



IBM Spectrum Scale On AWS
Version 1.1.0

*IBM Spectrum Scale
On AWS Guide*



IBM Spectrum Scale On AWS
Version 1.1.0

*IBM Spectrum Scale
On AWS Guide*



Note

Before using this information and the product it supports, read the information in “Notices” on page 41.

This edition applies to version 5 release 0 modification 1 of the following products, and to all subsequent releases and modifications until otherwise indicated in new editions:

- IBM Spectrum Scale ordered through Passport Advantage® (product number 5725-Q01)
- IBM Spectrum Scale ordered through AAS/eConfig (product number 5641-GPF)
- IBM Spectrum Scale for Linux on Z (product number 5725-S28)
- IBM Spectrum Scale for IBM ESS (product number 5765-ESS)

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 xxi. 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 IBM Corporation 2018.

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

Contents

Tables	v
---------------	----------

About this information	vii
Prerequisite and related information	xx
Conventions used in this information	xx
How to send your comments	xxi

Summary of changes	xxiii
---------------------------	--------------

Chapter 1. Introduction to IBM Spectrum Scale on AWS	1
AWS Services	1
Regions and availability zones	2
IBM Spectrum Scale instance types and operating system	3
IBM Spectrum Scale usage restrictions	3

Chapter 2. Setting up the IBM Spectrum Scale environment in the AWS Cloud	5
Optimal setup considerations	7

Chapter 3. Deploying IBM Spectrum Scale on AWS	9
Deployment options	12
Option 1: Deploying IBM Spectrum Scale on a new Amazon VPC with a single availability zone	13
Option 2: Deploying IBM Spectrum Scale on a new Amazon VPC with multiple availability zones	14
Option 3: Deploying IBM Spectrum Scale on an existing Amazon VPC	16

Chapter 4. Cleaning up the cluster and the stack	19
---	-----------

Chapter 5. Data security and AWS Identity and Access Management	21
--	-----------

Chapter 6. Cluster lifecycle management and debug data collection	23
mmaws utility	24

Lambda function to start nodes	28
--------------------------------	----

Chapter 7. Upgrading IBM Spectrum Scale	33
--	-----------

Chapter 8. Troubleshooting	35
CREATE_FAILED error with timeout message is encountered on launching AMI	35
Service.RequestLimitExceeded error in the cfn-init-cmd.log	35
Size limitation error on deploying the AWS CloudFormation templates	35
Stack creation failure message encountered	35

Chapter 9. Frequently Asked Questions	37
--	-----------

Accessibility features for IBM Spectrum Scale	39
Accessibility features	39
Keyboard navigation	39
IBM and accessibility	39

Notices	41
Trademarks	42
Terms and conditions for product documentation	43
IBM Online Privacy Statement	43

Glossary	45
-----------------	-----------

Index	51
--------------	-----------

Tables

1. IBM Spectrum Scale library information units	viii	15. Compute Node Configurations	15
2. Conventions	xx	16. Network Configuration	15
3. List of changes in documentation	xxix	17. Amazon EC2 Configuration	16
4. File System Configurations	13	18. Personal Configuration.	16
5. NSD Configurations	13	19. License Information.	16
6. Server Node Configurations	13	20. File System Configurations	17
7. Compute Node Configurations	13	21. NSD Configurations	17
8. Network Configuration	14	22. Server Node Configurations	17
9. Amazon EC2 Configuration	14	23. Compute Node Configurations	17
10. Personal Configuration.	14	24. Network Configuration	18
11. License Information.	14	25. Amazon EC2 Configuration	18
12. File System Configurations	14	26. Personal Configuration.	18
13. NSD Configurations	15	27. License Information.	18
14. Server Node Configurations	15		

About this information

This edition applies to IBM Spectrum Scale™ version 5.0.1 for AIX®, Linux, and Windows.

IBM Spectrum 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 Spectrum Scale is running on a particular AIX node, enter:

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

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

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

The IBM Spectrum Scale library consists of the information units listed in Table 1 on page viii.

To use these information units effectively, you must be familiar with IBM Spectrum 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 Spectrum Scale library information units

Information unit	Type of information	Intended users
<i>IBM Spectrum Scale: Concepts, Planning, and Installation Guide</i>	<p>This guide provides the following information:</p> <p>Product overview</p> <ul style="list-style-type: none"> • Overview of IBM Spectrum Scale • GPFS architecture • Protocols support overview: Integration of protocol access methods with GPFS • Active File Management • AFM-based Asynchronous Disaster Recovery (AFM DR) • Data protection and disaster recovery in IBM Spectrum Scale • Introduction to IBM Spectrum Scale GUI • IBM Spectrum Scale management API • Introduction to Cloud services • Introduction to file audit logging • IBM Spectrum Scale in an OpenStack cloud deployment • IBM Spectrum Scale product editions • IBM Spectrum Scale license designation • Capacity based licensing • IBM Spectrum Storage™ Suite <p>Planning</p> <ul style="list-style-type: none"> • Planning for GPFS • Planning for protocols • Planning for Cloud services • Firewall recommendations • Considerations for GPFS applications 	System administrators, analysts, installers, planners, and programmers of IBM Spectrum Scale clusters who are very experienced with the operating systems on which each IBM Spectrum Scale cluster is based

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

Information unit	Type of information	Intended users
<i>IBM Spectrum Scale: Concepts, Planning, and Installation Guide</i>	Installing <ul style="list-style-type: none"> • Steps for establishing and starting your IBM Spectrum Scale cluster • Installing IBM Spectrum Scale on Linux nodes and deploying protocols • Installing IBM Spectrum Scale on AIX nodes • Installing IBM Spectrum Scale on Windows nodes • Installing Cloud services on IBM Spectrum Scale nodes • Installing and configuring IBM Spectrum Scale management API • Installing Active File Management • Installing and upgrading AFM-based Disaster Recovery • Installing call home • Installing file audit logging • Steps to permanently uninstall GPFS and/or Protocols 	System administrators, analysts, installers, planners, and programmers of IBM Spectrum Scale clusters who are very experienced with the operating systems on which each IBM Spectrum Scale cluster is based

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

Information unit	Type of information	Intended users
IBM Spectrum Scale: Concepts, Planning, and Installation Guide	Upgrading <ul style="list-style-type: none"> • IBM Spectrum Scale supported upgrade paths • Upgrading to IBM Spectrum Scale 5.0.x from IBM Spectrum Scale 4.2.y • Upgrading to IBM Spectrum Scale 4.2.y from IBM Spectrum Scale 4.1.x • Upgrading to IBM Spectrum Scale 4.1.1.x from GPFS V4.1.0.x • Upgrading from GPFS 3.5 • Online upgrade support for protocols and performance monitoring • Upgrading AFM and AFM DR • Upgrading object packages • Upgrading NFS packages • Upgrading SMB packages • Upgrading call home configuration • Manually upgrading pmswift • Manually upgrading the performance monitoring tool • Manually upgrading the IBM Spectrum Scale management GUI • Upgrading Cloud services • Upgrading to IBM Cloud Object Storage software level 3.7.2 and above • Manually upgrading file audit logging • Upgrading IBM Spectrum Scale components with the installation toolkit • Migrating from Express Edition to Standard Edition • Completing the upgrade to a new level of IBM Spectrum Scale • Reverting to the previous level of IBM Spectrum Scale • Coexistence considerations • Compatibility considerations • Considerations for IBM Spectrum Protect™ for Space Management • GUI user role considerations • Applying maintenance to your GPFS system 	System administrators, analysts, installers, planners, and programmers of IBM Spectrum Scale clusters who are very experienced with the operating systems on which each IBM Spectrum Scale cluster is based

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

Information unit	Type of information	Intended users
IBM Spectrum Scale: Administration Guide	<p>This guide provides the following information:</p> <p>Configuring</p> <ul style="list-style-type: none"> • Configuring the GPFS cluster • Configuring the CES and protocol configuration • 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 file audit logging • Configuring Active File Management • Configuring AFM-based DR • Tuning for Kernel NFS backend on AFM and AFM DR <p>Administering</p> <ul style="list-style-type: none"> • Performing GPFS administration tasks • Verifying network operation with the mmnetverify command • Managing file systems • File system format changes between versions of IBM Spectrum Scale • Managing disks • Managing protocol services • Managing protocol user authentication • Managing protocol data exports • Managing object storage • Managing GPFS quotas • Managing GUI users • Managing GPFS access control lists • Considerations for GPFS applications • Accessing a remote GPFS file system 	System administrators or programmers of IBM Spectrum Scale systems

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

Information unit	Type of information	Intended users
IBM Spectrum Scale: Administration Guide	<ul style="list-style-type: none"> • Information lifecycle management for IBM Spectrum 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 • 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 • Highly-available write cache (HAWC) • Local read-only cache • Miscellaneous advanced administration • GUI limitations 	System administrators or programmers of IBM Spectrum Scale systems

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

Information unit	Type of information	Intended users
IBM Spectrum Scale: Problem Determination Guide	<p>This guide provides the following information:</p> <p>Monitoring</p> <ul style="list-style-type: none"> • Performance monitoring • Monitoring system health through the IBM Spectrum Scale GUI • Monitoring system health by using the mmhealth command • Monitoring events through callbacks • Monitoring capacity through GUI • Monitoring AFM and AFM DR • GPFS SNMP support • Monitoring the IBM Spectrum Scale system by using call home • Monitoring remote cluster through GUI • Monitoring file audit logging <p>Troubleshooting</p> <ul style="list-style-type: none"> • Best practices for troubleshooting • Understanding the system limitations • Collecting details of the issues • Managing deadlocks • Installation and configuration issues • Upgrade issues • Network issues • File system issues • Disk issues • Security issues • Protocol issues • Disaster recovery issues • Performance issues • GUI issues • AFM issues • AFM DR issues • Transparent cloud tiering issues • File audit logging issues • Maintenance procedures • Recovery procedures • Support for troubleshooting • References 	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 Spectrum Scale: Concepts, Planning, and Installation Guide</i>

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

Information unit	Type of information	Intended users
IBM Spectrum Scale: Command and Programming Reference	<p>This guide provides the following information:</p> <p>Command reference</p> <ul style="list-style-type: none"> • gpfs.snap command • mmaddcallback command • mmaddddisk command • mmaddnode command • mmadquery command • mmafmconfig command • mmafmctl command • mmafmlocal command • mmapplypolicy command • mmaudit command • mmauth command • mmbackup command • mmbackupconfig command • mmblock command • mmbuildgpl command • mmcachectl command • mmcallhome command • mmces command • mmcesdr 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 • mmchqos command • mmclidecode command • mmclone command • mmcloudgateway command • mmcrcluster command • mmcrfileset command • mmcrfs command • mmcrnodeclass command • mmcrnsd command • mmcrsnapshot command 	<ul style="list-style-type: none"> • System administrators of IBM Spectrum Scale systems • Application programmers who are experienced with IBM Spectrum Scale systems and familiar with the terminology and concepts in the XDSM standard

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

Information unit	Type of information	Intended users
IBM Spectrum Scale: Command and Programming Reference	<ul style="list-style-type: none"> • mmdefquota command • mmdefquotaoff command • mmdefquotaon command • mmdefragfs command • mmdelacl command • mmdelcallback command • mmdeldisk command • mmdelfileset command • mmdelfs command • mmdelnnode command • mmdelnnodeclass command • mmdelnsd command • mmdelsnapshot command • mmdf command • mmdiag command • mmdsh command • mmeditacl command • mmedquota command • mmexportfs command • mmfsck command • mmfsctl command • mmgetacl command • mmgetstate command • mmhadoopctl command • mmhealth command • mmimgbackup command • mmimgrestore command • mmimportfs command • mmkeyserv command • mmlinkfileset command • mmlsattr command • mmlscallback command • mmlscluster command • mmlsconfig command • mmlsdisk command 	<ul style="list-style-type: none"> • System administrators of IBM Spectrum Scale systems • Application programmers who are experienced with IBM Spectrum Scale systems and familiar with the terminology and concepts in the XDSM standard

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

Information unit	Type of information	Intended users
IBM Spectrum Scale: Command and Programming Reference	<ul style="list-style-type: none"> • mmlsfileset command • mmlsfs command • mmlslicense command • mmlsmgr command • mmlsmount command • mmlsnodeclass command • mmlsnsd command • mmlspolicy command • mmlspool command • mmlsqos command • mmlsquota command • mmlssnapshot command • mmmigratefs command • mmmount command • mmmmsgqueue command • mmnetverify command • mmnfs command • mmnsddiscover command • mmobj command • mmperfmon command • mmpmon command • mmprotocoltrace command • mm snapsnap command • mmputacl command • mmquotaoff command • mmquotaon command • mmremotefluster command • mmremotefs command • mmrepquota command • mmrestoreconfig command • mmrestorefs command • mmrestripefile command • mmrestripefs command • mmrpldisk command • mmsdrrestore command • mmsetquota command • mmshutdown command • mmsmb command • mmsnapdir command • mmstartup command • mmtracectl command • mmumount command • mmunlinkfileset command • mmuserauth command • mmwinservctl command • spectrumscale command 	<ul style="list-style-type: none"> • System administrators of IBM Spectrum Scale systems • Application programmers who are experienced with IBM Spectrum Scale systems and familiar with the terminology and concepts in the XDSM standard

