

IBM Spectrum Scale On AWS
Version 1.3.1

*IBM Spectrum Scale
On AWS Guide*



Note

Before using this information and the product it supports, read the information in [“Notices” on page 59.](#)

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

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

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 xxii.](#) When you send information to IBM, you grant IBM a nonexclusive right to use or distribute the information in any way it believes appropriate without incurring any obligation to you.

© **Copyright International Business Machines Corporation 2018, 2020.**

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.....	xxi
Conventions used in this information.....	xxi
How to send your comments.....	xxii
Summary of changes.....	xxv
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 systems.....	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.....	13
Option 1: Deploying into a new Amazon VPC using a single availability zone.....	13
Option 2: Deploying IBM Spectrum Scale on a new Amazon VPC with multiple Availability Zones	16
Option 3: Deploying IBM Spectrum Scale on an existing Amazon VPC.....	20
Chapter 4. Creating custom AMI.....	25
Chapter 5. Cluster lifecycle management	27
mmcloudworkflows utility.....	27
Chapter 6. Accessing IBM Spectrum Scale GUI in AWS.....	31
Chapter 7. Active file management on AWS.....	33
Preparing the environment for AFM.....	33
AFM cache modes.....	35
Deploying AFM on AWS.....	39
Configuration best practices.....	41
Limitations of AFM on AWS.....	41
Chapter 8. Upgrading IBM Spectrum Scale.....	43
Chapter 9. Cleaning up the cluster and the stack.....	45
Chapter 10. Data security and AWS Identity and Access Management.....	47
Chapter 11. Diagnosing and cleaning-up deployment failures.....	49
Chapter 12. Collecting debug data.....	51

Chapter 13. Troubleshooting	53
CREATE_FAILED error with timeout message is encountered on launching AMI.....	53
Service.RequestLimitExceeded error in the cfn-init-cmd.log	53
Size limitation error on deploying the AWS CloudFormation templates.....	53
Stack creation failure message encountered	53
Chapter 14. Frequently Asked Questions	55
Accessibility features for IBM Spectrum Scale	57
Accessibility features.....	57
Keyboard navigation.....	57
IBM and accessibility.....	57
Notices	59
Trademarks.....	60
Terms and conditions for product documentation.....	60
IBM Online Privacy Statement.....	61
Glossary	63
Index	71

Tables

- 1. IBM Spectrum Scale library information units..... viii
- 2. Conventions.....xxii
- 3. xxv
- 4. File system configurations..... 13
- 5. NSD Configurations..... 13
- 6. Server Node Configurations..... 14
- 7. Compute Node Configurations.....14
- 8. Network Configuration..... 15
- 9. Amazon EC2 Configuration..... 15
- 10. Personal Configuration.....16
- 11. License Information..... 16
- 12. File System Configurations..... 16
- 13. NSD Configurations..... 16
- 14. Server Node Configurations..... 17
- 15. Compute Node Configurations..... 18
- 16. Network Configuration..... 18
- 17. Amazon EC2 Configuration..... 19
- 18. Personal Configuration.....19
- 19. License Information..... 19
- 20. File System Configurations..... 20
- 21. NSD Configurations..... 20
- 22. Server Node Configurations..... 21
- 23. Compute Node Configurations..... 22

24. Network Configuration.....	22
25. Amazon EC2 Configuration.....	23
26. Personal Configuration.....	23
27. License Information.....	23
28. IBM Spectrum Scale BYOL AWS Marketplace Functional Support Matrix.....	55

About this information

This edition applies to IBM Spectrum Scale version 5.0.5 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
<p><i>IBM Spectrum Scale: Concepts, Planning, and Installation Guide</i></p>	<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 • Introduction to watch folder API • Introduction to clustered watch folder • 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 • Understanding call home <p>Planning</p> <ul style="list-style-type: none"> • Planning for GPFS • Planning for protocols • Planning for Cloud services • Planning for AFM • Planning for AFM DR • Firewall recommendations • Considerations for GPFS applications • Security-Enhanced Linux support • Space requirements for call home data upload 	<p>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</p>

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

Information unit	Type of information	Intended users
<p><i>IBM Spectrum Scale: Concepts, Planning, and Installation Guide</i></p>	<p>Installing</p> <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 • Installation of Active File Management (AFM) • Installing and upgrading AFM-based Disaster Recovery • Installing call home • Installing file audit logging • Installing watch folder API • Installing clustered watch folder • Steps to permanently uninstall GPFS <p>Upgrading</p> <ul style="list-style-type: none"> • IBM Spectrum Scale supported upgrade paths • Online upgrade support for protocols and performance monitoring 	<p>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</p>

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

Information unit	Type of information	Intended users
<p><i>IBM Spectrum Scale: Concepts, Planning, and Installation Guide</i></p>	<ul style="list-style-type: none"> • Upgrading IBM Spectrum® Scale non-protocol Linux nodes • Upgrading IBM Spectrum Scale protocol nodes • Upgrading AFM and AFM DR • Upgrading object packages • Upgrading SMB packages • Upgrading NFS packages • Upgrading call home • Manually upgrading the performance monitoring tool • Manually upgrading pmswift • Manually upgrading the IBM Spectrum Scale management GUI • Upgrading Cloud services • Upgrading to IBM Cloud Object Storage software level 3.7.2 and above • Upgrade paths and commands for file audit logging, watch folder API, and clustered watch folder • Upgrading with clustered watch folder enabled • Upgrading IBM Spectrum Scale components with the installation toolkit • Changing IBM Spectrum Scale product edition • Completing the upgrade to a new level of IBM Spectrum Scale • Reverting to the previous level of IBM Spectrum Scale 	<p>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</p>
<p><i>IBM Spectrum Scale: Concepts, Planning, and Installation Guide</i></p>	<ul style="list-style-type: none"> • Coexistence considerations • Compatibility considerations • Considerations for IBM Spectrum Protect for Space Management • Applying maintenance to your GPFS system • Guidance for upgrading the operating system on IBM Spectrum Scale nodes • Servicing IBM Spectrum Scale protocol nodes • Offline upgrade with complete cluster shutdown 	

