container using the **mmcloudgateway containerpairset update** command, after the upgrade.

After the upgrade, the previous sensors are no more valid and cloud services performance monitoring will fail on the GUI. Therefore, you must remove the old sensors and populate new ones. For more information, see [“Upgrading the cloud services sensors ” on page 718](#).

Upgrading to cloud services 1.1.5 from 1.1.4

This topic describes the procedure for upgrading to cloud services 1.1.5 from 1.1.4.

Before you begin

Ensure the following before you start the upgrade the process:

- Uninstalling the old rpm (**rpm -e** command) and then installing the new rpm (**rpm -i**) will not work. The upgrade must be done with the **rpm -Uvh** command only.

- Transparent cloud tiering service is stopped on all the server nodes by using the **mmcloudgateway service stop** command.
- IBM Storage Scale is at release 5.0.0. Migration might fail if your IBM Storage Scale release is different.
- File system version is upgraded to the latest (18.00 and beyond) by using the **mmchfs** command, to leverage the default transparent recall policy. If file system version is not upgraded to a given level before cloud services rpm upgrade, enabling transparent recalls for a container that is created with earlier cloud services versions will fail. If file system version is upgraded to 18.00 and beyond after upgrading the cloud services rpms, you must restart cloud services by using the **mmcloudgateway service** command to be able to effectuate transparent recalls for containers.
- No data migration or recall is in progress.

Note: Ensure that you uninstall transparent recall policy as you need to use transparent recall-enabled flag in place of the transparent recall policy.

About this task

Perform the following steps:

Procedure

1. Copy the `gpfs.tct.server-1.1.5.x86_64.rpm` file to each of the nodes that is specified as Transparent cloud tiering server nodes.
2. Run this command:

```
rpm -Uvh gpfs.tct.server-1.1.5.x86_64.rpm
```

3. To upgrade the client nodes, issue this command:

```
rpm -Uvh gpfs.tct.client-1.1.5*.rpm
```

4. To upgrade the Debian package, issue this command:

```
sudo dpkg -i gpfs.tct.client-1.1.5*.deb
```

5. To upgrade the SLES package, issue this command:

```
rpm -Uvh gpfs.tct.server-1.1.5*.rpm
```

6. Verify the version by using the **mmcloudgateway service version** command. The system displays output similar to this:

```
Cluster minReleaseLevel: 5.0.1

Node  Daemon node name          TCT Type      TCT Version    Equivalent Product
Version
-----
1    c350f1u1b7.pok.stglabs.ibm.com Server         1.1.5          5.0.1
```

Note: This upgrade takes longer than others because of the database schema change. The larger the database of migrated files, the longer this upgrade takes.

What to do next

After the upgrade, the previous sensors are no more valid and cloud services performance monitoring will fail on the GUI. Therefore, you must remove the old sensors and populate new ones. For more information, see [“Upgrading the cloud services sensors” on page 718](#).

Upgrading to cloud services 1.1.6 from 1.1.5

This topic describes the procedure for upgrading to cloud services 1.1.6 from 1.1.5.

Before you begin

Ensure the following before you start the upgrade the process:

- Uninstalling the old rpm (**rpm -e** command) and then installing the new rpm (**rpm -i**) will not work. The upgrade must be done with the **rpm -Uvh** command only.
- Transparent cloud tiering service is stopped on all the server nodes by using the **mmcloudgateway service stop** command.
- IBM Storage Scale is at release 5.0.2. Migration might fail if your IBM Storage Scale release is different.
- File system version is upgraded to the latest (18.00 and beyond) by using the **mmchfs** command, to leverage the default transparent recall policy. If file system version is not upgraded to a given level prior to cloud services rpm upgrade, enabling transparent recalls for a container that is created with earlier cloud services versions will fail. If file system version is upgraded to 18.00 and beyond after upgrading the cloud services rpms, you must restart cloud services by using the **mmcloudgateway service** command to be able to effectuate transparent recalls for containers.
- No data migration or recall is in progress.

About this task

Perform the following steps:

Procedure

1. Copy the `gpfs.tct.server-1.1.6.x86_64.rpm` file to each of the nodes that is specified as Transparent cloud tiering server nodes.
2. Run this command:

```
rpm -Uvh gpfs.tct.server-1.1.6.x86_64.rpm
```

3. To upgrade the client nodes, issue this command:

```
rpm -Uvh gpfs.tct.client-1.1.6*.rpm
```

4. To upgrade the Debian package, issue this command:

```
sudo dpkg -i gpfs.tct.client-1.1.6*.deb
```

5. To upgrade the SLES package, issue this command:

```
rpm -Uvh gpfs.tct.server-1.1.6*.rpm
```

6. Verify the version by using the **mmcloudgateway service version** command. The system displays output similar to this:

```
Cluster minReleaseLevel: 5.0.2

Node  Daemon node name          TCT Type      TCT Version   Equivalent Product
Version
-----
1    c350f1u1b7.pok.stglabs.ibm.com Server         1.1.6         5.0.2
```

Note: This upgrade takes longer than others because of the database schema change. The larger the database of migrated files, the longer this upgrade will take.

What to do next

After the upgrade, the previous sensors are no more valid and cloud services performance monitoring will fail on the GUI. Therefore, you must remove the old sensors and populate new ones. For more information, see [“Upgrading the cloud services sensors ” on page 718](#).

Upgrading the cloud services sensors

The sensors in earlier versions of the cloud services are no longer valid in the 1.1.4 version. Therefore, you must change the sensors after upgrading to cloud services version 1.1.4 from either 1.1.3.x or 1.1.2.x version as described below.

Before you begin

You must know which sensors are configured and active.

About this task

Perform the following steps to upgrade the sensors:

Procedure

1. Remove configuration of the old sensors by issuing the following commands:

- `mmperfmon config delete --sensors MCStoreGPFSSStats`
- `mmperfmon config delete --sensors MCStoreIcstoreStats`
- `mmperfmon config delete --sensors MCStoreLWESStats`

2. Remove the old sensors files from the `/opt/IBM/zimon` folder of every cloud services node:

- `rm MCStoreGPFSSStats.cfg`
- `rm MCStoreIcstoreStats.cfg`
- `rm MCStoreLWESStats.cfg`

3. Copy the following sensor files from the `/opt/ibm/MCStore/config/` folder to the `/opt/IBM/zimon` folder:

- `TCTDebugDbStats.cfg`
- `TCTDebugLweDestroyStats.cfg`
- `TCTFsetGpfsConnectorStats.cfg`
- `TCTFsetIcstoreStats.cfg`
- `TCTFsGpfsConnectorStats.cfg`
- `TCTFsIcstoreStats.cfg`

4. On each cloud services server node, create a new `MCStore-sensor-definition.cfg` file under `/opt/IBM/zimon/` and copy the following sensors text into it:

```
sensors=
{
  #Transparent cloud tiering statistics
  name = "TCTDebugDbStats"
  period = 10
  type = "Generic"
},
{
  #Transparent cloud tiering statistics
  name = "TCTDebugLweDestroyStats"
  period = 10
  type = "Generic"
```

```

},
{
    #Transparent cloud tiering statistics
    name = "TCTFsetGpfsConnectorStats"
    period = 10
    type = "Generic"
},

{
    #Transparent cloud tiering statistics
    name = "TCTFsetIcstoreStats"
    period = 10
    type = "Generic"
},
{
    #Transparent cloud tiering statistics
    name = "TCTFsGpfsConnectorStats"
    period = 10
    type = "Generic"
},

{
    #Transparent cloud tiering statistics
    name = "TCTFsIcstoreStats"
    period = 10
    type = "Generic"
}

```

5. Specify the following command to register the sensors:

```
mmperfmon config add --sensors MCStore-sensor-definition.cfg
```

6. Specify the following command to verify that all the sensors are added:

```
mmperfmon config show
```

The system displays output similar to this:

```

# This file has been generated automatically and SHOULD NOT
# be edited manually. It may be overwritten at any point
# in time.

cephMon = "/opt/IBM/zimon/CephMonProxy"
cephRados = "/opt/IBM/zimon/CephRadosProxy"
colCandidates = "jupiter-vm856.pok.stglabs.ibm.com", "jupiter-vm934.pok.stglabs.ibm.com"
colRedundancy = 2
collectors = {
    host = ""
    port = "4739"
}
config = "/opt/IBM/zimon/ZIMonSensors.cfg"
ctdbstat = ""
daemonize = T
hostname = ""
ipfixinterface = "0.0.0.0"
logfile = "/var/log/zimon/ZIMonSensors.log"
loglevel = "info"
mmcmd = "/opt/IBM/zimon/MMCmdProxy"
mmdfcmd = "/opt/IBM/zimon/MMDFProxy"
mmpmon = "/opt/IBM/zimon/MmpmonSockProxy"
piddir = "/var/run"
release = "6.0.1-x"
sensors = {
    name = "CPU"
    period = 1
},
{
    name = "Load"
    period = 1
},
{
    name = "Memory"
    period = 1
},
{
    name = "Network"
    period = 1
},
}

```

```

    {
        name = "Netstat"
        period = 10
    },
    {
        name = "Diskstat"
        period = 0
    },
    {
        name = "DiskFree"
        period = 600
    },
    {
        name = "Infiniband"
        period = 0
    },
    {
        name = "GPFSDisk"
        period = 0
    },
    {
        name = "GPFSFilesystem"
        period = 10
    },
    {
        name = "GPFSNSDDisk"
        period = 10
        restrict = "nsdNodes"
    },
    {
        name = "GPFSPoolIO"
        period = 0
    },
    {
        name = "GPFSVFS"
        period = 10
    },
    {
        name = "GPFSIOC"
        period = 0
    },
    {
        name = "GPFSVIO64"
        period = 0
    },
    {
        name = "GPFSPPDisk"
        period = 10
        restrict = "nsdNodes"
    },
    {
        name = "GPFSvFLUSH"
        period = 0
    },
    {
        name = "GPFSNode"
        period = 10
    },
    {
        name = "GPFSNodeAPI"
        period = 10
    },
    {
        name = "GPFSFilesystemAPI"
        period = 10
    },
    {
        name = "GPFSLROC"
        period = 0
    },
    {
        name = "GPFSCHMS"
        period = 0
    },
    {
        name = "GPFSAFM"
        period = 0
    },
    {
        name = "GPFSAFMFS"
        period = 0
    },
}

```

```

{
  name = "GPFSAFMFSET"
  period = 0
},
{
  name = "GPFSRPCS"
  period = 10
},
{
  name = "GPFSWaiters"
  period = 10
},
{
  name = "GPFSFilesetQuota"
  period = 3600
  restrict = "jupiter-vm856.pok.stglabs.ibm.com"
},
{
  name = "GPFSFileset"
  period = 300
  restrict = "jupiter-vm934.pok.stglabs.ibm.com"
},
{
  name = "GPFSPool"
  period = 300
  restrict = "jupiter-vm934.pok.stglabs.ibm.com"
},
{
  name = "GPFSDiskCap"
  period = 86400
  restrict = "jupiter-vm856.pok.stglabs.ibm.com"
},
{
  name = "TCTDebugDbStats"
  period = 10
  type = "Generic"
},
{
  name = "TCTDebugLweDestroyStats"
  period = 10
  type = "Generic"
},
{
  name = "TCTFsetGpfsConnectorStats"
  period = 10
  type = "Generic"
},
{
  name = "TCTFsetIcstoreStats"
  period = 10
  type = "Generic"
},
{
  name = "TCTFsGpfsConnectorStats"
  period = 10
  type = "Generic"
},
{
  name = "TCTFsIcstoreStats"
  period = 10
  type = "Generic"
}
smbstat = ""

```

7. Specify the following commands to restart your pmsensors and pmcollector on the nodes:

- `service pmsensors restart`
- `service pmcollector restart`

8. Verify the status of your pmsensors and pmcollector on the nodes:

- `service pmsensors status`
- `service pmcollector status`

Upgrading to IBM Cloud Object Storage software level 3.7.2 and above

If you use IBM Cloud Object Storage software level system (3.7.x and below) as the cloud storage tier for some time and then upgrade to IBM Cloud Object Storage software level 3.7.2 or above, you must perform certain configurations. Perform these configurations because of a change in computation of the entity tag (ETag) in IBM Cloud Object Storage software level 3.7.2.

About this task

Do the following steps to upgrade to IBM Cloud Object Storage software level 3.7.2 and above:

Procedure

1. To display the configured cloud account, run this command: **mmcloudgateway account list**
2. Edit `/var/MCStore/.mcstore_settings` and change the configuration property "`<cloudname>.provider=cleversafe`" to "`<cloudname>.provider=cleversafe-new`".
3. Restart the transparent cloud tiering service by running this command: **mmcloudgateway service restart**

Upgrade paths and commands for file audit logging and clustered watch folder

Use the following information to upgrade file audit logging and clustered watch folder in IBM Storage Scale.

File audit logging base functionality was introduced in IBM Storage Scale 5.0.0.

Clustered watch folder base functionality was introduced in IBM Storage Scale 5.0.3.

In IBM Storage Scale 5.1.0, there is a large architecture change that removes the dependency for the message queue and the `gpfs.kafka` package. The **mmmsgqueue config --remove-msgqueue** command automatically upgrades all components that are needed and disables then re-enables all of the existing healthy audits and watches. When you run this command, there is a small outage where no events are generated. The prerequisites to run this command are as follows:

- Run **mmchconfig release=LATEST** first.
- All file system versions where file audit logging or clustered watch folder are enabled or are planned to be enabled must set to at least IBM Storage Scale 5.1.0 or later. See **mmllsfs <fs> -V** to query the file system version.
- When upgrading file audit logging and clustered watch folder in remote cluster environments, you must ensure the following:
 - For file audit logging, ensure that your system meets the requirements for using file audit logging with remotely mounted file systems. For more information, see [“Requirements for using file audit logging with remotely mounted file systems”](#) on page 648.
 - To avoid version conflicts, it is required that you upgrade the accessing clusters before upgrading the owning cluster. Unexpected errors might occur on the remotely mounted file systems with file audit logging enabled if the remote cluster is at IBM Storage Scale 5.0.1 or earlier and the owning cluster was upgraded first.

Note: For more information about issues that might occur during an upgrade, see *File audit logging issues* in the *IBM Storage Scale: Problem Determination Guide*.

If file audit logging or clustered watch folder has never been enabled or your cluster was installed at 5.1.2 or later, the **mmmsgqueue config --remove-config** command does not need to run.

If you are upgrading from 5.0.5.X with **msgqueue** enabled, you must disable file audit logging and clustered watch folder and then run **mmmsgqueue config --remove** before the upgrade. File audit logging and clustered watch folder can be enabled after the upgrade completes.

If you are upgrading from 5.1.0.X/5.1.1.X with **msgqueue** enabled, you must run **mmmsgqueue config --remove-msgqueue** prior to the upgrade. This re-enables file audit logging and clustered watch folders automatically.

After a full cluster upgrade to 5.1.3 and above, regardless of file system level, the **mmrestorefs** command will not restore files that are located in the `.audit_log` fileset or configuration fileset. The current configuration will not be overwritten and audit records will not be removed or restored.

If running file audit logging or watch folder in a mixed release environment, issue the management commands from the lowest release level node.

Upgrading IBM Storage Scale components with the installation toolkit

You can use the installation toolkit to do an upgrade of almost all IBM Storage Scale components from the current version of IBM Storage Scale to the wanted version. These components include core GPFS, NFS, SMB, S3, HDFS, management GUI, file audit logging, and performance monitoring. The installation toolkit can be used to do an online upgrade, an entire cluster offline upgrade, or an upgrade with a subset of nodes offline or excluded to meet your environment needs.

Note:

- The installation toolkit detects and upgrades the IBM Storage Scale native REST API feature as part of the IBM Storage Scale upgrade from 5.2.3.1 version.
- Root and passwordless SSH access must be available on all cluster nodes to perform upgrade activities by using the installation toolkit.
- You can upgrade CES S3 protocol from IBM Storage Scale 5.2.1.1.

For more information about the upgrade phases and the flow of the upgrade procedure that is done by using the installation toolkit, see [“Upgrade process flow”](#) on page 725.

Important:

- Protocol upgrades are non-concurrent and they cause brief outages to user access. Schedule an outage window during the upgrade, which is approximately 12 minutes per protocol node.
 - Before upgrading from IBM Storage Scale 5.0.x to 5.1.x, remove the `gpfs.nfs-ganesha-selinux` package if it is present in the protocol node. If the `gpfs.nfs-ganesha-selinux` package is not removed, it might break the NFS Ganesha upgrade process.
 - Before upgrading from IBM Storage Scale 5.0.x to 5.1.x on the Ubuntu operating system, make sure that the upstream NFS Ganesha package is not installed or configured in the protocol nodes. If the NFS Ganesha package is installed or configured, uninstall the package by using the **apt purge nfs-ganesha** command.
 - For NFS upgrade, quiesce all I/O during the upgrade. During the NFS upgrade, I/O pauses occur and depending upon the client, NFS mounts and disconnects.
 - SMB requires quiescing all I/O during the upgrade. Due to the SMB clustering functions, differing SMB levels cannot co-exist within a cluster at the same time. This restriction requires a full outage of SMB during the upgrade.
- Before you proceed with the IBM Storage Scale upgrade, ensure that transparent cloud tiering is not active on the cluster. If it is active, shut down transparent cloud tiering services on all nodes by issuing the **mmcloudgateway service stop** command. If transparent cloud tiering is active, the IBM Storage Scale upgrade might fail.
- If you are upgrading from IBM Storage Scale 5.0.0 or earlier versions, and you store your LDAP, AD, or GUI SSL configuration in `/opt/ibm/wlp/usr/servers/gpfscli/server.xml` then you need to back up your settings. After the upgrade, recreate your configuration with **mkldap** for LDAP or AD

and with **sethttpskeystore** for GUI SSL certificates. For more information, see *Configuring external authentication for GUI users in IBM Storage Scale: Administration Guide*.

Before you proceed with upgrade by using the installation toolkit, ensure that you do the following steps:

- The prerequisite steps for using the installation toolkit are done. Doing these prerequisite steps are important if you are using the installation toolkit to upgrade a cluster that was created manually. For more information, see [“Preparing to use the installation toolkit” on page 528](#).
- Review limitations of the installation toolkit related to upgrade. For more information, see [“Limitations of the installation toolkit” on page 518](#).

For doing an upgrade by using the installation toolkit, a key requirement is the cluster definition file that you can create or update by using one of the following ways.

- Use the **./spectrumscale config populate** command to retrieve the cluster configuration and populate it in the cluster definition file.
- Manually enter the configuration into the cluster definition file by using the **./spectrumscale** command options.

For more information about obtaining the current cluster configuration, see [“Performing online upgrade by using the installation toolkit” on page 731](#).

Important: You can enable the prompt to administrators to shut down their workloads before starting the upgrade by using the following command.

```
./spectrumscale upgrade config workloadprompt -N Node1,Node2,...
```

The prompt is displayed on every specified node. You can also enable workload prompt on all nodes in the cluster definition file by using the **-N all** option of the command.

Note: You can skip package dependency check during the Toolkit upgrade precheck by using the following command.

```
./spectrumscale upgrade precheck --skip pkg-dependency-check
```

Upgrading by using the installation toolkit

- [“Performing online upgrade by using the installation toolkit” on page 731](#).
- [“Performing offline upgrade or excluding nodes from upgrade by using installation toolkit” on page 733](#).
- [“Upgrading IBM Storage Scale on IBM Storage Archive Enterprise Edition \(EE\) nodes by using the installation toolkit” on page 742](#).

For more information on upgrading HDFS, see [Upgrading CES HDFS Transparency in IBM Storage Scale: Big data and analytics support](#).

Upgrading nodes that have and interim fix installed

You can use the installation toolkit to upgrade nodes that have an interim fix (efix) that is installed.

Installing new functions after upgrade

You can use the installation toolkit to install new functions that are available in the new version that you upgraded to and that are supported by the installation toolkit. Use the following steps to do so:

1. Enable the functions in the cluster definition file by using the installation toolkit.
2. Issue the **./spectrumscale install** or the **./spectrumscale deploy** command.

Failure recovery

In case the upgrade does not complete successfully, certain steps need to be done to identify the issue and take appropriate action to recover from the upgrade failure. For more information, see *Upgrade issues* and *Upgrade recovery* in *IBM Storage Scale: Problem Determination Guide*.

Upgrade process flow

The upgrade procedure by using the installation toolkit is done in several phases. If you are doing a manual upgrade, you can use this process flow as a reference.

Important: Swift Object protocol is not supported from IBM Storage Scale 5.2.0 and later. If Swift Object is configured in the system, refer the 5.2.0 guide for the upgrade steps.

Upgrade phases

- [Upgrade precheck](#)
- [“Phase 1” on page 726](#): Upgrade of all non-protocol nodes
Note: If the nodes that are being upgraded are protocol nodes, this phase is not applicable.
- [“Phase 2” on page 727](#): Upgrade of protocol nodes (except SMB and CES components)
Note: If the nodes that are being upgraded are non-protocol nodes, this phase is not applicable.
- [“Phase 3” on page 729](#): SMB upgrade, and CES upgrade on protocol nodes; License package upgrade on all nodes; and File audit logging message queue enabling at cluster level.
Note: If the nodes that are being upgraded are non-protocol nodes, SMB upgrade, and CES upgrade tasks of this phase are not applicable.
- [Phase 4](#): CES HDFS upgrade on DataNodes and NameNodes
- [Upgrade post-check](#)


Upgrade precheck

Important: It is highly recommended to run the upgrade precheck while you are planning to upgrade, before the actual upgrade procedure. Doing this allows you adequate time to address any issues flagged during the upgrade precheck. The upgrade precheck can be done hours or days in advance and it does not have any impact on the cluster operations.

As part of the upgrade precheck, the following tasks are performed by the installation toolkit.

- Determine the installed IBM Storage Scale packages across the cluster on all nodes.
- Compare the versions of the installed packages with the versions in the repository of the packages that you want to upgrade to.
In a mixed operating system cluster, the comparison is done with the package repository applicable for the operating system running on the respective nodes.
- Determine which packages need to be upgraded on each node.
- For every package that needs to be upgraded or installed, dependency check is done to ensure that all dependencies are met. The upgrade procedure throws an error and exits, if any dependency is not met.
- Check whether passwordless SSH is set up.
- Check if supported OS is installed.
- Check whether upgrade path is valid.


Phase 1



1	Check the cluster health.
	Perform following steps on each non-protocol node, one node at a time.
2	Check the node health.
3	Unmount file systems and shut down GPFS.  Attention: At this point, I/O from this node stops.
4	Upgrade core GPFS.
5	Build a GPFS portability layer (GPL).
6	Upgrade file audit logging. Install if file audit logging is enabled in the cluster configuration but not installed.
7	Upgrade performance monitoring.
8	Upgrade GUI.

9	Start GPFS and mount the file systems.  Attention: At this point, I/O from this node resumes.
10	Check the node health.
	These steps are repeated on the remaining non-protocol nodes, one node at a time. After all of the non-protocol nodes are upgraded, phase 1 is complete.


Phase 2



	Perform following steps on each protocol node, one node at a time.
1	Check the node health.

2	Suspend CES on the node.  Attention: At this point, protocol I/O from this node stops.
3	Unmount file systems and shut down GPFS.  Attention: At this point, I/O from this node stops.
4	Upgrade core GPFS.
5	Build GPFS portability layer (GPL).
6	Upgrade file audit logging. Install if file audit logging is enabled in the cluster configuration but not installed.
8	Upgrade NFS.
9	Upgrade performance monitoring.
10	Upgrade GUI.
11	Start GPFS and mount the file systems.  Attention: At this point, I/O from this node resumes.
12	Resume CES on the node.  Attention: At this point, protocol I/O from this node resumes.
13	Check the node health.
	These steps are repeated on the remaining protocol nodes, one node at a time. After all of the protocol nodes are upgraded, phase 2 is complete.

Phase 3



2	Suspend CES on the first half of nodes.
3	Stop SMB on the 1st half of nodes.  Attention: At this point, SMB I/O on this half of nodes stops and CES IPs on these nodes are migrated to the nodes in the other half.
4	Upgrade SMB on the 1st half of nodes.
5	Suspend CES on the 2nd half of nodes.  Attention: At this point, SMB I/O stops on the whole cluster for a brief period.

6	Stop SMB on the 2nd half of nodes.
7	Start SMB on the 1st half of nodes.
8	Resume CES on the 1st half of nodes.  Attention: At this point, SMB I/O resumes on this half of nodes.
9	Upgrade SMB on the 2nd half of nodes.
10	Start SMB on the 2nd half of nodes.
11	Resume CES on the 2nd half of nodes.  Attention: At this point, SMB I/O resumes in the whole cluster.
12	Upgrade CES on all protocol nodes, one node at a time.
13	Upgrade license package on all nodes, one node at a time.
14	Enable file audit logging message queue on the cluster.
15	Check the cluster health.
	After these steps are done, the installation toolkit automatically runs upgrade post-check.

Phase 4

In this phase, as part of the CES HDFS upgrade, the `gpfs.hdfs-protocol` package is upgraded on all DataNodes and NameNodes. For more information, see [Upgrading CES HDFS Transparency in IBM Storage Scale: Big data and analytics support](#).

Upgrade post-check

As part of the upgrade post-check, the following tasks are performed by the installation toolkit.

- Check the health of the nodes.
- Compare the versions of the installed packages with the versions in the package repository.
In a mixed operating system cluster, the comparison is done with the package repository applicable for the operating system running on the respective nodes.
- If any package is not upgraded, the installation toolkit prompts you to rerun the upgrade procedure.

You must complete the upgrade to the new code level to take advantage of the new functionality. For more information, see [“Completing the upgrade to a new level of IBM Storage Scale” on page 753](#).

Related concepts

[Performing offline upgrade or excluding nodes from upgrade by using installation toolkit](#)

[Upgrade rerun after an upgrade failure](#)

The installation toolkit supports upgrade rerun if an upgrade fails.

Related tasks

[Performing online upgrade by using the installation toolkit](#)

Use the following procedure to do an online upgrade of IBM Storage Scale by using the installation toolkit.

[Upgrading Swift Object packages to IBM Storage Scale 5.1.8 or earlier by using the installation toolkit](#)

[Upgrading IBM Storage Scale on IBM Storage Archive Enterprise Edition \(EE\) nodes by using the installation toolkit](#)

Performing online upgrade by using the installation toolkit

Use the following procedure to do an online upgrade of IBM Storage Scale by using the installation toolkit.

Procedure

1. Download the new IBM Storage Scale self-extracting package by using the substeps and then place it on the installer node.
 - a. Go to the [IBM Storage Scale page on Fix Central](#), select the new package, and then click **Continue**.
 - b. Choose the download option **Download using Download Director** to download the new spectrumscale package and place it in the wanted location on the install node.