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

Information unit	Type of information	Intended users
<i>IBM Spectrum Scale: Command and Programming Reference</i>	Programming reference <ul style="list-style-type: none"> • IBM Spectrum Scale Data Management API for GPFS information • GPFS programming interfaces • GPFS user exits • IBM Spectrum Scale management API commands 	<ul style="list-style-type: none"> • System administrators of IBM Spectrum Scale systems • Application programmers who are experienced with IBM Spectrum Scale systems and familiar with the terminology and concepts in the XDSM standard

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

Information unit	Type of information	Intended users
IBM Spectrum Scale: Big Data and Analytics Guide	<p>This guide provides the following information:</p> <p>IBM Spectrum Scale support for Hadoop</p> <ul style="list-style-type: none"> • HDFS transparency • Supported IBM Spectrum Scale storage modes • Hadoop cluster planning • Installation and configuration of HDFS transparency • Application interaction with HDFS transparency • Upgrading the HDFS Transparency cluster • Rolling upgrade for HDFS Transparency • Security • Advanced features • Hadoop distribution support • Limitations and differences from native HDFS • Problem determination <p>BigInsights® 4.2.5 and Hortonworks Data Platform 2.6</p> <ul style="list-style-type: none"> • Planning <ul style="list-style-type: none"> – Hardware requirements – Preparing the environment – Preparing a stanza file • Installation <ul style="list-style-type: none"> – Set up – Installation of software stack – BigInsights value-add services on IBM Spectrum Scale • Upgrading software stack <ul style="list-style-type: none"> – Migrating from BI IOP to HDP – Upgrading IBM Spectrum Scale service MPack – Upgrading HDFS Transparency – Upgrading IBM Spectrum Scale file system – Upgrading to BI IOP 4.2.5 	<ul style="list-style-type: none"> • System administrators of IBM Spectrum Scale systems • Application programmers who are experienced with IBM Spectrum Scale systems and familiar with the terminology and concepts in the XD SM standard

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

Information unit	Type of information	Intended users
IBM Spectrum Scale: Big Data and Analytics Guide	<ul style="list-style-type: none"> • Configuration <ul style="list-style-type: none"> – Setting up High Availability [HA] – IBM Spectrum Scale configuration parameter checklist – Dual-network deployment – Manually starting services in Ambari – Setting up local repository – Configuring LogSearch – Setting IBM Spectrum Scale configuration for BigSQL – Hadoop Kafka/Zookeeper and IBM Spectrum Scale Kafka/Zookeeper • Administration <ul style="list-style-type: none"> – IBM Spectrum Scale-FPO deployment – Ranger – Kerberos – Short-circuit read (SSR) – Disabling short circuit write – IBM Spectrum Scale service management – Ambari node management – Restricting root access – IBM Spectrum Scale management GUI – IBM Spectrum Scale versus Native HDFS • Troubleshooting <ul style="list-style-type: none"> – Snap data collection • Limitations <ul style="list-style-type: none"> – Limitations and information • FAQ <ul style="list-style-type: none"> – General – Service fails to start – Service check failures 	<ul style="list-style-type: none"> • System administrators of IBM Spectrum Scale systems • Application programmers who are experienced with IBM Spectrum Scale systems and familiar with the terminology and concepts in the XDSM standard

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

Information unit	Type of information	Intended users
IBM Spectrum Scale on AWS Guide	IBM Spectrum Scale on AWS <ul style="list-style-type: none"> • Introduction to IBM Spectrum Scale on AWS • Setting up the IBM Spectrum Scale environment in the AWS Cloud • Deploying IBM Spectrum Scale on AWS • Cleaning up the cluster and the stack • Data security and AWS Identity and Access Management • Cluster lifecycle management and debug data collection • Upgrading IBM Spectrum Scale • Troubleshooting • Frequently Asked Questions 	<ul style="list-style-type: none"> • System administrators of IBM Spectrum Scale systems • Application programmers who are experienced with IBM Spectrum Scale systems and familiar with the terminology and concepts in the XDSM standard

Prerequisite and related information

For updates to this information, see IBM Spectrum Scale in IBM Knowledge Center (www.ibm.com/support/knowledgecenter/STXKQY/ibmspectrumscale_welcome.html).

For the latest support information, see the IBM Spectrum Scale FAQ in IBM Knowledge Center (www.ibm.com/support/knowledgecenter/STXKQY/gpfsclustersfaq.html).

Conventions used in this information

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

Note: Users of IBM Spectrum 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	<p>Bold words or characters represent system elements that you must use literally, such as commands, flags, values, and selected menu options.</p> <p>Depending on the context, bold typeface sometimes represents path names, directories, or file names.</p>
<u>bold underlined</u>	<u>bold underlined</u> keywords are defaults. These take effect if you do not specify a different keyword.
constant width	<p>Examples and information that the system displays appear in constant-width typeface.</p> <p>Depending on the context, constant-width typeface sometimes represents path names, directories, or file names.</p>

Table 2. Conventions (continued)

Convention	Usage
<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 Spectrum 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 Spectrum Scale development organization, send your comments to the following e-mail address:

gpfs@us.ibm.com

Summary of changes

This topic summarizes changes to the IBM Spectrum Scale licensed program and the IBM Spectrum 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 Spectrum Scale version 5.0.1 as updated, September 2018

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

IBM Spectrum Scale on AWS

IBM Spectrum Scale can now be deployed on AWS. For more information, see IBM Spectrum Scale on AWS.

AFM and AFM DR-related changes

- A secondary site can be converted into the primary site, if the primary goes down. This role can be reversed permanently. See *Role reversal* in *IBM Spectrum Scale: Concepts, Planning, and Installation Guide*.
- For enhanced data-in-flight security, a new configuration parameter added - **afmEnableNFSec**. See *Configuration parameters for AFM* and *Configuration parameters for AFM-based DR* in *IBM Spectrum Scale: Administration Guide*.
- An administrator can stop replication activity of filesets from cache to home during planned downtime. See *Stop and start replication on a fileset* in *IBM Spectrum Scale: Concepts, Planning, and Installation Guide*.

Authentication-related changes

The *mmuserauth* command can use passwords saved in a Stanza file.

Big data and analytics changes

For information on changes in IBM Spectrum Scale Big Data and Analytics support, see Big data and analytics - summary of changes.

Cloud services changes

Cloud services has the following updates:

- Support for backup and restore using SOBAR. For more information, see the *Scale out backup and restore (SOBAR) for Cloud services* topic in the *IBM Spectrum Scale: Administration Guide*.
- Support for automated Cloud services maintenance service for the following operations:
 - Background removal of deleted files from the object storage
 - Backing up the Cloud services full database to the cloud
 - Reconciling the Cloud services database
- Support for setting up a customized maintenance window, overriding the default values
- Support for multi-threading within a file in a single node to improve large file recall latency.

File audit logging updates

File audit logging has the following updates:

- If a file system with file audit logging is unmounted and then remounted, the producer reacquires the password. In addition, if a node mounts a file system after file audit logging has already been enabled, the producer on that node acquires the password. Previously, in both cases, the GPFS daemon would have to be restarted for the producer to acquire the new password.

- File audit logging authentication is changed from SASL PLAINTEXT to SASL SCRAM for improved security.
- If the Object protocol is enabled, no file system activity occurring within the primary object fileset is audited.
- Added **Services > File Auditing** page in the management GUI to enable monitoring file auditing through GUI.

File system core improvements

File systems: Integration with systemd is broader

You can now use systemd to manage IBM Spectrum Scale systemd services on configured systems. IBM Spectrum Scale automatically installs and configures GPFS as a suite of systemd services on systems that have systemd version 219 or later installed. Support for the IBM Spectrum Scale Cluster Configuration Repository (CCR) is included. For more information, see the topic *Planning for systemd* in the *IBM Spectrum Scale: Concepts, Planning, and Installation Guide*.

File systems: Traditional NSD nodes and servers can use checksums

NSD clients and servers that are configured with IBM Spectrum Scale can use checksums to verify data integrity and detect network corruption of file data that the client reads from or writes to the NSD server. For more information, see the **nsdChecksumTraditional** and **nsdDumpBuffersOnChecksumError** attributes in the topic *mmchconfig command* in the *IBM Spectrum Scale: Command and Programming Reference*.

File systems: Concurrent updates to small shared directories are faster

Fine-grained directory locking significantly improves the performance of concurrent updates to small directories that are accessed by more than one node. "Concurrent" updates means updates by multiple nodes within a 10-second interval. A "small" directory is one with fewer than 8 KiB entries.

File systems: NSD disk discovery on Linux now detects NVMe devices

The default script for NSD disk discovery on Linux, `/usr/lpp/mmfs/bin/mmdevdiscover`, now automatically detects NVM Express (NVMe) devices. It is no longer necessary to create an `nsddevises` user exit script to detect NVMe devices on a node. For more information, see the topic *NSD disk discovery* in the *IBM Spectrum Scale: Concepts, Planning, and Installation Guide* and the topic *nsddevises user exit* in the *IBM Spectrum Scale: Command and Programming Reference*.

mmapplypolicy command: New default values are available for the parameters **-N** (helper nodes) and **-g** (global work directory)

- If the **-N** parameter is not specified and the **defaultHelperNodes** attribute is not set, then the list of helper nodes defaults to the **managerNodes** node class. The target file system must be at format version 5.0.1 (format number 19.01) or later.
- If the **-g** parameter is not specified, then the global work directory defaults to the path (absolute or relative) that is stored in the new **sharedTmpDir** attribute. The target file system can be at any supported format version.

For more information, see the topics *mmapplypolicy command* and *mmchconfig command* in the *IBM Spectrum Scale: Command and Programming Reference*.

mmbackup command: A new default value is available for the **-g** (global work directory) parameter

If the **-g** parameter is not specified, then the global work directory defaults to the path (absolute or relative) that is stored in the new **sharedTmpDir** attribute. The target file system can be at any supported format version. For more information, see the topics *mmbackup command* and *mmchconfig command* in the *IBM Spectrum Scale: Command and Programming Reference*.

mmcachectl command: You can list the file and directory entries in the local page pool

You can display the number of bytes of file data that are stored in the local page pool for each file in a set of files, along with related information. You can display information for a single file, for the files in a fileset, for all the files in a file system, or for all the file systems that are mounted by the node. For more information, see *mmcachectl command* in the *IBM Spectrum Scale: Command and Programming Reference*.

IBM Spectrum Scale functionality to support GDPR requirements

To understand the requirements of EU General Data Protection Regulation (GDPR) compliance that are applicable to unstructured data storage and how IBM Spectrum Scale helps to address them, see the *IBM Spectrum Scale functionality to support GDPR requirements technote*.

IBM Spectrum Scale management API changes

Added the following API commands:

- GET /nodes/{name}/services
- GET /nodes/{name}/services/{serviceName}
- PUT /nodes/{name}/services/{serviceName}
- POST /filesystems/{filesystemName}/filesets/{filesetName}/afmctl
- GET /filesystems/{filesystemName}/policies
- PUT /filesystems/{filesystemName}/policies
- GET /perfmon/sensors
- GET /perfmon/sensors/{sensorName}
- PUT /perfmon/sensors/{sensorName}
- GET /cliauditlog

For more information on the API commands, see *IBM Spectrum Scale management API commands* in *IBM Spectrum Scale: Command and Programming Reference*. You can also access the documentation corresponding to each API command from the GUI itself. The API documentation is available in the GUI at: <https://<IP address or host name of API server>:<port>/ibm/api/explorer/>. For example: <https://scalegui.ibm.com:443/ibm/api/explorer/>.

IBM Spectrum Scale GUI changes

The following changes are made to the GUI:

- Added a new **Services** page that provides options to monitor, configure, and manage various services that are available in the IBM Spectrum Scale system. You can monitor and manage the following services from the **Services** page:
 - GPFS daemon
 - GUI
 - CES
 - CES network
 - Hadoop connector
 - Performance monitoring
 - NFS
 - SMB
 - Object
 - File auditing
 - Message queue
 - File authentication
 - Object authentication
- Added a new **Access > Command Audit Log** page that lists the various actions that are performed through the CLI. This page helps the system administrator to audit the commands

and tasks the users and administrators are performing. These logs can also be used to troubleshoot issues that are reported in the system.

- Moved the NFS Service, SMB Service, Object Service, and Object Administrator pages from the Settings menu to the newly created Services page.
- Removed GUI Preferences page and moved the options in that page to the GUI section of the Services page.
- New option is added in the GUI section of Services page to define session timeout for the GUI users.
- Support for creating and installing a self-signed or CA-certified SSL certificates is added in the GUI section of the Services page.
- Remote cluster monitoring capabilities are added. You can now create customized performance charts in the **Monitoring > Statistics** page and use them in the **Monitoring > Dashboard** page. If a file system is mounted on the remote cluster node, the performance of the remote node can be monitored through the detailed view of file systems in **Files > File Systems** page.
- Modified the **Files > Transparent Cloud Tiering** page to display details of the container pairs and cloud account.
- Added support for creating and modifying encryption rules in the **Files > Information Lifecycle** page. You can now create and manage the following types of encryption rules as well:
 - Encryption
 - Encryption specification
 - Encryption exclude
- Added ILM policy run settings in the **Files > Information Lifecycle** page.
- Added the **Provide Feedback** option in the user menu that is available at the upper right corner of the GUI.
- In the **Monitoring > Events** page, the events that are occurred multiple times are now grouped under the newly introduced **Event Groups** tab and the number of occurrences of the events are indicated in the **Occurrences** column. The **Individual Events** tab lists all the events irrespective of the multiple occurrences.
- Added a report for fileset growth and size distribution in the **Files > Filesets** page.
- Added the GPFSFilesystemAPI-sensor based performance monitoring in the File Systems and Nodes pages.
- NFS performance monitoring metrics are added in the detailed view of the **Files > Active File Management** page.