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

Information unit	Type of information	Intended users
<p><i>IBM Spectrum Scale: Administration Guide</i></p>	<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 IBM Power Systems for IBM Spectrum Scale • Configuring the message queue • Configuring file audit logging • Configuring clustered watch folder • Configuring Active File Management • Configuring AFM-based DR • Tuning for Kernel NFS backend on AFM and AFM DR • Configuring call home <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 	<p>System administrators or programmers of IBM Spectrum Scale systems</p>

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

Information unit	Type of information	Intended users
<p><i>IBM Spectrum Scale: Administration Guide</i></p>	<ul style="list-style-type: none"> • Native NFS and GPFS • Considerations for GPFS applications • Accessing a remote GPFS file system • 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 / RFC 2307 Attributes • Protocols cluster disaster recovery • File Placement Optimizer • Encryption • Managing certificates to secure communications between GUI web server and web browsers • Securing protocol data • Cloud services: Transparent cloud tiering and Cloud data sharing • Managing file audit logging • Performing a watch with watch folder API • RDMA tuning • Configuring Mellanox Memory Translation Table (MTT) for GPFS RDMA VERBS Operation • Administering AFM • Administering AFM DR • Highly-available write cache (HAWC) • Local read-only cache • Miscellaneous advanced administration • GUI limitations 	<p>System administrators or programmers of IBM Spectrum Scale systems</p>

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

Information unit	Type of information	Intended users
<p><i>IBM Spectrum Scale: Problem Determination Guide</i></p>	<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 the message queue • Monitoring file audit logging • Monitoring clustered watch <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 	<p>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></p>

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

Information unit	Type of information	Intended users
<i>IBM Spectrum Scale: Problem Determination Guide</i>	<ul style="list-style-type: none">• GUI issues• AFM issues• AFM DR issues• Transparent cloud tiering issues• File audit logging issues• Troubleshooting watch folder API• Troubleshooting mmwatch• Message queue issues• Maintenance procedures• Recovery procedures• Support for troubleshooting• References	

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

Information unit	Type of information	Intended users
<p><i>IBM Spectrum Scale: Command and Programming Reference</i></p>	<p>This guide provides the following information:</p> <p>Command reference</p> <ul style="list-style-type: none"> • gpfs.snap command • mmaddcallback command • mmadddisk 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 	<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
<p><i>IBM Spectrum Scale: Command and Programming Reference</i></p>	<ul style="list-style-type: none"> • mmclone command • mmcloudgateway command • mmcrcluster command • mmcrfileset command • mmcrfs command • mmcrnodeclass command • mmcrnsd command • mmcrsnapshot command • mmdefedquota command • mmdefquotaoff command • mmdefquotaon command • mmdefragfs command • mmdelacl command • mmdelcallback command • mmdeldisk command • mmdelfileset command • mmdelfs command • mmdelnode command • mmdelnodeclass 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 • mmhdfs command • mmhealth command • mmimgbackup command • mmimgrestore command • mmimportfs command • mmkeyserv 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
<p><i>IBM Spectrum Scale: Command and Programming Reference</i></p>	<ul style="list-style-type: none"> • mmlinkfileset command • mmlsattr command • mmlscallback command • mmlscluster command • mmlsconfig command • mmlsdisk command • mmlsfileset command • mmlsfs command • mmlslicense command • mmlsmgr command • mmlsmount command • mmlsnodeclass command • mmlsnsd command • mmlspolicy command • mmlspool command • mmlsqos command • mmlsquota command • mmlsnapshot command • mmmigratefs command • mmmount command • mmmmsgqueue command • mmnetverify command • mmnfs command • mmnsddiscover command • mmobj command • mmperfmon command • mmpmon command • mmprotocoltrace command • mmpsnap command • mmputacl command • mmquotaoff command • mmquotaon command • mmreclaimspace command • mmremotefilesystem command • mmremotefs command • mmrepquota command • mmrestoreconfig command • mmrestorefs command • mmrestripefile 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
<p><i>IBM Spectrum Scale: Command and Programming Reference</i></p>	<ul style="list-style-type: none"> • mmrestripefs command • mmrpldisk command • mmsdrrestore command • mmsetquota command • mmshutdown command • mmsmb command • mmsnapdir command • mmstartup command • mmtracectl command • mmumount command • mmunlinkfileset command • mmuserauth command • mmwatch command • mmwinservctl command • spectrumscale command <p>Programming reference</p> <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 • Watch folder API 	<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
<p><i>IBM Spectrum Scale: Big Data and Analytics Guide</i></p>	<p>This guide provides the following information:</p> <p>Summary of changes</p> <p>Hadoop Scale Storage Architecture</p> <ul style="list-style-type: none"> • Elastic Storage Server (ESS) • Erasure Code Edition • Share Storage (SAN-based storage) • File Placement Optimizer (FPO) • Deployment model • Additional supported features about storage <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 • CES HDFS • 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>IBM Spectrum Scale Hadoop performance tuning guide</p> <ul style="list-style-type: none"> • Overview • Performance overview • Hadoop Performance Planning over IBM Spectrum Scale • Performance guide 	<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
<p><i>IBM Spectrum Scale: Big Data and Analytics Guide</i></p>	<p>Hortonworks Data Platform 3.X</p> <ul style="list-style-type: none"> • Planning • Installation • Upgrading and uninstallation • Configuration • Administration • Limitations • Problem determination <p>Open Source Apache Hadoop</p> <ul style="list-style-type: none"> • Open Source Apache Hadoop without CES HDFS • Open Source Apache Hadoop with CES HDFS <p>BigInsights® 4.2.5 and Hortonworks Data Platform 2.6</p> <ul style="list-style-type: none"> • Planning • Installation • Upgrading software stack • Configuration • Administration • Troubleshooting • Limitations • FAQ 	<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 XDMS standard

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

Information unit	Type of information	Intended users
<i>IBM Spectrum Scale Erasure Code Edition Guide</i>	IBM Spectrum Scale Erasure Code Edition <ul style="list-style-type: none"> • Summary of changes • Introduction to IBM Spectrum Scale Erasure Code Edition • Planning for IBM Spectrum Scale Erasure Code Edition • Installing IBM Spectrum Scale Erasure Code Edition • Uninstalling IBM Spectrum Scale Erasure Code Edition • Incorporating IBM Spectrum Scale Erasure Code Edition in an Elastic Storage Server (ESS) cluster • Creating an IBM Spectrum Scale Erasure Code Edition storage environment • Upgrading IBM Spectrum Scale Erasure Code Edition • Administering IBM Spectrum Scale Erasure Code Edition • Troubleshooting • IBM Spectrum Scale RAID Administration 	<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 \(www.ibm.com/support/knowledgecenter/STXKQY/ibmspectrumscale_welcome.html\)](http://www.ibm.com/support/knowledgecenter/STXKQY/ibmspectrumscale_welcome.html).