Note: If you use **Download using your browser (HTTPS)**, do not click the down arrow to the left of the package name. Instead, right-click on the package name and select the **Save Link As** option. If you click the download arrow, the browser might hang.

2. Extract the new IBM Storage Scale self-extracting package by using the package name.
`/tmp/Spectrum_Scale_Standard-6.0.1.x_x86_64-Linux_install`

This step creates a new directory structure (`/usr/lpp/mmfs/6.0.1.x/`).

3. Accept the license agreement.

After the self-extracting package extraction is done, several next steps and instructions are displayed on the screen.

- Cluster installation and protocol deployment
- Upgrading an existing cluster by using the installation toolkit
- Adding nodes to an existing cluster by using the installation toolkit
- Adding NSDs or file systems to an existing cluster by using the installation toolkit

4. Depending on your setup, do one of the following sets of steps to obtain the current cluster configuration through the cluster definition file:

- If you are upgrading an IBM Storage Scale cluster that is created by using the installation toolkit, do these steps to upgrade by using the installation toolkit.

Note: If you used the installation toolkit for the initial installation and deployment, it is assumed that the prerequisites are already in place.

- a. If you want to change the installer node before the upgrade, issue the `spectrumscale setup -s InstallNodeIP` command to set up the new installer node.
- b. From the installer directory of the target version that you are upgrading to, populate the cluster definition file with the current cluster state by issuing the `./spectrumscale config populate` command. For more information, see [“Populating cluster definition file with current cluster state using the installation toolkit” on page 551](#) and [“Limitations of config populate option of the installation toolkit” on page 552](#).

If the command fails, update the cluster definition file with the cluster state change operations that were done after the initial installation and deployment by using the `./spectrumscale` commands. These cluster state change operations include:

- Node-related changes such as adding nodes
 - NSD-related changes such as adding or modifying NSDs
 - File system changes such as creating or modifying file systems
 - Protocols-related changes such as enabling or disabling protocols
- If you are upgrading a manually created IBM Storage Scale cluster, do these steps to upgrade by using the installation toolkit.
 - a. Ensure that the prerequisite steps for using the installation toolkit are completed. For more information, see [“Preparing to use the installation toolkit” on page 528](#).

b. Configure your cluster definition file as explained in the following topics:

- [“Setting up the installer node” on page 533](#)
- [“Defining the cluster topology for the installation toolkit” on page 534](#)

c. From the installer directory of the target version that you are upgrading to, populate the cluster definition file with the current cluster state by issuing the **./spectrumscale config populate** command. For more information, see [“Populating cluster definition file with current cluster state using the installation toolkit” on page 551](#) and [“Limitations of config populate option of the installation toolkit” on page 552](#).

If the config populate command fails, update the cluster definition file with the cluster state change operations that were done after the initial installation and deployment by using the **./spectrumscale** commands. These cluster state change operations include:

- Node-related changes such as adding nodes
- NSD-related changes such as adding or modifying NSDs
- File system changes such as creating or modifying file systems
- Protocols-related changes such as enabling or disabling protocols

5. Ensure that the cluster definition file is accurate and reflects the current state of your IBM Storage Scale cluster.

If you want to upgrade IBM Storage Scale on IBM Storage Archive Enterprise Edition (EE) nodes by using the installation toolkit, you must do some additional steps to manage IBM Spectrum Archive components. For more information, see [“Upgrading IBM Storage Scale on IBM Storage Archive Enterprise Edition \(EE\) nodes by using the installation toolkit” on page 742](#).

6. Run the upgrade precheck by using the **./spectrumscale upgrade precheck** command.

Important: It is highly recommended to run the upgrade precheck while you are planning to upgrade. Doing this step gives you adequate time to address any issues that are flagged during the upgrade precheck.

For any cluster with existing performance monitoring collectors, the upgrade prechecks might find that a reconfiguration is needed. To enable the installation toolkit to do the reconfiguration, issue the following command:

```
./spectrumscale config perfmon -r on
```

To keep your existing algorithms, issue the following command:

```
./spectrumscale config perfmon -r off
```

To do an offline upgrade based on the findings in the precheck results, do the steps in [“Designating nodes as offline in the upgrade configuration” on page 736](#).

7. Run the upgrade process by using the **./spectrumscale upgrade run** command.

Accept the warnings for system impact and interruption during the upgrade.

Note: On Ubuntu nodes, if the call home capability is enabled, after upgrading the call home packages to version 5.0.1 or later, issue the following commands from one of the nodes that is a part of a call home group.

```
mmcallhome capability disable  
mmcallhome capability enable accept
```

After the upgrade is completed, the installation toolkit automatically runs the upgrade post checks. You can manually run the upgrade post checks by issuing the **./spectrumscale upgrade postcheck** command.

8. If the fresh upgrade is not successful, do an upgrade rerun as needed. For more information, see [“Upgrade rerun after an upgrade failure” on page 741](#).

9. After the upgrade process is done, complete the upgrade to the new code level to take advantage of the new functions. For more information, see [“Completing the upgrade to a new level of IBM Storage Scale”](#) on page 753.
10. Do the protocols authentication configuration that is required after you upgrade protocol nodes. For more information, see [“Protocol authentication configuration changes during upgrade”](#) on page 744.

Related concepts

[Upgrade process flow](#)

[Performing offline upgrade or excluding nodes from upgrade by using installation toolkit](#)

[Upgrade rerun after an upgrade failure](#)

The installation toolkit supports upgrade rerun if an upgrade fails.

Related tasks

[Upgrading Swift Object packages to IBM Storage Scale 5.1.8 or earlier by using the installation toolkit](#)

[Upgrading IBM Storage Scale on IBM Storage Archive Enterprise Edition \(EE\) nodes by using the installation toolkit](#)

Performing offline upgrade or excluding nodes from upgrade by using installation toolkit

The installation toolkit can upgrade and tolerate nodes that are in an offline state, and it can exclude some nodes from the upgrade.

Note: Online upgrades can be done to the current release (N) from the previous release (N-1). For example, to release 5.1.y (N) from release 5.0.z (N-1). However, you can use the manual offline upgrade procedure to upgrade directly to the current release (N) from a previous release (N-2) or earlier. For example, upgrading to release 5.1.x (N) from release 4.2.x (N-2) or earlier. For more information, see [“Offline upgrade with complete cluster shutdown”](#) on page 769. You can also use the installation toolkit to do an offline upgrade from release 5.0.2 or later.

- [“Upgrade when nodes are unhealthy”](#) on page 735
- [“Designating nodes as offline in the upgrade configuration”](#) on page 736
- [“Upgrade when nodes are not reachable”](#) on page 737
- [“Excluding nodes from the upgrade configuration”](#) on page 737
- [“Upgrading the excluded nodes or offline designated nodes”](#) on page 737
- [“Parallel offline upgrade with the installation toolkit”](#) on page 738
- [“Upgrading mixed operating system cluster with Ubuntu nodes”](#) on page 739
- [“Upgrading mixed CPU architecture cluster”](#) on page 739
- [“Offline upgrade with the installation toolkit when OS upgrade is required”](#) on page 740
- [“Populating cluster configuration when nodes are designated as offline in the upgrade configuration”](#) on page 740
- [“Limitations of excluding nodes in the upgrade configuration”](#) on page 740

Tip: Before you start with any of the procedures included in the previous list, complete the procedure that is detailed in [Preparing to use the installation toolkit](#).

Preparing to use the installation toolkit

1. Download the new IBM Storage Scale self-extracting package by using the substeps and then place it on the installer node.
 - a. Go to the IBM Storage Scale page on Fix Central, select the new package, and then click **Continue**.
 - b. Choose the download option **Download using Download Director** to download the new spectrumscale package and place it in the wanted location on the install node.

Note: If you use **Download using your browser (HTTPS)**, do not click the down arrow to the left of the package name. Instead, right-click on the package name and select the **Save Link As** option. If you click the download arrow, the browser might hang.

2. Extract the new IBM Storage Scale self-extracting package by using the package name. `/tmp/Spectrum_Scale_Standard-6.0.1.x_x86_64-Linux_install`

This step creates a new directory structure (`/usr/lpp/mmfs/6.0.1.x/`).

3. Accept the license agreement.

After the self-extracting package extraction is done, several next steps and instructions are displayed on the screen.

- Cluster installation and protocol deployment
- Upgrading an existing cluster by using the installation toolkit
- Adding nodes to an existing cluster by using the installation toolkit
- Adding NSDs or file systems to an existing cluster by using the installation toolkit

4. Depending on your setup, do one of the following sets of steps to obtain the current cluster configuration through the cluster definition file.

If you are upgrading an IBM Storage Scale cluster that is created by using the installation toolkit, do these steps to upgrade by using the installation toolkit.

Note: If you used the installation toolkit for the initial installation and deployment, it is assumed that the prerequisites are already in place.

- a. If you want to change the installer node before the upgrade, issue the `spectrumscale setup -s InstallNodeIP` command to set up the new installer node.
- b. From the installer directory of the target version that you are upgrading to, populate the cluster definition file with the current cluster state by issuing the `./spectrumscale config populate` command. For more information, see [“Populating cluster definition file with current cluster state using the installation toolkit” on page 551](#) and [“Limitations of config populate option of the installation toolkit” on page 552](#).

If the command fails, update the cluster definition file with the cluster state change operations that were done after the initial installation and deployment by using the `./spectrumscale` commands. These cluster state change operations include:

- Node-related changes such as adding nodes
- NSD-related changes such as adding or modifying NSDs
- File system changes such as creating or modifying file systems
- Protocols-related changes such as enabling or disabling protocols

If you are upgrading a manually created IBM Storage Scale cluster, do these steps to upgrade by using the installation toolkit.

- a. Ensure that the prerequisite steps for using the installation toolkit are completed. For more information, see [“Preparing to use the installation toolkit” on page 528](#).
- b. Configure your cluster definition file as explained in the following topics:
 - [“Setting up the installer node” on page 533](#)
 - [“Defining the cluster topology for the installation toolkit” on page 534](#)
- c. From the installer directory of the target version that you are upgrading to, populate the cluster definition file with the current cluster state by issuing the `./spectrumscale config populate` command. For more information, see [“Populating cluster definition file with current cluster state using the installation toolkit” on page 551](#) and [“Limitations of config populate option of the installation toolkit” on page 552](#).

If the config populate command fails, update the cluster definition file with the cluster state change operations that were done after the initial installation and deployment by using the **./spectrumscale** commands. These cluster state change operations include:

- Node-related changes such as adding nodes
 - NSD-related changes such as adding or modifying NSDs
 - File system changes such as creating or modifying file systems
 - Protocols-related changes such as enabling or disabling protocols
5. Ensure that the cluster definition file is accurate and reflects the current state of your IBM Storage Scale cluster.

If you want to upgrade IBM Storage Scale on IBM Storage Archive Enterprise Edition (EE) nodes by using the installation toolkit, you must do some additional steps to manage IBM Spectrum Archive components. For more information, see [“Upgrading IBM Storage Scale on IBM Storage Archive Enterprise Edition \(EE\) nodes by using the installation toolkit” on page 742.](#)

Upgrade when nodes are unhealthy

By using the IBM Storage Scale offline upgrade, you can upgrade your cluster even if one or more nodes are unhealthy. A node is called unhealthy when the services are down but it is reachable through ping commands.

When you designate a node as offline in the cluster configuration, during the upgrade run, the installation toolkit upgrades all installed packages. However, no attempt is made to stop or restart the respective services. Manually restart the previously offline services by using these commands: **mmces service start** for protocol components and **mmstartup** for GPFS daemon.

For example, you try to upgrade a 5-node cluster whose nodes are node1, node2, node3, node4, and node5 (protocol node).

On upgrade precheck, you notice the following issues:

- node3 is reachable but NFS is down.
- node5 is reachable but SMB is down.
- node2 is reachable but all services, including GPFS, are down.

Here, you can do the following steps:

- Designate node3 as offline. During or after the upgrade, the installation toolkit does not restart any services, including NFS, on this node, but all installed packages, including NFS, are upgraded on node3.
- Designate node5 as offline. During or after the upgrade, the installation toolkit does not restart any services, including SMB, on this node, but all installed packages, including SMB, are upgraded on node5.
- Designate node2 as offline. All installed packages are upgraded, but none of the services are restarted.

Note:

- If you designate all nodes in a cluster as offline, then a full offline upgrade is done on all nodes. An upgrade of all installed packages is done without starting or stopping any services.
- If you try to designate a node that is already excluded as offline, then the exclude designation of the node is cleared, and the offline designation is added. For example,

```
./spectrumscale upgrade config offline -N vm1
```

```
[ INFO ] The node vm1.ibm.com was added in excluded list previously. Clearing  
this from excluded list. [ INFO ] Adding vm1.ibm.com as smb offline
```

- After an offline upgrade, you must ensure that all unhealthy services are manually started by using **mmces service start** for protocol components and **mmstartup** for GPFS.

Designating nodes as offline in the upgrade configuration

- To designate a node as offline, issue this command:

```
./spectrumscale upgrade config offline -N nodename
```

An offline upgrade is done on this node, which means that all installed packages are upgraded without restarting any services.

Important: Before you designate a node as offline, you must ensure that none of the components are active and if the node is a protocol node, then it must be suspended.

- To check the status of the GPFS daemon, issue the **mmgetstate** command.
- To stop the GPFS daemon, issue the **mmshutdown** command.
- To check the status of protocol components, issue the **mmces service list** command.
- To suspend the protocol node and stop the protocol services, issue the **mmces node suspend --stop** command.

If you are upgrading from IBM Storage Scale version 5.0.2.0 or earlier, issue the following commands to suspend the protocol node and stop the protocol services:

```
mmces node suspend
mmces service stop Protocol
```

- To designate all nodes as offline and do a full offline upgrade across the cluster, issue this command:

```
./spectrumscale upgrade config offline -N node1,node2,node3...,noden
```

All installed packages are upgraded on all the nodes in the cluster, but no services are restarted on any of the nodes.

- Clearing the offline designation.

- To clear all offline designations from a specific node, issue this command:

```
./spectrumscale upgrade config offline -N nodename --clear
```

- To clear all the offline designations from all the nodes, issue this command:

```
./spectrumscale upgrade config offline --clear
```

- To clear both the offline and the exclude configuration, issue this command:

```
./spectrumscale upgrade config clear
```

- To view all configuration that is done for offline upgrade, issue this command:

```
./spectrumscale upgrade config list
```

The output of this command includes the nodes that are excluded and the nodes where the components are designated as offline. An offline upgrade is initiated based on this configuration.

For example,

```
./spectrumscale upgrade config list
```

```
[ INFO ] GPFS Node SMB NFS OBJ GPFS
[ INFO ]
[ INFO ] Phase1: Non Protocol Nodes Upgrade
[ INFO ] nsd001st001 - - -
[ INFO ] nsd002st001 - - -
[ INFO ] nsd003st001 - - -
[ INFO ] nsd004st001 - - -
[ INFO ]
[ INFO ] Phase2: Protocol Nodes Upgrade
[ INFO ] prt002st001
[ INFO ] prt003st001
```

```
[ INFO ] prt004st001
[ INFO ] prt006st001
[ INFO ] prt008st001
[ INFO ] prt009st001
[ INFO ] prt011st001
[ INFO ]
[ INFO ] Excluded Nodes : prt007st001,prt001st001,prt010st001,prt005st001
[ INFO ]
```

Upgrade when nodes are not reachable

You can exclude one or more nodes from the current upgrade run, if the nodes are unreachable or if you want to upgrade them later. When you exclude a node from the upgrade configuration, no action is done on this node during the upgrade.

Note:

- It is not recommended to exclude a subset of protocol nodes. For example, if you have three protocol nodes, then you must exclude all three nodes together. It is not recommended to exclude only a subset (1 or 2) of nodes. For example,

```
./spectrumscale upgrade config exclude -N vm1
```

```
[ INFO ] Adding node vm1.ibm.com in excluded list.
[ WARN ] Protocol nodes should all be upgraded together if possible,
since mixed versions of the code are not allowed in CES components (SMB/OBJ).
You may add the remaining protocol node(s) : vm2.ibm.com in the excluded list or clear
node(s):
vm1.ibm.com with the ./spectrumscale config exclude --clear option so that no protocol nodes
are excluded.
```

- Ensure that not all admin nodes are excluded and that at least one admin node is available in the nonexcluded list. For example, if you have three admin nodes in a cluster that you want to upgrade, then you can exclude a maximum of two admin nodes only. If you have only one admin node, then it must not be excluded.

Excluding nodes from the upgrade configuration

- To exclude one or more nodes from the upgrade configuration, issue this command:

```
./spectrumscale upgrade config exclude -N node1,node2
```

This command ensures that the installation toolkit does not do any action on node1 and node2 during upgrade.

- Clearing the exclude designations.
 - To clear the exclude configuration from specific nodes, issue this command:

```
./spectrumscale upgrade config exclude -N node1,node2 --clear
```

Note: It is not recommended to clear only a subset of the protocol nodes that are designated as offline.

- To clear the exclude configuration from all nodes, issue this command:

```
./spectrumscale upgrade config exclude --clear
```

Upgrading the excluded nodes or offline designated nodes

1. Ensure that the nodes on which you want to do an offline upgrade are reachable through ping commands.
2. For nodes that are designated as excluded, clear the exclude designation of the nodes in the cluster definition file by using this command:

```
./spectrumscale upgrade config exclude -N node1,node2 --clear
```

- For nodes that are designated as excluded, designate them as offline if all services are not running by using this command:

```
./spectrumscale upgrade config offline -N node1,node2
```

- Run the upgrade procedure on the offline designated nodes by using this command:

```
./spectrumscale upgrade run
```

The installation toolkit upgrades the packages and restarts the services only for an online upgrade. For an offline upgrade, the installation toolkit only upgrades the packages that are currently installed on the offline designated nodes.

- After the upgrade procedure is completed, do the following:

- Restart the GPFS daemon by using the **mmstartup** command on each offline designated node.
- If the object protocol is configured, do the post-upgrade object configuration by using the following command from one of the protocol nodes.

```
mmobj config manage --version-sync
```

- Applicable for S3 protocol only. If a cluster is upgrading from an earlier release of IBM Storage Scale 5.2.3.0, issue the following command to disable the NooBaa service:

```
systemctl disable noobaa
```

Note: The NooBaa service does not start automatically after a reboot.

- Resume the protocol node and restart the protocol services by using the **mmces node resume --start** command for every offline designated node that is a protocol node.

If you are upgrading from IBM Storage Scale version 5.0.2.0 or earlier, issue the following commands to resume the protocol node and start the protocol services:

```
mmces node resume  
mmces service start Protocol
```

- Configure the protocols authentication that is required after you upgrade protocol nodes. For more information, see [“Protocol authentication configuration changes during upgrade”](#) on page 744.

- Upgrade the S3 protocol by running the following command on a node that is not excluded.

```
mms3 upgrade start [--upgrade_version] [-d]
```

Important: Make sure that you run this command on a healthy node where the S3 service RPM is upgraded, otherwise some **mms3** commands might not work as expected.

Parallel offline upgrade with the installation toolkit

Use the parallel offline upgrade option to upgrade all nodes in the cluster in parallel. By using parallel offline upgrade, you can reduce the overall upgrade duration compared to doing the upgrade in a serial manner.

Note:

- You can run parallel offline upgrade only:
 - If you can shut down IBM Storage Scale on all nodes in the cluster.
 - If you can suspend protocol services on all protocol nodes, if applicable, in the cluster.
 - If one of the nodes in the upgrade configuration is defined as an admin node.
- To run a parallel offline upgrade, all nodes in the cluster must be designated offline in the upgrade configuration. Otherwise, all the nodes are upgraded serially. If you want to use parallel offline upgrade

on a subset of nodes in the cluster, you must exclude the other nodes from the upgrade configuration by using the **./spectrumscale upgrade config exclude -N Node1, Node2,...** command. Only the nodes that are not designated as excluded in the upgrade configuration are upgraded.

Do a parallel offline upgrade as follows.

1. Shut down IBM Storage Scale on all nodes in the cluster.

```
# mmsshutdown -a
```

2. If your cluster comprises protocol nodes, suspend protocol services on those nodes.

```
# mmces node suspend --stop -a
```

3. Designate all nodes in the cluster as offline in the upgrade configuration of the installation toolkit.

```
# ./spectrumscale upgrade config offline -N all
```

4. Run the upgrade procedure.

```
# ./spectrumscale upgrade run
```

In the 1st phase, the installation toolkit upgrades core components such as GPFS, file audit logging, performance monitoring, and GUI on all nodes in parallel. In the 2nd phase, the installation toolkit upgrades all protocol components on all protocol nodes in parallel.

5. Upgrade the S3 protocol by running the following command on a node that is not excluded.

```
mms3 upgrade start [--upgrade_version] [-d]
```

Important: Make sure that you run this command on a healthy node where the S3 service RPM is upgraded, otherwise some **mms3** commands might not work as expected.

Upgrading mixed operating system cluster with Ubuntu nodes

On Ubuntu nodes in a mixed operating system cluster, upgrades by using the installation toolkit are not supported. You can upgrade a mixed operating system cluster with Ubuntu nodes with the installation toolkit in two hops. You can do these hops in any order.

1. Upgrade the non-Ubuntu nodes as follows.
 - a. Extract the installation image on the non-Ubuntu installer node.
 - b. Run the config populate command.
 - c. Exclude all Ubuntu nodes from the upgrade configuration.
 - d. Run the upgrade command.
2. Upgrade the Ubuntu nodes as follows.
 - a. Extract the installation image on the Ubuntu installer node.
 - b. Run the config populate command.
 - c. Exclude all non-Ubuntu nodes from the upgrade configuration.
 - d. Run the upgrade command.

Upgrading mixed CPU architecture cluster

You can upgrade a mixed CPU architecture cluster with the installation toolkit in two hops. You can do these hops in any order.

1. Upgrade the first set of nodes that have the same CPU architecture as follows.
 - a. Extract the installation image on the installer node of the same CPU architecture.
 - b. Run the config populate command.

The cluster configuration is populated for the nodes that have the same CPU architecture as the installer node.

- c. Run the upgrade command.
2. Upgrade the second set of nodes that have the other type of CPU architecture as follows.
 - a. Extract the installation image on the installer node for the other type of CPU architecture.
 - b. Run the config populate command.

The cluster configuration is populated for the nodes that have the same CPU architecture as the installer node.

- c. Run the upgrade command.

Offline upgrade with the installation toolkit when OS upgrade is required

If you need to upgrade the OS during an offline upgrade of IBM Storage Scale with the installation toolkit, use the following high-level procedure as guidance.

1. Run the config populate command.
2. Stop the services and suspend the protocol nodes.
3. Do the OS upgrade.
4. Download and extract the IBM Storage Scale installation package of the version that you want to upgrade to.
5. Copy the newly generated cluster definition file to the extracted installation package location.
6. Designate the nodes that you want to upgrade as offline in the upgrade configuration.
7. Run the upgrade pre-check procedure.
8. Run the upgrade procedure.
9. Start the services and resume the protocol nodes.

Populating cluster configuration when nodes are designated as offline in the upgrade configuration

You cannot use the config populate function if one or more nodes in the cluster are designated as offline in the upgrade configuration. Use the following steps to populate the cluster configuration in scenarios in which you plan to designate one or more nodes as offline in the upgrade configuration.

1. Extract the installation image that you want to use for doing the upgrade.
2. Use the **./spectrumscale config populate** command or create a cluster definition file by using the **./spectrumscale** command.
3. Shut down one or more nodes.
4. Use the **./spectrumscale upgrade config** command to designate the nodes as offline in the upgrade configuration.

Limitations of excluding nodes in the upgrade configuration

You cannot exclude a node if SMB service is running on it. This restriction is not applicable for NFS. For example, when you try to exclude from the upgrade configuration the vm1 node that is running SMB, an error occurs.

```
./spectrumscale upgrade config exclude -N vm1
```

```
[ FATAL ] In order to exclude a protocol node running SMB from the current upgrade, the SMB
service must first be stopped on that node.
Please stop SMB using the mmces command and retry.
```


Related concepts

[Upgrade process flow](#)

[Upgrade rerun after an upgrade failure](#)

The installation toolkit supports upgrade rerun if an upgrade fails.

Related tasks

[Performing online upgrade by using the installation toolkit](#)

Use the following procedure to do an online upgrade of IBM Storage Scale by using the installation toolkit.

[Upgrading Swift Object packages to IBM Storage Scale 5.1.8 or earlier by using the installation toolkit](#)

[Upgrading IBM Storage Scale on IBM Storage Archive Enterprise Edition \(EE\) nodes by using the installation toolkit](#)

Upgrade rerun after an upgrade failure

The installation toolkit supports upgrade rerun if an upgrade fails.

Do an upgrade rerun if the upgrade fails due to the following reasons:

- A component stops during the upgrade process and it is in an unhealthy state
- Network issue during the upgrade
- Mixed mode packages

You can do a rerun by using the regular upgrade (**./spectrumscale upgrade run**) command. The installation toolkit automatically identifies whether this command is run for a fresh upgrade or for a rerun. During a rerun, the installation toolkit does an upgrade by ignoring all unrecoverable errors that you received during the upgrade.

When an upgrade rerun is done:

- The installation toolkit tries to do the upgrade even if some components become unhealthy.
- The installation toolkit displays warnings instead of unrecoverable messages if some components are in a healthy state on a node.
- The installation toolkit does not try to stop the services if they are already stopped. If the services are already running, then the installation toolkit does not try to start the services again.
- The installation toolkit tries to start the services if they are stopped.
- The installation toolkit does not try to start or stop services on nodes where the services are designated as offline and on the nodes that are designated as excluded.

Related concepts

[Upgrade process flow](#)

[Performing offline upgrade or excluding nodes from upgrade by using installation toolkit](#)

Related tasks

[Performing online upgrade by using the installation toolkit](#)

Use the following procedure to do an online upgrade of IBM Storage Scale by using the installation toolkit.

[Upgrading Swift Object packages to IBM Storage Scale 5.1.8 or earlier by using the installation toolkit](#)

[Upgrading IBM Storage Scale on IBM Storage Archive Enterprise Edition \(EE\) nodes by using the installation toolkit](#)

Upgrading Swift Object packages to IBM Storage Scale 5.1.8 or earlier by using the installation toolkit

Important:

- CES Swift Object protocol feature is not supported from IBM Storage Scale 5.2.0 onwards.
- IBM Storage Scale 5.1.8 is the last release that has CES Swift Object protocol.
- IBM Storage Scale 5.2.0 will tolerate the update of a CES node from IBM Storage Scale 5.1.8.