Installation toolkit changes

- The installation toolkit supports the installation and the deployment of IBM Spectrum Scale on Red Hat Enterprise Linux 7.5 on x86_64, PPC64, and PPC64LE.
- The installation toolkit supports the installation and the deployment of IBM Spectrum Scale on Ubuntu 16.04.4 and 18.04 (x86_64).
- The installation toolkit config populate option now supports call home and file audit logging.
- The installation toolkit performance monitoring configuration for protocols sensors has been improved.

mmhealth command: Enhancements

New options have been added to the **mmhealth node show** and **mmhealth cluster show** commands. For more information, see the topic *mmhealth command* in the *IBM Spectrum Scale: Command and Programming Reference*.

Object changes

- In Cluster Export Services (CES), the Pike release of OpenStack is used for Swift, Keystone, and their dependent packages.

NFS changes

- NFS Daemon rename - improved compatibility with Red Hat Selinux environment.
- CES NFS logs its crash stack trace, in case of an abnormal termination.

SMB changes

The following enhancements are done:

- **winbind** queuing improvements under high load
- **winbind** enhancements to re-establish domain controller connection on reboot
- Support for hardware encryption on Power®
- Enhanced debug messages
- Enhanced share mode handling with Microsoft Excel temporary files
- IBM Spectrum Scale by default does no longer allow anonymous access via SMB. With this change, it is no longer possible to enumerate local users without credentials. If this change causes issues with customer application requiring anonymous access, the original behavior can be restored by issuing the mmsmb config change --option "restrict anonymous = 0" command.

System Health changes

- Improved configuration options for performance monitoring tools
- Enhanced checking of Infiniband port state: monitoring of port speed and width
- Upgrade: better shutdown/unmount orchestration to avoid upgrade problems
- Recovery: Automatic restart of nfs-ganesha
- Added performance data to call home package
- CES improvements
- Network monitoring for CES IPs when using node affinity
- Identification of unassigned CES IPs
- Improve error messages for easier problem determination
- Enhanced reporting for CES IP moves and rebalancing

Introduction of global group for call home

Call home now uses a global group as a default group that contains the global values that are applied to all groups. For more information on this group and its uses, see the *Understanding call home* topic in the *IBM Spectrum Scale: Problem Determination Guide*.

Automatic single node assignment for performance monitoring

A single node is automatically selected by the system to run certain sensors. If the selected node becomes unresponsive, the system reconfigures a healthy node to act as the singleton sensor node. For more information on assigning single node sensors, see the *Automatic assignment of single node sensors* topic in the *IBM Spectrum Scale: Problem Determination Guide*.

Upgrades to call home configuration

Upgrading call home to a higher version requires specific steps. For more information, see the *Upgrading call home configuration* topic in the *IBM Spectrum Scale: Concepts, Planning, and Installation Guide*.

New options to monitor node health

The CESIP option can be used to monitor a cluster manager node. For more information, see the *Monitoring the health of a node* topic in the *IBM Spectrum Scale: Problem Determination Guide*.

Documented commands, structures, and subroutines

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

New commands

The following command is new in this release:

- **mmcachectl**

New structures

There are no new structures.

New subroutines

There are no new subroutines.

Changed commands

The following commands were changed:

- **mmapplypolicy**
- **mmbackup**
- **mmces**
- **mmcallhome**
- **mmchconfig**
- **mmchfileset**
- **mmcloudgateway**
- **mmfsck**
- **mmhealth**
- **mmisqos**
- **mmisnapshot**
- **mmnfs**
- **mm smb**
- **mmuserauth**
- **spectrum scale**

Changed structures

There are no changed structures.

Changed subroutines

There are no changed subroutines.

Deleted commands

There are no deleted commands.

Deleted structures

There are no deleted structures.

Deleted subroutines

There are no deleted subroutines.

Messages

The following are the new, changed, and deleted messages:

New messages

6027-307, 6027-2402, 6027-2403, 6027-2404, 6027-2405, 6027-2406, 6027-2407, and 6027-3932.

Changed messages

None.

Deleted messages

None.

Changes in documentation

List of documentation changes in product guides and respective Knowledge Center sections

The following is a list of documentation changes including changes in topic titles, changes in placement of topics, and deleted topics:

Table 3. List of changes in documentation

Guide	Knowledge center section	List of changes
Concepts, Planning, and Installation Guide	Installing	<ul style="list-style-type: none"> Upgrading-related information removed from the <i>Installing AFM</i> topic in <i>IBM Spectrum Scale: Concepts, Planning, and Installation Guide</i>. Use case for arping binary for Ubuntu removed from the <i>Software requirements</i> topic in <i>IBM Spectrum Scale: Concepts, Planning, and Installation Guide</i>.
	Upgrading	<ul style="list-style-type: none"> Upgrading-related information removed from the <i>Installing AFM</i> topic added to a new topic <i>Upgrading AFM and AFM DR</i> in <i>IBM Spectrum Scale: Concepts, Planning, and Installation Guide</i>. The following object upgrade subtopics are removed and the object upgrade information is consolidated in <i>Upgrading object packages</i> in <i>IBM Spectrum Scale: Concepts, Planning, and Installation Guide</i> <ul style="list-style-type: none"> Upgrading object packages to 5.0.x from 4.2.x Upgrading object packages to version 4.2.3.x from 4.2.2.x Upgrading object packages to version 4.2.2.x from 4.2.1.x
Administration Guide	Administering	<ul style="list-style-type: none"> Removed the following topic from <i>Administering files for Transparent cloud tiering</i> section: <ul style="list-style-type: none"> "Restoring Transparent cloud tiering service on a backup cluster" Changed the titles of the following topics: <ul style="list-style-type: none"> <i>Configuring encryption with SKLM v2.7 or later</i> in the <i>IBM Spectrum Scale: Administration Guide</i>. <i>Installing Windows IDMU</i> in the <i>IBM Spectrum Scale: Administration Guide</i>. <i>Configuring ID mappings in IDMU</i> in the <i>IBM Spectrum Scale: Administration Guide</i>.
Problem Determination Guide	Monitoring	The <i>Threshold monitoring use cases</i> is now a sub-topic of the <i>System health monitoring use cases</i> section in the <i>IBM Spectrum Scale: Problem Determination Guide</i> .

Chapter 1. Introduction to IBM Spectrum Scale on AWS

Amazon Web Services (AWS) provides an easy way to create and manage a collection of related AWS resources, provisioning and updating them in an orderly and predictable fashion.

IBM Spectrum Scale addresses the needs of applications for which performance or performance-to-capacity ratio demands cannot be met by traditional scale-up storage systems.

IBM Spectrum Scale is a high-performance, highly available, clustered file system and associated management software, available on a variety of platforms. IBM Spectrum Scale can scale in several dimensions, including performance (bandwidth and IOPS), capacity, and number of nodes or instances that can mount the file system.

For information on how to deploy a highly available IBM Spectrum Scale cluster on the Amazon Web Services (AWS) Cloud into a configuration of your choice, see [IBM Spectrum Scale on AWS](#).

Note: IBM Spectrum Scale is not itself an application, but instead provides the storage infrastructure for the applications.

IBM Spectrum Scale is deployed for many I/O-demanding enterprise applications that require high performance or scale. IBM Spectrum Scale provides various configuration options and access methods including traditional POSIX-based file access, and many features such as snapshots, compression, and encryption. This offering automates the deployment of IBM Spectrum Scale on AWS for users who require highly available access to a shared name space across multiple instances with good performance, without requiring an in-depth knowledge of IBM Spectrum Scale.

It is recommended that the IBM Spectrum Scale users subscribe to the IBM notifications for important updates on issues such as security vulnerabilities and other IBM Spectrum Scale fixes. You can subscribe to the IBM Spectrum Scale notifications from the [My Notifications](#) page.

Note: In IBM Spectrum 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 allows you 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 lets you provision a private, isolated section of the AWS Cloud, where you can launch 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 enables you to launch virtual machine instances with a variety of

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 down times. 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 a storage for the 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 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 lets you run code without provisioning or managing servers. For more information, see [AWS Lambda documentation](#).

Regions and availability zones

IBM Spectrum 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 Spectrum Scale cluster nodes across two or more availability zones within an AWS region by default. Within each region there can be two or more availability zones.

Regions are completely independent of each other, but each availability zone, although isolated, is connected through stable low-latency links. You can only view resources that are associated with the region that you specify, because regions are isolated from each other, and resources are automatically replicated across regions. Reading and writing data to an IBM Spectrum Scale file system can cause data to be sent between instances in different availability zones, which results in Amazon charging per GB data transfer charges. You can avoid these cross availability zone data movement charges by deploying IBM Spectrum Scale resources into a single availability zone.

Traffic to and from an Amazon EC2 instance in the same or different availability zones within a region is limited by the network bandwidth of the instance types. For instances that have more than a high network bandwidth, there are additional considerations related to whether two communicating nodes are in the same placement group. The following additional bandwidth limitations apply to nodes with a greater network bandwidth than high:

- Up to 5 Gbps of bandwidth for single-flow traffic

- Either 25 Gbps or 10 Gbps of bandwidth for multi-flow traffic by using private IPv4 or IPv6 addresses based on instance type, enhanced networking

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](#).

When you deploy the multi-availability zone template configuration for an IBM Spectrum Scale cluster with replication, each element of the data and the metadata is replicated in a separate availability zone 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 Spectrum 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 Spectrum Scale creates multiple copies of each data element.

IBM Spectrum Scale instance types and operating system

The IBM Spectrum Scale AMI supports a large selection of EC2 instance types for the IBM Spectrum Scale cluster instances.

It is recommended that you benchmark the environment to make sure you achieve the level of performance you need before starting a production deployment. The deployment launches an EC2 instance running RHEL 7.5. For more information on EC2 instances, see [Amazon EC2 Instance Types](#).

IBM Spectrum Scale usage restrictions

This version does not support the following features of IBM Spectrum 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)
- Active File Management (AFM)
- 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 Spectrum Scale file system from one IBM Spectrum Scale cluster to another IBM Spectrum Scale
- IBM Spectrum Scale Graphical User Interface
- User name space management and quota management
- Snapshots and clones

Additional limitations include the following:

- Using EBS volume encryption for IBM Spectrum Scale file systems is not supported.
- The archiving and restoring of IBM Spectrum Scale data through the use of AWS services is not supported.

Chapter 2. Setting up the IBM Spectrum Scale environment in the AWS Cloud

Setting up a VPC with the default parameters builds the IBM Spectrum Scale environment in the AWS Cloud with the following properties.

Note: You can choose to create a new VPC for the IBM Spectrum 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:

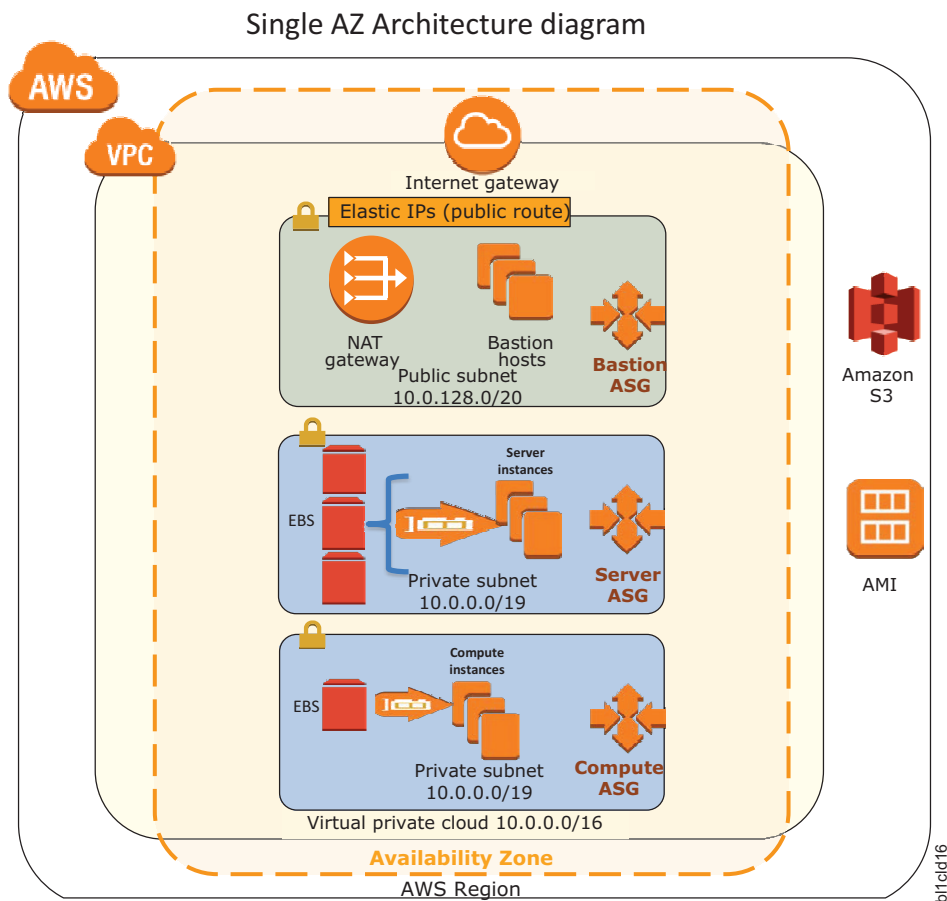


Figure 1. Single availability zone architecture diagram

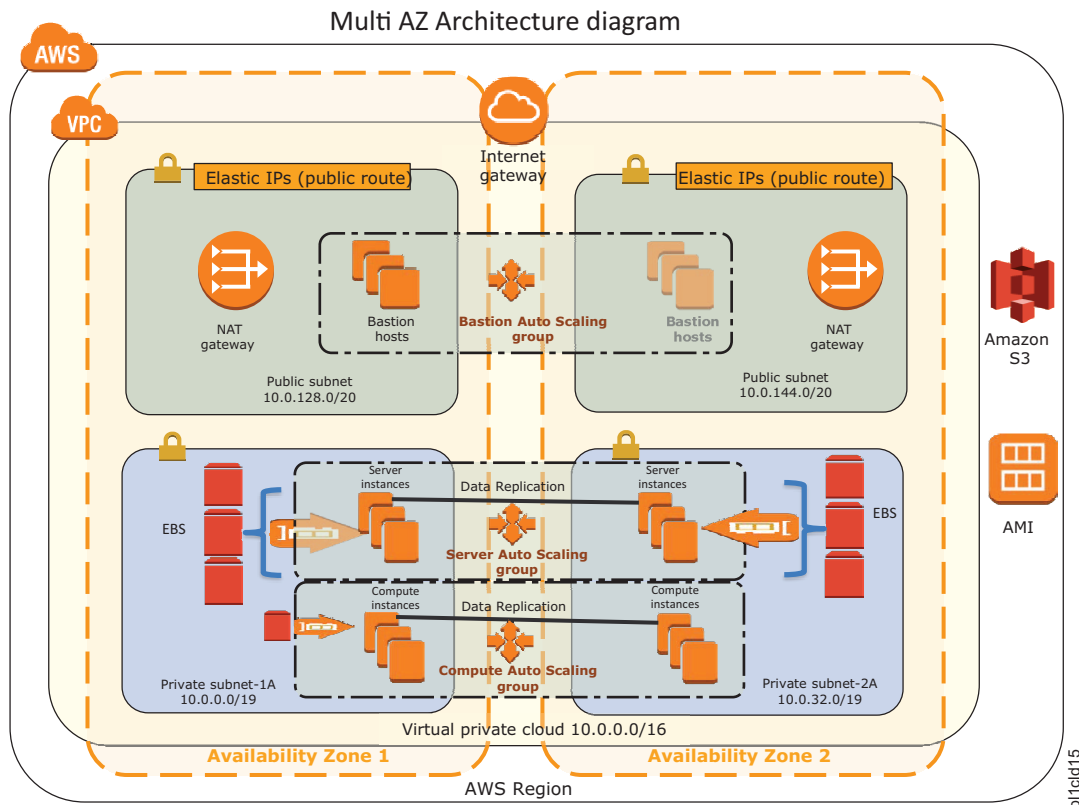


Figure 2. Multi availability zone architecture diagram

- A VPC that spans two availability zones and includes two public and two private subnets, for security and high availability. The number of private subnets is configurable. For more information, see “Optimal setup considerations” on page 7.
- 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. For more information, see “Optimal setup considerations” on page 7.
- A bastion host in a public subnet provides Secure Shell (SSH) access to the IBM Spectrum Scale cluster. The bastion host instance is managed by an auto scaling group of one. This 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 Spectrum Scale daemon.

Each IBM Network Shared Disk (NSD) storage server instance, referred to as an NSD server, deployed has:

- A 100-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 Spectrum Scale compute instance has:

- A 100-GB EBS volume for the root device.
- An extra 5-GB EBS volume attached to one of the compute instances in the cluster. This disk is added to improve the resiliency of the IBM Spectrum Scale cluster, and to avoid loss of file system access after a disk failure.

Note: In IBM Spectrum Scale terminology, the disk is called a descriptor quorum disk, and it is added to avoid a file system descriptor quorum loss.

Optimal setup considerations

IBM Spectrum Scale cluster nodes can be deployed across a single availability zone and multiple availability zones within an AWS region.

It is recommended to use the multiple availability zone architecture to ensure optimal high availability. It deploys an IBM Spectrum Scale cluster with replication, each element of the data and the metadata is replicated in a separate availability zone to avoid the loss of data when hardware failures occur in a single availability zone.

In a single availability zone deployment, IBM Spectrum Scale stores one copy of data and replicates two copies of each metadata element within the single availability zone.

Note: In case of a single availability zone deployment, all IBM Spectrum Scale server and compute instances are deployed into a single availability zone, and IBM Spectrum Scale creates only one copy of each data element. This approach avoids data movement charges across availability zone, but compromises on the high availability and the level of data protection of the solution. If you want to ensure optimal data protection, it is recommended that you choose the multiple availability zone deployment option, which creates two copies of each data element, with each copy in a separate availability zone.

The Amazon VPC service creates a logically isolated networking environment that you can connect to your on-premises data centers, or use as a stand-alone environment. It is recommended that you carefully consider the environment into which you are deploying IBM Spectrum Scale. You can deploy the AMI into a new VPC in which all the IBM Spectrum Scale cluster instances are in private subnets and the bastion host instance is the only host that has direct access to the Internet. For more information on deploying IBM Spectrum Scale, see “Deployment options” on page 12.

Ensure that it is similarly set up with NAT gateways, has at least two private subnets to deploy the IBM Spectrum Scale instances, and has bastion hosts for secure inbound access.

The following figure provides a high-level view of the IBM Spectrum Scale architecture that includes compute nodes and NSD servers, equally split between two availability zones by default, to form one IBM Spectrum Scale cluster.

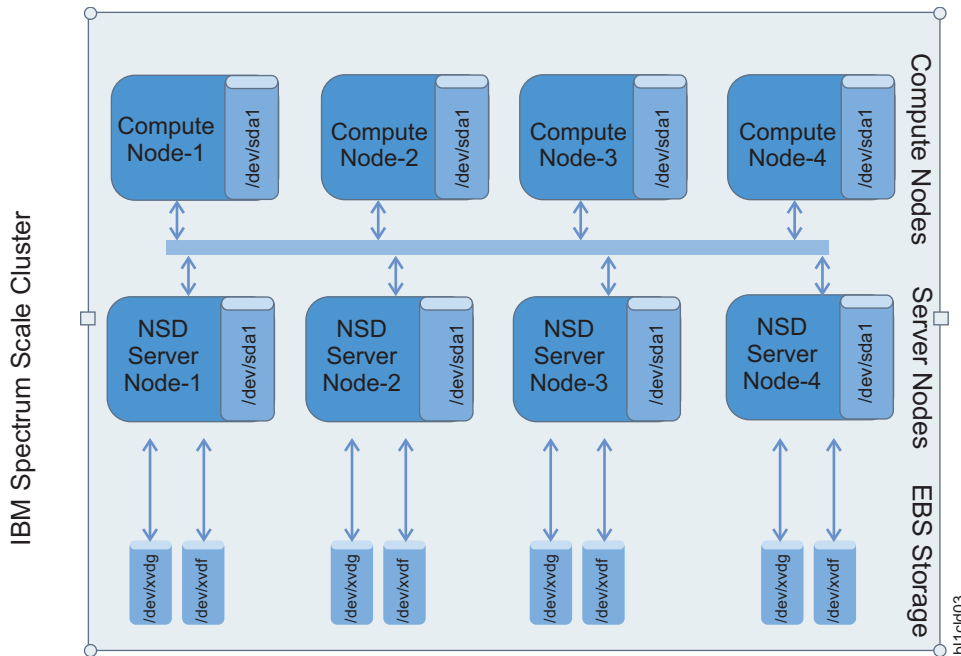


Figure 3. High-level IBM Spectrum Scale cluster architecture on AWS

The compute nodes and NSD nodes are all part of the IBM Spectrum Scale cluster and mount the shared file system as shown. The bastion host is not part of the IBM Spectrum Scale cluster and does not mount the IBM Spectrum Scale shared file system.

Chapter 3. Deploying IBM Spectrum Scale on AWS

Follow these given steps to deploy IBM Spectrum 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 Spectrum Scale on AWS.
 - c. Create a key pair in your preferred region. For information on how to create a key pair, see [Amazon EC2 Key Pairs](#)
 - d. If necessary, [request a service limit increase](#) for the EC2 instance types that you intend to deploy. To do this, in the AWS Support Center, choose **Create Case** > **Service Limit Increase** > **EC2 instances**, and then complete the fields in the **Limit Increase** form.
2. Launch the IBM Spectrum 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 Spectrum Scale does not replicate any of the file system data. It is recommended that this option only be used for cases in which a higher probability of data loss is acceptable in order 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 Spectrum Scale replicates the data such 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 given point. It is recommended that users carefully read Amazon's statement regarding the durability of EBS volumes, particularly the annual failure rates, in order to better assess the possibility of potential data loss resulting from EBS volume failures. For more information, see [Amazon EBS features](#).

Note: This is where the network infrastructure for IBM Spectrum Scale 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 check box 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 Spectrum 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 Spectrum 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 Spectrum Scale cluster. Deleting the stack created by the main template deletes the entire IBM Spectrum Scale deployment stack. However, you still need to delete CloudWatch alerts manually.

3. Connect to the IBM Spectrum Scale cluster.

When the AWS CloudFormation template has successfully created the stack, all instances of compute and NSD servers launched by the AWS set up is up and running with the IBM Spectrum Scale file system mounted on it.

To connect to the IBM Spectrum Scale cluster take the following steps:

- a. Open the Amazon EC2 console at <https://console.aws.amazon.com/ec2/>.
- b. In the navigation pane, under **Instances**, check for the public DNS (IPv4) value for the instance named `LinuxBastion`.
- c. Use your AWS private key to connect to the bastion host using SSH.

Note: This is the key that you specify in the **Key Name** parameter of the AWS CloudFormation template during deployment.

- d. From the bastion host, use the SSH agent to log in to any of the compute instances or NSD server instances that were launched by the AWS CloudFormation templates.

Note: 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.

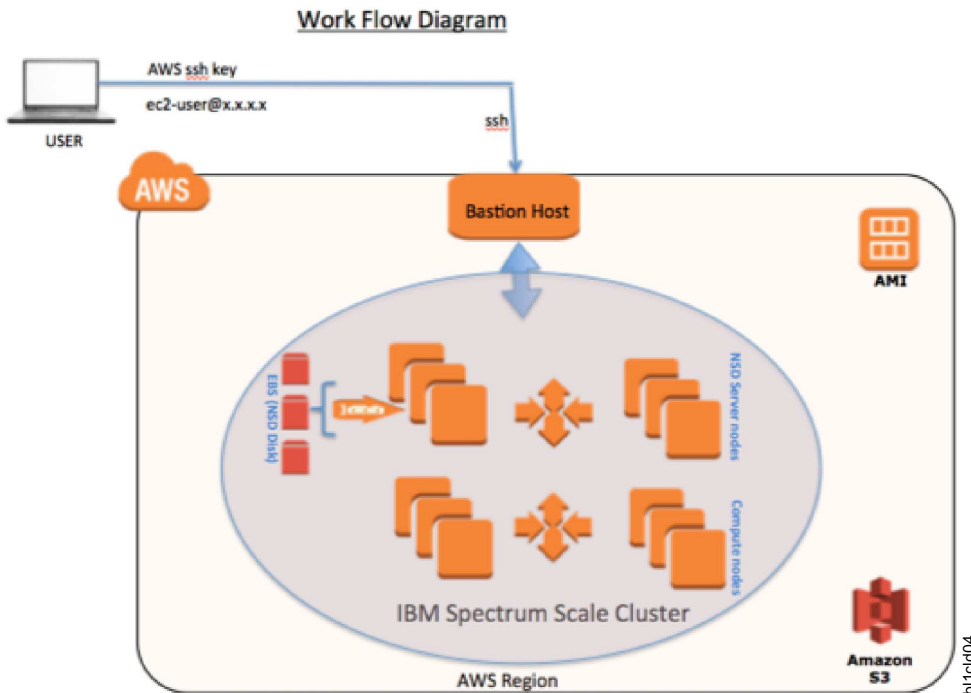


Figure 4. High-level IBM Spectrum Scale cluster architecture to connect from host

4. Test the deployment using IBM Spectrum Scale commands

After you log in to a compute or NSD server instance, you can administer IBM Spectrum Scale. For more information on administering IBM Spectrum Scale, see the *IBM Spectrum Scale: Administration Guide*.

Note: You must be a root user to run IBM Spectrum Scale administration commands.

The **mmfsccluster** command displays the details of the IBM Spectrum Scale cluster. The command gives an output similar to the following:

```
[ec2-user@ip-10-0-1-110]$ /usr/lpp/mmfs/bin/mmfscluster
```

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	Daemon node name	IP address	Admin node name	Designation
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	quorum-manager-perfmon
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	quorum-manager-perfmon
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	quorum-perfmon
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	perfmon

For more information on the **mmfsccluster** command, see the *IBM Spectrum Scale: Command and Programming Reference* guide in the IBM Spectrum Scale Knowledge Centre.