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

Conventions used in this information

Table 2 on page xxii 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>	<p><u>bold underlined</u> keywords are defaults. These take effect if you do not specify a different keyword.</p>
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>
<i>italic</i>	<p><i>Italic</i> words or characters represent variable values that you must supply.</p> <p><i>Italics</i> are also used for information unit titles, for the first use of a glossary term, and for general emphasis in text.</p>
< <i>key</i> >	<p>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>.</p>
\	<p>In command examples, a backslash indicates that the command or coding example continues on the next line. For example:</p> <pre>mkcondition -r IBM.FileSystem -e "PercentTotUsed > 90" \ -E "PercentTotUsed < 85" -m p "FileSystem space used"</pre>
{ <i>item</i> }	<p>Braces enclose a list from which you must choose an item in format and syntax descriptions.</p>
[<i>item</i>]	<p>Brackets enclose optional items in format and syntax descriptions.</p>
<Ctrl- <i>x</i> >	<p>The notation <Ctrl-<i>x</i>> indicates a control character sequence. For example, <Ctrl-<i>c</i>> means that you hold down the control key while pressing <<i>c</i>>.</p>
<i>item</i> ...	<p>Ellipses indicate that you can repeat the preceding item one or more times.</p>
	<p>In <i>synopsis</i> statements, vertical lines separate a list of choices. In other words, a vertical line means <i>Or</i>.</p> <p>In the left margin of the document, vertical lines indicate technical changes to the information.</p>

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:

`scale@us.ibm.com`

Summary of changes

This topic summarizes the changes to the IBM Spectrum Scale AWS Marketplace offering support section.

Summary of changes for IBM Spectrum Scale AWS Marketplace offering version 1.3.1 as updated on October 2020

This release of the AWS Marketplace offering includes the following improvements. All improvements are available after an upgrade, unless otherwise specified.

Changes in Operating System support

Customers have the ability to choose from any of the following supported operating systems:

- RHEL 7.7
- RHEL 7.8
- RHEL 8.1
- RHEL 8.2

Changes in features

Customers can install IBM Spectrum Scale on a customer built AMI.

Customers can install and configure IBM Spectrum Scale GUI.

IBM Spectrum Scale version 5.0.5.3 is supported for AWS 1.3.1

Dedicated controller node for IBM Spectrum Scale cluster deployment and cluster lifecycle management.

The **mmaws utility** is replaced with the **mmcloudworkflows utility**.

Documentation updates

The titles of the following sections have changed:

<i>Table 3.</i>	
Old Title	New Title
<i>Cluster lifecycle management and debug data collection</i>	<i>Cluster lifecycle management</i>
<i>Collecting AWS-specific debug data</i>	<i>Collecting debug data in the event of initial deployment failure</i>
<i>Diagnosing deployment failures</i>	<i>Diagnosing and cleaning-up deployment failures</i>
<i>Collecting debug data if initial deployment fails</i>	<i>Collecting debug data</i>

Summary of changes for IBM Spectrum Scale AWS Marketplace offering version 1.2.0 as updated on June 2019

This release of the AWS Marketplace offering includes the following improvements. All improvements are available after an upgrade, unless otherwise specified.

Changes in Operating System

Upgraded to RHEL 7.6.

Changes in features

Added manual AFM setup and support.

Fixes in features

Fixed entitlement check issue.

Documentation updates

Added documentation for AFM support.

Minor change in lambda function usage.

Summary of changes**for AWS Marketplace offering version 1.1.0
as updated on January 2019**

This release of the AWS Marketplace offering includes the following improvements. All improvements are available after an upgrade, unless otherwise specified.

Changes in Operating System

Upgraded to RHEL 7.5.

Changes in features

Added entitlement checks.

Fixes in features

Upgrade of IBM Spectrum Scale made available locally through IBM Spectrum Scale install toolkit.

Documentation updates

Added documentation for IBM Spectrum Scale upgrade.

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 the 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 namespace across multiple instances with high 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 is used to create and manage a collection of related AWS resources, and provision and update them in an orderly and predictable way. You use a template to describe all the AWS resources that are needed, while the AWS CloudFormation creates and configures the resources, and figures out dependencies. For more information, see [AWS CloudFormation Documentation](#).

Amazon VPC

The Amazon VPC service is used to provision a private, isolated section of the AWS Cloud, where you can start AWS services and other resources in a virtual network that you define. You have complete control over your virtual networking environment, including selection of your own IP address range, subnet creation, and configuration of route tables and network gateways. For more information, see [Amazon Virtual Private Cloud Documentation](#).

Amazon EC2

The Amazon EC2 service is used to start virtual machine instances with various operating systems. You can choose from the existing AMIs or import your own virtual machine images. For more information, see [Amazon Elastic Compute Cloud Documentation](#).

Amazon Auto Scaling

The Amazon Auto Scaling service helps maintain high availability and manage capacity by automatically increasing or decreasing the EC2 instance fleet. You can use auto scaling to run your

fleet at optimal utilization by increasing instance capacity during demand spikes and decreasing capacity during downtimes. For more information, see [AWS Auto Scaling Documentation](#).

Amazon EBS

Amazon Elastic Block Store (Amazon EBS) provides persistent block storage volumes for use with Amazon EC2 instances in the AWS Cloud. For more information, see [Amazon Elastic Block Store Documentation](#).