– *Tolerate* means:

- The CES node will be updated to 5.2.0.
- Swift Object support will not be updated as part of the 5.2.0 update.
- You may continue to use the version of Swift Object protocol that was provided in IBM Storage Scale 5.1.8 on the CES 5.2.0 node.
- IBM will provide usage and known defect support for the version of Swift Object that was provided in IBM Storage Scale 5.1.8 until you migrate to a supported object solution that IBM Storage Scale provides.
- CES Swift Object is replaced with IBM Storage Scale S3. For more details, refer the *S3 support overview* section the *IBM Storage Scale: Concepts, Planning, and Installation Guide*.
- For more information about Swift Object in IBM Storage Scale, refer to the [IBM Storage Scale 5.2.0 documentation](#).

- Contact IBM for further details and migration planning.

Related concepts

[Upgrade process flow](#)

[Performing offline upgrade or excluding nodes from upgrade by using installation toolkit](#)

[Upgrade rerun after an upgrade failure](#)

The installation toolkit supports upgrade rerun if an upgrade fails.

Related tasks

[Performing online upgrade by using the installation toolkit](#)

Use the following procedure to do an online upgrade of IBM Storage Scale by using the installation toolkit.

[Upgrading IBM Storage Scale on IBM Storage Archive Enterprise Edition \(EE\) nodes by using the installation toolkit](#)

Upgrading IBM Storage Scale on IBM Storage Archive Enterprise Edition (EE) nodes by using the installation toolkit

On an IBM Storage Archive Enterprise Edition (EE) node, if you want to upgrade IBM Storage Scale, you need to do some steps before and after the upgrade to manage the IBM Spectrum Archive related components. Upgrade IBM Storage Scale on IBM Storage Archive Enterprise Edition (EE) nodes by using the installation toolkit.

Before you begin

Note: For more information about IBM Storage Archive Enterprise Edition (EE) commands, see [IBM Storage Archive Enterprise Edition \(EE\) documentation](#).

Procedure

1. Do the following steps that must be carried out on IBM Storage Archive Enterprise Edition (EE) components before an IBM Storage Scale upgrade.
 - a) Stop IBM Storage Archive Enterprise Edition (EE) by issuing the following command from one of the IBM Storage Archive Enterprise Edition (EE) nodes.

```
eeadm cluster stop
```

- b) Deactivate failover operations by issuing the following command on all IBM Storage Archive Enterprise Edition (EE) nodes.

```
dsmmigfs disablefailover
```

- c) Stop the IBM Storage Protect for Space Management daemons by issuing the following command on all IBM Storage Archive Enterprise Edition (EE) nodes.

```
dsmmigfs stop
```

- d) Stop the IBM Storage Protect for Space Management service by issuing the following command on all IBM Storage Archive Enterprise Edition (EE) nodes.

```
systemctl stop hsm.service
```

After these steps, proceed with downloading, extracting, and doing other steps before an installation toolkit upgrade precheck. For more information, see [“Upgrading IBM Storage Scale components with the installation toolkit”](#) on page 723.

2. Set the environment variable to bypass the installation toolkit RPM dependency check.

```
export SSFEATURE_OVERRIDE_EXT_PKG_DEPS=true
```

3. Run the installation toolkit upgrade precheck by issuing the following commands.

```
cd /usr/lpp/mmfs/6.0.x.x/ansible-toolkit  
./spectrumscale upgrade precheck
```

Where 6.0.x.x is the release-specific directory.

This precheck fails and lists all RPMs that are failing the dependency check. Ensure that the RPMs listed are what is expected when IBM Storage Archive Enterprise Edition (EE) is installed.

4. Rerun the installation toolkit upgrade precheck by issuing the following command.

```
./spectrumscale upgrade precheck
```

This precheck passes if no other problems are detected.

5. Run the installation toolkit upgrade by issuing the following command.

```
./spectrumscale upgrade run
```

6. Upgrade IBM Storage Archive Enterprise Edition (EE), if needed.
7. Do the following steps that must be carried out on IBM Storage Archive Enterprise Edition (EE) components after a successful IBM Storage Scale upgrade.
 - a) Start the IBM Storage Protect for Space Management service by issuing the following command on all IBM Storage Archive Enterprise Edition (EE) nodes.

```
systemctl start hsm.service
```

- b) Start the IBM Storage Protect for Space Management daemons by issuing the following command on all IBM Storage Archive Enterprise Edition (EE) nodes.

```
dsmmigfs start
```

- c) Activate failover operations by issuing the following command on all IBM Storage Archive Enterprise Edition (EE) nodes.

```
dsmmigfs enablefailover
```

- d) Start IBM Storage Archive Enterprise Edition (EE) by issuing the following command from one of the IBM Storage Archive Enterprise Edition (EE) nodes.

```
eeadm cluster start
```

What to do next

After these steps, you must complete the upgrade to a new code level to take advantage of the new functions. For more information, see [“Completing the upgrade to a new level of IBM Storage Scale”](#) on page 753.

Related concepts

[Upgrade process flow](#)

[Performing offline upgrade or excluding nodes from upgrade by using installation toolkit](#)

[Upgrade rerun after an upgrade failure](#)

The installation toolkit supports upgrade rerun if an upgrade fails.

Related tasks

[Performing online upgrade by using the installation toolkit](#)

Use the following procedure to do an online upgrade of IBM Storage Scale by using the installation toolkit.

[Upgrading Swift Object packages to IBM Storage Scale 5.1.8 or earlier by using the installation toolkit](#)

Upgrading IBM Storage Scale native REST API

You can upgrade IBM Storage Scale native REST API manually or by using the installation toolkit.

Prerequisite

- Root access must be available on all cluster nodes to perform upgrade activities.
- When using the installation toolkit, passwordless SSH must be configured across all nodes.
- IBM Storage Scale native REST API does not support the mmdsh function for general use. Use tools such as SSH or SCP to perform remote operations across nodes.

Note: The IBM Storage Scale native REST API feature is upgraded as part of the IBM Storage Scale upgrade process.

Manual upgrade

Perform the IBM Storage Scale upgrade on each node by following the documented manual upgrade procedure. Ensure that the IBM Storage Scale native REST API packages are upgraded as part of the overall upgrade. For more information, see [“Upgrading IBM Storage Scale protocol nodes”](#) on page 682 or [“Upgrading IBM Storage Scale non-protocol nodes that run on Linux”](#) on page 677.

Upgrade using installation toolkit

The installation toolkit detects and upgrades the IBM Storage Scale native REST API feature as part of the IBM Storage Scale upgrade. Make yourself familiar with the installation toolkit procedure. For more information, see [“Overview of the installation toolkit”](#) on page 509 and [“Preparing to use the installation toolkit”](#) on page 528.

Procedure

1. Ensure that the root passwordless SSH is enabled on all relevant nodes.
2. Perform the IBM Storage Scale upgrade procedure by using the installation toolkit. For more information, see [“Upgrading IBM Storage Scale components with the installation toolkit”](#) on page 723.
3. Disable root and passwordless SSH access after the upgrade is complete if it is no longer required.

Protocol authentication configuration changes during upgrade

During IBM Storage Scale protocol nodes upgrade, do protocol authentication-related configuration depending on your authentication setup.

- [“Identify the current authentication scheme configured for file protocols”](#) on page 745
- [“Upgrade authentication for file protocols set up with the LDAP-based authentication scheme”](#) on page 745
- [“Upgrade authentication for file protocols set up with the LDAP-based authentication scheme and communication with the LDAP server secured by TLS”](#) on page 746
- [“Upgrade authentication for file protocols set up with AD or LDAP file authentication with Kerberos”](#) on page 746

- [“Resolve file protocols authentication scheme configuration command failure” on page 747](#)

Identify the current authentication scheme configured for file protocols

To identify the current authentication scheme that is configured for file protocols, issue the following command.

```
# mmuserauth service list --data-access-method file
```

You identify authentication scheme that is configured for file protocols with the value of the field `FILE protocols` is configured for in the command output.

Upgrade authentication for file protocols set up with the LDAP-based authentication scheme

If file protocols are set up with the LDAP-based authentication scheme, complete the following steps. These steps are applicable for any variation of the LDAP-based authentication scheme.

Note: Issue the following commands when upgrade steps for the first protocol node in the cluster are being done. Do not repeat the following steps for the remaining protocol nodes.

1. Install the `sssd-tools` package on all the protocol nodes.
2. Obscure the password that is stored in the SSSD configuration file by doing the following steps.
 - a. Copy the current SSSD configuration file to the `/tmp` path.

```
# /bin/cp /etc/sss/sss.conf /tmp/sss_update.conf
```

- b. Store the secret of the LDAP user that is used to integrate with the LDAP server in the current session.

```
# secret="$(/usr/lpp/mmfs/bin/mmgetconfdata -f /tmp/sss_update.conf -s "domain/LDAPDOMAIN" -a "ldap_default_authtok")"
```

- c. From IBM Storage Scale 5.1.0 and higher, obscure the secret of the user that is used to communicate with the LDAP server by issuing the following command.

```
# echo -n "${secret}" | /usr/sbin/sss_obfuscate -d LDAPDOMAIN -f /tmp/sss_update.conf --stdin
```

- d. Clear the secret from the current session by issuing the following command.

```
# unset secret
```

- e. Publish the updated file to the CCR by issuing the following command.

```
# /usr/lpp/mmfs/bin/mmccr fput SSSD_CONF /tmp/sss_update.conf
```

- f. Delete the SSSD configuration file from the `/tmp` path by issuing the following command.

```
# /bin/rm /tmp/sss_update.conf
```

3. On each protocol node, publish the updated SSSD configuration file as follows.
 - a. Publish the new SSSD configuration file from IBM Storage Scale configuration repository to the node.

```
# /usr/lpp/mmfs/bin/mmccr fget SSSD_CONF /etc/sss/sss.conf
```

- b. Restart the SSSD service to reflect the change.

```
# systemctl restart sssd
```

4. Validate that the users from the LDAP server can be successfully resolved on all the protocol nodes.

Upgrade authentication for file protocols set up with the LDAP-based authentication scheme and communication with the LDAP server secured by TLS

If file protocols are set up with the LDAP-based authentication scheme and the communication with the LDAP server is secured by using the TLS protocol, do the following steps. These steps are applicable for any variation of the LDAP-based authentication scheme that is secured by TLS.

Note:

- You can confirm that you are using LDAP secured with TLS, if `ENABLE_SERVER_TLS` is `true` in the output of the **mmuserauth service list** command.
- Issue the following commands when upgrade steps for the first protocol node in the cluster are being done. Do not repeat the following steps for the remaining protocol nodes.

1. Fetch the configuration file from the CCR.

```
#/usr/lpp/mmfs/bin/mmccr fget LDAP_CONF /tmp/ldap_conf.from.ccr
```

2. Update the temporary file for the following changes.

- a. Delete the entry for the `TLS_CIPHER_SUITE` setting from the file.
- b. Add the following new entry to the file.

```
TLS_PROTOCOL_MIN 3.3
```

- c. Add the following new entry to the file based on an operating system (OS) of the protocol node.

If the OS of the protocol node is RHEL or SLES, add the following entry.

```
TLS_CIPHER_SUITE DEFAULT:!SSLv3:!TLSv1:!TLSv1.1:@STRENGTH
```

If the OS of the protocol node is Ubuntu, add the following entry.

```
TLS_CIPHER_SUITE NORMAL:-VERS-ALL:+VERS-TLS1.3:+VERS-TLS1.2:-AES-128-CBC:-AES-256-CBC
```

- d. Save the file.

3. Publish the updated file to the CCR by issuing the following command.

```
# /usr/lpp/mmfs/bin/mmccr fput LDAP_CONF /tmp/ldap_conf.from.ccr
```

4. Propagate the changes to all the protocol nodes of the cluster by issuing the following command.

```
# mmuserauth service check --data-access-method file --rectify
```

5. Delete the temporary file.

```
# /bin/rm /tmp/ldap_conf.from.ccr
```

Upgrade authentication for file protocols set up with AD or LDAP file authentication with Kerberos

Important: This section applies only if you are upgrading from a version earlier than IBM Storage Scale 5.1.2.

From IBM Storage Scale 5.1.2, the system file `/etc/krb5.keytab` is no longer for Kerberos-based file authentication. IBM Storage Scale uses a custom keytab file instead. If you are upgrading from version 5.1.2 or later, skip this step.

If you are using AD or LDAP file authentication with Kerberos, do the following steps to upgrade the authentication configuration.

Note: You can confirm that you are using AD or LDAP file authentication with Kerberos, if either `ENABLE_NFS_KERBEROS` or `ENABLE_KERBEROS` is `true` in the output of the **mmuserauth service list** command.

1. Complete the upgrade steps on protocol nodes by using the installation toolkit or the manual upgrade procedure.
2. Run the `/usr/lpp/mmfs/bin/mmktupgrade` script.

The script creates a custom keytab file `/var/mmfs/etc/krb5_scale.keytab` and does the necessary NFS and SMB configuration changes.

3. Restart NFS and SMB on protocol nodes for the configuration changes to take effect.

Note: To prevent downtime, you can restart services on these nodes in a phased manner.

After these steps, IBM Storage Scale uses the custom keytab file for Kerberos authentication.

Resolve file protocols authentication scheme configuration command failure

The file protocols authentication scheme configuration command might fail with the following error.

```
mmuserauth: [E] CCR command failed: service
Failed to read value for variable LDAPMAP_DOMAINS from CCR.
mmuserauth service list: Command failed. Examine previous error messages to determine cause.
```

If this error occurs, set the `LDAPMAP_DOMAINS` variable to none in the cluster configuration repository (CCR) by doing the following steps from any protocol node.

1. Edit the `/tmp/authccr` file and set `LDAPMAP_DOMAINS=none`. If the `LDAPMAP_DOMAINS` entry is not present, add the entry and set it to none.
2. Issue this command.

```
usr/lpp/mmfs/bin/mmccr fget authccr /tmp/authccr
```



Warning: The `mmccr` command is an IBM Storage Scale internal component and it must be used only under the guidance of IBM Support. If the `mmccr` command is used incorrectly, the cluster can become nonoperational.

Changing the IBM Storage Scale product edition

You can change the IBM Storage Scale product edition from your currently installed product edition either by using the installation toolkit or manually.

Before you begin

The available product edition change paths¹ are listed next.

1. IBM Storage Scale Standard Edition to IBM Storage Scale Data Access Edition
2. IBM Storage Scale Standard Edition to IBM Storage Scale Advanced Edition
3. IBM Storage Scale Standard Edition to IBM Storage Scale Data Management Edition
4. IBM Storage Scale Advanced Edition to IBM Storage Scale Data Management Edition
5. IBM Storage Scale Data Access Edition to IBM Storage Scale Data Management Edition

The methods that can be used to change the product edition are listed next.

1. Installation toolkit (cluster online)²
2. Manual node by node (cluster online)³
3. Manual all nodes (cluster offline)

Notes for paths and methods:

1. These paths are the only edition changes that are supported.
2. If you are planning to use the installation toolkit for changing the product edition, be aware that the product edition change by using the installation toolkit can be performed only when you are upgrading

to a newer version of IBM Storage Scale. If you do not want to upgrade to a newer version of IBM Storage Scale, you must change the product edition manually.

3. You must change the whole cluster to the same product edition. If you are using the manual node by node method, make sure that the product edition is changed on each node in the cluster.

About this task

You can use one of the following sets of steps to change the IBM Storage Scale product edition.

Important: If the cluster where you plan to make a supported edition change has an interim fix (efix) applied to its IBM Storage Scale version, you must either upgrade to a version with no interim fix applied, or contact IBM Storage Scale support team.

Changing the product edition by using the installation toolkit

You can change the IBM Storage Scale product edition from your currently installed product edition by using the installation toolkit.

About this task

Use the following steps for changing the IBM Storage Scale product edition by using the installation toolkit.

Important: If the cluster where you plan to make a supported edition change has an interim fix (efix) applied to its IBM Storage Scale version, you must either upgrade to a version with no interim fix applied, or contact IBM Storage Scale support team.

Procedure

1. Download from IBM Support Fix Central the installation image of the IBM Storage Scale edition that you are planning to change to.
2. Extract the installation image.

For performing an upgrade by using the installation toolkit, a key requirement is the cluster definition file.

3. Use the **./spectrumscale config populate** command to traverse the existing cluster and populate the cluster definition file.

For more information about obtaining the current cluster configuration, see Step 4 of [“Performing online upgrade by using the installation toolkit”](#) on page 731.

4. Run the upgrade precheck procedure of the installation toolkit.

```
cd /usr/lpp/mmfs/x.x.x.x/ansible-toolkit
./spectrumscale upgrade precheck
```

Note: Replace x.x.x.x with the IBM Storage Scale version.

5. Run the upgrade procedure of the installation toolkit.

```
cd /usr/lpp/mmfs/x.x.x.x/ansible-toolkit
./spectrumscale upgrade run
```

Results

With these steps, the product edition is changed, and the applicable packages are installed on all nodes that are specified in the cluster definition file.

Changing the product edition manually, one node at a time

You can manually change the IBM Storage Scale product edition from your currently installed product edition, one node at a time.

About this task

Use the following steps to change the IBM Storage Scale product edition manually, one node at a time.

Important: If the cluster where you plan to make a supported edition change has an interim fix (efix) applied to its IBM Storage Scale version, you must either upgrade to a version with no interim fix applied, or contact IBM Storage Scale support team.

Procedure

Tip: In the following steps, **rpm** and **dpkg** commands are used for installing and uninstalling packages. If the respective repository is set up in your environment, you can also use package manager commands such as **yum**, **zypper**, and **apt-get** for installing and uninstalling packages.

Remember: According to the edition that you are changing to, replace *License_Package_Name* and `gpfs.license.x` in the following procedure. Use the next table as reference.

Edition	License package name
IBM Storage Scale Data Access Edition	<code>gpfs.license.da*</code>
IBM Storage Scale Data Management Edition	<code>gpfs.license.dm*</code>
IBM Storage Scale Advanced Edition	<code>gpfs.license.adv*</code>

1. Download from IBM Support Fix Central the installation image of the IBM Storage Scale edition that you are planning to change to.
2. Extract the installation image.

```
mmumount all
```

3. Unmount all file systems on the node.
4. Verify that file systems are unmounted.

```
mmismount all -L -N NodeName
```

5. Shut down GPFS on the node.

```
mmshutdown
```

6. Uninstall the current license package by using one of the following commands depending on your product edition change path and the operating system.

RHEL and SLES:

```
rpm -e License_Package_Name
```

You can obtain the license package name by using the **rpm -qa | grep gpfs.license** command.

Ubuntu:

```
dpkg -P License_Package_Name
```

You can obtain the license package name by using the **dpkg -l | grep gpfs.license** command.

7. Install the new license package by using one of the following sets of commands depending on your product edition change path and the operating system.

RHEL and SLES:

```
cd /usr/lpp/mmfs/x.x.x.x/gpfs_rpms
rpm -ivh gpfs.license.x*.rpm
```

Ubuntu:

```
cd /usr/lpp/mmfs/x.x.x.x/gpfs_debs
dpkg -i gpfs.license.x*.deb
```

8. Start GPFS on the node.

```
mmstartup
```

9. Repeat the previous steps on all nodes in the cluster.

Changing the product edition manually, all nodes at the same time

You can manually change the IBM Storage Scale product edition from your currently installed product edition, all nodes at the same time.

About this task

Use the following steps to change the IBM Storage Scale product edition manually, all nodes at the same time.

Important: If the cluster where you plan to make a supported edition change has an interim fix (efix) applied to its IBM Storage Scale version, you must either upgrade to a version with no interim fix applied, or contact IBM Storage Scale support team.

Procedure

Tip: In the following steps, **rpm** and **dpkg** commands are used for installing and uninstalling packages. If the respective repository is set up in your environment, you can also use package manager commands such as **yum**, **zypper**, and **apt-get** for installing and uninstalling packages.

Remember: According to the edition that you are changing to, replace *License_Package_Name* and *gpfs.license.x* in the following procedure. Use the next table as reference.

Edition	License package name
IBM Storage Scale Data Access Edition	gpfs.license.da*
IBM Storage Scale Data Management Edition	gpfs.license.dm*
IBM Storage Scale Advanced Edition	gpfs.license.adv*

1. Download from IBM Support Fix Central the installation image of the IBM Storage Scale edition that you are planning to change to.
2. Extract the installation image.
3. Unmount all file systems.

```
mmumount all -a
```

4. Verify that file systems are unmounted.

```
mmismount all -L
```

5. Shut down GPFS on all nodes.

```
mmshutdown -a
```

6. Verify that GPFS was shut down on all nodes.

```
mmgetstate -a
```

7. Uninstall the current license package by using one of the following commands, depending on your product edition change path and the operating system.

RHEL and SLES:

```
mmdsh -N all rpm -e License_Package_Name
```

You can obtain the license package name by using the **rpm -qa | grep gpfs.license** command.

Ubuntu:

```
mmdsh -N all dpkg -P License_Package_Name
```

You can obtain the license package name by using the **dpkg -l | grep gpfs.license** command.

8. Copy `gpfs.adv`, `gpfs.crypto`, and `gpfs.license.x` to all the other nodes in the cluster.

Make sure to create the *same directory structure* in the nodes, so that you can copy these files to the same path within each node.

9. Install the `gpfs.adv`, `gpfs.crypto`, and new license packages.

- a) Use one of the following sets of commands, depending on your product edition change path and the operating system.

RHEL and SLES:

```
mmdsh -N all rpm -ivh /usr/lpp/mmfs/x.x.x.x/gpfs_rpms/gpfs.adv*.rpm
mmdsh -N all rpm -ivh /usr/lpp/mmfs/x.x.x.x/gpfs_rpms/gpfs.crypto*.rpm
```

Ubuntu:

```
mmdsh -N all dpkg -i /usr/lpp/mmfs/x.x.x.x/gpfs_debs/gpfs.adv*.deb
mmdsh -N all dpkg -i /usr/lpp/mmfs/x.x.x.x/gpfs_debs/gpfs.crypto*.deb
```

- b) Verify that all the nodes have a copy of the `gpfs.adv`, `gpfs.crypto`, and `gpfs.license.x` packages.

In RHEL and SLES, use the **rpm -qa | grep gpfs.license** command.

In Ubuntu, use the **dpkg -l | grep gpfs.license** command.

10. Install the new license package.

Use one of the following commands, depending on your product edition change path and the operating system.

RHEL and SLES:

```
mmdsh -N all rpm -ivh /usr/lpp/mmfs/x.x.x.x/gpfs_rpms/gpfs.license.x*.rpm
```

Ubuntu:

```
mmdsh -N all dpkg -i /usr/lpp/mmfs/x.x.x.x/gpfs_debs/gpfs.license.x*.deb
```

11. Start GPFS on all nodes.

```
mmstartup -a
```

12. Make sure that all the nodes are online.

```
mmgetsate -a
```

13. Verify that the edition change was successful.

```
mmllslicense -L
```

14. Check that all your file systems are mounted and running.

```
mmlsmount all -L
```

Changing Express Edition to Standard Edition

Starting with release 4.2.3, IBM Storage Scale Express Edition is no longer available. Use this information for migrating from IBM Storage Scale Express Edition 4.2.2.x or earlier to IBM Storage Scale Standard Edition 4.2.3.x or later.

Before you begin

Before you begin, extract the IBM Storage Scale software as described in the [“Extracting the IBM Storage Scale software on Linux nodes”](#) on page 473 topic. For more information about the location of extracted installation images, see [“Location of extracted packages”](#) on page 480.

About this task

Use one of the following sets of steps depending on the operating system and the IBM Storage Scale version.

Procedure

Red Hat Enterprise Linux and SLES

- To migrate from 4.1.1.x or 4.2.0.x or 4.2.1.x Express Edition to 4.2.3 Standard Edition or later on Red Hat Enterprise Linux and SLES, update the existing RPMs and install new RPMs by issuing this command.

```
rpm -Uvh gpfs*rpm
```

- To migrate from 4.2.2.x Express Edition to 4.2.3 Standard Edition or later on Red Hat Enterprise Linux and SLES, use these steps.
 - a) Uninstall the Express Edition license RPM by issuing the following command.

```
rpm -e gpfs.license.exp*.rpm
```

- b) Update the existing RPMs and install new RPMs by issuing the following command.

```
rpm -Uvh gpfs*rpm
rpm -ivh gpfs.compression*.rpm
rpm -ivh gpfs.license.std*.rpm
```

Debian and Ubuntu

- To migrate from 4.1.1.x or 4.2.0.x or 4.2.1.x Express Edition to 4.2.3 Standard Edition or later on Debian and Ubuntu, update the existing packages and install new packages by issuing the following command.

```
dpkg -i gpfs*deb
```

- To migrate from 4.2.2.x Express Edition to the 4.2.3 Standard Edition or later on Debian and Ubuntu, use these steps.
 - a) Uninstall the Express Edition license package by issuing the following command.

```
dpkg -r gpfs.license.exp*deb
```

- b) Update the existing packages and install new packages by issuing the following command.

```
dpkg -i gpfs*deb
```

AIX

- Depending on whether you want to be able to revert to the previous version without reinstalling or uninstalling or not, do one of the following sets of steps.
- If you want to be able to revert to previous version without reinstalling or uninstalling, do one of the following sets of steps:
 - To migrate from 4.2.0.x or 4.2.1.x Express Edition to 4.2.3 Standard Edition or later on AIX, use these steps.
 1. Update existing packages to 4.2.3 level or a later version by using the PTF images on IBM Fix Central.
 2. Install additional packages from CD images.
 - To migrate from 4.2.2.x Express Edition to 4.2.3 Standard Edition on AIX, use these steps.
 1. Update existing packages to 4.2.3 level or a later version by using the PTF images on IBM Fix Central.
 2. Install more packages from CD images.
 3. After a successful update, uninstall the `gpfs.license.exp` package.

To revert to the previous version, do the following steps:

1. Uninstall all additional packages installed from CD images.
 2. Reject the PTF images.
 3. If you are reverting to 4.2.2.x and you had uninstalled the `gpfs.license.exp` package for migrating to Standard Edition, install the `gpfs.license.exp` package.
- If you can reinstall or uninstall to revert to the previous version, do one of the following sets of steps:
 - To migrate from 4.2.0.x or 4.2.1.x Express Edition to 4.2.3 Standard Edition or later on AIX, install 4.2.3 packages from CD images.
 - To migrate from 4.2.2.x Express Edition to 4.2.3 Standard Edition or later on AIX, use these steps.
 1. Install 4.2.3 or a later version packages from CD images.
 2. After a successful update, uninstall the `gpfs.license.exp` package.

Windows

- To migrate from 4.2.2.x or earlier Express Edition to Standard Edition 4.2.3 or later on Windows, use these steps.
 - a) Uninstall the current IBM Storage Scale packages (For example, `gpfs.base-4.x.x.0-Windows.msi`, `gpfs.gskit-8.0.x.x.msi`, and `gpfs.x.license.msi`).
 - b) Restart the system.
 - c) Install `gpfs.ext-4.x.x.-Windows-license.msi`, `gpfs.ext-4.x.x.0-Windows.msi`, and `gpfs.gskit-8.0.x.x.msi` packages.