The **mmfssnsd** command displays the NSD server information. The command gives an output similar to the following:

```
[ec2-user@ip-10-0-1-110]$ /usr/lpp/mmfs/bin/mmfsnsd
```

```
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 **mm1nsd** command, see the *IBM Spectrum Scale: Command and Programming Reference* guide in the IBM Spectrum Scale Knowledge Centre.

The **mm1sdisk** command displays disk details. The command gives an output similar to the following:

```
[ec2-user@ip-10-0-1-110]$ mm1sdisk fs1 -L
disk      driver  sector  failure holds  holds  status  availability disk id pool  remarks
name      type    size    group metadata data
-----
0a9503b2cf264bf08 nsd      512      1 Yes    Yes    ready    up      1 system  desc
0d7e3d725140a0cdb nsd      512      2 Yes    Yes    ready    up      2 system  desc
0fb10f7d9a981f39e nsd      512      3 No     No     ready    up      3 system  desc
Number of quorum disks: 3
Read quorum value:      2
Write quorum value:     2
```

For more information on the **mm1sdisk** command, see the *IBM Spectrum Scale: Command and Programming Reference* guide in the IBM Spectrum Scale Knowledge Centre.

The **mmddf** command displays the fs1 filesystem information. The command gives an output similar to the following:

```
[ec2-user@ip-10-0-1-110]$ mmddf fs1
disk      disk size  failure holds  holds  free in KB  free in KB
name      in KB    group metadata data    in full blocks  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 **mmddf** command, see the *IBM Spectrum Scale: Command and Programming Reference* guide in the IBM Spectrum Scale Knowledge Centre.

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 Spectrum Scale: Command and Programming Reference* guide in the IBM Spectrum Scale Knowledge Centre.

Deployment options

The following three deployment options are available to users.

- **Deploy IBM Spectrum Scale into a new VPC with a single availability zone (end-to-end deployment)** : This option builds a new AWS environment consisting of the VPC, subnets, security groups, bastion hosts, and other infrastructure components, and then deploys IBM Spectrum Scale into this new VPC with a single availability zone .

- **Deploy IBM Spectrum Scale into a new VPC with multiple availability zones (end-to-end deployment):** This option builds a new AWS environment consisting of the VPC, subnets, security groups, bastion hosts, and other infrastructure components, and then deploys IBM Spectrum Scale into this new VPC with multiple availability zones.
- **Deploy IBM Spectrum Scale into an existing VPC:** This option provisions IBM Spectrum Scale in your existing AWS infrastructure.

Option 1: Deploying IBM Spectrum 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 4. File System Configurations

Parameter label (name)	Default	Description
Block Size (BlockSize)	4M	The file system block size. You can choose a value from 256 KiB to 16MiB.
GPFS Mount Point (GpfsMountPoint)	/gpfs/fs1	The mount point for the IBM Spectrum Scale file system. Note: fs1 is the file system name that is created by default, and is mounted on location /gpfs/fs1. For example, changing the mount path to /gpfs/fs2 does not create filesystem with name fs2. Instead, the file system fs1 is mounted on /gpfs/fs2.

Table 5. NSD Configurations

Parameter label (name)	Default	Description
EBS Type (EBSType)	gp2	The EBS volume type for IBM Spectrum Scale storage attached to each NSD server node. Options are: General Purpose SSD (gp2), Provisioned IOPS SSD (io1), Cold HDD (sc1), Throughput Optimized HDD (st1), and EBS Magnetic (standard). For more information about EBS volume type choices, see Amazon EBS.
Disk Per Node (DiskPerNode)	1	The number of NSD volumes to attach to each NSD server node. You can choose 1-5 disks.
Disk Size (DiskSize)	500	The disk size of NSD volume(s) attached to each NSD server node, in GiBs. Supported disk sizes are 10-16,384 GiB.

Table 6. Server Node Configurations

Parameter label (name)	Default	Description
Server Node Count (ServerNodeCount)	2	The number of EC2 instances to launch for the NSD server on the GPFS cluster. You can select 2-32 instances.
Server Instance Type (ServerInstanceType)	t2.medium	The instance type to use for the NSD server node instances.

Note: By default, all server nodes are assigned with an IBM Spectrum Scale server license.

Table 7. Compute Node Configurations

Parameter label (name)	Default	Description
Compute Node Count (ComputeNodeCount)	2	The number of IBM Spectrum 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 instances.

Table 8. 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 (PrivateSubnetCIDR)	10.0.1.0/24	The CIDR block for the private subnet 1A located in availability zone 1.
Public Subnet CIDR (PublicSubnetCIDR)	10.0.3.0/24	The CIDR block for the public DMZ subnet 1 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, e.g., <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 9. Amazon EC2 Configuration

Parameter label (name)	Default	Description
Key Pair Name (KeyPairName)	requires input	A public/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 you created in your preferred region.
Bastion AMI OS (BastionAMIOS)	Amazon-Linux- HVM	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	TheEC2 instance type for the bastion host instances.

Table 10. Personal Configuration

Parameter label (name)	Default	Description
Spectrum S3 Bucket (SpectrumS3Bucket)	requires input	An optional parameter that defines the S3 bucket name used for shared object storage among IBM Spectrum Scale nodes. An IAM role is set up to allow read/write operations on the specified bucket. You can provide the name of a pre-existing bucket you own or specify the name of a new bucket to create. Deployment logs are stored in this bucket in case stack creation fails.
Operator Email (OperatorEmail)	requires input	The email address to which notifications of any scaling operations is sent to.

Table 11. License Information

Parameter label (name)	Default	Description
IBM Customer Number (IBMCustomerNumber)	requires input	Required 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 Spectrum 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 12. File System Configurations

Parameter label (name)	Default	Description
------------------------	---------	-------------

Table 12. File System Configurations (continued)

Block Size (BlockSize)	4M	The file system block size. You can choose a value from 256 KiB to 16MiB.
GPFS Mount Point (GpfsMountPoint)	/gpfs/fs1	The mount point for the IBM Spectrum Scale volume. Note: fs1 is the file system name that is created by default, and is mounted on location /gpfs/fs1. For example, changing the mount path to /gpfs/fs2 does not create filesystem with name fs2. Instead, the file system fs1 is mounted on /gpfs/fs2.

Table 13. NSD Configurations

Parameter label (name)	Default	Description
EBS Type (EBSType)	gp2	The EBS volume type for IBM Spectrum Scale storage attached to each NSD server node. Options are: General Purpose SSD (gp2), Provisioned IOPS SSD (io1), Cold HDD (sc1), Throughput Optimized HDD (st1), and EBS Magnetic (standard). For more information about EBS volume type choices, see Amazon EBS.
Disk Per Node (DiskPerNode)	1	The number of NSD volumes to attach to each NSD server node. You can choose 1-5 disks.
Disk Size (DiskSize)	500	The disk size of NSD volume(s) attached to each NSD server node, in GiBs. Supported disk sizes are 10-16,384 GiB.

Table 14. Server Node Configurations

Parameter label (name)	Default	Description
Server Node Count (ServerNodeCount)	2	The number of EC2 instances to launch for the NSD server on the GPFS cluster. You can select 2-32 instances.
Server Instance Type (ServerInstanceType)	t2.medium	The instance type to use for the NSD server node instances.

Note: By default, all server nodes are assigned with an IBM Spectrum Scale server license.

Table 15. Compute Node Configurations

Parameter label (name)	Default	Description
Compute Node Count (ComputeNodeCount)	2	The number of IBM Spectrum 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 instances.

Table 16. 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.
Public Subnet 1 CIDR (PublicSubnet1CIDR)	10.0.0.0/24	The CIDR block for the public subnet located in availability zone 1.

Table 16. Network Configuration (continued)

Public Subnet 2 CIDR (PublicSubnet2CIDR)	10.0.2.0/24	The CIDR block for the public subnet located in availability zone 2.
Allowed External Access CIDR (RemoteAccessCIDR)	requires input	The CIDR block that is allowed external SSH access to the bastion hosts, e.g., <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 17. Amazon EC2 Configuration

Parameter label (name)	Default	Description
Key Pair Name (KeyPairName)	requires input	A public/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 you created in your preferred region.
Bastion AMI OS (BastionAMIOS)	Amazon-Linux- HVM	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 18. Personal Configuration

Parameter label (name)	Default	Description
Spectrum S3 Bucket (SpectrumS3Bucket)	requires input	An optional parameter that defines the S3 bucket name used for shared object storage among IBM Spectrum Scale nodes. An IAM role is set up to allow read/write operations on the specified bucket. You can provide the name of a pre-existing bucket you own or specify the name of a new bucket to create. Deployment logs are stored in this bucket in case stack creation fails.
Operator Email (OperatorEmail)	requires input	The email address to which notifications of any scaling operations is sent to.

Table 19. License Information

Parameter label (name)	Default	Description
IBM Customer Number (IBMCustomerNumber)	requires input	Required 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 Spectrum Scale on an existing Amazon VPC

The following requirements must be met for IBM Spectrum Scale to be deployed on an existing VPC.

- One private and one public network per availability zone. A minimum of two availability zones are required, so two private and two public subnets are needed.
- Private subnet requirements:
 1. Internet gateway (IGW) is not configured in the route table.
 2. Auto-assign public IPv4 address is disabled. The IBM Spectrum Scale nodes with public IP are not supported due to security issues.
 3. Configure the Network Address Translation (NAT) for this subnet. The IBM Spectrum Scale stack creation requires access to AWS S3 and RHUI servers.

- Public subnets are required for bastion nodes. It is recommended to use Amazon templates to create bastion nodes. The IBM Spectrum Scale stack creation does not have any requirement for public subnets.

Note: It is recommended to use the following:

1. Dedicated private subnet for the IBM Spectrum Scale cluster
 2. Subnet masks that are large enough to support maximum number of nodes
 3. Same subnet mask for all private subnets
 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 20. File System Configurations

Parameter label (name)	Default	Description
Block Size (BlockSize)	4M	The file system block size. You can choose a value from 256 KiB to 16MiB.
GPFS Mount Point(GpfsMountPoint)	/gpfs/fs1	The mount point for the IBM Spectrum Scale volume.

Table 21. NSD Configurations

Parameter label (name)	Default	Description
EBS Type (EBSType)	gp2	The EBS volume type for IBM Spectrum Scale storage attached to each NSD server node. Options are: General Purpose SSD (gp2), Provisioned IOPS SSD (io1), Cold HDD (sc1), Throughput Optimized HDD (st1), and EBS Magnetic (standard). For more information about EBS volume type choices, see Amazon EBS.
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 volume(s) attached to each NSD server node, in GiBs. Supported disk sizes are 10-16,384 GiB.

Table 22. Server Node Configurations

Parameter label (name)	Default	Description
Server Node Count (ServerNodeCount)	2	The number of EC2 instances to launch for the NSD server on the GPFS cluster. You can select 2-64 instances.
Server Instance Type(ServerInstanceType)	t2.medium	The instance type to use for the NSD server node instances.

Table 23. Compute Node Configurations

Parameter label (name)	Default	Description
Compute Node Count (ComputeNodeCount)	2	The number of IBM Spectrum 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 instances.

Table 24. 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.
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.
Allowed External Access CIDR (RemoteAccessCIDR)	requires input	The CIDR block that is allowed external SSH access to the bastion hosts, e.g., <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 25. Amazon EC2 Configuration

Parameter label (name)	Default	Description
Key Pair Name (KeyPairName)	requires input	A public/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 you created in your preferred region.
Bastion AMI OS (BastionAMIOS)	Amazon-Linux- HVM	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	TheEC2 instance type for the bastion host instances.

Table 26. Personal Configuration

Parameter label (name)	Default	Description
Spectrum S3 Bucket (SpectrumS3Bucket)	requires input	An optional parameter that defines the S3 bucket name used for shared object storage among IBM Spectrum Scale nodes. An IAM role is set up to allow read/write operations on the specified bucket. You can provide the name of a pre-existing bucket you own or specify the name of a new bucket to create.
Operator Email (OperatorEmail)	requires input	The email address to which notifications of any scaling operations is sent to.

Table 27. 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 .

Chapter 4. Cleaning up the cluster and the stack

To cleanup the stack, you must delete the stack created by the main template. This also deletes the IBM Spectrum Scale cluster.

Important: You can delete the stack, and clean up the alarm and Lambda console functions to free up memory space. This is a non-reversible action. Stacks once deleted cannot be accessed.

To delete the stack created by the main template follow these steps:

1. Open the CloudFormation console, and select the root stack to be deleted.
2. Go to **Actions > Delete Stack**.

When the stack is deleted, the Config Status field for CloudWatch alarms displays an Invalid notification message. The state of alarms which initially displays OK change to INSUFFICIENT_DATA after 5 minutes. Take the following steps to delete the alarms:

1. Open the CloudWatch console at <https://console.aws.amazon.com/cloudwatch/>.
2. In the navigation pane, go to **Alarms > Insufficient**.
3. Select the alarms corresponding to the stack, and click **Delete**.

The deletion operation on a stack does not delete the EBS volumes used for NSD servers. If they are no longer needed, the EBS volumes can be deleted by switching to AWS EBS console and deleting them manually. The EBS volumes have the following naming convention: <StackName>-ClusterStack-<ID>.

If the start_nodes lambda function has been created using the `mmaws create_lambda_function` command, it needs to be deleted using the following steps:

1. Open the lambda console at <https://console.aws.amazon.com/lambda/>.
2. In the **Functions** pane, select the **start_nodes** function name, and select **Actions > Delete**.

Chapter 5. Data security and AWS Identity and Access Management

The AWS cloud provides a scalable, highly reliable platform that helps customers deploy applications and data quickly and securely.