Amazon S3

Amazon Simple Storage Service (Amazon S3) is 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 the AWS services and resources for your users. With IAM, you can manage users, security credentials such as access keys, and permissions that control which AWS resources users can access, from a central location. For more information, see [AWS Identity and Access Management Documentation](#).

AWS Lambda function

AWS Lambda is a stateless compute service that lets you run code without provisioning or managing servers. For more information, see [AWS Lambda documentation](#).

AWS Systems Manager

AWS Systems Manager is a service that is used to view and control AWS infrastructure. The Systems Manager enables the viewing of operational data from other AWS services, and the automation of operational tasks. In addition, the Systems Manager can also be used to store and retrieve configuration data. For more information, see [AWS Systems Manager](#).

Regions and Availability Zones

IBM 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 three Availability Zones within an AWS region.

Regions are 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 that are 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 a single-flow traffic
- Either 25 Gbps or 10 Gbps of bandwidth for a multi-flow traffic by using a private IPv4 or IPv6 address 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. The replication is done to avoid the loss of data when hardware failures occur in one of the Availability Zones. By default, instances are distributed evenly between the Availability Zones. Therefore, each availability zone has one private subnet. However, these zones remain a single or logical cluster.

- If you choose to deploy a `single_availability_zone` configuration, all the server and compute instances are deployed into a single availability zone, and IBM 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 two copies of each data element.

IBM Spectrum Scale instance types and operating systems

The IBM Spectrum Scale deployment on AWS supports a large selection of EC2 instance types for the IBM Spectrum Scale cluster nodes. The instance type must be chosen after a detailed evaluation of the IBM Spectrum Scale performance and resource requirements, and also the resource requirements for any application workloads that might be running on these nodes. For more information, see [Amazon EC2 Instance Types](#).

IBM Spectrum Scale on AWS BYOL 1.3.1 offers support for the following Operating Systems versions:

- Red Hat® Enterprise Linux (RHEL) versions 7.7 and 7.8
- Red Hat Enterprise Linux (RHEL) versions 8.1 and 8.2

Customers must specify an appropriate AMI running a supported operating system. Red Hat Enterprise Linux AMIs are available from the AWS Marketplace and the AWS Community. Additionally, customers can choose to create their own custom AMI. For more information, see [Chapter 4, “Creating custom AMI,” on page 25](#).

Important: The customer is responsible to ensure that they have the necessary support agreements in place for the Red Hat Enterprise Linux Operating Systems AMIs they specify to facilitate resolution with Red Hat support for any issues.

IBM 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)
- 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
- 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:

- A VPC that spans either a single Availability Zone or three Availability Zones. A single Availability Zone configuration that includes one public and one private subnet. A multi Availability Zone configuration includes three public and three private subnets that are spread across each Availability Zone. 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, which ensures that there is always at least one host available.
- An AWS Identity and Access Management instance role with fine-grained permissions for access to AWS services necessary for the deployment process.
- Appropriate security groups for each instance or function to restrict access to only necessary protocols and ports. The AWS offering opens only ports for SSH and the IBM Spectrum Scale daemon.