Completing the upgrade to a new level of IBM Storage Scale

It is a good idea to use the cluster for a while with the new level of IBM Storage Scale installed, until you are sure that you are ready to permanently upgrade the cluster to the new level.

Before you begin

When you are ready to permanently upgrade the cluster, follow the steps in this topic to complete the upgrade. If you decide not to complete the upgrade, you can revert to the previous level of IBM Storage Scale. For more information, see [“Reverting to the previous level of IBM Storage Scale” on page 758](#).

Before you begin this task, verify that you have upgraded all the nodes in the cluster to the latest licensed version of IBM Storage Scale.

About this task

When you run `mmchconfig release=LATEST` in Step 2 of these directions, you can add other parameters and their values to the command line:

```
mmchconfig release=LATEST,<parameterN=value>,<parameterN+1=value>...
```

The following table describes the options and parameters that are referred to in this topic.

Table 80. Other options and parameters with <code>mmchconfig release=LATEST</code>		
Option/Parameter	Purpose	Comment
<code>--accept-empty-cipherlist-security</code>	Specify this option if you want to continue running the cluster with the lowest level of security for communications between nodes or with other clusters. That is, you want the <code>cipherList</code> attribute to remain set to <code>EMPTY</code> or undefined. For more information, see the topic <i>Security mode</i> in the <i>IBM Storage Scale: Administration Guide</i>	It is a good idea to have some level of security for cluster communications: - Set <code>cipherList</code> to <code>AUTHONLY</code> or to a supported cipher: <pre>mmchconfig cipherList=AUTHONLY</pre> - Do not include the parameter <code>--accept-empty-cipherlist-security</code> when you run <code>mmchconfig release=LATEST</code>. For more information, see the topic <i>Security mode</i> in the <i>IBM Storage Scale: Administration Guide</i>
<code>--accept-no-compliance-to-nist-standards</code>	Specify this option if you do not want to continue running the cluster with the security transport that follows the NIST SP800-131A recommendations. For more information, see the topic <i>NIST compliance</i> in the <i>IBM Storage Scale: Administration Guide</i>	It is recommended to run the security transport that follows the NIST SP800-131A recommendations: - Set <code>nistCompliance</code> attribute to <code>SP800-131A</code> <pre>mmchconfig nistCompliance=SP800-131A</pre> - Set <code>nistCompliance</code> attribute to <code>SP800-131A</code> at the same time when you upgrade IBM Storage Scale to a new level. <pre>mmchconfig release=LATEST,nistCompliance=SP800-131A</pre>

Table 80. Other options and parameters with mmchconfig release=LATEST (continued)		
Option/Parameter	Purpose	Comment
mmfsLogTimeStampISO8601={yes no}	Setting this parameter to no allows the cluster to continue running with the earlier log time stamp format. For more information, see the topic <i>Security mode</i> in the <i>IBM Storage Scale: Administration Guide</i>	- Set mmfsLogTimeStampISO8601 to no if you save log information and you are not yet ready to switch to the new log time stamp format. After you complete the upgrade, you can change the log time stamp format at any time with the mmchconfig command. - Omit this parameter if you are ready to switch to the new format. The default value is yes
tscCmdAllowRemoteConnections	Setting this attribute to no forces the ts* commands to communicate with the local mmfsd daemon only over a UNIX domain socket (UDS), and to communicate with mmfsd daemons running on other nodes over the RPC communication framework used by IBM Storage Scale.	After a cluster is upgraded, the tscCmdAllowRemoteConnections parameter value is set to yes by default. Setting the tscCmdAllowRemoteConnections configuration to no requires a cluster minimum release level of 5.1.3 or later. System administrators can set the tscCmdAllowRemoteConnections value to no after setting the minReleaseLevel=LATEST, which is 5.1.3 or later. Setting tscCmdAllowRemoteConnections to no in a multi-cluster setting must be done after ensuring that the remote clusters are running IBM Storage Scale version 5.1.3 or later. For more information, see <i>mmchconfig command</i> in <i>IBM Storage Scale: Command and Programming Reference Guide</i>.

Note: It is not recommended to use the --accept-no-compliance-to-nist-standards option and this option might not be available in the subsequent releases.

Procedure

1. Verify that the SHA message digest and the cipherList configuration variable are set to valid values.

Note:

- The SHA message digest is a hash result that is generated by a cryptographic hash function.
- The cipherList variable specifies the security mode for communications among nodes in the cluster and with nodes in other clusters.

Follow these steps:

- a) Display the current values by entering the following command. The listing shows both the command and example output:

```
# mmauth show .
```

```
Cluster name: zounds.cluster (this cluster) Cipher list: (none specified)
SHA digest: (undefined) File system access: (all rw)
```

- b) If the value for the SHA digest is (undefined), follow these steps:
 - i) Enter the following command to generate a public/private key pair and an SHA message digest:

```
mmauth genkey new
```

- ii) Enter `mmauth show .` again and verify that the value for SHA digest is no longer (undefined).

- c) If the value for cipherList is (none specified) or EMPTY, do one of the following actions:

- If you want a level of security in communications between nodes and with other clusters, follow these steps:

- i) Set cipherList to AUTHONLY or to a supported cipher:

```
mmauth update . -l AUTHONLY
```

- ii) Enter `mmauth show .` again and verify that the value for cipherList is no longer (none specified) or EMPTY.

- If you do not want a level of security cluster communications, let cipherList remain set to (none specified) or EMPTY.

2. Enter the following command to upgrade the cluster configuration data and enable new functionality. You can add the parameters that are described in [Table 80 on page 754](#) to the command line:

```
mmchconfig release=LATEST
```

Note: Until you run the **mmchconfig release=LATEST** command, the management GUI might not be fully operational at the new code level.

Important: If the `mmchconfig` command detects any nodes that are not available or cannot be reached, it lists the names of those nodes. If any such nodes are listed, correct the problem and run the command again until it verifies all the nodes and completes successfully.

Note: If the `mmchconfig` command fails with an error message that indicates that cipherlist is set to EMPTY, do one of the following actions:

- If you want the cluster to run with a higher security mode than EMPTY, set cipherList to AUTHONLY or to a supported cipher:

```
mmauth update . -l AUTHONLY
```

Return to the first part of Step 2 and run the `mmchconfig` command as before.

- If you want the cluster to continue with the security mode set to EMPTY, return to the first part of Step 2 and run the `mmchconfig` command with the additional parameter `--accept-empty-cipherlist-security`.

3. If you have not already done so, assign an appropriate IBM Storage Scale license to each of the nodes in the cluster.

See “IBM Storage Scale license designation” on page 300 for a detailed description of the IBM Storage Scale license types. To see what the minimum required IBM Storage Scale license is for each of the nodes in the cluster, enter the following command:

```
mmllslicense -L
```

To assign an IBM Storage Scale server license to the nodes that require it, enter the following command:

```
mmchlicense server -N NodeList
```

To assign an IBM Storage Scale client license to the nodes that require it, enter:

```
mmchlicense client -N NodeList
```

4. Enable backward compatible format changes or upgrade all file systems to the latest metadata format changes.



Attention: Before you continue with this step, it is important to understand the differences between `mmchfs -V compat` and `mmchfs -V full`:

- If you enter `mmchfs -V compat`, only changes that are backward compatible with the previous file system version are enabled. Nodes in remote clusters that are running at the previous IBM Storage Scale version or later will still be able to mount the file system. Nodes running an IBM Storage Scale version earlier than the previous version will no longer be able to mount the file system. For example, if you are upgrading from IBM Storage Scale 5.0.2.x to 5.1.0.x, a file system of version corresponding to IBM Storage Scale 5.0.2.x will remain mountable from a remote cluster running IBM Storage Scale 5.0.2.x or later. However, a file system of version corresponding to IBM Storage Scale 5.0.0.0 will remain mountable from a remote cluster running IBM Storage Scale 5.0.0.0 or later.

Running in the compatibility mode might prevent enablement of new functionality that relies on disk structures available only with the latest version.

- If you enter `mmchfs -V full`, all new functions that require different on-disk data structures are enabled. Nodes in remote clusters that run an older IBM Storage Scale version will no longer be able to mount the file system. For example, remote container native clusters. If any nodes that run an older IBM Storage Scale version have mounted the file system at the time this command is entered, the `mmchfs` command fails. This consideration might also apply to minor releases within the same major release. For example, release 5.0.2.x and release 5.1.0.x might have different metadata formats.

To enable backward-compatible format changes, enter the following command:

```
mmchfs FileSystem -V compat
```

To upgrade the desired file system to the latest metadata format changes, enter the following command:

```
mmchfs FileSystem -V full
```

Certain new file system features might require more processing that cannot be handled by the `mmchfs -V` command alone. To fully activate such features, in addition to `mmchfs -V`, you must also run the `mmigratefs` command.

Note: The first mount of a file system after you run `mmchfs -V` might fail with a no-disk-space error. This situation might occur if the file system is relatively low on metadata disk space (10% or less free). If so, enter the command again. Typically the file system is mounted without a problem after the initial failed mount.

After completing the upgrade, you might need to enable cluster configuration repository (CCR) and fast extended attributes (`fastea`).

5. Enable cluster configuration repository (CCR) and fast extended attributes (fastea) as follows, if required.

- a) Enable the CCR in the cluster as follows.

```
mmchcluster --ccr-enable
```

- b) Enable fastea in the cluster as follows.

- i) Check if fastea is enabled.

```
mmfsfs FileSystemName --fastea
```

A sample output is as follows.

```
flag value description -----  
----- --fastea Yes Fast external  
attributes enabled?
```

- ii) If fastea is not enabled, issue the following command to enable it.

```
mmmigratefs FileSystemName --fastea
```

For more information, see *mmmigratefs* command in *IBM Storage Scale: Command and Programming Reference Guide*.

Note: The following features are a few among several IBM Storage Scale features that require fast extended attributes (fastea) to be enabled to work.

- Clones
- Independent filesets
- Encryption
- Active file management (AFM)

6. If you have file audit logging or clustered watch folder enabled on any file systems, follow the steps in [“Upgrade paths and commands for file audit logging and clustered watch folder” on page 722](#).

Note: Ensure that you include the `gpfs.librdkafka` package as part of the upgrade process if those packages are currently installed.

7. Set the value of the `tscCmdAllowRemoteConnections` attribute to `no`.

The `tscCmdAllowRemoteConnections` attribute specifies whether the `ts*` commands are allowed to use the remote TCP/IP connections when communicating with the local or other `mmfsd` daemons. The scope of the `tscCmdAllowRemoteConnections` attribute is local to a cluster and it has the same value on all the nodes in the cluster.



Warning: Make sure that all remote clusters must be at 5.1.3 or later. Setting the `tscCmdAllowRemoteConnections` parameter to `no` impacts the remote cluster functions, hence the recommendation is to check the version of the remote clusters. If any remote cluster is running a version of IBM Storage Scale older than 5.1.3, setting the `tscCmdAllowRemoteConnections` parameter value to `no` in the home cluster can generate `Operation not permitted` errors. For more information, see *mmchconfig* command in *IBM Storage Scale: Command and Programming Reference Guide*.

Reverting to the previous level of IBM Storage Scale

If you decide not to continue the upgrade to the latest level of GPFS, and you have not yet issued the `mmchfs -V` command, you can reinstall the earlier level of GPFS.

About this task

Important: Once a file system are upgraded explicitly by issuing the `mmchfs -V full` command, the disk images can no longer be read by a prior version of GPFS. Re-create the file system from the backup

media and restore the content if you choose to go back after this command is issued. The same rules apply for file systems that are newly created with GPFS 5.0.

You can revert to GPFS 4.2.x.

If you have performed backups with the `mmbbackup` command by using the 5.0 version and decide to revert to an earlier version, you must rebuild the `mmbbackup` shadow database by using the `mmbbackup` command with either the `-q` or `--rebuild` option.

The procedure differs depending on whether you have issued the `mmchconfig release=LATEST` command or not.

Reverting to a previous level of GPFS when you have not issued `mmchconfig release=LATEST`

If you have not issued the `mmchconfig release=LATEST` command, perform these steps.

About this task

1. Stop all user activity in the file systems.
2. Cleanly unmount all mounted GPFS file systems. Do not use force unmount.
3. Stop GPFS on all nodes in the cluster:

```
mmshutdown -a
```

4. Run the appropriate uninstallation program to remove GPFS from each node in the cluster. For example:

- For Linux nodes (this example is only applicable on the IBM Storage Scale Standard Edition):

```
rpm -e gpfs.gpl gpfs.license.std gpfs.base gpfs.docs gpfs.gskit gpfs.msg.en_US
```

- For AIX nodes:

```
installp -u gpfs
```

- For Windows nodes, open the Programs and Features control panel and remove IBM General Parallel File System.

For more information about the remaining steps, see *IBM Storage Scale* in *IBM Documentation*, and search for the appropriate *IBM Storage Scale: Concepts, Planning, and Installation Guide* for your release.

5. Copy the installation images of the previous GPFS licensed program on all affected nodes.
6. Install the original install images and all required PTFs.
7. For Linux nodes running GPFS, you must rebuild the GPFS portability layer.
8. Restart all nodes.

Related tasks

[Reverting to a previous level of GPFS when you have issued `mmchconfig release=LATEST`](#)

If you *have* issued the `mmchconfig release=LATEST` command, you must rebuild the cluster. Perform these steps.

Reverting to a previous level of GPFS when you *have* issued `mmchconfig release=LATEST`

If you *have* issued the `mmchconfig release=LATEST` command, you must rebuild the cluster. Perform these steps.

About this task

1. Stop all user activity in the file systems.
2. Cleanly unmount all mounted GPFS file systems. Do not use force unmount.
3. Stop GPFS on all nodes in the cluster:

```
mmshutdown -a
```

4. Export the GPFS file systems by issuing the `mmexportfs` command:

```
mmexportfs all -o exportDataFile
```

5. Delete the cluster:

```
mmdelnode -a
```

6. Run the appropriate uninstallation command to remove GPFS from each node in the cluster. For example:

- For Linux nodes:

```
rpm -e gpfs.gpl gpfs.license.xx gpfs.base gpfs.docs gpfs.msg.en_US
```

- For AIX nodes:

```
installp -u gpfs
```

- For Windows nodes, open the Programs and Features control panel and remove IBM General Parallel File System.

For the remaining steps, see *IBM Storage Scale* in IBM Documentation, and search for the appropriate *IBM Storage Scale: Concepts, Planning, and Installation Guide* for your release.

7. Copy the installation images of the previous GPFS licensed program on all affected nodes.
8. Install the original installation images and all required PTFs.
9. For Linux nodes running GPFS, you must rebuild the GPFS portability layer.
10. Restart all nodes.
11. Recreate your original cluster by using the `mmcrcluster` command. Run the `mmchlicense` command to set the appropriate licenses after the cluster is created.
12. Use the `mmchconfig` command to restore any previously set configuration settings that are compatible with GPFS 4.2 or below.
13. Import the file system information using the `mmimportfs` command. Specify the file created by the `mmexportfs` command from Step “4” on page 760:

```
mmimportfs all -i exportDataFile
```

14. Start GPFS on all nodes in the cluster, issue:

```
mmstartup -a
```

15. Mount the file systems if this is not done automatically when the GPFS daemon starts.

Related tasks

Reverting to a previous level of GPFS when you have not issued `mmchconfig release=LATEST`
If you have not issued the `mmchconfig release=LATEST` command, perform these steps.

Coexistence considerations

Each GPFS cluster can have multiple GPFS file systems that coexist on the cluster, but function independently of each other. In addition, each file system might have different data management programs.

Note: The GPFS Data Management API (DMAPI) and GPFS file system snapshots can coexist; however, access to the files in a snapshot using DMAPI is restricted. For more information, see *IBM Storage Scale: Command and Programming Reference Guide*.

Compatibility considerations

All applications that ran on the previous release of GPFS will run on the new level of GPFS. File systems that were created under the previous release of GPFS can be used under the new level of GPFS.

Considerations for IBM Storage Protect for Space Management

Upgrading to GPFS 3.3 or beyond requires consideration for IBM Storage Protect for Space Management. IBM Storage Protect for Space Management requires that all nodes in a cluster are configured with a DMAPI file handle size of 32 bytes.

During upgrade, it is possible to have older versions of GPFS that have a DMAPI file handle size of 16 bytes. Until all nodes in the cluster are updated to the latest release and the DMAPI file handle size is changed to 32 bytes, IBM Storage Protect for Space Management must be disabled. Run `mmfsconfig dmapifileHandleSize` to see what value is set.

After all nodes in the cluster are upgraded to the latest release and you change the DMAPI file handle size to 32 bytes, you can enable IBM Storage Protect for Space Management.

The DMAPI file handle size is configured with the **dmapifileHandleSize** option. For more information about this option, see *GPFS configuration attributes for DMAPI* in *IBM Storage Scale: Command and Programming Reference Guide*.

Applying maintenance to your IBM Storage Scale system

Before applying maintenance to your IBM Storage Scale system, there are several things you must consider.

About this task

Remember that:

1. There is limited interoperability between IBM Storage Scale 5.1.x nodes and nodes running IBM Storage Scale 5.0.y. This function is intended for short-term use. You do not get the full functions of IBM Storage Scale 5.1.x until all nodes are using IBM Storage Scale 5.1.x.
2. Interoperability between maintenance levels is specified on a per-maintenance-level basis. See the [IBM Storage Scale FAQ](#) in [IBM Documentation](#).
3. When applying maintenance, on the node being upgraded, the GPFS file systems need to be unmounted and GPFS must be shut down. Any applications by using a GPFS file system must be shut down before applying maintenance to avoid application errors or possible system halts.

Note: If the AIX Alternate Disk Install process is used, the new image is placed on a different disk. The running install image is not changed, and GPFS is not shut down. The new image is not used until the node is rebooted.

4. For the latest service information, see the [IBM Storage Scale FAQ](#) in [IBM Documentation](#).

To download fixes, go to the IBM Support Portal: Downloads for General Parallel File System (www.ibm.com/support/entry/portal/Downloads/Software/Cluster_software/General_Parallel_File_System) and obtain the fixes for your hardware and operating system.

To apply the latest fixes:

- For Linux nodes, see [Chapter 4, “Installing IBM Storage Scale on Linux nodes and deploying protocols,” on page 469.](#)
- For AIX nodes, see [Chapter 5, “Installing IBM Storage Scale on AIX nodes,” on page 575.](#)
- For Windows nodes, see [Chapter 6, “Installing IBM Storage Scale on Windows nodes,” on page 579.](#)

Guidance for upgrading the operating system on IBM Storage Scale nodes

This is a high-level guidance for upgrading the operating system (OS) on your IBM Storage Scale nodes with focus on steps to be performed from the IBM Storage Scale side.

For OS specific upgrade steps, refer to the respective OS documentation.



Warning: If upgrading from an older version of the OS to a newer version requires a reinstallation of the OS, the devices that are used for the install disk must be reviewed carefully before accepting the disk presented by the installer. This is required to avoid situations in which the installer chooses the IBM Storage Scale NSDs and causes file system corruption. Ensure that IBM Storage Scale NSDs are not altered during OS re-installation or upgrade.

Order of upgrade

Important: Before upgrading the OS on a node, ensure that the new OS version that you are planning to upgrade to is supported on the current IBM Storage Scale version. For information on the supported OS versions, see [Functional Support Matrices in IBM Storage Scale FAQ.](#)

- If the new OS version that you are planning to upgrade to is supported on the current IBM Storage Scale version, upgrading IBM Storage Scale is not mandatory. Proceed with upgrading the OS.
- If the new OS version that you are planning to upgrade to is not supported on the current IBM Storage Scale version, do the following:
 1. Determine the IBM Storage Scale version on which the new OS version that you are planning to upgrade to is supported by referring to [Functional Support Matrices in IBM Storage Scale FAQ.](#)
 2. Upgrade IBM Storage Scale.
 3. After the IBM Storage Scale upgrade is completed, upgrade to the new OS version.

Prerequisites for upgrading Linux distributions

- Download and mount the operating system installation image on each node that is being upgraded.
- Create an OS base repository that points to the installation image on each node that is being upgraded.
 - If the repository name is the same as an old OS repository, clean the repository cache.
 - Verify that the repository works.

- If required, download OFED drivers on each node that is being upgraded.
- Note the current OS and OFED settings. You might need to restore them after the OS upgrade.

It is assumed that mixed level of OS and drivers across nodes are supported by vendors.

- Note the current SELinux and firewall settings. Although an OS upgrade typically does not change these settings, it is recommended to note these settings.
- Check the file system space and ensure that /boot has adequate free space.
- If you are upgrading protocol nodes to Red Hat Enterprise Linux and Swift Object is configured then you must perform a cleanup of Swift Object services and related packages before upgrading. For more

information, see [“Instructions for removing object protocol packages when upgrading protocol nodes to Red Hat Enterprise Linux 8.x”](#) on page 767.

- If you are upgrading perfmon pmcollector nodes and the upgrade requires a reinstallation of the OS, backup and migrate the perfmon data of each pmcollector node as described in the *Migrating the pmcollector* section in the *IBM Storage Scale: Problem Determination Guide*.

OS upgrade flow

Do the following steps on each node on which the OS is being upgraded.

1. If a protocol node is being upgraded, do the following:

- a. Suspend the protocol node and stop the protocol services running on the node by issuing the following command.

```
mmces node suspend --stop
```

- b. Wait for the protocol service monitor to change to the suspended state.

mmhealth node show ces	Component	Status	Status Change	Reasons
CES		SUSPENDED	30 min. ago	-
AUTH		SUSPENDED	30 min. ago	-
AUTH_OBJ		SUSPENDED	30 min. ago	-
BLOCK		SUSPENDED	30 min. ago	-
CESNETWORK		SUSPENDED	30 min. ago	-
NFS		SUSPENDED	30 min. ago	-
SMB		SUSPENDED	30 min. ago	-

2. Disable the GPFS autoloading setting that causes GPFS to be started whenever the operating system reboots by issuing the following command.

```
mmchconfig autoloading=no -N NodeName
```

3. Unmount GPFS file systems on the node by issuing the following command.

```
mmumount all
```

4. Stop GPFS on the node by issuing the following command.

```
mmshutdown
```

5. Check if there are any residual GPFS processes running and kill them, if needed.
6. If required, uninstall the OFED driver using the bundled uninstallation script in the OFED package.
7. Perform the OS upgrade steps according to the respective documentation.
8. Reboot the node when the OS upgrade steps are completed.
9. If required, install the new OFED driver and reboot the node after OFED driver installation.
10. Restore any OS and OFED settings (if applicable) that have been overwritten during the upgrade.
11. If a leap upgrade from Red Hat Enterprise Linux 7.x to 8.x is performed, make sure to restore some of the removed packages and follow the details in [“Guidance for Red Hat Enterprise Linux 8.x on IBM Storage Scale nodes”](#) on page 764.
12. Rebuild the GPFS portability layer. For more information, see [“Building the GPFS portability layer on Linux nodes”](#) on page 483.
13. Verify that no conflicting services have been started or enabled during the OS upgrade. For example, some versions of RHEL autostart a non-IBM version of NFS after upgrade and it conflicts with the CES NFS service of IBM Storage Scale.
14. Verify SELinux and firewall settings.

15. Enable the GPFS autoloading setting, if it was disabled earlier in this procedure, by issuing the following command.

```
mmchconfig autoloading=yes -N nodeName
```

16. Restart GPFS on the node by issuing the following command.

```
mmstartup
```

17. Verify that all file systems are mounted on the node by issuing the following command.

```
mmismount all
```

18. If a protocol node is being upgraded, do the following:

- a. Resume the protocol node and start the protocol services on the node by issuing the following command.

```
mmces node resume --start
```

- b. Ensure that CES IPs are hosted by this node and that all services are healthy by issuing the following commands.

```
mmces address list  
mmces service list -a  
mmces node list  
mmhealth node show CES  
mmhealth cluster show
```

Repeat these steps on all nodes on which the OS needs to be upgraded.

Related tasks

[“Upgrading the SMB package after upgrading OS” on page 699](#)

If you have upgraded the operating system on your protocol nodes, ensure that the SMB package is upgraded to the version required for the new OS.

[“Upgrading the performance monitoring packages after upgrading Ubuntu OS” on page 705](#)

If you have upgraded the Ubuntu operating system on your IBM Storage Scale nodes, more steps are required to ensure that performance monitoring packages are upgraded to the version required for the latest version of Ubuntu.

[“Servicing IBM Storage Scale protocol nodes” on page 768](#)

Use these steps to upgrade the OS and to install kernel updates, firmware updates, and driver updates on IBM Storage Scale protocol nodes.

Guidance for Red Hat Enterprise Linux 8.x on IBM Storage Scale nodes

Red Hat Enterprise Linux 8.x is supported starting with IBM Storage Scale release 5.0.4.

Remember: Red Hat Enterprise Linux 8.0 is not an Extended Update Support (EUS) release, so the errata stream stops when Red Hat Enterprise Linux 8.1 is released. Do not use Red Hat Enterprise Linux 8.0 in production. Upgrade Red Hat Enterprise Linux to 8.1 to stay within the IBM Storage Scale support matrix.

The following list contains some general guidance when installing Red Hat Enterprise Linux 8.x on IBM Storage Scale nodes.

- After Red Hat Enterprise Linux 8.x is installed, IBM Storage Scale on x86 requires `elfutils-devel` to be installed for successfully building the GPFS portability layer (GPL) by using the `mmbuildgp1` command.
- Review the Red Hat Enterprise Linux 8.x documentation topic *Considerations for adopting RHEL8* for all hardware and driver support that was previously supported but has since been removed.
- Review the [IBM Storage Scale FAQ in IBM Documentation](#) for more information and for the supported kernel versions.

Red Hat provides a `leapp` upgrade utility to upgrade the OS from Red Hat Enterprise Linux 7.x to 8.x but it has some limitations. It is recommended to provision a new Red Hat Enterprise Linux 8.x node instance instead of using a `leapp` upgrade if possible. If you are using the `leapp` upgrade utility, there are some details to be aware of that are as follows.

- A `leapp` upgrade might upgrade the IBM Storage Scale node to the latest release of Red Hat Enterprise Linux 8.x. Ensure that `leapp` upgrade does not upgrade the node to an unsupported OS version. For supported OS and kernel versions, see [IBM Storage Scale FAQ in IBM Documentation](#).

Note: On s390x, a `leapp` upgrade from Red Hat Enterprise 7.x to 8.x is not supported. Using the `leapp` utility for the upgrade can lead to an unbootable system as described in the Red Hat bugzillas RH1890215 and RH1881428. Contact IBM support if you plan to upgrade from Red Hat Enterprise 7.x to 8.x on s390x.