Security responsibilities are shared between IBM and AWS. AWS operates, manages, and controls the components from the host operating system and virtualization layer down to the physical security of the facilities in which the services operate. IBM assumes the responsibility and management of the guest operating system, including updates and security patches, other associated applications, as well as the configuration of the AWS-provided security group firewall. For more information about security on AWS, visit the [AWS Cloud Security Center](#)

IBM Spectrum Scale offering on AWS launches its instances in a logically isolated virtual network by using Amazon VPC. It supports launching the instances on a new VPC as well as on existing VPC.

Access to the IBM Spectrum Scale cluster is only allowed via Bastion host. It is a special purpose server instance designed to be the primary access point from the Internet and acts as a proxy to the other IBM Spectrum Scale instances. In case of deployment in a new VPC, the Bastion host is launched automatically. In case of deployment in existing VPC, it is a pre-requisite to have a bastion host configured.

This setup enforces login to the instances only through an EC2-user that has non-root privileges.

AWS Identity and Access Management (IAM)

This setup leverages an IAM role with the least privileged access. It is not necessary or recommended to store SSH keys, secret keys, or access keys on the provisioned instances.

The root user of the instances in a cluster can be accessed only by using the SSH key specified during the deployment process. AWS does not store these SSH keys, so loss of the SSH key can lead to loss of access to these instances. Updating operating system patches are the user's responsibility, and should be performed on a periodic basis.

A security group acts as a firewall that controls the traffic to one or more instances. When you launch an instance, you associate one or more security groups with the instance. You can add rules to each security group that allow traffic to or from its associated instances. You can modify the rules for a security group at any time. The new rules are automatically applied to all instances that are associated with the security group.

The security groups created and assigned to the individual instances as part of this solution are restricted as much as possible while allowing access to the various functions needed by IBM Spectrum Scale. It is recommended to review the security groups to further restrict access as needed once the cluster is up and running. The initial set up creates the following security groups for IBM Spectrum Scale:

BastionSecurityGroup

This security group is created by the nested Linux bastion stack when you deploy the AWS in a new VPC. It enables SSH access to the Linux bastion hosts.

ServerSecurityGroup

This group is for IBM Spectrum Scale NSD server instances. It allows SSH access for BastionSecurityGroup and enables communication between compute instances and NSD server instances.

ComputeSecurityGroup

This group is for IBM Spectrum Scale compute instances. It allows SSH access for BastionSecurityGroup and enables communication between compute instances and NSD server instances.

Chapter 6. Cluster lifecycle management and debug data collection

IBM Spectrum Scale on AWS provides convenient utilities and functions for cluster lifecycle management and debug data collection, such as expansion and contraction of cluster and starting and stopping of nodes.

The utility creates informational system snapshot at a single point in time to enforce debug data collection.

mmaaws utility

The **mmaaws** utility manages the IBM Spectrum Scale cluster lifecycle and workflows on AWS.

Synopsis

```
mmaaws add_nodes --node-type {compute,server} --node-instances NUM_INSTANCES
```

or

```
mmaaws remove_nodes(--ip-addresses IP_ADDRESSES [IP_ADDRESSES...] | --hostnames HOSTNAMES [HOSTNAMES...])
```

or

```
mmaaws list_instances
```

or

```
mmaaws start_nodes (--ip-addresses IP_ADDRESSES [IP_ADDRESSES...] |  
                  --instance-ids INSTANCE_IDS [INSTANCE_IDS...])
```

or

```
mmaaws stop_nodes (--all | --compute |  
                  --ip-addresses IP_ADDRESSES |  
                  --hostnames HOSTNAMES [HOSTNAMES...] |  
                  --nodeclass NODECLASS [NODECLASS...])
```

or

```
mmaaws create_lambda_functions [ --function-path FUNCTION_PATH ] --iam-role IAM_ROLE
```

or

```
mmaaws collect_debug_data
```

Availability

Available with IBM Spectrum Scale on AWS.

Description

Use the **mmaaws** utility to manage IBM Spectrum Scale cluster lifecycle and workflows on AWS.

Important: AWS credentials must be configured on the node by executing the **aws configure** command. For more information on the **aws configure** command, see [aws configure CLI command configuration](#). To obtain or manage AWS access keys, see [Managing Access Keys for IAM Users](#).

Parameters

add_nodes

Adds the compute or server nodes to the IBM Spectrum Scale cluster:

--node-type {compute,server}

Specifies the IBM Spectrum Scale node type.

--num-instances NUM_INSTANCES

Specifies the number of nodes to be added.

remove_nodes

Removes the compute or server nodes from the IBM Spectrum Scale cluster:

--ip-addresses IP_ADDRESSES [IP_ADDRESSES...]

Specifies the IP addresses of the nodes to be removed.

--hostnames *HOSTNAMES* [*HOSTNAMES...*]
Specifies the host names of the nodes to be removed.

list_instances

Lists the instances in the VPC.

start_nodes

Starts the specified nodes in an IBM Spectrum Scale cluster.

Important: This can be used only if at least one node in the cluster is alive. In case of a situation where all nodes are in a **Stopped** state, you can create a lambda function to the start nodes of the stack using the `create_lambda_functions` option.

--ip-addresses *IP_ADDRESSES* [*IP_ADDRESSES...*]
Specifies the IP addresses of the nodes to be started.

--instance-ids *INSTANCE_IDS* [*INSTANCE_IDS...*]
Specifies the instance ids to be started.

stop_nodes

Stops the specified nodes in an IBM Spectrum Scale cluster:

--all
Stops all the nodes in the cluster.

--compute
Stops all the nodes within a compute group.

--ip-addresses *IP_ADDRESSES* [*IP_ADDRESSES...*]
Specifies the IP addresses of the nodes to be stopped.

--hostnames *HOSTNAMES* [*HOSTNAMES...*]
Specifies the hostnames of the nodes to be stopped.

--nodeclass *NODECLASS* [*NODECLASS...*]
Specifies the node classes to be stopped.

Note: Server nodes will not be stopped by using the `--compute`, `--ip-addresses`, `--hostnames` or `--nodeclass` options. However, all the nodes can be stopped by using the `--all` option.

create_lambda_functions

Creates a lambda function.

--function-path *FUNCTION_PATH*
Specifies the specific lambda function call to be created. The default value for this is set to **all**.

--iam-role *IAM_ROLE*
Specifies the IAM role to be associated with the lambda functions. The default value for this is set to **None**.

Note: The iam-role provided for this utility must contain the `AWSLambdaBasicExecutionRole` policy

collect_debug_data

Collects debug data.

Note: This creates a snapshot at a single point in time. This system snapshot consists of the following information:

- RHEL version
- Kernel version
- AWS auto scaling group properties
- Deployment logs

- CloudFormation stack events
- AWS - Spectrum Scale NSD mappings
- AWS vs. Spectrum Scale node mapping
- AWS instances root device properties

The information gathered is stored at `/var/adm/ras/aws_scale_logs.<date_timestamp>.tar.gz`. This option does not perform `gpfs.snap` operation. The `gpfs.snap` operations must be done separately.

Exit status

0 Successful completion.

nonzero

A failure has occurred.

Security

You must have root authority to run the **mmaws** utility.

Examples

1. To add compute nodes:

```
# mmaws add_nodes --node-type compute --num-instance 1
```

The system displays output similar to this:

```
2018-07-31 08:33:38,084 - mm_aws_add_rm_nodes - INFO - Logging in to file: /var/adm/ras/aws_scale_logs/mm_aws_add_rm_nodes.log_2018-Jul-31_08-33-38
2018-07-31 08:33:38,084 - mm_aws_add_rm_nodes - INFO - A. Performing prerequisite check
2018-07-31 08:33:38,427 - mm_aws_add_rm_nodes - INFO - B. Performing stack resource identification
.....
2018-07-31 08:33:45,817 - mm_aws_add_rm_nodes - INFO - Created new instance with ID: i-0e6a246636eb25e39
2018-07-31 08:33:45,818 - mm_aws_add_rm_nodes - INFO - Waiting to obtain 'ok' status for instance id(s): ['i-0e6a246636eb25e39']
.....
2018-07-31 08:40:17,972 - mm_aws_add_rm_nodes - INFO - Operation (mmmount all -N 10.0.1.149) completed successfully.
2018-07-31 08:40:17,972 - mm_aws_add_rm_nodes - INFO - *****
2018-07-31 08:40:17,972 - mm_aws_add_rm_nodes - INFO - Adding new compute instance(s) to cluster completed successfully.
2018-07-31 08:40:17,973 - mm_aws_add_rm_nodes - INFO - *****
```

2. To add server nodes:

```
# mmaws add_nodes --node-type server --num-instance 2
```

The system displays output similar to this:

```
2018-07-31 08:52:24,957 - mm_aws_add_rm_nodes - INFO - Logging in to file: /var/adm/ras/aws_scale_logs/mm_aws_add_rm_nodes.log_2018-Jul-31_08-52-24
2018-07-31 08:52:24,958 - mm_aws_add_rm_nodes - INFO - A. Performing prerequisite check
2018-07-31 08:52:25,305 - mm_aws_add_rm_nodes - INFO - B. Performing stack resource identification
.....
2018-07-31 08:52:29,728 - mm_aws_add_rm_nodes - INFO - Created new instance with ID: i-05a64309d0b2cd5ae
2018-07-31 08:52:30,703 - mm_aws_add_rm_nodes - INFO - Created new instance with ID: i-01eb8be276a344971
.....
2018-07-31 08:59:12,549 - mm_aws_add_rm_nodes - INFO - Operation (mmmount all -N 10.0.3.47,10.0.1.235) completed successfully.
2018-07-31 08:59:12,549 - mm_aws_add_rm_nodes - INFO - *****
2018-07-31 08:59:12,549 - mm_aws_add_rm_nodes - INFO - Adding new server instance(s) to cluster completed successfully.
2018-07-31 08:59:12,549 - mm_aws_add_rm_nodes - INFO - *****
```

3. To remove nodes that are using a specific IP address:

```
# mmaws remove_nodes --ip-addresses 10.0.3.47 10.0.1.235 10.0.1.149
```

The system displays output similar to this:

```
2018-07-31 09:05:33,844 - mm_aws_add_rm_nodes - INFO - Logging in to file: /var/adm/ras/aws_scale_logs/mm_aws_add_rm_nodes.log_2018-Jul-31_09-05-33
2018-07-31 09:05:33,844 - mm_aws_add_rm_nodes - INFO - A. Performing prerequisite check
2018-07-31 09:05:34,203 - mm_aws_add_rm_nodes - INFO - B. Performing stack resource identification
.....
2018-07-31 09:06:09,966 - ibm_aws_scale_utils.autoscaling_utils - INFO - Terminating instance (i-0e6a246636eb25e39) from autoscaling group completed successfully.
2018-07-31 09:06:10,126 - ibm_aws_scale_utils.autoscaling_utils - INFO - Updating max size of autoscaling group completed successfully.
2018-07-31 09:06:10,127 - mm_aws_add_rm_nodes - INFO - *****
2018-07-31 09:06:10,127 - mm_aws_add_rm_nodes - INFO - Removal of specified compute instance(s) from cluster completed successfully.
2018-07-31 09:06:10,127 - mm_aws_add_rm_nodes - INFO - *****
```

```

.....
2018-07-31 09:06:09,760 - ibm_aws_scale_utils.autoscaling_utils - INFO - Terminating instance (i-05a64309d0b2cd5ae) from autoscaling group completed successfully.
2018-07-31 09:06:09,966 - ibm_aws_scale_utils.autoscaling_utils - INFO - Terminating instance (i-01eb8be276a344971) from autoscaling group completed successfully.
2018-07-31 09:06:10,126 - ibm_aws_scale_utils.autoscaling_utils - INFO - Updating max size of autoscaling group completed successfully.
2018-07-31 09:06:10,127 - mm_aws_add_rm_nodes - INFO - *****
2018-07-31 09:06:10,127 - mm_aws_add_rm_nodes - INFO - Removal of specified server instance(s) from cluster completed successfully.
2018-07-31 09:06:10,127 - mm_aws_add_rm_nodes - INFO - *****

```

4. To list the nodes:

```
# mmaws list_nodes
```

The system displays output similar to this:

Compute Node/Instances List:

Instance Id	Private IP	Instance Type	AvailabilityZone	State
i-0601f80f96812e70f	10.0.1.237	t2.micro	us-east-2a	running
i-0b25af79f2f1fd4ff	10.0.3.96	t2.micro	us-east-2b	running

NSD Server Node/Instances List:

Instance Id	Private IP	Instance Type	AvailabilityZone	State
i-0693826b9e5648e58	10.0.3.175	t2.micro	us-east-2b	running
i-0979e758a545da2ed	10.0.1.226	t2.micro	us-east-2a	running

5. To stop nodes that are using a specific IP address:

```
# mmaws stop_nodes --ip-addresses 10.0.1.237
```

The system displays output similar to this:

```

2018-07-31 08:28:19,351 - mm_aws_stop_nodes - INFO - Logging in to file: /var/adm/ras/aws_scale_logs/mm_aws_stop_nodes.log_2018-Jul-31_08-28-19
2018-07-31 08:28:19,351 - mm_aws_stop_nodes - INFO - A. Performing prerequisite check
2018-07-31 08:28:19,351 - mm_aws_stop_nodes - INFO - B. Performing stack resource identification
2018-07-31 08:28:20,043 - mm_aws_stop_nodes - INFO - C. Performing stop operation
2018-07-31 08:28:20,377 - mm_aws_stop_nodes - INFO - Unmounting filesystem on nodes (['10.0.1.237'])
2018-07-31 08:28:22,492 - mm_aws_stop_nodes - INFO - Shutting down gpfs on nodes (['10.0.1.237'])
2018-07-31 08:28:37,582 - mm_aws_stop_nodes - INFO - Status of node(s):
2018-07-31 08:28:37,582 - mm_aws_stop_nodes - INFO - i-0601f80f96812e70f : stopping