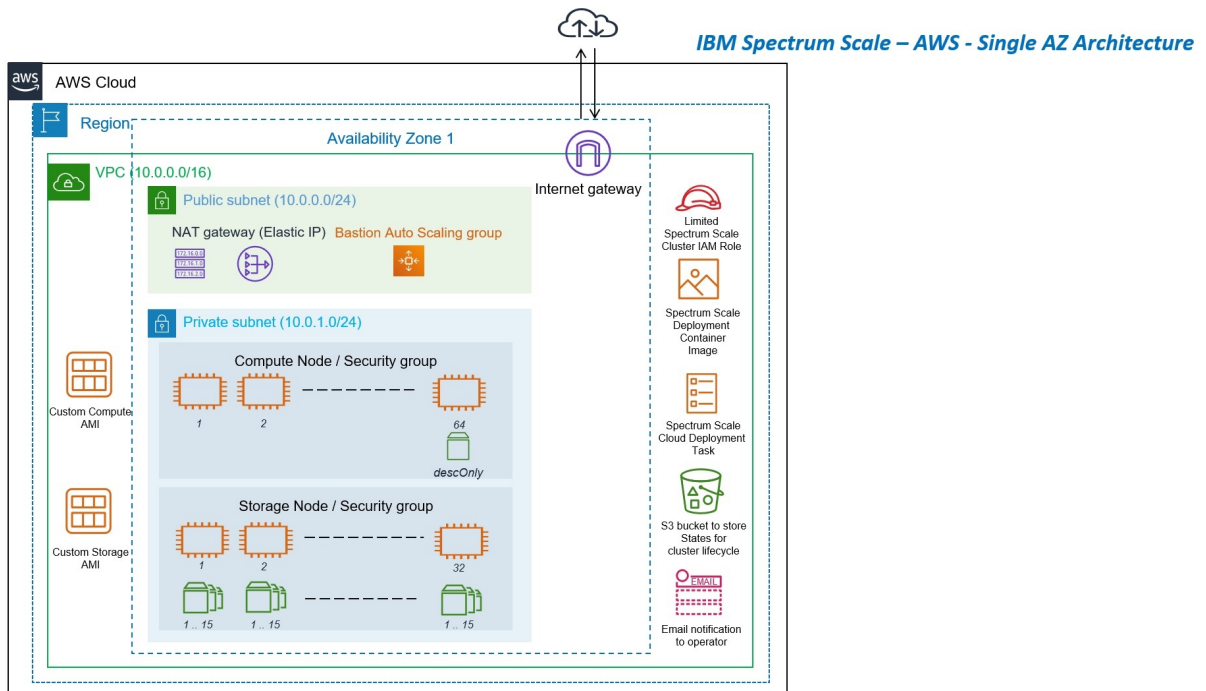


Figure 1. Single availability zone architecture diagram

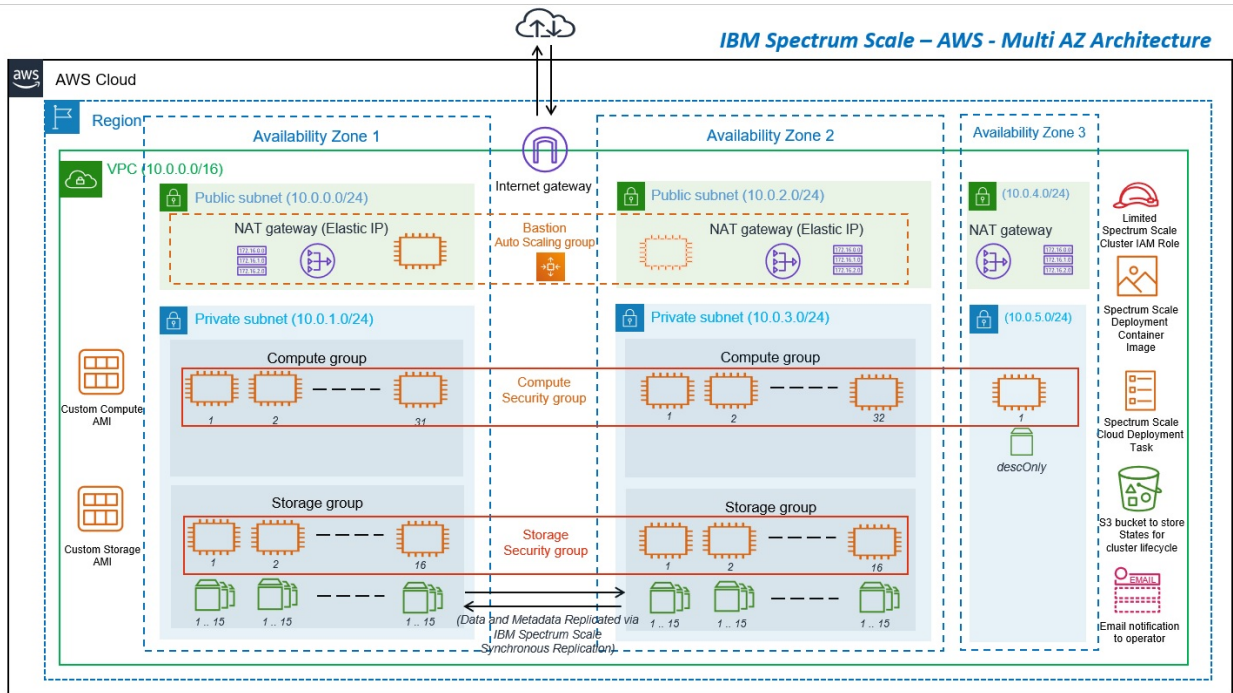


Figure 2. Multi availability zone architecture diagram

Each IBM Spectrum Scale storage node, also referred to as the IBM Network Shared Disk storage server or NSD server, 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.
- If a multi Availability Zone configuration is defined, an extra IBM Spectrum Scale node is created to improve the resiliency of the IBM Spectrum Scale cluster. The extra node is created to avoid loss of file system access after a disk failure. This node has an extra 5 GB EBS volume that is attached to it, and fulfills the role of a tiebreaker for the node and the file system descriptor quorum decisions. For more information, see the *Synchronous mirroring with GPFS replication* section in the IBM Spectrum Scale documentation.

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.

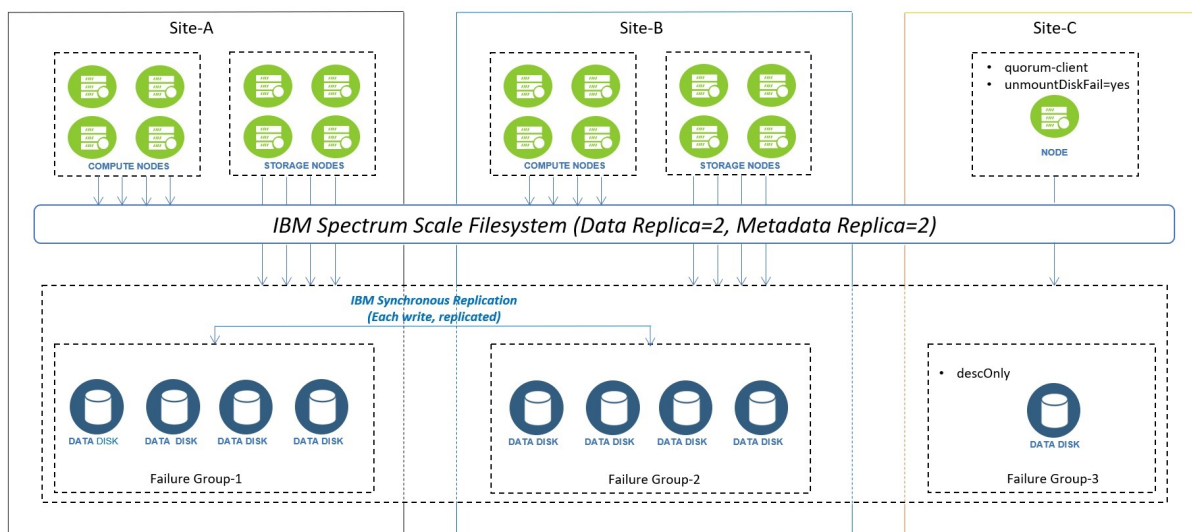


Figure 3. Synchronous mirroring by using GPFS replication

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 choose to deploy the IBM Spectrum Scale cluster into a new VPC that is created as part of the IBM Spectrum Scale deployment. Here, the IBM Spectrum Scale cluster instances are located 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 13](#).

If the IBM Spectrum Scale cluster needs to be deployed within a VPC that already exists, ensure that the following conditions are met:

- The VPC is set up with NAT gateways.
- The VPC has at least one private subnet or three private subnets to deploy the IBM Spectrum Scale instances.

- The VPC 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.