- Using the `leapp` upgrade utility to upgrade from Red Hat Enterprise Linux 7.6 to 8.x with IBM Storage Scale 5.0.4.x or later installed might result in the removal of some IBM Storage Scale packages. This package removal might potentially happen with any third-party software. It is highly recommended to use `leapp-0.8.1-1` or later. The `leapp` utility is supported by Red Hat and any issues that are encountered with the `leapp` upgrade utility must be reported to Red Hat. Before you run `leapp` upgrade, save the list of IBM Storage Scale packages that are installed on the system by using the `rpm -qa | grep gpfs` command and then compare it to the packages installed after the upgrade. The `/var/log/leapp` directory also contains the packages that are to be removed.
- If Cluster Export Services (CES) are enabled on an IBM Storage Scale node running on a version earlier than 5.1.0 and you run a `leapp` upgrade from Red Hat Enterprise Linux 7.x to 8.x. Also, if Swift Object is configured in the system then you must first stop Swift Object and remove its associated packages. For more information, see [“Instructions for removing object protocol packages when upgrading protocol nodes to Red Hat Enterprise Linux 8.x” on page 767](#).

Note: With certain versions of IBM Storage Scale, the Swift Object protocol is not supported on Red Hat Enterprise Linux 8.x. Refer to IBM Storage Scale FAQ in IBM Documentation for determining the IBM Storage Scale version on which the object protocol is supported on Red Hat Enterprise Linux 8.x and if you can reinstall object or not.

- If Cluster Export Services (CES) are enabled on an IBM Storage Scale 5.0.4.x or later node and you run a `leapp` upgrade from Red Hat Enterprise Linux 7.x to 8.x, the NFS el7 packages are removed. The NFS el8 packages (`gpfs.nfs-ganesha-*.rpm`, `gpfs.nfs-ganesha-gpfs-*.rpm`, `gpfs.nfs-ganesha-utils-*.rpm`) need to be installed after the `leapp` upgrade. The NFS el8 packages are located in the `/usr/lpp/mmfs/5.x.x.x/ganesha_rpms/rhel8` directory. You can use `yum localinstall` or `rpm -ivh` commands to install these packages.
- If Cluster Export Services (CES) are enabled on an IBM Storage Scale 5.0.4.x or later node and you run a `leapp` upgrade from Red Hat Enterprise Linux 7.x to 8.x, the SMB el7 package is removed. The SMB el8 package needs to be installed after the `leapp` upgrade. The SMB el8 package is located in the `/usr/lpp/mmfs/5.x.x.x/smb_rpms/rhel8` directory. You can use `yum localinstall` or `rpm -ivh` commands to install this package.
- If performance collection or GUI functionalists are enabled on an IBM Storage Scale 5.0.4.x or later node and you run a `leapp` upgrade from Red Hat Enterprise Linux 7.x to 8.x, the associated `gpfs.gss.pmcollector` or `gpfs.gui` packages are removed and they need to be reinstalled after the `leapp` upgrade. For more information, see [“Reinstalling performance monitoring collector on Red Hat Enterprise Linux 8.x nodes” on page 705](#).
 - Performance collection package can be installed from: `/usr/lpp/mmfs/5.x.x.x/zimon_rpms/rhel8/gpfs.gss.pmcollector`
 - If NFS is also installed, the NFS Ganesha performance monitoring package can be installed from: `/usr/lpp/mmfs/5.x.x.x/zimon_rpms/rhel8/gpfs.pm-ganesha`
 - Management GUI package can be installed from: `/usr/lpp/mmfs/5.x.x.x/gpfs_rpms/gpfs.gui*`

After installing the management GUI package, start the management GUI by using the **`systemctl start gpfsgui`** command.

- A leapp upgrade from Red Hat Enterprise Linux 7.x to 8.x removes the python-ldap package, which is prerequisite for the **mmadquery** command. Ensure that you install python3-ldap after leapp upgrade, if you plan to use **mmadquery** on Red Hat Enterprise Linux 8.x.
- The CLI utility **mmadquery** that is bundled with IBM Storage Scale 5.1.0 for Red Hat Enterprise Linux 7.x is based on python2. The CLI utility **mmadquery** that is bundled with IBM Storage Scale 5.1.0 for Red Hat Enterprise Linux 8.x is based on python3. When you install IBM Storage Scale 5.1.0 on Red Hat Enterprise Linux 7.x and upgrade an operating system from Red Hat Enterprise Linux 7.x to Red Hat Enterprise Linux 8.x, **mmadquery** still points to python2.

To point to the correct **mmadquery**, which is based on python3, perform the following steps:

1. Remove existing /usr/lpp/mmfs/bin/mmadquery.
 2. Copy /usr/lpp/mmfs/bin/mmadquery_py3 to /usr/lpp/mmfs/bin/mmadquery.
- After all packages that were removed by leapp upgrade are reinstalled as described in the preceding text, complete the steps in [“Guidance for upgrading the operating system on IBM Storage Scale nodes” on page 762](#) to finish bringing the node online after the OS is upgraded to Red Hat Enterprise Linux 8.x.
 - If Cluster Export Services (CES) are enabled on an IBM Storage Scale node, an in-place upgrade from Red Hat Enterprise Linux 8.x by using leapp upgrade utility might fail.

The leapp utility does not support upgrading Red Hat Enterprise Linux 8.x on the file protocols nodes because of the winbind strings in the /etc/nsswitch.conf file. When either SMB or NFS or both SMB and NFS protocols are configured for Active Directory (AD) based authentication schemes by using the **mmuserauth** command, this command adds winbind in the /etc/nsswitch.conf file of each protocol node. The in-place upgrade with winbind strings in the /etc/nsswitch.conf is not allowed because of a known OS upgrade limitation.

To do the in-place upgrade, you need to complete the steps from 1 through 4 in the [Troubleshooting known issues](#) documentation for upgrading Red Hat Enterprise Linux 8.x. After the OS upgrade, when you resume the protocol nodes, the IBM Storage Scale adds the removed winbind entries in the /etc/nsswitch.conf file on each protocol node. You can verify the successful addition of the removed winbind entries for one of the CES nodes by issuing the following command:

```
# mmuserauth service check -N cesNodes
```

After the OS upgrade, do the following steps to confirm correctness of the /etc/nsswitch.conf file on the nodes:

1. Check whether the current configuration is valid by issuing the following command:

```
# authselect check
```

2. Check whether the /etc/nsswitch.conf file is linked to the current authselect profile nsswitch.conf file by issuing the following command:

```
# ls -l /etc/nsswitch.conf
```

3. If either the current configuration is incorrect or the /etc/nsswitch.conf file is not linked to the current authselect profile, apply forcefully the current authselect profile on the nodes.

- a. Check the current authselect profile by issuing the following command:

```
# authselect current
```

- b. Apply forcefully the authselect profile again by issuing the following command:

```
# authselect select ProfileNameIdentifiedFromPreviousCommand --force
```

The packages that leapp upgrade might remove can vary depending on packages that are installed and dependencies that might exist. It might be required to reinstall any removed packages, which can typically be determined by viewing the log /var/log/leapp/leapp-upgrade.log and looking for *Removing dependent packages* after the leapp upgrade completes and the node restarts.

Instructions for removing object protocol packages when upgrading protocol nodes to Red Hat Enterprise Linux 8.x

The presence of the object packages and their dependencies can prevent upgrade to Red Hat Enterprise Linux® 8.x. Therefore, it is necessary to remove the associated packages from the system when you are upgrading protocol nodes to Red Hat Enterprise Linux® 8.x. After the upgrade is complete, the object packages can be reinstalled.

About this task

Important:

- CES Swift Object protocol feature is not supported from IBM Storage Scale 5.2.0 onwards.
- IBM Storage Scale 5.1.8 is the last release that has CES Swift Object protocol.
- IBM Storage Scale 5.2.0 will tolerate the update of a CES node from IBM Storage Scale 5.1.8.
 - *Tolerate* means:
 - The CES node will be updated to 5.2.0.
 - Swift Object support will not be updated as part of the 5.2.0 update.
 - You may continue to use the version of Swift Object protocol that was provided in IBM Storage Scale 5.1.8 on the CES 5.2.0 node.
 - IBM will provide usage and known defect support for the version of Swift Object that was provided in IBM Storage Scale 5.1.8 until you migrate to a supported object solution that IBM Storage Scale provides.
 - CES Swift Object is replaced with IBM Storage Scale S3. For more details, refer the *S3 support overview* section the *IBM Storage Scale: Concepts, Planning, and Installation Guide*.
 - For more information about Swift Object in IBM Storage Scale, refer to the [IBM Storage Scale 5.2.0 documentation](#).
- Contact IBM for further details and migration planning.

The object protocol is only supported on Red Hat Enterprise Linux® 8.x with IBM Storage Scale version 5.1.0.1 and later. Systems at earlier versions of IBM Storage Scale with object protocol need to upgrade to version 5.1.0.1 or later after the system upgrade to Red Hat Enterprise Linux® 8.x.



Attention: The object protocol is not supported in IBM Storage Scale 5.1.0.0. If you want to deploy the object, upgrade to the IBM Storage Scale 5.1.0.1 or a later release.

Procedure

1. After the object protocol is stopped, certain packages need to be cleaned up to prevent conflicts during the upgrade to Red Hat Enterprise Linux 8.x.

Run the following command to uninstall these packages and their dependencies:

```
yum erase isa-l librasurecode openstack-keystone openstack-selinux openstack-swift-account
openstack-swift-container openstack-swift-object
openstack-swift-proxy python2-bcrypt python2-cinderclient python2-dogpile-cache python2-
glanceclient python2-keystoneauth1 python2-keystoneclient python2-keystonemiddleware
python2-neutronclient python2-novaclient python2-openstackclient python2-osc-lib python2-os-
client-config python2-oslo-cache python2-oslo-concurrency python2-oslo-config
python2-oslo-context python2-oslo-db python2-oslo-i18n python2-oslo-middleware python2-oslo-
service python2-pbr python2-pysocks python2-repoze-lru python2-requests
python2-requestsexceptions python2-scrypt python-argparse python-cadf python-cliff-tablib
python-debtcollector python-dns python-funcsigs python-greenlet python-httplib2
python-jinja2 python-keystone python-msgpack python-openstackclient-lang python-oslo-cache-
lang python-oslo-concurrency-lang python-oslo-db-lang python-oslo-i18n-lang
python-oslo-middleware-lang python-pip python-posix_ipc python-pyeclib python-simplejson
python-sqlalchemy python-swift python-tablib python-testtools python-wrapit
python-wsgiref python-zope-event swift3 python-keyring python-idna python-PyMySQL
```

2. After the packages that are mentioned in the preceding step are removed, the following packages must be downgraded to the level of the system Red Hat Enterprise Linux packages by running the following command:

```
yum downgrade pyOpenSSL python2-cryptography
```

If these packages are not removed or downgraded, the upgrade to Red Hat Enterprise Linux 8.x using the **leapp** utility might fail with a message such as the following message:

```
[ERROR] Actor: prepare_upgrade_transaction Message: A Leapp Command Error occurred
```

The **leapp** command with the **--verbose** option might display messages such as the following message:

```
file *** from install of package ### conflicts with file from package ###
```

If the messages are related to the packages associated with the object protocol, performing the cleanup steps eliminates these errors.

What to do next

After the object protocol is stopped and the associated packages are removed, proceed with the Red Hat Enterprise Linux 8.x upgrade.

Note: Due to the object packages being removed for the leapp upgrade, these special steps need to be done when you are using the installation toolkit to upgrade the object protocol to version 5.1.x after the system upgrade is complete:

1. Use the installation toolkit offline upgrade for the protocol nodes.
2. After the installation toolkit has successfully completed the upgrade, manually migrate the object configuration data to version 5.1 by using the following command:

```
mmobj config manage --version-sync
```

For more information, see [“Upgrading Swift Object packages to IBM Storage Scale 5.1.8 or earlier by using the installation toolkit” on page 741.](#)

Servicing IBM Storage Scale protocol nodes

Use these steps to upgrade the OS and to install kernel updates, firmware updates, and driver updates on IBM Storage Scale protocol nodes.

About this task

It is advisable to match OS, kernel, firmware, and driver across all nodes within a cluster to help with performance and to ease debugging. As a part of this procedure, ensure the following:

- Before upgrading OS or installing kernel updates, refer to the IBM Storage Scale FAQ to ensure that the version of IBM Storage Scale currently running is supported on the version of OS that you are planning to upgrade to or the kernel update that you are planning to install: [IBM Storage Scale FAQ in IBM Documentation](#).

It is possible that an IBM Storage Scale upgrade might be required to get to a version of IBM Storage Scale that supports the newly wanted OS. If so, upgrade IBM Storage Scale first and then the OS.

- If kernel update is for a new OS (RHEL 7.3 vs RHEL 7.2, for example), always update the OS before applying the kernel update.
- When taking nodes offline to upgrade the OS or installing kernel updates, firmware updates, or driver updates, ensure the following:
 - The Quorum does not break.
 - Enough NSD nodes remain up to access NSDs.

- The remaining nodes can handle the desired workload.

Procedure

1. Uninstall the applicable drivers as follows.

- a) Obtain and extract the applicable drivers.
- b) Suspend CES on the node being upgraded.

```
mmces node suspend -N NodeBeingUpgraded --stop
```

- c) Shut down GPFS on the node being upgraded.

```
mmshutdown -N NodeBeingUpgraded
```

- d) Find the uninstallation script within the drivers package and execute it on the node being upgraded.
2. Create a local repository for the OS upgrade.

A repository must be created so that the OS can be upgraded. This repository can be DVD or ISO based. Make sure that you remove any repositories pointing to old OS versions.

3. Upgrade the OS and reboot the node after the OS upgrade.

For more information, see [“Guidance for upgrading the operating system on IBM Storage Scale nodes” on page 762](#).

4. Install the kernel update and reboot the node after the kernel update.

5. Rebuild the GPFS portability layer using the **mmbuildgpl** command.

For more information, see [“Building the GPFS portability layer on Linux nodes” on page 483](#).

6. Install the required firmware updates and reboot the node after the firmware update.

7. Install the latest drivers, as applicable and reboot the node after the driver update.

8. Verify that GPFS is active on the node and then resume CES.

```
mmgetstate -a
mmces node resume -N NodeBeingUpgraded --start
mmces node list
mmces service list -a
mmces address list
```

9. Repeat the preceding steps on each protocol node that needs to be serviced, one node at a time.

Offline upgrade with complete cluster shutdown

You can use the following steps for an offline upgrade.

Important: Online upgrade from an unsupported version of IBM Storage Scale or GPFS is not supported. However, you can use the following manual offline upgrade procedure to upgrade to a supported version of IBM Storage Scale.

1. Stop all file system operations on the cluster and any external access to the cluster. For example, protocol clients.
2. If you have protocol nodes, then stop all CES operations.

```
mmces service stop Protocol -a
```

3. Unmount all file systems.

```
mmumount all -a
```

4. Shut down GPFS on all nodes.

```
mmshutdown -a
```

Refer to the [IBM Storage Scale FAQ in IBM Documentation](#) for the supported operating system versions that apply to the IBM Storage Scale release that you are upgrading to.

5. Verify that the GPFS daemon has terminated and that the kernel extensions are unloaded.

```
mmfsenv -u
```

If the **mmfsenv -u** command reports that it cannot unload the kernel extensions because they are busy, then the installation can proceed, but the node must be rebooted after the installation. Kernel extensions are busy means that a process has a current directory in some GPFS file system directory or it has an open file descriptor. The Linux utility **lsof** can identify the process and that process can then be killed. Retry **mmfsenv -u** after killing the process and if the command succeeds, then a reboot of the node can be avoided.

Note:

- The **mmfsenv -u** command is normally used with assistance from support or for cases where its usage is specifically documented such as in this scenario.
 - The **mmfsenv -u** command only checks the local node from which it is run, not all the nodes in the cluster.
6. If the operating system needs to be upgraded, then complete that upgrade.
 7. Extract the IBM Storage Scale installation image so that all installation files (rpm, deb, bff, msi, and so on) are available.
 8. Run the necessary platform-specific commands to upgrade GPFS.
- SLES and RHEL Linux:

```
rpm -Fvh /usr/lpp/mmfs/5.x.y.z/gpfs_rpms/  
<OS>{gpfs.adv*.rpm,gpfs.base*.rpm,gpfs.crypto*.rpm,  
gpfs.docs*.rpm,gpfs.gpl*.rpm,gpfs.gskit*.rpm,gpfs.msg*.rpm,  
gpfs.license.XX*.rpm,gpfs.compression*.rpm}  
rpm -Fvh /usr/lpp/mmfs/5.x.y.z/zimon_rpms/  
<OS>{gpfs.gss.pmsensors*.rpm,gpfs.gss.pmc collector*.rpm}
```

Where **<OS>** is the operating system specific directory such as **rhel17** or **sles15**.

- The **gpfs-adv** and **gpfs-crypto** packages are available only with IBM Storage Scale Advanced Edition, IBM Storage Scale Data Management Edition, or IBM Storage Scale Developer Edition.
 - Depending on the IBM Storage Scale edition that you are installing, install the appropriate **gpfs.license** package.
 - IBM Storage Scale Developer Edition is available only on Red Hat Enterprise Linux on x86_64.
- Ubuntu:

```
dpkg -i /usr/lpp/mmfs/5.x.y.z/gpfs_debs/  
<OS>{gpfs.adv*.deb,gpfs.base*.deb,gpfs.crypto*.deb,  
gpfs.docs*.deb,gpfs.gpl*.deb,gpfs.gskit*.deb,gpfs.msg*.deb,  
gpfs.license.XX*.deb,gpfs.compression*.deb}  
  
dpkg -i /usr/lpp/mmfs/5.x.y.z/zimon_debs/  
<OS>{gpfs.gss.pmsensors*.deb,gpfs.gss.pmc collector*.deb}
```

Where **<OS>** is the operating system specific directory such as **ubuntu18**.

Note:

- The **gpfs-crypto** and the **gpfs-adv** packages are available only with IBM Storage Scale Advanced Edition and IBM Storage Scale Data Management Edition.
 - Depending on the IBM Storage Scale edition that you are installing, install the appropriate **gpfs.license** package.
- AIX:

```
installp -agXYd . gpfs
```

- Windows:
 - a. Uninstall the version of GPFS or IBM Storage Scale that you are upgrading from and restart.
 - b. Uninstall the license of GPFS or IBM Storage Scale version that you are upgrading from.
 - c. Uninstall Open Secure Shell for GPFS on Windows, if applicable.
 - d. Disable any SUA daemon, such as OpenSSH, that might be configured. Do not uninstall SUA yet, or you may lose GPFS configuration information.
 - e. Install the 64-bit version of Cygwin. For more information, see [“Installing Cygwin” on page 588](#).
 - f. Install `gpfs.base-5.1.x.x-Windows-license.msi`.
 - g. Install `gpfs.base-5.1.x.x-Windows.msi`.
 - h. Install IBM® GSKit for GPFS.
 - i. Uninstall SUA completely.
 - j. Configure passwordless SSH in Cygwin for mixed (Linux, AIX, Windows) clusters or use **mmwinservctl** on a Windows only cluster.
- 9. On Linux systems, rebuild the GPFS portability layer (GPL) with the **mmbuildgpl** command. For more information, see [“Building the GPFS portability layer on Linux nodes” on page 483](#).
- 10. Restart GPFS by issuing the **mmstartup** command and remount file systems by issuing the **mmmout** command, as necessary. If you have protocol nodes, start the CES operations by issuing the **mmces service start Protocol -a** command.

Note:

- For completing S3 protocol upgrade, you need to run the **mms3 upgrade start [--upgrade_version] [-d]** command from any one node.
 - At this point, all code is at the upgraded level, but the cluster is not upgraded and new functions cannot be configured until the upgrade is completed.
11. Complete the upgrade as documented in [“Completing the upgrade to a new level of IBM Storage Scale” on page 753](#).

Related concepts

[“IBM Storage Scale supported upgrade paths” on page 672](#)

Use this information to understand the online upgrade paths that are supported for IBM Storage Scale.

[“Performing offline upgrade or excluding nodes from upgrade by using installation toolkit” on page 733](#)

Accessibility features for IBM Storage Scale

Accessibility features help users who have a disability, such as restricted mobility or limited vision, to use information technology products successfully.

Accessibility features

The following list includes the major accessibility features in IBM Storage Scale:

- Keyboard-only operation
- Interfaces that are commonly used by screen readers
- Keys that are discernible by touch but do not activate just by touching them
- Industry-standard devices for ports and connectors
- The attachment of alternative input and output devices

IBM Documentation, and its related publications, are accessibility-enabled.

Keyboard navigation

This product uses standard Microsoft Windows navigation keys.

IBM and accessibility

See the [IBM Human Ability and Accessibility Center \(www.ibm.com/able\)](http://www.ibm.com/able) for more information about the commitment that IBM has to accessibility.

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, MD-NC119 Armonk, NY 10504-1785 US

For license inquiries regarding double-byte character set (DBCS) information, contact the IBM Intellectual Property Department in your country or send inquiries, in writing, to:

Intellectual Property Licensing Legal and Intellectual Property Law IBM Japan Ltd. 19-21, Nihonbashi-Hakozakicho, Chuo-ku Tokyo 103-8510, Japan

INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some jurisdictions do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM websites are provided for convenience only and do not in any manner serve as an endorsement of those websites. The materials at those websites are not part of the materials for this IBM product and use of those websites is at your own risk.

IBM may use or distribute any of the information you provide in any way it believes appropriate without incurring any obligation to you.

Licensees of this program who wish to have information about it for the purpose of enabling: (i) the exchange of information between independently created programs and other programs (including this one) and (ii) the mutual use of the information which has been exchanged, should contact:

IBM Director of Licensing IBM Corporation North Castle Drive, MD-NC119 Armonk, NY 10504-1785 US

Such information may be available, subject to appropriate terms and conditions, including in some cases, payment of a fee.

The licensed program described in this document and all licensed material available for it are provided by IBM under terms of the IBM Customer Agreement, IBM International Program License Agreement or any equivalent agreement between us.

The performance data discussed herein is presented as derived under specific operating conditions. Actual results may vary.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and

cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

Statements regarding IBM's future direction or intent are subject to change or withdrawal without notice, and represent goals and objectives only.

All IBM prices shown are IBM's suggested retail prices, are current and are subject to change without notice. Dealer prices may vary.

This information is for planning purposes only. The information herein is subject to change before the products described become available.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to actual people or business enterprises 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.

Each copy or any portion of these sample programs or any derivative work must include a copyright notice as follows:

© (your company name) (year).

Portions of this code are derived from IBM Corp.

Sample Programs. © Copyright IBM Corp. _enter the year or years_.

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](http://www.ibm.com/legal/copytrade.shtml).

Intel is a trademark 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.

The registered trademark Linux is used pursuant to a sublicense from the Linux Foundation, the exclusive licensee of Linus Torvalds, owner of the mark on a worldwide basis.

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

Red Hat, OpenShift®, and Ansible are trademarks or registered trademarks of Red Hat, Inc. or its subsidiaries in the United States and other countries.

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

Terms and conditions for product documentation

Permissions for the use of these publications are granted subject to the following terms and conditions.

IBM Privacy Policy

At IBM we recognize the importance of protecting your personal information and are committed to processing it responsibly and in compliance with applicable data protection laws in all countries in which IBM operates.

Visit the IBM Privacy Policy for additional information on this topic at <https://www.ibm.com/privacy/details/us/en/>.

Applicability

These terms and conditions are in addition to any terms of use for the IBM website.

Personal use

You can reproduce these publications for your personal, noncommercial use provided that all proprietary notices are preserved. You cannot distribute, display, or make derivative work of these publications, or any portion thereof, without the express consent of IBM.

Commercial use

You can reproduce, distribute, and display these publications solely within your enterprise provided that all proprietary notices are preserved. You cannot make derivative works of these publications, or reproduce, distribute, or display these publications or any portion thereof outside your enterprise, without the express consent of IBM.

Rights

Except as expressly granted in this permission, no other permissions, licenses, or rights are granted, either express or implied, to the Publications or any information, data, software or other intellectual property contained therein.

IBM reserves the right to withdraw the permissions that are granted herein whenever, in its discretion, the use of the publications is detrimental to its interest or as determined by IBM, the above instructions are not being properly followed.

You cannot download, export, or reexport this information except in full compliance with all applicable laws and regulations, including all United States export laws and regulations.

IBM MAKES NO GUARANTEE ABOUT THE CONTENT OF THESE PUBLICATIONS. THE PUBLICATIONS ARE PROVIDED "AS-IS" AND WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESSED OR IMPLIED, INCLUDING BUT NOT LIMITED TO IMPLIED WARRANTIES OF MERCHANTABILITY, NON-INFRINGEMENT, AND FITNESS FOR A PARTICULAR PURPOSE.

Glossary

This glossary provides terms and definitions for IBM Storage Scale.

The following cross-references are used in this glossary:

- *See* refers you from a nonpreferred term to the preferred term or from an abbreviation to the spelled-out form.
- *See also* refers you to a related or contrasting term.

For other terms and definitions, see the [IBM Terminology website \(www.ibm.com/software/globalization/terminology\)](http://www.ibm.com/software/globalization/terminology) (opens in new window).

B

block utilization

The measurement of the percentage of used subblocks per allocated blocks.

C

cluster

A loosely coupled collection of independent systems (nodes) organized into a network for the purpose of sharing resources and communicating with each other. See also *GPFS cluster*.

cluster configuration data

The configuration data that is stored on the cluster configuration servers.

Cluster Export Services (CES) nodes

A subset of nodes configured within a cluster to provide a solution for exporting GPFS file systems by using the Network File System (NFS), Server Message Block (SMB), and S3 protocols.

cluster manager

The node that monitors node status using disk leases, detects failures, drives recovery, and selects file system managers. The cluster manager must be a quorum node. The selection of the cluster manager node favors the quorum-manager node with the lowest node number among the nodes that are operating at that particular time.

Note: The cluster manager role is not moved to another node when a node with a lower node number becomes active.

clustered watch folder

Provides a scalable and fault-tolerant method for file system activity within an IBM Storage Scale file system. A clustered watch folder can watch file system activity on a fileset, inode space, or an entire file system. Events are streamed to an external Kafka sink cluster in an easy-to-parse JSON format. For more information, see the *mmwatch* command in the *IBM Storage Scale: Command and Programming Reference Guide*.

control data structures

Data structures needed to manage file data and metadata cached in memory. Control data structures include hash tables and link pointers for finding cached data; lock states and tokens to implement distributed locking; and various flags and sequence numbers to keep track of updates to the cached data.