```

6. To start the nodes that are using a specific IP address:

```
# mmaws start_nodes --ip-addresses 10.0.1.237
```

The system displays output similar to this:

```

2018-07-31 08:22:30,759 - mm_aws_start_nodes - INFO - Logging in to file: /var/adm/ras/aws_scale_logs/mm_aws_start_nodes.log_2018-Jul-31_08-22-30
2018-07-31 08:22:30,759 - mm_aws_start_nodes - INFO - A. Performing prerequisite check
2018-07-31 08:22:30,759 - mm_aws_start_nodes - INFO - B. Performing stack resource identification
2018-07-31 08:22:32,051 - mm_aws_start_nodes - INFO - C. Performing start operation
2018-07-31 08:22:32,319 - mm_aws_start_nodes - INFO - Status of node(s):
2018-07-31 08:22:32,319 - mm_aws_start_nodes - INFO - 10.0.1.237 : pending

```

7. To create a lambda function that can start nodes even if no node is available:

```
# mmaws create_lambda_func --function-path /usr/lpp/mmfs/bin/ibm_aws_workflows/ibm_aws_lambda_functions/mm_aws_lambda_start_nodes
--iam-role lambda_role
```

The system displays output similar to this:

```

2018-07-31 08:32:02,023 - mm_aws_create_lambda_functions - INFO - Logging in to file: /var/adm/ras/aws_scale_logs/mm_aws_create_lambda_functions.log_2018-Jul-31_08-32-02
2018-07-31 08:32:02,024 - mm_aws_create_lambda_functions - INFO - A. Performing prerequisite check
2018-07-31 08:32:02,024 - mm_aws_create_lambda_functions - INFO - B. Performing current stack resource identification
2018-07-31 08:32:02,742 - mm_aws_create_lambda_functions - INFO - Provided IAM role (lambda_role) contains required "AWSLambdaBasicExecutionRole" policy.
2018-07-31 08:32:02,742 - mm_aws_create_lambda_functions - INFO - C. Performing Lambda function vs. filepath mapping
2018-07-31 08:32:02,742 - mm_aws_create_lambda_functions - INFO - Identified Lambda function(s) vs. filepath:
{'start_nodes': '/usr/lpp/mmfs/bin/ibm_aws_workflows/ibm_aws_lambda_functions/mm_aws_lambda_start_nodes.py'}
2018-07-31 08:32:02,743 - mm_aws_create_lambda_functions - INFO - D. Performing Lambda function creation
2018-07-31 08:32:03,176 - mm_aws_create_lambda_functions - INFO - Lambda function (start_nodes) created successfully.

```

8. To collect debug data:

```
# mmaws collect_debug_data
```

The system displays output similar to this:

```

2018-07-31 08:11:36,636 - mm_aws_collect_data - INFO - Logging in to file: /var/adm/ras/aws_scale_logs/mm_aws_collect_data.log_2018-Jul-31_08-11-36
2018-07-31 08:11:36,636 - mm_aws_collect_data - INFO - A. Performing prerequisite check
2018-07-31 08:11:37,005 - mm_aws_collect_data - INFO - B. Collecting RHEL / OS version, kernel Details
2018-07-31 08:11:37,931 - mm_aws_collect_data - INFO - C. Collecting AutoScaling group properties
2018-07-31 08:11:38,152 - mm_aws_collect_data - INFO - D. Collecting AWS - Spectrum Scale Deployment logs, events
2018-07-31 08:11:38,849 - mm_aws_collect_data - INFO - E. Collecting AWS - Spectrum Scale NSD mapping

```

```

2018-07-31 08:11:46,939 - mm_aws_collect_data - INFO - F. Identifying AWS vs. Spectrum Scale node mapping
2018-07-31 08:11:47,541 - mm_aws_collect_data - INFO - G. Identifying AWS instances root device properties
2018-07-31 08:11:50,202 - mm_aws_collect_data - INFO - H. Compressing all debug data related data
2018-07-31 08:11:50,219 - mm_aws_collect_data - INFO - *****
2018-07-31 08:11:50,219 - mm_aws_collect_data - INFO - Compressed Log is available at /var/adm/ras/aws_scale_logs.2018-Jul-31_08-11-50.tar.gz
2018-07-31 08:11:50,219 - mm_aws_collect_data - INFO - *****

```

Location

/usr/lpp/mmfs/bin

Lambda function to start nodes

AWS Lambda is a stateless compute service that lets you run code without provisioning or managing servers.

IBM Spectrum Scale stack leverages the Lambda function to start compute or server nodes. The **mmaws create_lambda_functions** command create a Lambda function which helps to start all the nodes, even if all or some of the nodes in the cluster are in a **Stopped** state.

Consider the following before creating a Lambda function:

- The **mmaws create_lambda_functions** command creates a Lambda function **start_nodes** in the region where the IBM Spectrum Scale stack is deployed. This function needs to be created prior to stopping all nodes.
- The Lambda function creation requires the presence of a role with **AWSLambdaBasicExecutionRole** policy as prerequisite. For more information on IAM role creation, see [Create the Execution Role \(IAM Role\)](#)
- The Lambda function is created with the following parameters:
 - **handler:** lambda_handler
 - **timeout:** 300
 - **memory_size:** 128
- AWS charges only for the compute time or Lambda function execution time, and there is no charge when this function is not running.

```

# mmaws create_lambda_functions --function-path /usr/lpp/mmfs/bin/ibm_aws_workflows/ibm_aws_lambda_functions/mm_aws_lambda_start_nodes.py --iam-role lambda_role
2018-07-31 08:32:02,023 - mm_aws_create_lambda_functions - INFO - Logging in to file: /var/adm/ras/aws_scale_logs/mm_aws_create_lambda_functions.log_2018-Jul-31_08-32-02
2018-07-31 08:32:02,024 - mm_aws_create_lambda_functions - INFO - A. Performing prerequisite check
2018-07-31 08:32:02,024 - mm_aws_create_lambda_functions - INFO - B. Performing current stack resource identification
2018-07-31 08:32:02,742 - mm_aws_create_lambda_functions - INFO - Provided IAM role (lambda_role) contains required "AWSLambdaBasicExecutionRole" policy.
2018-07-31 08:32:02,742 - mm_aws_create_lambda_functions - INFO - C. Performing Lambda function vs. filepath mapping
2018-07-31 08:32:02,742 - mm_aws_create_lambda_functions - INFO - Identified Lambda function(s) vs. filepath:
{'start_nodes': '/usr/lpp/mmfs/bin/ibm_aws_workflows/ibm_aws_lambda_functions/mm_aws_lambda_start_nodes.py'}
2018-07-31 08:32:02,743 - mm_aws_create_lambda_functions - INFO - D. Performing Lambda function creation
2018-07-31 08:32:03,176 - mm_aws_create_lambda_functions - INFO - Lambda function (start_nodes) created successfully.

```

Note: You can also run this CLI without a function-path as follows:

```
mmaws create_lambda_func --iam-role lambda_role
```

This creates a Lambda function which can be viewed from the Lambda console.

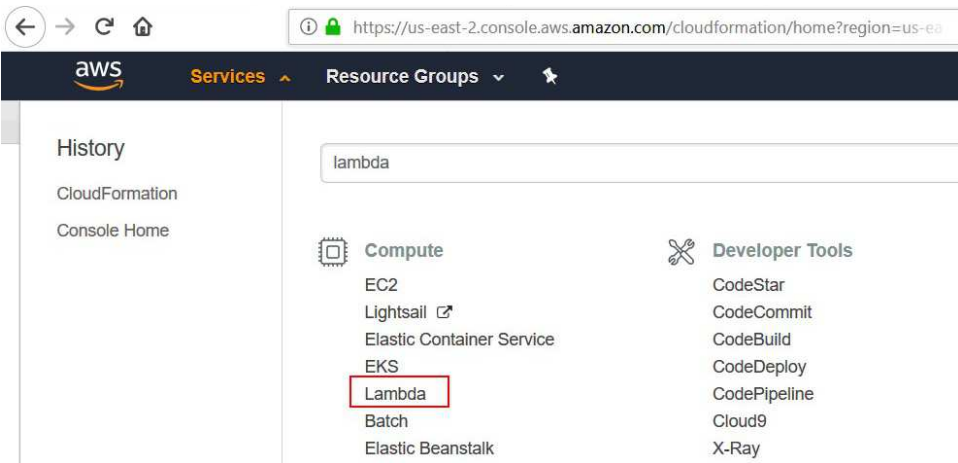


Figure 5. Lambda console

The `start_nodes` is the new Lambda function created, and its function details can be viewed by clicking on **start_nodes**.

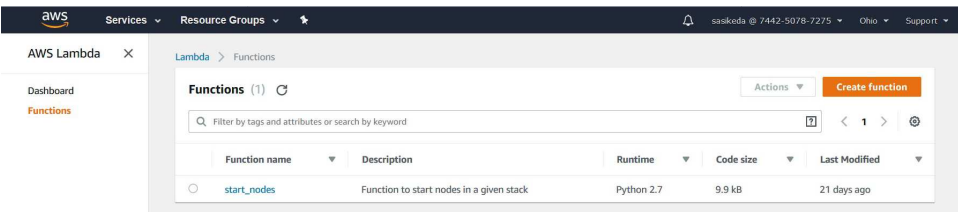


Figure 6. Viewing Lambda function details

The `start_nodes` function code can be viewed under the **Function code** tab.

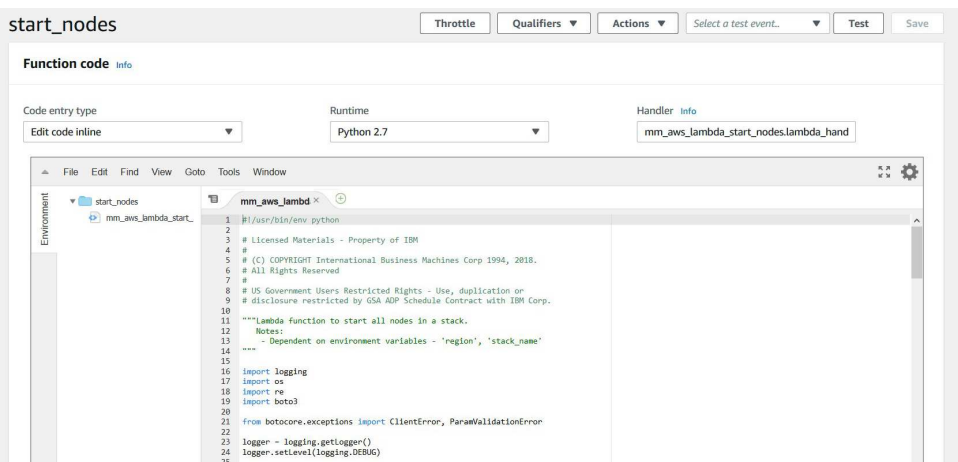


Figure 7. Function code information

The `start_nodes` function is populated with environment variables such as `region` and `stack_para_ids`. These variables are needed for the execution of the function.

start_nodes Throttle Qualifiers Actions Select a test event... Test Save

Environment variables

You can define Environment Variables as key-value pairs that are accessible from your function code. These are useful to store configuration settings without the need to change function code. [Learn more.](#)

region	us-east-2	Remove
stack_para_id1	87bc9280	Remove
stack_para_id2	947f	Remove
stack_para_id3	11e8	Remove
stack_para_id4	9c17	Remove
stack_para_id5	0a7dfa8cb208	Remove
Key	Value	Remove

► Encryption configuration

Figure 8. Environment variables

The `start_nodes` function is populated with role, memory, and timeout information.

start_nodes Throttle Qualifiers Actions Select a test event... Test Save

Execution role

Defines the permissions of your function. Note that new roles may not be available for a few minutes after creation. [Learn more](#) about Lambda execution roles.

Choose an existing role ▼

Existing role
You may use an existing role with this function. Note that the role must be assumable by Lambda and must have Cloudwatch Logs permissions.

lambda_role ▼

Basic settings

Description
Function to start nodes in a given stack

Memory (MB) [Info](#)
Your function is allocated CPU proportional to the memory configured.
128 MB

Timeout [Info](#)
5 min 0 sec

Figure 9. Environment variables

Note: This function can be executed by creating a test event.

Follow these steps to execute the Lambda function:

1. Click on **Test**.

Lambda > Functions > start_nodes ARN - arn:aws:lambda:us-east-2:744250787275:function:start_nodes

start_nodes Throttle Qualifiers Actions Select a test event... **Test** Save

Figure 10. Execute Lambda function

2. In the **Configure test event** page, enter the event name of your choice and remove the existing key values shown in the curly braces. Click **Create**.

Configure test event



A function can have up to 10 test events. The events are persisted so you can switch to another computer or web browser and test your function with the same events.

- ☒ Create new test event
- ☐ Edit saved test events

Event template

Hello World ▼

Event name

myTestEvent

```
1 {  
2  
3 }
```

bl1dd13

Figure 11. Execute Lambda function

Note: The event created in the previous step can be found in the **Select test event** box.



Figure 12. Select test event

Important: The `mmaws start_nodes` command can also be used to start nodes. However, at least one node of the cluster must be running in order to start the other stopped nodes of the cluster using the `mmaws start_nodes` command.