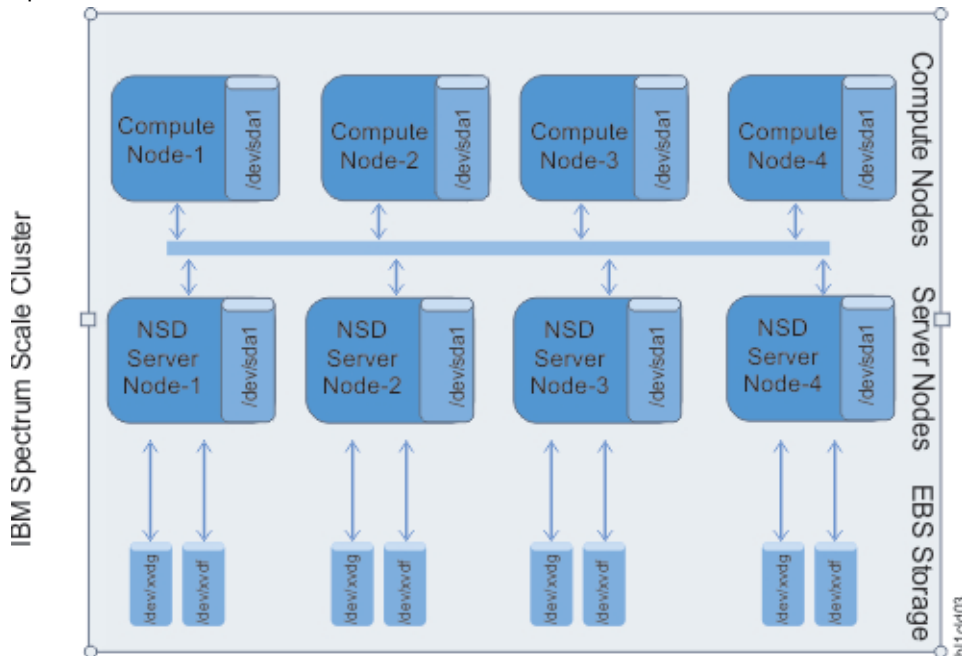


Figure 4. High-level IBM Spectrum Scale cluster architecture on AWS

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 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 more information, see [Amazon EC2 Key Pairs](#)
- d. Verify the desired AWS Region and Availability Zone supports the EC2 instance types that are chosen to be provisioned for the IBM Spectrum Scale storage and compute nodes. You can verify the AWS EC2 instance type from **AWS console > EC2 Service Page > Instance Type**.
- e. Verify the AWS EC2 quota limits, and ensure that enough quotas exist for the EC2 instance types that you intend to deploy. You can verify the AWS EC2 quota limits from **AWS console > EC2 Service Page > Limits**.

If necessary, [request a service limit increase](#) for the EC2 instance types that you intend to deploy. To request a [service limit increase](#), in the AWS Support Center, choose **Create Case > Service Limit Increase > EC2 instances**, and then complete the fields in the **Limit Increase** form.



Attention: The IBM Spectrum Scale deployment on AWS uses CloudFormation templates. These CloudFormation templates rely on the creation and use of IAM Roles to facilitate seamless access to required AWS resources. Therefore, it is important to ensure that while launching the CloudFormation template, the user has sufficient permissions to create and manage IAM Roles.

Note: If you intend to provision a cluster with a large number of storage and compute nodes, it is recommended to increase the AWS Systems Manager's Parameter Store throughput limit. You can change the throughput limit from **AWS Console > Systems Manager > Parameter store > Setting > Parameter Store throughput**.

2. Launch the IBM 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 is used only for cases in which a higher probability of data loss is acceptable to potentially improve performance, and save on data movement costs. For example, in a scratch file system scenario.

To better tolerate EBS failures, it is recommended to use the multiple availability zone option, in which IBM Spectrum Scale replicates the data so that there are two total copies of each data element. In this case, the replication is intended to deal with a loss of a volume. However, the probability of data loss depends on how the failure rates impact the total number of volumes that might fail at any point. It is recommended that users carefully read Amazon's statement about the durability of EBS volumes, particularly the annual failure rates, to better assess the possibility of potential data loss that results from EBS volume failures. For more information, see [Amazon EBS features](#).

Note: The region in which the IBM Spectrum Scale CloudFormation template is launched is where the network infrastructure for the IBM Spectrum Scale cluster is built. The template is launched in the US East (Ohio) region by default. You can change the region.

- b. On the **Select Template** page, keep the default setting for the **Template URL**, and then choose **Next**.
- c. On the **Specify Details** page, change the stack name if needed. Review the parameters for the template.

Note: Provide values for the parameters that require input. For all other parameters, review the default settings and customize them as necessary.

- d. In the **Options** page, you can specify tags or key-value pairs for resources in your stack, and set **Advanced options**.

Note: For more information on how to specify tags, see [AWS CloudFormation Resource Tags Type](#). For setting stack options in the **Advanced options** section, see [Setting AWS CloudFormation Stack Options](#).

- e. In the **Review** page, review and confirm the template settings. Under **Capabilities**, select the checkbox to acknowledge that the template creates IAM resources.
- f. Click **Create** to deploy the stack.
- g. Monitor the status of the stack. When the status displays CREATE_COMPLETE, the IBM 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 that is 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 successfully creates the stack, all instances of compute and NSD servers launched by the AWS setup is up and running with the IBM Spectrum Scale file system that is mounted on it.

Follow these steps to connect to the IBM Spectrum Scale cluster:

- 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 by 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 or storage nodes that were launched by the AWS CloudFormation templates.

For more information about using an SSH agent to forward your private key on connection, see [Using SSH agent forwarding](#).

Important: Do not copy your private key to the bastion host instance.

Note: The IBM Spectrum Scale controller node is not part of the IBM Spectrum Scale cluster, and is used for cluster lifecycle management. For more information, see [Chapter 5, “Cluster lifecycle management,”](#) on page 27.