D

Data Management Application Program Interface (DMAPI)

The interface defined by the Open Group's XDSM standard as described in the publication *System Management: Data Storage Management (XDSM) API Common Application Environment (CAE) Specification C429*, The Open Group ISBN 1-85912-190-X.

deadman switch timer

A kernel timer that works on a node that has lost its disk lease and has outstanding I/O requests. This timer ensures that the node cannot complete the outstanding I/O requests (which would risk causing file system corruption), by causing a panic in the kernel.

dependent fileset

A fileset that shares the inode space of an existing independent fileset.

disk descriptor

A definition of the type of data that the disk contains and the failure group to which this disk belongs. See also *failure group*.

disk leasing

A method for controlling access to storage devices from multiple host systems. Any host that wants to access a storage device configured to use disk leasing registers for a lease; in the event of a perceived failure, a host system can deny access, preventing I/O operations with the storage device until the preempted system has reregistered.

disposition

The session to which a data management event is delivered. An individual disposition is set for each type of event from each file system.

domain

A logical grouping of resources in a network for the purpose of common management and administration.

E**ECKD**

See *extended count key data (ECKD)*.

ECKD device

See *extended count key data device (ECKD device)*.

encryption key

A mathematical value that allows components to verify that they are in communication with the expected server. Encryption keys are based on a public or private key pair that is created during the installation process. See also *file encryption key*, *master encryption key*.

extended count key data (ECKD)

An extension of the count-key-data (CKD) architecture. It includes additional commands that can be used to improve performance.

extended count key data device (ECKD device)

A disk storage device that has a data transfer rate faster than some processors can utilize and that is connected to the processor through use of a speed matching buffer. A specialized channel program is needed to communicate with such a device. See also *fixed-block architecture disk device*.

F**failback**

Cluster recovery from failover following repair. See also *failover*.

failover

(1) The assumption of file system duties by another node when a node fails. (2) The process of transferring all control of the ESS to a single cluster in the ESS when the other clusters in the ESS fails. See also *cluster*. (3) The routing of all transactions to a second controller when the first controller fails. See also *cluster*.

failure group

A collection of disks that share common access paths or adapter connections, and could all become unavailable through a single hardware failure.

FEK

See *file encryption key*.

fileset

A hierarchical grouping of files managed as a unit for balancing workload across a cluster. See also *dependent fileset*, *independent fileset*.

fileset snapshot

A snapshot of an independent fileset plus all dependent filesets.

file audit logging

Provides the ability to monitor user activity of IBM Storage Scale file systems and store events related to the user activity in a security-enhanced fileset. Events are stored in an easy-to-parse JSON format. For more information, see the *mmaudit* command in the *IBM Storage Scale: Command and Programming Reference Guide*.

file clone

A writable snapshot of an individual file.

file encryption key (FEK)

A key used to encrypt sectors of an individual file. See also *encryption key*.

file-management policy

A set of rules defined in a policy file that GPFS uses to manage file migration and file deletion. See also *policy*.

file-placement policy

A set of rules defined in a policy file that GPFS uses to manage the initial placement of a newly created file. See also *policy*.

file system descriptor

A data structure containing key information about a file system. This information includes the disks assigned to the file system (*stripe group*), the current state of the file system, and pointers to key files such as quota files and log files.

file system descriptor quorum

The number of disks needed in order to write the file system descriptor correctly.

file system manager

The provider of services for all the nodes using a single file system. A file system manager processes changes to the state or description of the file system, controls the regions of disks that are allocated to each node, and controls token management and quota management.

fixed-block architecture disk device (FBA disk device)

A disk device that stores data in blocks of fixed size. These blocks are addressed by block number relative to the beginning of the file. See also *extended count key data device*.

fragment

The space allocated for an amount of data too small to require a full block. A fragment consists of one or more subblocks.

G**GPUDirect Storage**

IBM Storage Scale's support for NVIDIA's GPUDirect Storage (GDS) enables a direct path between GPU memory and storage. File system storage is directly connected to the GPU buffers to reduce latency and load on CPU. Data is read directly from an NSD server's pagepool and it is sent to the GPU buffer of the IBM Storage Scale clients by using RDMA.

global snapshot

A snapshot of an entire GPFS file system.

GPFS cluster

A cluster of nodes defined as being available for use by GPFS file systems.

GPFS portability layer

The interface module that each installation must build for its specific hardware platform and Linux distribution.

GPFS recovery log

A file that contains a record of metadata activity and exists for each node of a cluster. In the event of a node failure, the recovery log for the failed node is replayed, restoring the file system to a consistent state and allowing other nodes to continue working.

I

ill-placed file

A file assigned to one storage pool but having some or all of its data in a different storage pool.

ill-replicated file

A file with contents that are not correctly replicated according to the desired setting for that file. This situation occurs in the interval between a change in the file's replication settings or suspending one of its disks, and the restripe of the file.

independent filesset

A filesset that has its own inode space.

indirect block

A block containing pointers to other blocks.

inode

The internal structure that describes the individual files in the file system. There is one inode for each file.

inode space

A collection of inode number ranges reserved for an independent filesset, which enables more efficient per-filesset functions.

ISKLM

IBM Security Key Lifecycle Manager. For GPFS encryption, the ISKLM is used as an RKM server to store MEKs.

J

journaled file system (JFS)

A technology designed for high-throughput server environments, which are important for running intranet and other high-performance e-business file servers.

junction

A special directory entry that connects a name in a directory of one filesset to the root directory of another filesset.

K

kernel

The part of an operating system that contains programs for such tasks as input/output, management and control of hardware, and the scheduling of user tasks.

M

master encryption key (MEK)

A key used to encrypt other keys. See also *encryption key*.

MEK

See *master encryption key*.

metadata

Data structures that contain information that is needed to access file data. Metadata includes inodes, indirect blocks, and directories. Metadata is not accessible to user applications.

metanode

The one node per open file that is responsible for maintaining file metadata integrity. In most cases, the node that has had the file open for the longest period of continuous time is the metanode.

mirroring

The process of writing the same data to multiple disks at the same time. The mirroring of data protects it against data loss within the database or within the recovery log.

Microsoft Management Console (MMC)

A Windows tool that can be used to do basic configuration tasks on an SMB server. These tasks include administrative tasks such as listing or closing the connected users and open files, and creating and manipulating SMB shares.

multi-tailed

A disk connected to multiple nodes.

N**namespace**

Space reserved by a file system to contain the names of its objects.

Network File System (NFS)

A protocol, developed by Sun Microsystems, Incorporated, that allows any host in a network to gain access to another host or netgroup and their file directories.

Network Shared Disk (NSD)

A component for cluster-wide disk naming and access.

NSD volume ID

A unique 16-digit hex number that is used to identify and access all NSDs.

node

An individual operating-system image within a cluster. Depending on the way in which the computer system is partitioned, it may contain one or more nodes.

node descriptor

A definition that indicates how GPFS uses a node. Possible functions include: manager node, client node, quorum node, and nonquorum node.

node number

A number that is generated and maintained by GPFS as the cluster is created, and as nodes are added to or deleted from the cluster.

node quorum

The minimum number of nodes that must be running in order for the daemon to start.

node quorum with tiebreaker disks

A form of quorum that allows GPFS to run with as little as one quorum node available, as long as there is access to a majority of the quorum disks.

non-quorum node

A node in a cluster that is not counted for the purposes of quorum determination.

Non-Volatile Memory Express (NVMe)

An interface specification that allows host software to communicate with non-volatile memory storage media.

P**policy**

A list of file-placement, service-class, and encryption rules that define characteristics and placement of files. Several policies can be defined within the configuration, but only one policy set is active at one time.

policy rule

A programming statement within a policy that defines a specific action to be performed.

pool

A group of resources with similar characteristics and attributes.

portability

The ability of a programming language to compile successfully on different operating systems without requiring changes to the source code.

primary GPFS cluster configuration server

In a GPFS cluster, the node chosen to maintain the GPFS cluster configuration data.

private IP address

An IP address used to communicate on a private network.

public IP address

An IP address used to communicate on a public network.

Q**quorum node**

A node in the cluster that is counted to determine whether a quorum exists.

quota

The amount of disk space and number of inodes assigned as upper limits for a specified user, group of users, or fileset.

quota management

The allocation of disk blocks to the other nodes writing to the file system, and comparison of the allocated space to quota limits at regular intervals.

R**Redundant Array of Independent Disks (RAID)**

A collection of two or more disk physical drives that present to the host an image of one or more logical disk drives. In the event of a single physical device failure, the data can be read or regenerated from the other disk drives in the array due to data redundancy.

recovery

The process of restoring access to file system data when a failure has occurred. Recovery can involve reconstructing data or providing alternative routing through a different server.

remote key management server (RKM server)

A server that is used to store master encryption keys.

replication

The process of maintaining a defined set of data in more than one location. Replication consists of copying designated changes for one location (a source) to another (a target) and synchronizing the data in both locations.

RKM server

See *remote key management server*.

rule

A list of conditions and actions that are triggered when certain conditions are met. Conditions include attributes about an object (file name, type or extension, dates, owner, and groups), the requesting client, and the container name associated with the object.

S**SAN-attached**

Disks that are physically attached to all nodes in the cluster using Serial Storage Architecture (SSA) connections or using Fibre Channel switches.

Scale Out Backup and Restore (SOBAR)

A specialized mechanism for data protection against disaster only for GPFS file systems that are managed by IBM Storage Protect for Space Management.

secondary GPFS cluster configuration server

In a GPFS cluster, the node chosen to maintain the GPFS cluster configuration data in the event that the primary GPFS cluster configuration server fails or becomes unavailable.

Secure Hash Algorithm digest (SHA digest)

A character string used to identify a GPFS security key.

session failure

The loss of all resources of a data management session due to the failure of the daemon on the session node.

session node

The node on which a data management session was created.

Small Computer System Interface (SCSI)

An ANSI-standard electronic interface that allows personal computers to communicate with peripheral hardware, such as disk drives, tape drives, CD-ROM drives, printers, and scanners faster and more flexibly than previous interfaces.

snapshot

An exact copy of changed data in the active files and directories of a file system or fileset at a single point in time. See also *fileset snapshot*, *global snapshot*.

source node

The node on which a data management event is generated.

stand-alone client

The node in a one-node cluster.

storage area network (SAN)

A dedicated storage network tailored to a specific environment, combining servers, storage products, networking products, software, and services.

storage pool

A grouping of storage space consisting of volumes, logical unit numbers (LUNs), or addresses that share a common set of administrative characteristics.

stripe group

The set of disks comprising the storage assigned to a file system.

striping

A storage process in which information is split into blocks (a fixed amount of data) and the blocks are written to (or read from) a series of disks in parallel.

subblock

The smallest unit of data accessible in an I/O operation, equal to one thirty-second of a data block.

system storage pool

A storage pool containing file system control structures, reserved files, directories, symbolic links, special devices, as well as the metadata associated with regular files, including indirect blocks and extended attributes. The system storage pool can also contain user data.

T**token management**

A system for controlling file access in which each application performing a read or write operation is granted some form of access to a specific block of file data. Token management provides data consistency and controls conflicts. Token management has two components: the token management server, and the token management function.

token management function

A component of token management that requests tokens from the token management server. The token management function is located on each cluster node.

token management server

A component of token management that controls tokens relating to the operation of the file system. The token management server is located at the file system manager node.

transparent cloud tiering (TCT)

A separately installable add-on feature of IBM Storage Scale that provides a native cloud storage tier. It allows data center administrators to free up on-premise storage capacity, by moving out cooler data to the cloud storage, thereby reducing capital and operational expenditures.

twin-tailed

A disk connected to two nodes.

U**user storage pool**

A storage pool containing the blocks of data that make up user files.

V**VFS**

See *virtual file system*.

virtual file system (VFS)

A remote file system that has been mounted so that it is accessible to the local user.

virtual node (vnode)

The structure that contains information about a file system object in a virtual file system (VFS).

W**watch folder API**

Provides a programming interface where a custom C program can be written that incorporates the ability to monitor inode spaces, filesets, or directories for specific user activity-related events within IBM Storage Scale file systems. For more information, a sample program is provided in the following directory on IBM Storage Scale nodes: `/usr/lpp/mmfs/samples/util` called `tswf` that can be modified according to the user's needs.

Index

Special Characters

- ?
 - deleting old sensors [718](#)
- .afm
 - internal directory [69](#)
- .pconflicts
 - internal directory [69](#)
- .ptrash
 - internal directory [69](#)
- (excluding a node
 - offline upgrade [733](#)
- (marking a component offline(
 - during upgrade [733](#)
- /tmp/mmfs
 - collecting problem determination data in [468](#)

Numerics

- 2 K metadata [175](#)

A

- access control [47](#)
- access control lists (ACLs)
 - file system authorization [354](#)
- access control on GPFS file systems
 - Windows [586](#)
- access to file systems
 - simultaneous [2](#)
- accessibility features for IBM Storage Scale [773](#)
- active file management
 - deployment [623](#)
- Active File Management
 - failover mapping [125](#)
- Active SMB connections
 - determining factors [399](#)
- AD [220](#)
- adapter
 - invariant address requirement [309](#)
- adding cloud servicesnode
 - existing cloud servicescluster [630](#)
- adding LTFS nodes [572](#)
- adding protocol nodes
 - with toolkit [567](#)
- administration commands
 - GPFS [4](#), [5](#), [19](#)
- administration daemon [218](#), [227](#), [233](#)
- Administration daemon [216](#), [228](#)
- AFM
 - AFM gateway nodes requirements on the cache and home clusters [453](#)
 - AFM to COS replication [168](#)
 - asynchronous auto prefetch [126](#)
 - cache eviction [94](#)
 - caching modes [56](#)
 - disable online [115](#)

AFM (continued)

- encryption [113](#)
 - failover mapping [125](#)
 - Fast create [74](#)
 - fast remove [121](#)
 - fileset disabling [114](#)
 - gateway node [453](#)
 - global namespace [61](#)
 - independent-writer (IW) [56](#)
 - install [641](#)
 - Limitations [128](#), [144](#)
 - Local Update mode [123](#)
 - local-update (LU) [56](#)
 - NFS backend filesets [126](#)
 - operation with disconnected home [98](#)
 - outband updates [122](#)
 - Parallel data transfer using multiple remote mounts [80](#)
 - Parallel data transfers [77](#)
 - partial file caching [81](#)
 - peer snapshot [90](#)
 - planned maintenance
 - IW cache [111](#)
 - planning [452](#)
 - prefetch [82](#)
 - primary gateway [59](#)
 - psnap [90](#)
 - read-only (RO) [56](#)
 - resync on SW filesets [103](#)
 - single-writer (SW) [56](#)
 - Synchronous read mount [104](#)
 - UID and GID requirements on the cache and home clusters [452](#)
 - upgrade [692](#)
 - WAN latency [54](#)
 - workerThreads on cache cluster [452](#)
 - writable snapshot [123](#)
- AFM directory [175](#)
- AFM DR
- best practices [152](#)
 - cache cluster
 - UID/GID requirements [455](#)
 - changing
 - secondary site [140](#)
 - characteristics
 - active-passive relationships [147](#)
 - independent fileset [147](#)
 - NFD exports [148](#)
 - NFSv3 [148](#)
 - one-on-one relationships [147](#)
 - failback
 - new primary site [139](#)
 - old primary site [138](#)
 - failover
 - secondary site [138](#)
 - features [133](#)
 - Immutability and appendOnly [133](#)

- AFM DR (*continued*)
 - installing
 - AFM-based disaster recovery [643](#)
 - Limitations [128](#), [144](#)
 - planning [455](#)
 - role reversal [136](#)
 - RPO snapshots [134](#)
 - secondary cluster
 - NFS setup [455](#)
 - trucking
 - inband [148](#)
 - outband [148](#)
 - worker1Threads on primary cluster [455](#)
- AFM DR characteristics
 - failback [150](#)
 - failover [150](#)
 - trucking
 - inband [149](#)
 - outband [149](#)
- AFM gateway node [12](#), [453](#)
- AFM gateway nodes requirements on the cache and home clusters [453](#)
- AFM mode
 - operations [64](#)
- AFM node [12](#)
- AFM recovery improvements
 - remove operations [108](#)
 - rename operations [108](#)
- AFM to cloud object services replication
 - Limitations [184](#)
- AFM to cloud object storage [175](#), [179](#)
- afmAsyncPrefetchInterval parameter [126](#)
- afmFailoverMap [125](#)
- afmFastRemove parameter [121](#)
- afmRecoveryVer2 [108](#)
- afmSyncReadMount [104](#)
- AIX
 - electronic license agreement [576](#)
 - installation instructions for GPFS [575](#)
 - installing GPFS [576](#)
 - prerequisite software [575](#)
- AIX node
 - upgrade [674](#)
- allocation map
 - block [14](#)
 - inode [14](#)
 - logging of [15](#)
- allowing the GPFS administrative account to run as a service, Windows [593](#)
- Amazon Auto Scaling [186](#)
- Amazon CloudWatch [186](#)
- Amazon EBS [186](#)
- Amazon EC2 [186](#)
- Amazon IAM [186](#)
- Amazon S3 [186](#)
- Amazon VPC [186](#)
- Amazon web services
 - IBM Storage Scale [186–188](#), [450](#), [608](#), [612](#), [615](#), [619](#), [623](#), [625](#), [690](#)
- Amazon web services)IBM Storage Scale [625](#)
- antivirus software
 - Windows [584](#)
- API
 - configuring [633](#)

- API (*continued*)
 - fields parameter [201](#)
 - filter parameter [203](#)
 - installing [633](#)
 - paging [205](#)
- API nodes [219](#)
- application programs
 - communicating with GPFS [19](#)
- applying maintenance levels to GPFS [761](#)
- architecture, GPFS [9](#)
- assigning a static IP address
 - Windows [588](#)
- asynchronous
 - delay [64](#)
 - operations [64](#)
- asynchronous auto prefetch
 - AFM [126](#)
- asynchronous delay
 - AFM [64](#)
- atime value [353](#)
- audit logging
 - installation [546](#)
- authentication
 - basic concepts [381](#)
 - protocol user authentication [381](#)
- Authentication for IBM Storage Scale native REST API [220](#)
- authentication planning
 - file access [384](#)
 - protocols [381](#), [383](#)
- autoload attribute [327](#)
- Automatic installation [550](#)
- automatic mount
 - shared file system access [3](#)
- availability zone [619](#)
- Availability Zone
 - new Amazon VPC
 - multiple Availability Zone [615](#)
- Availability Zones [187](#)
- AWS
 - AFM
 - best practices [625](#)
 - AFM configuration best practices [625](#)
 - AFM limitations [625](#)
- AWS Cloud
 - AFM
 - configuration [625](#)
 - deployment [623](#)
 - limitations [625](#)
 - availability zones [187](#)
 - Data security [187](#)
 - Deployment [608](#)
 - IBM Storage Scaleupgrade [690](#)
 - Instance types [187](#)
 - Operating system [187](#)
 - regions [187](#)
 - usage restrictions [188](#)
- AWS CloudFormation [186](#)
- AWS Lambda function [186](#)

B

- Backup considerations [417](#)
- backup planning [362](#)
- bandwidth

- bandwidth (*continued*)
 - increasing aggregate [2](#)
- Base64 encoded [232](#)
- best practices [146](#), [147](#)
- bin directory [472](#)
- block
 - allocation map [14](#)
 - size [350](#)
- block allocation map [354](#)
- block placement [306](#)
- Bridge for Grafana installation [505](#), [550](#)
- bucket policy [47](#)

C

- cache [16](#)
- cache and home
 - AFM [56](#)
- cache cluster
 - AFM [55](#)
 - UID/GID requirements [455](#)
- cache eviction [94](#)
- cache fileset
 - AFM LU mode [123](#)
- cache synchronization
 - AFM [122](#)
- cache warming
 - AFM [122](#)
- cached
 - files [63](#)
- cached and uncached files
 - AFM [63](#)
- caching modes [56](#)
- call home
 - configuration [548](#)
 - event based uploads [264](#)
 - heartbeats [270](#)
 - installation [548](#)
 - mmcallhome [264](#), [270](#), [292](#)
 - monitoring IBM Storage Scale system remotely [264](#), [270](#), [292](#)
- Call home
 - Upgrade
 - v4.2.0. to v4.2.1 or later [703](#)
 - v4.2.1. to v5.0.0 [702](#)
- case sensitivity
 - Windows [583](#)
- CCR [27](#)
- CES
 - data uploads [188](#)
 - mmces command [188](#)
 - overview [31](#)
 - protocol [188](#)
 - SMB limitations
 - exports [394](#)
- CES HDFS support
 - planning [407](#)
- CES IPs [543](#)
- CES NFS [34](#), [35](#)
- CES NFS limitations [391](#)
- CES NFS Linux limitations [391](#)
- CES NFS support
 - overview [33](#)

- CES shared root file system [542](#)
- changing
 - secondary site [140](#)
- characteristics [146](#), [147](#)
- clean up
 - cloud services [668](#)
- clean up the environment [668](#)
- cleanup procedures for installation [660](#)
- cloud data sharing
 - creating a node class [627](#)
 - security considerations [416](#)
- Cloud data sharing
 - operation [243](#)
- cloud providers
 - cloud services [246](#)
 - transparent cloud tiering [246](#)
- cloud services
 - creating a ndoe class [627](#)
 - modifying [630](#)
 - planning [407](#)
 - security considerations [416](#)
 - software requirements [407](#)
 - upgrade [710](#)
 - upgrade to 1.1.2.1 [711](#)
 - upgrade to 1.1.3 [712](#)
 - upgrade to 1.1.4 [713](#), [714](#)
 - upgrade to 1.1.5 [715](#)
 - upgrade to 1.1.6 [717](#)
- cloudkit [441](#), [688](#)
- cloudkit on IBM Cloud [441](#)
- cloudkit upgrade cluster [688](#)
- cluster configuration data files
 - /var/mmfs/gen/mmsdrfs file [26](#)
 - content [26](#)
- Cluster Export Services
 - overview [31](#)
- cluster manager
 - description [10](#)
 - initialization of GPFS [20](#)
- cluster node
 - considerations [409](#)
- Clustered configuration repository [27](#)
- clustered watch folder
 - AFM [126](#)
 - events [257](#)
 - introduction [256](#)
 - JSON [258](#)
 - limitations [651](#)
 - remotely mounted file systems [652](#)
 - requirements [651](#)
- coexistence considerations [761](#)
- Coherency option [401](#)
- collectgpfsrtraceonhang [36](#)
- collecting details [261](#)
- collecting problem determination data [468](#)
- collectnfsclientperf [35](#)
- collectnfscoreonhang [36](#)
- collectnfsperf [35](#)
- commands
 - cloudkit [688](#)
 - description of GPFS commands [4](#)
 - failure of [322](#)
 - mmbackup [27](#)

- commands (*continued*)
 - mmcallhome [702, 703](#)
 - mmchcluster [323, 324](#)
 - mmchconfig [16, 18](#)
 - mmchdisk [23](#)
 - mmcheckquota [20](#)
 - mmchfs [345, 355, 357, 391](#)
 - mmcrcluster [18, 319, 323, 324, 575](#)
 - mmcrfs [345, 355, 357, 391](#)
 - mmcrnsd [329, 330](#)
 - mmdefedquota [358](#)
 - mmdefquotaon [358](#)
 - mmdelnsd [329, 330](#)
 - mmedquota [358](#)
 - mmfsck [15, 20, 24](#)
 - mmlsdisk [23, 341](#)
 - mmlsfs [15, 391](#)
 - mmlsquota [357](#)
 - mmmount [20](#)
 - mmrepquota [357](#)
 - mmstartup [327](#)
 - mmwinservctl [323, 324](#)
 - operating system [19](#)
 - processing [23, 24](#)
 - remote file copy
 - rcp [324](#)
 - scp [324](#)
 - remote shell
 - rsh [323](#)
 - ssh [323](#)
- communication
 - cache and home [56](#)
 - GPFS daemon to daemon [321](#)
 - invariant address requirement [309](#)
- comparison
 - live system backups [368](#)
 - snapshot based backups [368](#)
- compatibility considerations [761](#)
- concepts
 - AFM DR [133](#)
- configuration
 - files [26](#)
 - flexibility in your GPFS cluster [3](#)
 - of a GPFS cluster [319](#)
- configuration and tuning settings
 - configuration file [327](#)
 - default values [327](#)
 - GPFS files [4](#)
- configuration of protocol nodes [573](#)
- configuration tasks
 - SMB export limitations [394](#)
- configuring
 - api [633](#)
- configuring a mixed Windows and UNIX cluster [592](#)
- configuring a Windows HPC server [595](#)
- configuring GPFS [534](#)
- configuring quota
 - transparent cloud tiering [418](#)
- configuring system for transparent cloud tiering [410](#)
- configuring the GPFS Administration service, Windows [593](#)
- configuring Windows [588](#)
- considerations for backup
 - Transparent Cloud Tiering [417](#)
- considerations for GPFS applications

- considerations for GPFS applications (*continued*)
 - exceptions to Open Group technical standards [463](#)
 - NFSv4 ACL [463](#)
 - stat() system call [463](#)
- controlling the order in which file systems are mounted [361](#)
- conversion
 - mode [69](#)
- conversion of mode
 - AFM [69](#)
- CPU workload [49](#)
- created files (maximum number) [360](#)
- creating a node class
 - transparent cloud tiering [627](#)
- creating GPFS directory
 - /tmp/gpfs1pp on AIX nodes [576](#)
- creating the GPFS administrative account, Windows [593](#)
- creating the GPFS directory
 - /tmp/gpfs1pp on Linux nodes [473](#)

D

- daemon
 - communication [321](#)
 - description of the GPFS daemon [5](#)
 - memory [15](#)
 - quorum requirement [10](#)
 - starting [327](#)
- DASD for NSDs, preparing [341](#)
- data
 - availability [2](#)
 - consistency of [2](#)
 - guarding against failure of a path to a disk [317](#)
 - maximum replicas [356](#)
 - recoverability [313](#)
 - replication [355, 356](#)
- data blocks
 - logging of [15](#)
 - recovery of [15](#)
- Data Management API (DMAPI)
 - enabling [359](#)
- data privacy
 - methodical command [293](#)
- data protection
 - backing up data [191](#)
 - commands [193](#)
 - create and maintain snapshots [192](#)
 - data mirroring [192](#)
 - fileset backup [191](#)
 - restoring a file system from a snapshot [192](#)
 - restoring data [192](#)
- data recovery
 - commands [193](#)
- data tiering
 - overview [306](#)
 - planning [464](#)
- data upload
 - scheduled upload [270](#)
- default data replication [356](#)
- default metadata replication [355](#)
- default quotas
 - description [358](#)
 - files [15](#)