Chapter 7. Upgrading IBM Spectrum Scale

Follow these steps to upgrade your version of IBM Spectrum Scale.

1. Download the latest IBM Spectrum Scale Data Management bundle from IBM Fix Central and upload it to the IBM Spectrum Scale EC2 instance or an S3 bucket.
2. Download the latest workflow rpm from Developer Works and install it in the existing cluster.
3. Upgrade IBM Spectrum Scale by running the **aws_upgrade_scale** command on any of the AWS EC2 instances:

```
aws_upgrade_scale [-h][-v] [-s S3BUCKET | -l LOCALPATH] [-b BUILDFILE] [-V VERSION]
```

The command has the following options:

- h** Displays the options for this script.
- s** Specifies the s3bucket that contains IBM Spectrum Scale Data Management bundle.
- l** Specifies the local path of the directory or folder that contains IBM Spectrum Scale Data Management bundle.
- b** Specifies the build file name to be upgraded.
- V** Provides the version information to be upgraded to.
- v** Displays the verbose output.
- f** Specifies the force option to silently accept the IBM Spectrum Scale electronic license agreement.

You can upload the bundle into the AWS using one of the following methods:

- Download the IBM Spectrum Scale bundle from S3 using the **-s** option.
- Copy the IBM Spectrum Scale bundle manually using the **-l** option.

The following code displays an example of the command to upgrade from IBM Spectrum Scale version 5.0.1.1 to IBM Spectrum Scale version 5.0.1.2:

```
[ec2-user@ip-10-0-1-187]$/usr/lpp/mmfs/bin/ibm_aws_workflows/aws_upgrade_scale_1 -s s3_bucket/5.0.1.2_PTF -V 5.0.1.2 -b Spectrum_Scale_Protocols_Data_Management-5.0.1.2-x86_64-Linux-install -v
```

The following code displays an example of the command to upgrade from IBM Spectrum Scale version 5.0.1.1 to IBM Spectrum Scale version 5.0.2:

```
[ec2-user@ip-10-0-1-187]$/usr/lpp/mmfs/bin/ibm_aws_workflows/aws_upgrade_scale_1 -s s3_bucket/5.0.2_GA -V 5.0.2.0 -b Spectrum_Scale_Data_Management-5.0.2.0-x86_64-Linux-install -v
```

Note: Ensure that the installer directory `/usr/lpp/mmfs/<version>/installer` exists.

Chapter 8. Troubleshooting

The following issues have been encountered in AWS cloud:

CREATE_FAILED error with timeout message is encountered on launching AMI

If AWS CloudFormation fails to create the stack, it is recommended that you relaunch the template with **Rollback on failure** set to **No**.

This setting is under the **Advanced** section in the **AWS CloudFormation** console under the **Options** page. With this setting, the stack's state is retained and the instance are left running, so you can troubleshoot the issue. For details about the stack, look at the log files in `/var/adm/ras/aws_scale_logs/cfn-init.log` and the details log file in `/var/adm/ras/aws_scale_logs/aws_bootstrap_setup.log`.

Note: When you set the **Rollback on failure** to **No**, you continue to incur AWS charges for this stack. Ensure that the stack is deleted when you have finished troubleshooting.

For more information, see *Troubleshooting AWS CloudFormation* on the AWS website, or the *AWSQuick Start Discussion Forum*.

Service.RequestLimitExceeded error in the cfn-init-cmd.log

Stack creation fails if you encounter a `Service.RequestLimitExceeded` error in the `cfn-init-cmd.log`.

You might encounter this error if you try to deploy a large cluster that exceeds your account's limits. To address this problem, request a service limit increase for the EC2 instance types that you intend to deploy. To do this in the AWS Support Center, click **Create Case > Service Limit Increase > EC2 instances**, and then complete the fields in the **Limit increase** form.

Size limitation error on deploying the AWS CloudFormation templates

If you deploy the templates from a local copy on your computer or from a non-S3 location, you might encounter template size limitations when you create the stack.

It is recommended that you launch the AWS CloudFormation templates from the location that is provided or from another S3 bucket. For more information about AWS CloudFormation limits, see the *AWS CloudFormation Limits*.

Stack creation failure message encountered

In the AWS CloudFormation **Events** tab, the reason for stack creation failure is described as follows: The maximum number of addresses has been reached.

You might encounter this error if you try to deploy two clusters in an AWS region by using the default root or IAM user account settings, but you have insufficient elastic IP addresses available for the AMI deployment. The IBM Spectrum Scale deployment requires three elastic IP addresses. You can deploy the IBM Spectrum Scale offering on AWS in a different AWS Region where you are not using elastic IP addresses, or you can use one of the following approaches:

- For the region in which you are hitting a failure to allocate elastic IPs, request a limit increase for your elastic IP limit.

- If you have elastic IP addresses allocated that you do not need, delete some of these addresses. Deleting unneeded addresses frees up the elastic IP addresses available in the region where you intend to launch IBM Spectrum Scale.

Chapter 9. Frequently Asked Questions

This page details some frequently asked questions.

How to create an IAM user account that can satisfy all the requirements needed to deploy the IBM Spectrum Scale on AWS?

Ensure that any sub-account you use to launch the IBM Spectrum Scale has the required policies enabled. For instructions on creating an IAM user account with the policies needed to deploy the IBM Spectrum Scale, see [IAM user to deploy IBM Spectrum Scale on AWS](#).

How to resolve issues encountered while running IBM Spectrum Scale?

IBM provides support for IBM Spectrum Scale issues through the [IBM Spectrum Scale forum](#). You can also get support by sending a mail to the IBM Spectrum Scale mailing list at scale@us.ibm.com.

How can I find out which version of the template and IBM Spectrum Scale code I deployed?

After the deployment is complete, click the **Outputs** tab for the stack created in the AWS CloudFormation console. This tab displays the version of the CloudFormation template, the version of IBM Spectrum Scale that was deployed, and the name of the optional S3 bucket used by the IBM Spectrum Scale, which is set by the **SpectrumS3bucket** parameter.

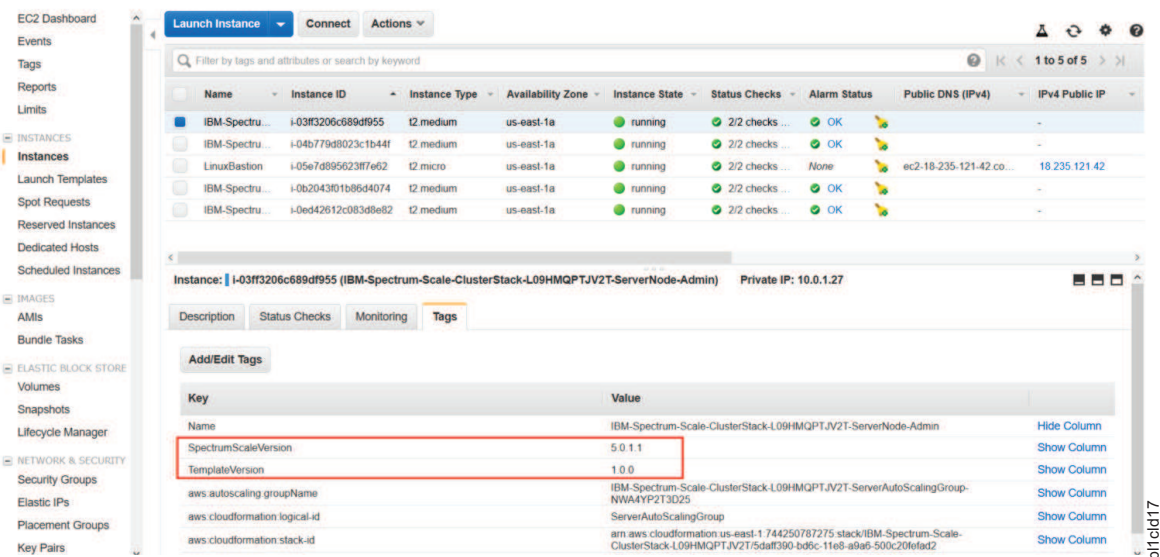


Figure 13. IBM Spectrum Scale stack outputs

Accessibility features for IBM Spectrum 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 Spectrum 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 Knowledge Center, and its related publications, are accessibility-enabled. The accessibility features are described in IBM Knowledge Center (www.ibm.com/support/knowledgecenter).

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) 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.

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.

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

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

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

Terms and conditions for product documentation

Permissions for the use of these publications are granted subject to the following terms and conditions.

Applicability

These terms and conditions are in addition to any terms of use for the IBM website.

Personal use

You may reproduce these publications for your personal, noncommercial use provided that all proprietary notices are preserved. You may not distribute, display or make derivative work of these publications, or any portion thereof, without the express consent of IBM.

Commercial use

You may reproduce, distribute and display these publications solely within your enterprise provided that all proprietary notices are preserved. You may not 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 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 may not download, export or re-export 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.

IBM Online Privacy Statement

IBM Software products, including software as a service solutions, ("Software Offerings") may use cookies or other technologies to collect product usage information, to help improve the end user experience, to tailor interactions with the end user or for other purposes. In many cases no personally identifiable information is collected by the Software Offerings. Some of our Software Offerings can help enable you to

collect personally identifiable information. If this Software Offering uses cookies to collect personally identifiable information, specific information about this offering's use of cookies is set forth below.

This Software Offering does not use cookies or other technologies to collect personally identifiable information.

If the configurations deployed for this Software Offering provide you as customer the ability to collect personally identifiable information from end users via cookies and other technologies, you should seek your own legal advice about any laws applicable to such data collection, including any requirements for notice and consent.

For more information about the use of various technologies, including cookies, for these purposes, See IBM's Privacy Policy at <http://www.ibm.com/privacy> and IBM's Online Privacy Statement at <http://www.ibm.com/privacy/details> the section entitled "Cookies, Web Beacons and Other Technologies" and the "IBM Software Products and Software-as-a-Service Privacy Statement" at <http://www.ibm.com/software/info/product-privacy>.

Glossary

This glossary provides terms and definitions for IBM Spectrum 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) (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 Object 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.

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 connection, 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 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

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 fileset

A fileset 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 fileset, which enables more efficient per-fileset 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 fileset to the root directory of another fileset.

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.

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

A IP address used to communicate on a private network.

public IP address

A 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 involves 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 Spectrum Protect Hierarchical Storage Management (HSM).

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 Spectrum 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).

Index

A

- accessibility features for IBM Spectrum Scale 39
- Amazon Auto Scaling 1
- Amazon CloudWatch 1
- Amazon EBS 1
- Amazon EC2 1
- Amazon IAM 1
- Amazon S3 1
- Amazon VPC 1
- Amazon web services
 - IBM Spectrum Scale 1, 2, 3, 5, 7, 9, 12, 13, 14, 16, 19, 21, 23, 24, 28, 33, 35, 37
- availability zone 16
 - new Amazon VPC
 - multiple availability zone 14
- availability zones 2
- AWS cloud 19
- AWS Cloud 1, 2, 3, 5, 7, 9, 12, 13, 14, 21, 23, 24, 28, 35, 37
 - availability zones 2, 7
 - Bastion host 21
 - Data security 2, 21
 - deployment
 - deployment options 12
 - Deployment 9
 - EC2-user 21
 - Frequently Asked Questions 37
 - IBM Spectrum Scaleupgrade 33
 - Instance types 3
 - Operating system 3
 - regions 2
 - Setup
 - Optimal setup 7
 - Troubleshooting 35
 - usage restrictions 3
- AWS CloudFormation 1
- AWS CloudFormation templates deployment 35
- AWS Lambda
 - Lambda function 28
- AWS Lambda function 1

B

- Bastion host 21, 23

C

- command
 - mmaws 19

D

- Data security 21, 23
- Delete
 - cluster 19
 - stack 19
- Deployment 9
 - existing VPC 16
 - multiple availability zone 14
 - new VPC 14

- Deployment (*continued*)
 - single availability zone 13
- deployment options
 - existing Amazon VPC 9, 12
 - new Amazon VPC
 - multiple availability zone 9, 12
 - single availability zone 9, 12

E

- EC2-user 21, 23
- existing Amazon VPC 16

F

- Frequently Asked Questions 37

G

- GPFS
 - mmaws utility 24

I

- IBM Spectrum Scale information units vii
- IBM Spectrum Scaleon AWS 1, 2, 3, 5, 7, 9, 12, 13, 14, 16, 19, 21, 23, 24, 28, 33, 35, 37
- Identity and Access Management
 - IAM 21
- Instance types 3

L

- Lambda function 19
 - execute 28

M

- mmaws 24

N

- new Amazon VPC
 - multiple availability zone 14
 - single availability zone 13

O

- Operating system 3

R

- regions 2

S

- Service.RequestLimitExceeded error 35
- setup
 - optimal setup 7
- Setup 5
- size limitation error 35
- Stack creation failure 35
- Stack creation failure message 35

T

- Troubleshooting
 - AWS CloudFormation templates deployment 35
 - AWS CloudIBM Spectrum Scale 35
 - CREATE_FAILED error 35
 - no timeout message 35
 - Service.RequestLimitExceeded error 35
 - size limitation error 35
 - Stack creation failure 35
 - Stack creation failure message 35

U

- usage restrictions 3
- utility
 - mmaws 24



Product Number: 5725-Q01
5641-GPF
5725-S28
5765-ESS

Printed in USA

SC27-9283-01