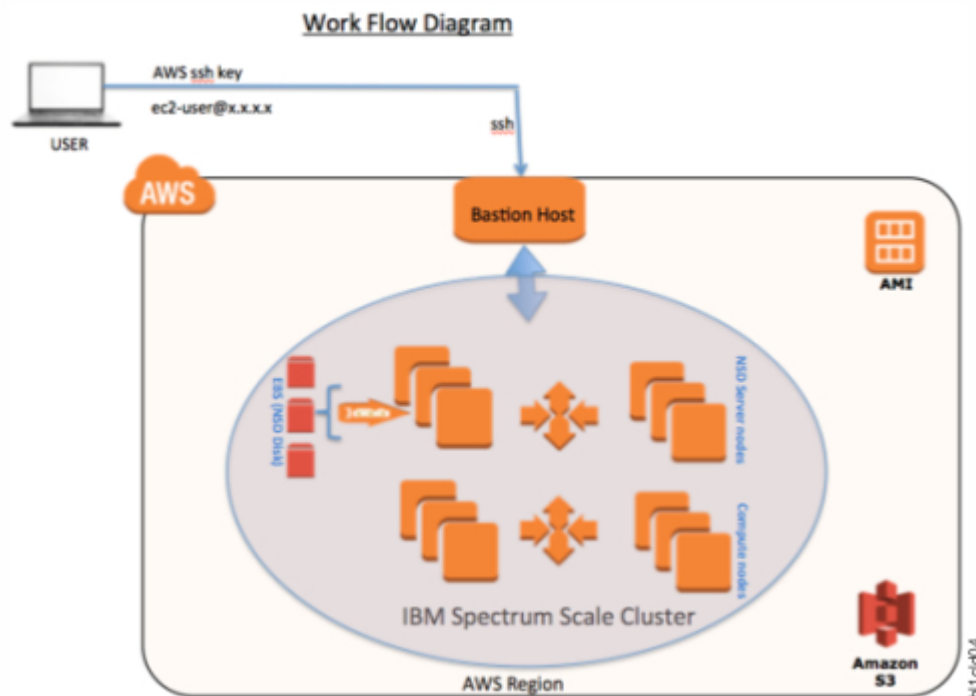


Figure 5. High-level IBM Spectrum Scale cluster architecture to connect from host

4. Test the deployment by using IBM Spectrum Scale commands.

After you log in to a compute or storage nodes, 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 controller is not part of the IBM Spectrum Scale cluster, and therefore cannot run any of the following administrative commands. The controller node can be used only for Lifecycle management commands. For more information, see Chapter 5, “Cluster lifecycle management,” on page 27.

The **mmfscluster** command displays the details of the IBM Spectrum Scale cluster. The command gives the following output:

```
[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 Designation Daemon node name IP address Admin node name
-----
1 quorum-manager-perfmon ip-10-0-1-110.ap-south-1.compute.internal 10.0.1.110 ip-10-0-1-110.ap-south-1.compute.internal
2 quorum-manager-perfmon ip-10-0-3-42.ap-south-1.compute.internal 10.0.3.42 ip-10-0-3-42.ap-south-1.compute.internal
3 quorum-manager-perfmon ip-10-0-1-72.ap-south-1.compute.internal 10.0.1.72 ip-10-0-1-72.ap-south-1.compute.internal
4 perfmon ip-10-0-3-82.ap-south-1.compute.internal 10.0.3.82 ip-10-0-3-82.ap-south-1.compute.internal
```

For more information on the **mmfscluster** command, see the *IBM Spectrum Scale: Command and Programming Reference* guide in the [IBM Spectrum Scale Knowledge Centre](#).

The **mmlsnsd** command displays the NSD server information. The command gives the following output:

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

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 **mmlsnsd** command, see the *IBM Spectrum Scale: Command and Programming Reference* guide in the [IBM Spectrum Scale Knowledge Centre](#).

The **mmlsdisk** command displays disk details. The command gives the following output:

```
[ec2-user@ip-10-0-1-110]$ mmlsdisk fs1 -L
```

disk name	driver type	sector size	failure group	holds metadata	holds data	status	availability	storage disk id	pool
0a9503b2cf264bf08	nsd system desc	512	1	Yes	Yes	ready	up	1	
0d7e3d725140a0cdb	nsd system desc	512	2	Yes	Yes	ready	up	2	
0fb10f7d9a981f39e	nsd system desc	512	3	No	No	ready	up	3	

Number of quorum disks: 3
Read quorum value: 2
Write quorum value: 2

For more information on the **mmlsdisk** command, see the *IBM Spectrum Scale: Command and Programming Reference* guide in the [IBM Spectrum Scale Knowledge Centre](#).

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

```
[ec2-user@ip-10-0-1-110]$ mmdf fs1
```

disk name	disk size in KB	failure group	holds metadata	holds data	free in KB in full blocks	free in KB in fragments
Disks in storage pool: system (Maximum disk size allowed is 3.97 TB)						
0a9503b2cf264bf08	524288000	1	Yes	Yes	520880128 (99%)	8792 (0%)
0d7e3d725140a0cdb	524288000	2	Yes	Yes	520880128 (99%)	8792 (0%)
0fb10f7d9a981f39e	5242880	3	No	No	0 (0%)	0 (0%)
(pool total)	1053818880				1041760256 (99%)	17584 (0%)
(total)	1053818880				1041760256 (99%)	17584 (0%)

Inode Information

Number of used inodes:	4038
Number of free inodes:	497722
Number of allocated inodes:	501760
Maximum number of inodes:	1025024

For more information on the **mmdf** command, see the *IBM 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.

Parameter label (name)	Default	Description
Block Size (BlockSize)	1 M	The file system block size. You can choose a value from 256 KiB - 16 MiB.
GPFS Mount Point (GpfsMountPoint)	/gpfs/fs1	The mount point for the IBM Spectrum Scale file system. Note: fs1 is the file system name that is created by default, and is mounted on location /gpfs/fs1. Changing the mount path to /gpfs/fs2 creates a file system with name fs2, and the file system is mounted on /gpfs/fs2.

Parameter label (name)	Default	Description
EBS Type (EBSType)	gp2	The EBS volume type for IBM Spectrum Scale storage that is attached to each NSD server node. The following options are available: <ul style="list-style-type: none"> • General Purpose SSD (gp2) • Provisioned IOPS SSD (io1) • Cold HDD (sc1) • Throughput Optimized HDD (st1) • EBS Magnetic (standard) For more information about EBS volume type choices, see Amazon EBS . To maximize the benefit of io1 EBS volume type, it is recommended to use the appropriate EBS optimized instances. For more information, see the AWS documentation.