- deleting authentication [389](#)
- deleting ID mapping [389](#)
- deleting old sensors
 - ? [718](#)
- deploying protocols
 - existing cluster [567](#)
- deploying protocols on Linux nodes
 - procedure for [541](#), [544](#), [545](#), [550](#), [551](#), [723](#), [731](#)
- Deployment
 - existing VPC [619](#)
 - multiple Availability Zone [615](#)
 - new VPC
 - single Availability Zone [612](#)
- deployment considerations [146](#), [147](#), [151](#)
- deployment options
 - existing Amazon VPC [608](#)
 - new Amazon VPC
 - multiple availability zone [608](#)
 - single availability zone [608](#)
- df command (specifying whether it will report numbers based on quotas for the fileset) [359](#)
- diagnosing errors [570](#)
- diagnostic data [34](#)
- differences between GPFS and NTFS
 - Windows [585](#)
- direct access storage devices (DASD) for NSDs, preparing [341](#)
- directory deletion
 - AFM [121](#)
- directory, bin [472](#)
- Disable gpfs leases [399](#)
- disable leases [399](#)
- disabling protocols
 - authentication considerations [383](#)
- disabling the Windows firewall [588](#)
- disaster recovery
 - data mirroring [191](#)
 - protocol cluster disaster recovery [191](#)
 - SOBAR [191](#)
 - use of GPFS replication and failure groups [3](#)
- disk descriptor replica [341](#)
- disk usage
 - verifying [359](#)
- disks
 - considerations [328](#)
 - failure [316](#)
 - file system descriptor [13](#)
 - media failure [24](#)
 - mmcrfs command [349](#)
 - recovery [23](#)
 - releasing blocks [24](#)
 - stanza files [330](#)
 - state of [23](#)
 - storage area network [328](#)
 - stripe group [13](#)
- DMAPI
 - coexistence considerations [761](#)
 - considerations for IBM Storage Protect for Space Management [761](#)
- DMAPI file handle size considerations
 - for IBM Storage Protect for Space Management [761](#)
- documentation
 - installing man pages on AIX nodes [577](#)
 - installing man pages on Linux nodes [476](#)

- domain
 - Active Directory [588](#)
 - Windows [588](#)

E

- ECKD devices, preparing environment for [342](#)
- electronic license agreement
 - AIX nodes [576](#)
 - Linux nodes [473](#)
- Enable gpfs leases [399](#)
- enable leases [399](#)
- enable share modes [400](#)
- enabling
 - protocols [493](#)
- enabling file system features [359](#)
- enabling protocols
 - authentication considerations [383](#)
 - existing cluster [570](#)
- environment for ECKD devices, preparing [342](#)
- environment, preparing [472](#)
- ESS
 - adding protocol nodes [571](#)
 - deploying protocols [571](#)
- estimated node count [357](#)
- event based uploads
 - FTDC2CallHome [264](#)
- events
 - file audit logging [252](#)
- exceptions
 - CES NFS Linux [391](#)
- existing Amazon VPC [619](#)
- expiring a disconnected RO cache [100](#)
- explanations and examples, installing GPFS and protocols [532](#), [533](#)
- Express Edition
 - migrating to Standard [752](#)
- Extended attributes (EAs)
 - planning [376](#)
- external Kafka sink [256](#)
- extracting GPFS patches [476](#)
- extracting the IBM Storage Scale software [473](#)
- extracting update SUSE Linux Enterprise Server and Red Hat Enterprise Linux RPMs [476](#)
- extraction, packaging overview and [473](#)

F

- Fail-over scenarios [394](#)
- failback
 - new primary site [139](#)
 - old primary site [138](#)
- failed fresh upgrade
 - upgrade rerun [741](#)
- failover
 - secondary site [138](#)
- failover in AFM [101](#)
- failure
 - disk [316](#)
 - Network Shared Disk server [316](#)
 - node [313](#), [314](#), [316](#)
- failure groups
 - definition of [2](#)

- failure groups (*continued*)
 - loss of [341](#)
 - preventing loss of data access [316](#)
 - use of [340](#)
- fast tier pool
 - requirements [464](#)
- features
 - AFM DR [133](#)
- Fields [230](#)
- file audit logging
 - attributes [252](#)
 - events [252](#)
 - fileset [251](#)
 - installing [647](#)
 - limitations [647](#)
 - producers [250](#)
 - records [251](#)
 - requirements [647](#)
- file authentication [384](#)
- File client limitations [394](#)
- file name considerations
 - Windows [583](#)
- file system creation [354](#)
- file system descriptor
 - failure groups [340](#)
 - inaccessible [341](#)
 - quorum [340](#)
- file system features
 - enabling [359](#)
- file system level backups [369](#)
- file system manager
 - command processing [23](#), [24](#)
 - description [10](#)
 - internal log file [355](#)
 - mount of a file system [20](#)
 - NSD creation considerations [339](#)
 - quota management function [11](#)
 - selection of [12](#)
 - token management function [11](#)
 - Windows drive letter [357](#)
- file system name considerations
 - Windows [582](#)
- file systems
 - administrative state of [4](#), [26](#)
 - AFM node [12](#)
 - authorization [354](#)
 - block size [350](#)
 - creating [345](#)
 - descriptor [13](#)
 - device name [349](#)
 - disk descriptor [349](#)
 - enabling DMAPI [359](#)
 - interacting with a GPFS file system [19](#)
 - internal log file [355](#)
 - last time accessed [353](#)
 - list of disk descriptors [355](#)
 - maximum number of [14](#)
 - maximum number of files [360](#)
 - maximum number of mounted files [14](#)
 - maximum size [14](#)
 - maximum size supported [14](#)
 - metadata [13](#)
 - metadata integrity [12](#)
 - mount options [357](#)

- file systems (*continued*)
 - mounting [3](#), [20](#), [350](#)
 - mountpoint [357](#)
 - NFS export [391](#)
 - NFSv4 export [391](#)
 - number of nodes mounted by [357](#)
 - opening a file [20](#)
 - quotas [358](#)
 - reading a file [21](#)
 - recoverability parameters [355](#), [356](#)
 - repairing [24](#)
 - sample creation [361](#)
 - shared access among clusters [1](#)
 - simultaneous access [2](#)
 - sizing [345](#)
 - stripe group [13](#)
 - time last modified [354](#)
 - Windows drive letter [357](#)
 - writing to a file [22](#), [23](#)
- file systems (controlling the order in which they are mounted) [361](#)
- files
 - .toc [577](#)
 - /etc/filesystems [26](#), [27](#)
 - /etc/fstab [26](#), [27](#)
 - /var/mmfs/etc/mmfs.cfg [27](#)
 - /var/mmfs/gen/mmsdrfs [26](#), [27](#)
 - consistency of data [2](#)
 - GPFS recovery logs [15](#)
 - inode [14](#)
 - installation on AIX nodes [575](#)
 - maximum number of [14](#), [360](#)
 - mmfslinux [6](#)
 - structure within GPFS [13](#)
- files that can be created, maximum number of [360](#)
- fileset backups [369](#)
- fileset disabling
 - AFM [114](#)
- fileset snapshots
 - writable [123](#)
- fileset to home
 - AFM [68](#)
- filesetdf option [359](#)
- filesets
 - AFM
 - asynchronous prefetch [126](#)
- firewall
 - Windows, disabling [588](#)
- firewall recommendations
 - transparent cloud tiering [412](#)
- First time data capture [34](#)
- Force flushing contents before Async Delay [77](#)
- fresh transparent cloud tiering cluster
 - setting up [668](#)
- fresh upgrade failure
 - solution [741](#)
- FTDC [34](#)
- Functional overview
 - REST API [198](#)

G

- Ganesha core dump [36](#)

Ganesha limitations [391](#)

ganesha_stats [35](#)

gateway

node failure [91](#)

recovery [91](#)

GDS [28](#), [343](#), [639](#), [692](#)

global namespace

AFM [61](#)

Google Cloud Platform [179](#)

GPFS

adding file systems [568](#)

adding LTFS nodes [572](#)

adding nodes [568](#), [571](#), [572](#)

adding nodes with toolkit [567](#)

adding NSDs [568](#)

adding protocol nodes [571](#)

adding protocols with toolkit [567](#)

administration commands [5](#)

AFM gateway node [12](#)

application interaction [19–24](#)

architecture [9](#), [10](#), [13](#), [15](#), [17](#), [19](#)

backup data [27](#)

basic structure [5](#), [6](#)

CES [31](#)

CES NFS support [33](#)

CES node [12](#)

cluster configuration data files [26](#)

cluster configurations [6](#)

cluster creation [319–327](#)

Cluster Export Services [31](#)

cluster manager [10](#)

considerations for applications [463](#)

daemon [5](#)

daemon communication [17](#), [18](#)

deploying protocols [541](#), [544](#), [545](#), [550–552](#), [567](#), [723](#), [725](#), [731](#)

diagnosing errors [570](#)

disk considerations [329](#), [339–342](#)

disk storage use [13](#)

enabling protocols [570](#)

enabling protocols on Linux [493](#)

Express Edition [752](#)

failure recovery processing [26](#)

file consistency [2](#)

file structure [13](#), [15](#)

file system creation [345](#), [349](#), [350](#), [353–361](#)

file system manager [10](#)

for improved performance [2](#)

hardware requirements [309](#)

increased data availability [2](#)

installing [467](#), [469](#), [470](#), [472](#), [473](#), [475–477](#), [483–485](#), [493–495](#), [497](#), [509](#), [516](#), [518](#), [532–534](#), [539](#), [541](#), [544](#), [545](#), [550–552](#), [573](#), [575–588](#), [590–592](#), [595](#), [628](#), [706](#), [723](#), [731](#)

installing CES on Linux [491](#)

installing on AIX nodes [575–578](#)

installing on Linux [486](#)

installing on Linux nodes [469](#), [470](#), [472](#), [473](#), [476](#), [477](#), [493–495](#), [497](#), [509](#), [516](#), [532–534](#), [539](#), [573](#), [706](#)

installing on Windows nodes [579–588](#), [590–592](#), [595](#)

installing packages on Linux [487](#), [490](#)

installing protocols on Linux [491](#)

installing with toolkit [509](#)

IP address usage [18](#)

GPFS (*continued*)

kernel extensions [5](#)

management functions [10](#), [12](#)

memory usage [15](#), [16](#)

metanode [12](#)

migration [671](#), [758–761](#)

network communication [17](#), [19](#)

NSD disk discovery [24](#)

OS upgrade [762](#), [767](#)

pinned memory [16](#)

planning [300](#), [309](#), [310](#), [313](#), [316](#), [318](#), [319](#), [345](#), [362](#), [363](#), [365](#), [368](#), [369](#), [372](#), [373](#), [380](#), [381](#), [383](#), [384](#), [389](#), [390](#), [403](#)

planning for transparent cloud tiering [408](#), [409](#)

portability layer [6](#)

product editions [298](#), [304](#)

product overview [240](#)

protocol node [12](#)

protocols support [30](#), [31](#), [33](#), [37](#), [52](#)

quota files [15](#)

recoverability [313](#), [314](#), [316](#), [318](#)

recovery logs [15](#)

RHEL 8.x upgrade [764](#)

servicing protocol nodes [768](#)

simplified administration [4](#)

simplified storage management [3](#)

SMB support [37](#)

software requirements [310](#)

special management functions [10](#)

strengths [1–4](#)

stretch cluster use case [556](#)

Swift Object storage support [52](#)

system flexibility [3](#)

thin provisioned disks [330](#)

uninstalling [659](#), [660](#), [666](#)

upgrade [674](#), [676](#), [677](#), [682](#), [753](#)

upgrading [671–674](#), [676](#), [677](#), [682](#), [695](#), [697](#), [699](#), [700](#), [704](#), [705](#), [725](#), [753](#), [758–761](#)

upgrading from 3.5 [769](#)

user interaction [19–24](#)

GPFS administration commands [17](#)

GPFS administrative adapter port name [320](#)

GPFS architecture [9](#)

GPFS clusters

administration adapter port name [320](#)

configuration data files [4](#)

configuration file [327](#)

configuration servers [322](#)

creating [319](#)

daemon

starting [327](#)

daemon communication [17](#)

establishing [467](#)

introduction [1](#)

naming [325](#)

nodes in the cluster [321–327](#)

operating environment [6](#)

planning nodes [319](#)

portability layer [483–485](#)

recovery logs

creation of [15](#)

unavailable [24](#)

server nodes [319](#)

starting [467](#)

- GPFS clusters (*continued*)
 - starting the GPFS daemon [10](#), [327](#)
 - user ID domain [326](#)
- GPFS communications adapter port name [320](#)
- GPFS daemon communications [17](#)
- GPFS for Linux on Z, running [477](#)
- GPFS for Windows Multiplatform
 - overview [579](#)
- GPFS introduction [1](#)
- GPFS limitations on Windows [581](#)
- GPFS Metrics Plugin [305](#)
- GPFS patches, extracting [476](#)
- GPFS product structure
 - capacity based licensing [304](#)
- GPFS strengths [1](#)
- GPFS trace [36](#)
- GPFS, configuring [534](#)
- GPFS, installing [516](#)
- GPFS, installing over a network [577](#)
- GPFS, planning for [309](#)
- GPFS, reverting to a previous level [759](#), [760](#)
- GPG key
 - verifying [475](#)
- GPU [28](#)
- GPU workload [49](#)
- GPUDirect RDMA [49](#)
- GPUDirect Storage
 - installing [639](#)
 - planning [343](#)
 - upgrading [692](#)
- Grafana
 - with Installation Toolkit [550](#)
 - with RPM [505](#)
- GSKit [1](#)
- GUI
 - overview [194](#)
- GUI without root privileges [505](#)
- guidelines for maintenance tasks [416](#)

H

- hang scenario [36](#)
- hard limit, quotas [358](#)
- hardware requirements [309](#)
- Hardware requirements
 - transparent cloud tiering [407](#)
- heartbeats
 - call home events [270](#)
 - mmhealth command [270](#)
- home
 - AFM [55](#)
- home cluster
 - UID/GID requirements [455](#)
- home file system
 - snapshot [123](#)
- How WORM works [246](#)
- HPC server (Windows), configuring [595](#)

I

- IAM users [47](#)
- IBM Cloud

- IBM Cloud (*continued*)
 - DNS [441](#)
 - VPC [441](#)
- IBM Cloud Object Storage migration
 - transparent cloud tiering [722](#)
- IBM Cloud Object Storage considerations [410](#)
- IBM Global Security Kit (GSKit) [1](#)
- IBM Spectrum Archive
 - adding nodes [572](#)
- IBM Spectrum Scale
 - AFM
 - .pconflicts [69](#)
- IBM Storage Protect
 - backup planning [362](#), [363](#), [365](#), [368](#), [369](#), [372](#), [373](#)
 - file data storage [363](#)
 - fileset backups [372](#)
 - identify backup candidates [365](#)
 - metadata storage [363](#)
 - provisioning [362](#)
- IBM Storage Protect backup planning [362](#), [365](#)
- IBM Storage Protect for Space Management
 - DMAPI file handle size considerations [761](#)
 - file system backups [373](#)
- IBM Storage Protect space managed file system backups [373](#)
- IBM Storage Scale
 - 6.0.0 [674](#)
 - 6.0.0.0 [674](#)
 - 6.0.0.1 [674](#)
 - 6.0.0.2 [674](#)
 - 6.0.1.0 [676](#)
 - 6.0.1.x [676](#)
 - Active File Management [54–57](#), [59](#), [61–64](#), [68](#), [69](#), [74](#), [77](#), [80–82](#), [90](#), [91](#), [94](#), [98](#), [100](#), [101](#), [103](#), [108](#), [111–114](#), [154](#), [168](#), [452](#), [453](#), [641](#), [692](#)
 - Active File Management - Disaster Recovery [455](#)
 - Active File Management DR [130](#), [132–134](#), [138–140](#), [455](#), [643](#)
 - adding file systems [568](#)
 - adding nodes [568](#), [571](#), [572](#)
 - adding nodes with toolkit [567](#)
 - adding NSDs [568](#)
 - adding protocols with toolkit [567](#)
 - administration commands [5](#)
 - AFM
 - .afm [69](#)
 - .ptrash [69](#)
 - AFM gateway nodes requirements on the cache and home clusters [453](#)
 - AFM to Cloud Object Server replication
 - Eviction feature [168](#)
 - AFM to cloud object storage [154](#)
 - asynchronous delay [64](#)
 - cache cluster [55](#)
 - cache eviction [94](#)
 - cached and uncached files [63](#)
 - caching modes [56](#)
 - encryption [113](#)
 - expiring a disconnected RO cache [100](#)
 - Failover [101](#)
 - Fast creates [74](#)
 - fileset disabling [114](#)
 - fileset to home [68](#)
 - Force flushing contents before Async Delay [77](#)

IBM Storage Scale (continued)

AFM (continued)

- gateway node [453](#)
- gateway node failure [91](#)
- gateway recovery [91](#)
- global namespace [61](#)
- home [55](#)
- IBM Storage Protect [108](#)
- Inode limits to set at cache and home [452](#)
- install [641](#)
- iw cache maintenance [112](#)
- NFS backend protocol [57](#)
- operation with disconnected home [98](#)
- Parallel data transfer using multiple remote mounts [80](#)
- Parallel data transfers [77](#)
- partial file caching [81](#)
- peer snapshot [90](#)
- planned maintenance [111](#)
- planning [452](#)
- prefetch [82](#)
- primary gateway [59](#)
- psnap [90](#)
- resync on SW filesets [103](#)
- revalidation parameters [62](#)
- synchronous or asynchronous operations [64](#)
- UID and GID requirements on the cache and home clusters [452](#)
- unplanned maintenance [112](#)
- upgrade [692](#)
- using mmbbackup [114](#)
- viewing snapshots at home [100](#)
- WAN latency [54](#)
- workerThreads on cache cluster [452](#)

AFM communication

- cache and home [56](#)

AFM DR

- concepts [133](#)
- failback [138](#), [139](#)
- failover [138](#)
- installing [643](#)
- introduction [130](#)
- modes [133](#)
- planning [455](#)
- recovery time objective [132](#)
- role reversal [136](#)
- RPO snapshots [134](#)
- secondary cluster [455](#)
- secondary site [140](#)
- UID/GID requirements [455](#)
- worker1Threads on primary cluster [455](#)

AFM inode limits [452](#)

AFM mode

- operations [64](#)

AFM monitoring

- cache and home [56](#)

AFM-based Asynchronous Disaster Recovery

- role reversal [136](#)

- application interaction [19–24](#)

- architecture [9](#), [10](#), [13](#), [15](#), [17](#), [19](#)

- backup data [27](#)

- basic structure [4–6](#)

- call home [263](#), [270](#), [293](#), [294](#), [645](#)

- call home data upload [263](#), [270](#), [293](#)

IBM Storage Scale (continued)

- call home update PTF [294](#)

- capacity based licensing [304](#)

- CES [31](#)

- CES HDFS planning [407](#)

- CES NFS support [33](#)

- cloudkit [441](#)

- cluster configuration data files [26](#)

- cluster configurations [6](#)

- cluster creation [319–327](#)

- Cluster Export Services [31](#)

- cluster manager [10](#)

- configuring and tuning system

- transparent cloud tiering [410](#)

- considerations for applications [463](#)

- daemon [5](#)

- daemon communication [17](#), [18](#)

- data protection [193](#)

- data tiering

- overview [306](#)

- planning [464](#)

- deploying protocols [541](#), [544](#), [545](#), [550–552](#), [567](#), [723](#), [731](#)

- diagnosing errors [570](#)

- disk considerations [328](#), [329](#), [339–342](#)

- enabling protocols [570](#)

- enabling protocols on Linux [493](#)

- Express Edition [752](#)

- failure recovery processing [26](#)

- file audit logging [647](#)

- file consistency [2](#)

- file structure [13](#), [15](#)

- file system creation [345](#), [349](#), [350](#), [353–361](#)

- for improved performance [2](#)

- GUI [194](#)

- hardware requirements [309](#)

- Immutability and appendOnly

- IAM modes [133](#)

- increased data availability [2](#)

- inspect call home [289](#)

- inspect call home data upload [289](#)

- installation toolkit [526](#), [528](#), [546](#), [548](#), [555](#)

- installing [467](#), [469](#), [470](#), [472](#), [473](#), [475–477](#), [483–485](#), [493–495](#), [497](#), [509](#), [516](#), [518](#), [532–534](#), [539](#), [541](#), [544](#), [545](#), [550–552](#), [573](#), [575–588](#), [590–592](#), [595](#), [628](#), [706](#), [723](#), [731](#)

- installing call home [645](#)

- installing CES on Linux [491](#)

- installing GUI [497](#)

- installing on AIX nodes [575–578](#)

- installing on Linux [486](#)

- installing on Linux nodes [469](#), [470](#), [472](#), [473](#), [475–477](#), [493–495](#), [497](#), [509](#), [516](#), [532–534](#), [539](#), [573](#), [706](#)

- installing on Windows nodes [579–588](#), [590–592](#), [595](#)

- installing packages on Linux [487](#), [490](#)

- installing protocols on Linux [491](#)

- installing with toolkit [509](#)

- IP address usage [18](#)

- kernel extensions [5](#)

- key features [1](#)

- licensing [298](#)

- management functions [10](#), [12](#)

- management GUI

- manual installation [497](#)

IBM Storage Scale (continued)

- management GUI (continued)
 - node classes [667](#)
- manual installation [477](#), [493–495](#), [706](#)
- manual upgrade [704](#), [705](#)
- memory usage [15](#), [16](#)
- migration [671](#), [695](#), [697](#), [699](#), [700](#), [758–761](#)
- network communication [17](#), [19](#)
- node failure [313](#), [314](#), [316](#)
- non-protocol nodes [677](#)
- NSD disk discovery [24](#)
- offline upgrade [733](#)
- OpenStack cloud deployment [295](#)
- OS upgrade [762](#), [767](#)
- pinned memory [16](#)
- planning [300](#), [309](#), [310](#), [313](#), [316](#), [318](#), [319](#), [328](#), [345](#), [362](#), [363](#), [365](#), [368](#), [369](#), [372](#), [373](#), [380](#), [381](#), [383](#), [384](#), [389](#), [390](#), [403](#), [413](#)
- planning for transparent cloud tiering [408](#), [409](#)
- portability layer [6](#)
- prerequisites for call home [645](#)
- product edition [747–750](#)
- product editions [298](#), [304](#)
- product overview [240](#)
- product structure [298](#)
- protocol nodes [682](#)
- protocol nodes authentication [744](#)
- protocols prerequisites [470](#)
- protocols support [30](#), [31](#), [33](#), [37](#), [52](#), [407](#)
- recoverability [313](#), [314](#), [316](#), [318](#)
- RHEL 8.x upgrade [764](#)
- SELinux
 - enforcing mode [463](#)
 - permissive mode [463](#)
- servicing protocol nodes [768](#)
- shared file system access [1](#)
- simplified administration [4](#)
- simplified storage management [3](#)
- SMB support [37](#)
- software requirements [310](#)
- special management functions [10](#)
- strengths [1–4](#)
- stretch cluster use case [556](#)
- supported cloud providers [246](#)
- supported NFS clients [391](#)
- supported upgrade paths [672](#)
- Swift Object storage support [52](#)
- system flexibility [3](#)
- thin provisioned disks [330](#)
- transparent cloud tiering
 - how it works [242](#)
 - overview [240](#)
- uninstalling [659](#), [660](#), [666](#)
- upgrade
 - non-protocol nodes [677](#)
 - protocol nodes [682](#)
- upgrade considerations [673](#)
- upgrade process flow [725](#)
- upgrading
 - non-protocol nodes [677](#)
 - protocol nodes [682](#)
 - protocol nodes authentication [744](#)
- Upgrading [674](#), [676](#), [677](#), [682](#), [744](#)
- upgrading from 3.5 [769](#)

IBM Storage Scale (continued)

- upgrading LTFS nodes [742](#)
- user interaction [19–24](#)
- WORM solutions [246](#)
- IBM Storage Scale AFM node [12](#)
- IBM Storage Scale Bridge
 - installing [505](#), [550](#)
- IBM Storage Scale CES node [12](#)
- IBM Storage Scale Client license [300](#)
- IBM Storage Scale deployment
 - cloudkit [441](#)
- IBM Storage Scale disk storage use [13](#)
- IBM Storage Scale file system
 - migrating SMB data [397](#)
- IBM Storage Scale file system manager [10](#)
- IBM Storage Scale for object storage
 - upgrade to 5.0.x [695](#)
 - upgrading [695](#)
- IBM Storage Scale FPO license [300](#)
- IBM Storage Scale GUI
 - upgrading to the latest version [708](#)
- IBM Storage Scale information units [xix](#)
- IBM Storage Scale introduction [1](#)
- IBM Storage Scale license designation [300](#)
- IBM Storage Scale management API
 - configuring [633](#)
 - fields parameter [201](#)
 - filter parameter [203](#)
 - Functional overview [198](#)
 - installing [633](#)
 - paging [205](#)
- IBM Storage Scale metanode [12](#)
- IBM Storage Scale NFS
 - upgrading [699](#)
- IBM Storage Scale on AWS [187](#), [450](#), [615](#)
- IBM Storage Scale overview [1](#)
- IBM Storage Scale protocol node [12](#)
- IBM Storage Scale protocols
 - planning [380](#), [381](#), [383](#), [384](#), [389](#), [390](#)
- IBM Storage Scale quota files [15](#)
- IBM Storage Scale recovery logs [15](#)
- IBM Storage Scale S3
 - upgrading [700](#)
- IBM Storage Scale Server license [300](#)
- IBM Storage Scale Swift Object storage
 - overview [52](#)
- IBM Storage Scale System 3000 awareness [555](#)
- IBM Storage Scale System awareness [555](#)
- IBM Storage Scale on AWS [186–188](#), [608](#), [612](#), [619](#), [623](#), [625](#), [690](#)
- IBM Z, DASD tested with Linux on [341](#)
- IBM Z, running GPFS for Linux on [477](#)
- Identity and Access Management [47](#)
- IDMU (Identity Mapping for Unix (IDMU) / RFC 2307 Attributes) [592](#)
- Independent writer
 - AFM mode [64](#)
- independent-writer (IW) [56](#)
- indirect blocks [13](#), [15](#)
- indirection level [13](#)
- InfiniBand [49](#)
- initialization of the GPFS daemon [20](#)
- inode
 - allocation file [14](#)

- inode (*continued*)
 - allocation map [14](#)
 - cache [16](#)
 - logging of [15](#)
 - usage [13](#), [23](#)
 - Inode limits to set at cache and home [452](#)
 - install
 - query
 - uninstall [628](#)
 - Install toolkit [635](#)
 - installation cleanup procedures [660](#)
 - installation toolkit
 - audit logging [546](#)
 - call home configuration [548](#)
 - IBM Storage Scale System 3000 awareness [555](#)
 - IBM Storage Scale System awareness [555](#)
 - limitations [518](#), [552](#)
 - mixed operating system support [526](#)
 - populating cluster definition file [552](#)
 - preparing [528](#)
 - product edition change [747–750](#)
 - stretch cluster [556](#)
 - upgrade flow chart [725](#)
 - upgrade process flow [725](#)
 - upgrading LTFS nodes [742](#)
 - installation toolkit(
 - offline upgrade [733](#)
 - installing
 - AFM-based disaster recovery [643](#)
 - api [633](#)
 - CES [491](#)
 - cloud data sharing [628](#)
 - cloud services [628](#)
 - clustered watch folder [651](#)
 - GPUDirect Storage [639](#)
 - manually [486](#)
 - transparent cloud tiering
 - on GPFS nodes [628](#)
 - installing and configuring OpenSSH, Windows [594](#)
 - installing Cygwin [588](#)
 - installing GPFS
 - on Windows nodes [590](#), [591](#)
 - over a network [577](#)
 - installing GPFS (Ubuntu)
 - verifying the GPFS installation [493](#)
 - installing GPFS and protocols, examples [532](#), [533](#), [541](#), [545](#)
 - installing GPFS on AIX nodes
 - creating the GPFS directory [576](#)
 - directions [577](#)
 - electronic license agreement [576](#)
 - files used during [575](#)
 - man pages [577](#)
 - procedures for [576](#)
 - table of contents file [577](#)
 - verifying the GPFS installation [578](#)
 - what to do before you install GPFS [575](#)
 - installing GPFS on Linux nodes
 - building the GPFS portability layer
 - using the Autoconfig tool [485](#)
 - using the mmbuildgpl command [484](#), [485](#)
 - electronic license agreement [473](#)
 - installing the software packages [477](#), [493–495](#), [497](#), [706](#)
 - man pages [476](#)
 - installing GPFS on Linux nodes (*continued*)
 - procedure for [469](#), [470](#), [473](#), [476](#), [477](#), [493–495](#), [497](#), [509](#), [516](#), [532–534](#), [539](#), [573](#), [706](#)
 - verifying the GPFS installation [494](#)
 - what to do before you install GPFS [469](#)
 - installing GPFS on Windows nodes [579](#)
 - installing GPFS prerequisites [587](#)
 - Installing IBM Storage Scale Bridge for Grafana with the installation toolkit [550](#)
 - installing IBM Storage Scale on Linux nodes
 - creating the GPFS directory [473](#)
 - License Acceptance Process (LAP) tool [473](#)
 - installing packages
 - manually [487](#)
 - NSDs [490](#)
 - with toolkit [509](#)
 - installing protocols [469](#), [470](#), [473](#), [509](#), [550](#)
 - installing Tracefmt [588](#)
 - installing Tracelog [588](#)
 - Instance types [187](#)
 - inter-protocol level locking [400](#)
 - Internet Protocol Version 6 [119](#), [120](#)
 - interoperability of
 - transparent cloud tiering [247](#)
 - introduction
 - AFM DR [130](#)
 - file audit logging [250](#)
 - transparent cloud tiering [240](#)
 - invariant address adapter
 - requirement [309](#)
 - IO Throughput [399](#)
 - IOPS [399](#)
 - IP address
 - private [18](#)
 - public [18](#)
 - iptables rules
 - migration [504](#)
- ## J
- joining an Active Directory domain [588](#)
 - JSON
 - attributes [252](#)
 - events [252](#)
 - JSON attributes [258](#)
- ## K
- Kafka
 - external sink [256](#)
 - kernel extensions [5](#)
 - kernel memory [15](#)
- ## L
- latest level of file system
 - migrating [359](#)
 - LDAP [220](#)
 - librdkafka [722](#)
 - License Acceptance Process (LAP) tool [473](#)
 - license designation [300](#)
 - limitations
 - CES NFS Linux [391](#)

- limitations (*continued*)
 - clustered watch folder [651](#)
 - NFS protocol nodes [391](#)
- Limitations
 - SMB clients limitations [394](#)
 - SMB share limitations [394](#)
- limitations of
 - transparent cloud tiering [247](#)
- limitations of cloudkit
 - IBM Cloud [441](#)
- Limitations-Native-Rest-API [238](#)
- Linux
 - building the GPFS portability layer
 - using the mmbuildgpl command [484](#), [485](#)
 - installation instructions for GPFS [469](#)
 - installing GPFS [473](#)
 - kernel requirement [310](#)
 - prerequisite software [472](#)
- Linux on Z, DASD tested with [341](#)
- Linux on Z, running GPFS [477](#)
- load balancing across disks [2](#)
- Local updates
 - AFM mode [64](#)
- local-update (LU) [56](#)
- Long-running operations [225](#)
- LRO [225](#)

M

- maintenance
 - iw cache [112](#)
- maintenance levels of GPFS, applying [761](#)
- maintenance tasks
 - planning [416](#)
- man pages
 - installing on AIX nodes [577](#)
 - installing on Linux nodes [476](#)
- management API
 - fields parameter [201](#)
 - filter parameter [203](#)
 - paging [205](#)
- management GUI
 - installing [550](#)
 - node classes [667](#)
 - nodeclasses [497](#)
 - uninstalling [666](#)
- management GUI node classes
 - removing nodes [667](#)
- Manual install [635](#)
- manual installation [477](#), [486](#), [487](#), [490](#), [491](#), [493–495](#), [706](#)
- Manual installation [505](#)
- manual upgrade
 - performance monitoring
 - Ubuntu 18.04 [705](#)
- Manually installing IBM Storage Scale Bridge for Grafana with the RPM [505](#)
- maxFilesToCache parameter
 - memory usage [16](#)
- maximum data replicas [356](#)
- maximum metadata replicas [356](#)
- maximum number of files [360](#)
- maximum number of files that can be created [360](#)
- maxStatCache parameter
 - memory usage [16](#)

- memory
 - non-pinned [16](#)
 - pinned [16](#)
 - usage [15](#)
- metadata
 - default [355](#)
 - maximum replicas [356](#)
 - replication [355](#)
- metanode [12](#)
- migrating
 - reverting to the previous level of GPFS [758–760](#)
- migrating file system format to the latest level [359](#)
- migrating IBM Cloud Object Storage [722](#)
- Migrating SMB data from traditional filer [397](#)
- migrating SMB data from traditional NAS [397](#)
- migrating SMB data to IBM Storage Scale [397](#)
- mixed Windows and UNIX cluster, configuring [592](#)
- mmapfctl
 - outbandupdate [122](#)
- mmbackup command [27](#), [193](#)
- mmbackupconfig command [193](#)
- mmcesdr command [193](#)
- mmchcluster command [323](#), [324](#)
- mmchconfig command
 - defining a subset of nodes [5](#), [19](#)
- mmchdisk command [23](#)
- mmcheckquota command [20](#)
- mmchfs [391](#)
- mmchfs command [345](#), [355](#), [357](#)
- mmcrcluster command [18](#), [319](#), [323](#), [324](#), [575](#)
- mmcrfs [391](#)
- mmcrfs command [345](#), [355](#), [357](#)
- mmcrnsd command [329](#)
- mmdefedquota command [358](#)
- mmdefquotaon command [358](#)
- mmdeinsd command [329](#)
- mmquota command [358](#)
- mmfsck command [15](#), [20](#), [24](#), [193](#)
- mmimagebackup command [193](#)
- mmimgrestore command [193](#)
- mmlsdisk command [23](#), [341](#)
- mmlsfs [391](#)
- mmlsfs command [15](#)
- mmlsquota command [357](#)
- mmmount command [20](#)
- mmrepquota command [357](#)
- mmrestoreconfig command [193](#)
- mmstorefs command [193](#)
- mmstartup command [327](#)
- mmwinservctl command [323](#), [324](#)
- modes
 - AFM DR [133](#)
- monitoring
 - cache and home [56](#)
- mount options [357](#)
- mount-priority option [361](#)
- mounting a file system [20](#), [350](#), [357](#)
- mounting of file systems, controlling the order of the [361](#)
- mountpoint [357](#)
- moving SMB data from a traditional filer [397](#)
- mtime values [354](#)
- multi-node cloud services
 - set up [629](#)
- multi-node setup for cloud services [629](#)

Multiple Path I/O (MPIO)
utilizing [317](#)

N

NAS protocols [399](#)
Native REST API [458](#), [635](#)
Native Scale API [216](#)
network
 communication within your cluster [2](#)
network communication
 administration commands [19](#)
network considerations [408](#)
Network File System (NFS)
 access control lists [354](#)
 deny-write open lock
 [349](#)
network installing GPFS [577](#)
Network Shared Disk (NSD)
 creation of [329](#)
 disk discovery [24](#)
 server disk considerations [328](#)
 server failure [316](#)
 server node considerations [339](#)
new Amazon VPC
 multiple Availability Zone [615](#)
 single Availability Zone [612](#)
NFS
 backend protocol [57](#)
NFS backend [59](#)
NFS planning
 file system considerations [390](#)
NFS ports [59](#)
NFS statistics collection [35](#)
NFS target mapping [125](#)
NFS unresponsive [36](#)
NFS-Ganesha [34](#), [35](#)
NFS-supported clients [391](#)
nfsdebug.snap [36](#)
node quorum
 definition of [314](#)
 selecting nodes [316](#)
node quorum with tiebreaker disks
 definition of [314](#)
 selecting nodes [316](#)
nodes
 acting as special managers [10](#)
 cluster manager [10](#)
 descriptor form [321](#)
 designation as manager or client [321](#)
 estimating the number of [357](#)
 failure [313](#), [314](#), [316](#)
 file of nodes in the cluster for installation [575](#)
 file system manager [10](#)
 file system manager selection [12](#)
 in a GPFS cluster [319](#)
 in the GPFS cluster [321](#)
 quorum [321](#)
nofilesetdf option [359](#)
Non API nodes [219](#)
non-pinned memory [16](#)
non-standard ports [59](#)
NSD
 backend protocol [57](#)

number of files that can be created, maximum [360](#)
NVIDIA
 Nsight Systems [305](#)

O

Object
 upgrade to 5.0.x [695](#)
object storage [49](#)
OIDC [220](#)
Open Secure Socket Layer [1](#)
OpenSSL [1](#)
operating system
 calls [20–23](#)
 commands [19](#), [20](#)
Operating system [187](#)
opportunistic locks [399](#)
order in which file systems are mounted, controlling the [361](#)
outbandupdate command [122](#)
overview
 of GPFS for Windows Multiplatform [579](#)
 transparent cloud tiering [240](#)

P

packaging overview and extraction [473](#)
pagepool parameter
 affect on performance [22](#)
 in support of I/O [16](#)
 memory usage [16](#)
Pagination [228](#)
PAM [220](#)
Parallel data transfer [121](#)
Parallel data transfer using multiple remote mounts [80](#)
Parallel data transfers [77](#)
partial file caching [81](#)
patches (GPFS), extracting [476](#)
patches (GPFS), extracting SUSE Linux Enterprise Server and
Red Hat Enterprise Linux [476](#)
PATH environment variable [575](#)
peer snapshot [90](#)
perf_collection_interval [35](#)
performance
 pagepool parameter [22](#)
 use of GPFS to improve [2](#)
 use of pagepool [16](#)
performance considerations
 cloud services [413](#)
 transparent cloud tiering [413](#)
performance monitoring
 Limitations [462](#)
 mmpfmon command [190](#)
 performance monitoring tool [190](#)
performance monitoring tool
 manual installation [495](#)
Performance monitoring tool
 pmcollector memory consumption [461](#)
Performance Monitoring tool
 manual installation [706](#)
 pmswift [706](#)
 uninstalling [666](#)
performance optimization
 AFM [121](#)

- Persistent Reserve
 - reduced recovery time [318](#)
- pinned memory [16](#)
- planned maintenance
 - IW cache [111](#)
- planning
 - data tiering [464](#)
 - GPUDirect Storage [343](#)
 - IBM Storage Scale [416](#)
 - NFS [389](#), [390](#)
 - SMB
 - SMB fail-over scenarios [394](#)
 - SMB upgrades [394](#)
 - systemd [378](#)
 - transparent cloud tiering [407](#)
- Planning
 - extended attributes (EAs) [376](#)
 - Quality of service for I/O operations (QoS) [374](#)
- planning considerations
 - cluster creation [319](#)
 - disks [328](#)
 - file system creation [345](#)
 - hardware requirements [309](#)
 - IBM Storage Scale license designation [300](#)
 - product structure [298](#), [304](#)
 - recoverability [313](#)
 - software requirements [310](#)
- planning for
 - cloud data sharing [407](#)
 - IBM Storage Scale [413](#)
- planning for IBM Cloud [441](#)
- planning for IBM Storage Scale [407](#)
- planning for maintenance activities
 - cloud services [416](#)
- policies [3](#)
- populating cluster definition file
 - limitations [552](#)
- port configuration [59](#)
- port forwarding [505](#)
- portability layer
 - building
 - using the mmbuildgpl command [484](#), [485](#)
 - description [6](#)
- Posix locking [400](#)
- pre-installation steps
 - transparent cloud tiering [627](#)
- prefetch interval
 - AFM [126](#)
- preparing direct access storage devices (DASD) for NSDs [341](#)
- preparing environment for ECKD devices [342](#)
- preparing the environment [472](#)
- prerequisites
 - for Windows [587](#)
- prerequisites for installing protocols [470](#), [473](#)
- primary gateway [59](#)
- private IP address [18](#)
- producers
 - clustered watch folder [256](#)
 - lightweight events [250](#)
- product edition
 - changing [747–750](#)
- programming specifications
 - AIX prerequisite software [575](#)

- programming specifications (*continued*)
 - Linux prerequisite software [472](#)
 - verifying prerequisite software [472](#), [575](#)
- protocol exports
 - fileset considerations [403](#)
- protocol node configuration [573](#)
- protocol node hardware [399](#)
- protocol nodes
 - driver updates [768](#)
 - firmware updates [768](#)
 - kernel updates [768](#)
 - servicing [768](#)
 - upgrading [677](#), [682](#)
- protocol nodes authentication
 - upgrading [744](#)
- protocols prerequisites [470](#), [473](#)
- protocols support
 - overview [30](#)
- protocols, deploying [541](#), [545](#)
- protocols, installing [469](#), [550](#)
- psnap [90](#)
- PTF support [761](#)
- public IP address [18](#)

Q

- Quality of service for I/O operations (QoS)
 - planning [374](#)
- quorum
 - definition of [314](#)
 - during node failure [314](#)
 - enforcement [10](#)
 - file system descriptor [340](#)
 - initialization of GPFS [20](#)
 - node [314](#)
 - selecting nodes [316](#)
- quota support
 - Cloud services [418](#)
- quotas
 - default quotas [358](#)
 - description [358](#)
 - files [15](#)
 - in a replicated system [358](#)
 - role of file system manager node [11](#)
 - values reported in a replicated file system [357](#)

R

- RBAC [223](#)
- rcp command [324](#)
- Read only
 - AFM mode [64](#)
- read operation
 - buffer available [21](#)
 - buffer not available [21](#)
 - requirements [21](#)
 - token management [22](#)
- read-only (RO) [56](#)
- recommendations for maintenance tasks
 - cloud services [416](#)
- recoverability
 - disk failure [316](#)

- recoverability (*continued*)
 - disks [23](#)
 - features of GPFS [2, 26](#)
 - file systems [24](#)
 - node failure [313](#)
 - parameters [313](#)
- recovery time
 - reducing with Persistent Reserve [318](#)
- recovery time objective
 - AFM DR [132](#)
- Red Hat Enterprise Linux 8.x
 - removing object protocol [767](#)
- Red hat RHEL 8.x [458](#)
- reduced recovery time using Persistent Reserve [318](#)
- Redundant Array of Independent Disks (RAID)
 - preventing loss of data access [316](#)
- regions [187](#)
- remote
 - file systems [255](#)
- remote command environment
 - rcp [324](#)
 - rsh [323](#)
 - scp [324](#)
 - ssh [323](#)
- remotely
 - mounted [255](#)
- remotely mounted
 - file system [648](#)
- remotely mounted file systems [255](#)
- removing GPFS, uninstalling [659](#)
- repairing a file system [24](#)
- replication
 - affect on quotas [357](#)
 - description of [2](#)
 - preventing loss of data access [316](#)
- reporting numbers based on quotas for the fileset [359](#)
- Requests
 - REST API [199, 206](#)
- requirements
 - clustered watch folder [651](#)
 - hardware [309, 647](#)
 - OS [647](#)
 - RPM [647](#)
 - software [310](#)
 - system [647](#)
- Response time [399](#)
- Responses
 - REST API [204](#)
- REST API
 - configuring [633](#)
 - Functional overview [198](#)
 - installing [633](#)
 - Requests [199, 206](#)
 - Responses [204](#)
- resync on SW filesets [103](#)
- revalidation
 - AFM [62](#)
 - parameters [62](#)
- reverting to a previous level [758–760](#)
- RHEL 9
 - GUI installation [504](#)
- RoCE [49](#)
- role reversal
 - AFM DR [136](#)

- Role-based access control [223](#)
- RPM database
 - query
 - RPM commands [628](#)
- RPM package
 - covert
 - deb package [628](#)
- RPMs (update), extracting SUSE Linux Enterprise Server and Linux [476](#)
- RPO snapshots [134](#)
- rsh command [323](#)
- running GPFS for Linux on Z [477](#)

S

- S3 [47](#)
- S3 accounts [47](#)
- S3 over RDMA [49](#)
- scaleadm [218, 238](#)
- scaleapi rpm [218](#)
- scalectl --url [219](#)
- scalegmt user [505](#)
- scaling
 - basic concepts [380](#)
 - nodes scaling [380](#)
- scaling planning
 - protocols [380](#)
- scp command [324](#)
- secondary cluster
 - NFS setup [455](#)
- secure communication
 - establishing [53](#)
- security
 - shared file system access [1](#)
- security considerations
 - transparent cloud tiering [416](#)
- self-extracting package [473](#)
- servicing protocol nodes [768](#)
- servicing your GPFS system [761](#)
- set up multi-nodes [629](#)
- setting quota limit
 - transparent cloud tiering [418](#)
- setting up a multi-node [629](#)
- Setup [450](#)
- share modes [400](#)
- shared file system access [1](#)
- shared segments [16](#)
- shell PATH [472](#)
- Single writer
 - AFM mode [64](#)
- single-writer (SW) [56](#)
- sizing file systems [345](#)
- SMB
 - planning [394](#)
 - update to 4.2.3.x [697](#)
 - update to 5.0.x [697](#)
 - upgrading [697, 699](#)
- SMB Client [400, 403](#)
- SMB connections
 - SMB active connections [394](#)
- SMB data migration [397](#)
- SMB data migration using Robocopy [397](#)
- SMB export [400](#)
- SMB migration prerequisites [397](#)

- SMB planning [393](#)
- SMB protocol [400](#)
- SMB share limitations [394](#)
- SMB support
 - overview [37](#)
- snapshot based backups [368](#)
- snapshots
 - coexistence considerations with DMAPAPI [761](#)
 - gpfs.snap command [188](#)
 - writable [123](#)
- SOBAR [193](#)
- soft limit, quotas [358](#)
- softcopy documentation [476](#), [577](#)
- software requirements [310](#)
- Software requirements
 - transparent cloud tiering [407](#)
- Specifying whether the df command will report numbers based on quotas for the fileset [359](#)
- spectrum scale API [197](#)
- spectrumscale [509](#)
- spectrumscale installation toolkit
 - adding file systems [568](#)
 - adding nodes [568](#)
 - adding NSD nodes [536](#)
 - adding NSDs [568](#)
 - configuration options [534](#)
 - creating file systems [537](#)
 - debugging [570](#)
 - deploying protocols
 - debugging [545](#)
 - logging [545](#)
 - diagnosing errors [570](#)
 - enabling protocols [570](#)
 - installing GPFS [539](#)
 - installing IBM Storage Scale [539](#), [541](#), [544](#), [545](#)
 - installing management GUI [550](#)
 - limitations [516](#), [518](#), [552](#)
 - options [516](#), [518](#)
 - overview [509](#)
 - populating cluster definition file [551](#), [552](#)
 - preparing [528](#)
 - setting installer node [533](#)
 - upgrade [723](#), [731](#)
 - upgrade flow chart [725](#)
 - upgrade process flow [725](#)
 - upgrading [723](#), [731](#)
- ssh command [323](#)
- start replication
 - AFM [118](#)
- starting GPFS [327](#)
- stat cache [16](#)
- stat() system call [16](#), [23](#)
- stop replication
 - AFM [118](#)
- Storage Area Network (SAN)
 - disk considerations [328](#)
- Storage configuration [399](#)
- storage disk subsystems [399](#)
- storage management
 - filesets [3](#)
 - policies [3](#)
 - storage pools [3](#)
- storage pools [3](#)
- Storage Scale [216](#)
- storage scale installation toolkit
 - adding node definitions [534](#)
- storage system [399](#)
- storage tiers [306](#)
- stretch cluster [556](#)
- strict replication [355](#)
- structure of IBM Storage Scale [4](#)
- substitution variables [400](#)
- support
 - failover [2](#)
- support and limitation for Windows [580](#)
- supported NFS clients in IBM Storage Scale [391](#)
- SUSE Linux Enterprise Server and Linux RPMs (update), extracting [476](#)
- Swift Object heatmap data tiering [53](#)
- Swift Object storage
 - overview [52](#)
- Swift Object support
 - overview [52](#)
- synchronous
 - operations [64](#)
- synchronous or asynchronous operations
 - AFM [64](#)
- system calls
 - open [20](#)
 - read [21](#)
 - stat() [23](#)
 - write [22](#)
- system health
 - events [188](#)
 - mmhealth command [188](#)
 - thresholdmonitoring [188](#)
- systemd
 - planning [378](#)

T

- Target nodes/node classes [227](#)
- Thanks [699](#)
- Thin provision
 - TRIM-supported NVMe SSDs [338](#)
 - TRIM-supported SSDs [336](#)
- Thin provisioned disks)
 - creation of [330](#)
- Thin provisioning
 - features [331](#)
- tiebreaker disks [314](#)
- token management
 - description [11](#)
 - large clusters [2](#)
 - system calls [20](#)
 - use of [2](#)
- toolkit installation [509](#), [567](#)
- Tracefmt program, installing [588](#)
- Tracelog program, installing [588](#)
- tracing
 - mmprotocoltrace command [188](#)
- transparent cloud tiering
 - considerations [416](#)
 - creating a node class [627](#)
 - firewall recommendations [412](#)
 - hardware requirements [407](#)
 - IBM Cloud Object Storage considerations [722](#)
 - installing

- transparent cloud tiering (*continued*)
 - installing (*continued*)
 - on GPFS nodes [628](#)
 - interoperability and limitations [247](#)
 - overview [242](#)
 - planning [407](#), [413](#)
 - security [416](#)
 - software requirements [407](#)
 - tuning and configuring [410](#)
 - upgrade to 1.1.2 [710](#)
 - upgrade to 1.1.2.1 [711](#)
 - upgrade to 1.1.3 [712](#)
 - upgrade to 1.1.4 [713](#), [714](#)
 - upgrade to 1.1.5 [715](#)
 - upgrade to 1.1.6 [717](#)
 - working mechanism [242](#)

- Transparent Cloud Tiering

- backup considerations [417](#)

- tuning system for transparent cloud tiering [410](#)

U

- Ubuntu Linux packages (update), extracting [476](#)

- UID and GID requirements on the cache and home clusters [452](#)

- uncached

- files [63](#)

- uninstall

- GPFS permanently [659](#)

- transparent cloud tiering [668](#)

- UNIX and Windows (mixed) cluster, configuring [592](#)

- UNIX, Identity Mapping for Unix (IDMU) / RFC 2307

- Attributes [592](#)

- unplanned maintenance

- IW cache [112](#)

- Update mask [233](#)

- update PTF notification [294](#)

- update SUSE Linux Enterprise Server and Red hat Enterprise Linux RPMs, extracting [476](#)

- update Ubuntu Linux packages, extracting [476](#)

- upgrade

- cloud services [710](#)

- Upgrade

- v4.2.0. to v4.2.1 or later [703](#)

- v4.2.1. to v5.0.0 [702](#)

- upgrade considerations

- performance monitoring [673](#)

- protocols [673](#)

- upgrade from 1.1.1 to 1.1.2

- transparent cloud tiering [710](#)

- upgrade rerun

- installation toolkit [741](#)

- upgrade to Red Hat Enterprise Linux 8.x [767](#)

- Upgrades [394](#)

- upgrading

- completing the upgrade [753](#)

- GPUDirect Storage [692](#)

- offline upgrade [672](#)

- online upgrade [672](#), [673](#)

- performance monitoring [673](#)

- protocols [673](#)

- supported paths [672](#)

- toolkit [723](#), [731](#)

- upgrading from version 3.5 [769](#)

- upgrading IBM Storage Scale GUI [708](#)

- upgrading operating system [762](#), [764](#)

- upgrading OS [762](#), [764](#)

- usage restrictions [188](#)

- user-defined node class

- cloud services [627](#)

- using CES [31](#)

- using mmbbackup [114](#)

V

- verifying

- GPFS for AIX installation [578](#)

- GPFS for Linux installation [494](#)

- GPFS installation (Ubuntu) [493](#)

- prerequisite software for AIX nodes [575](#)

- prerequisite software for Linux nodes [472](#)

- verifying disk usage [359](#)

- verifying signature [475](#)

- viewing snapshots at home [100](#)

- Views [230](#)

W

- WAN

- considerations [408](#)

- WAN considerations [408](#)

- watch folder

- AFM [126](#)

- watch folder API [258](#), [722](#)

- Windows

- access control on GPFS file systems [586](#)

- allowing the GPFS administrative account to run as a service [593](#)

- antivirus software [584](#)

- assigning a static IP address [588](#)

- case sensitivity [583](#)

- configuring [588](#)

- configuring a mixed Windows and UNIX cluster [592](#)

- configuring the GPFS Administration service [593](#)

- creating the GPFS administrative account [593](#)

- differences between GPFS and NTFS [585](#)

- disabling the firewall [588](#)

- drive letter [357](#)

- file name considerations [583](#)

- file system name considerations [582](#)

- GPFS limitations [581](#)

- Identity Mapping for Unix (IDMU) / RFC 2307 Attributes [592](#)

- installation procedure [579](#)

- installing and configuring OpenSSH [594](#)

- installing Cygwin [588](#)

- installing GPFS on Windows nodes [590](#), [591](#)

- joining an Active Directory domain [588](#)

- overview [579](#)

- prerequisites [587](#)

- static IP address, Windows [588](#)

- support and limitations [580](#)

- Windows HPC server, configuring [595](#)

- Windows node

- upgrade [676](#)

- worker1Threads on primary cluster [455](#)

- workerThreads on cache cluster [452](#)

Working of

Cloud data sharing [243](#)

workload profiling [399](#)

writable snapshot

concept [123](#)

Write Once Read Many solution [246](#)

write operation

buffer available [22](#)

buffer not available [22](#)

token management [23](#)

X

x86 Architecture [458](#)



Product Number: 5641-DM1
5641-DM3
5641-DM5
5641-DA1
5641-DA3
5641-DA5
5737-F34
5737-I39
5765-DME
5765-DAE

SC28-3600-00