Disk Per Node (DiskPerNode)	1	The number of NSD volumes to attach to each NSD server node. You can choose 1 - 15 disks.
Disk Size (DiskSize)	500	The disk size of NSD volumes that are attached to each NSD server node, in GiBs. Supported disk sizes are 500 - 16,384 GiB.

Parameter label (name)	Default	Description
Storage Node Count (StorageNodeCount)	2	The number of EC2 instances used for the NSD server on the GPFS cluster. You can select 2 - 64 instances.
Storage Instance Type (StorageInstanceType)	t2.medium	The instance type to use for the storage node. A minimum t2.medium instance type is needed. However, other types are available, and a value can be selected based on your requirements. Ensure that the chosen instance type is supported in the desired region.
StorageAMIID		<p>AMI-ID to be used for the storage node. Only Red Hat Enterprise Linux Operating system version 7.7, 7.8, 8.1 and 8.2 are supported. For more information, see Amazon Machine Images (AMI).</p> <p>Customers must specify an appropriate AMI running a supported Operating Systems. Red Hat Enterprise Linux AMIs are available from the AWS Marketplace and the AWS Community. Additionally, customers can choose to create their own custom AMI. For more information on how to build custom AMI, see Chapter 4, “Creating custom AMI,” on page 25.</p> <p>Important: The customer is responsible to ensure that they have the necessary support agreements in place for the Red Hat Enterprise Linux Operating Systems AMIs they specify. The support agreement is needed to facilitate resolutions with Red Hat support for any issues.</p>

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

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. A minimum t2.medium instance type is needed. However, other types are available, and a value can be selected based on your requirements. Ensure that the chosen instance type is supported in the desired region.

Table 7. Compute Node Configurations (continued)

ComputeAMIID		<p>AMI-ID to be used for the compute node. Only Red Hat Enterprise Linux Operating system version 7.7, 7.8, 8.1, and 8.2 are supported. For more information, see Amazon Machine Images (AMI).</p> <p>Customers must specify an appropriate AMI running a supported Operating Systems. Red Hat Enterprise Linux AMIs are available from the AWS Marketplace and the AWS Community. Additionally, customers can choose to create their own custom AMI. For more information on how to build custom AMI, see Chapter 4, “Creating custom AMI,” on page 25.</p> <p>Important: The customer is responsible to ensure that they have the necessary support agreements in place for the Red Hat Enterprise Linux Operating Systems AMIs they specify. The support agreement is needed to facilitate resolutions with Red Hat support for any issues.</p>
--------------	--	---

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 (PrivateSubnet1CIDR)	10.0.1.0 /24	The CIDR block for the private subnet located in Availability Zone 1.
Public Subnet CIDR (PublicSubnet1CIDR)	10.0.0.0 /24	The CIDR block for the public subnet located in Availability Zone 1.
Allowed External Access CIDR (RemoteAccessCIDR)	Requires input	The CIDR block that is allowed external SSH access to the bastion hosts. For example, x.x.x.x/16 - 28. It is recommended that you set this value to a trusted CIDR block. For example, you might want to restrict access to your corporate network.

Table 9. Amazon EC2 Configuration

Parameter label (name)	Default	Description
Key Pair Name (KeyPairName)	Requires input	A public or private key pair, which enables you to connect securely to your instance after it is launched. When you created an AWS account, this is the key pair that you created in your preferred region.
Bastion AMI OS (BastionAMIOS)	Amazon-Linux2-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.

<i>Table 9. Amazon EC2 Configuration (continued)</i>		
Controller Instance Type	t2.micro	The instance type to use for the controller node. The drop-down lists all the available instance types. The instance type of the controller node affects the deployment time of the IBM Spectrum Scale cluster. The larger the chosen cluster, the larger the controller instance type must be for optimal deployment time. The default value is usually preferred only for proof-of-concept deployments. For IBM Spectrum Scale clusters with larger than 30 nodes, a value with a large instance type must be specified for the controller instance type.

<i>Table 10. Personal Configuration</i>		
Parameter label (name)	Default	Description
Spectrum S3 Bucket (SpectrumS3Bucket)	Requires input	S3 bucket name to be created and used for storing state files. Ensure that the S3 bucket does not exist. You can review the rules for naming a bucket at Bucket Restriction .
Operator Email (OperatorEmail)	Requires input	The email address to which the notifications about scaling operations are sent.

<i>Table 11. License Information</i>		
Parameter label (name)	Default	Description
IBM Customer Number (IBMCustomerNumber)	Requires input	Needed to validate the customer entitlement for BYOL offering.
License Agreement Terms (LicenseAgreementTerms)	Requires input	Review the licensing terms at Licence Information , and if you agree to the terms, choose Accept .

Option 2: Deploying IBM 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.

<i>Table 12. File System Configurations</i>		
Parameter label (name)	Default	Description
Block Size (BlockSize)	1 M	The file system block size. You can choose a value from 256 KiB - 16 MiB.
GPFS Mount Point (GpfsMountPoint)	/gpfs/fs1	The mount point for the IBM Spectrum Scale volume. Note: fs1 is the file system name that is created by default, and is mounted on location /gpfs/fs1. Changing the mount path to /gpfs/fs2 creates a file system with name fs2, and the file system is mounted on /gpfs/fs2.

<i>Table 13. NSD Configurations</i>		
Parameter label (name)	Default	Description

EBS Type (EBSType)	gp2	<p>The EBS volume type for IBM Spectrum Scale storage that is attached to each NSD server node. The following options are available:</p> <ul style="list-style-type: none"> • General Purpose SSD (gp2) • Provisioned IOPS SSD (io1) • Cold HDD (sc1) • Throughput Optimized HDD (st1) • EBS Magnetic (standard) <p>For more information about EBS volume type choices, see Amazon EBS.</p> <p>To maximize the benefit of io1 EBS volume type, it is recommended to use the appropriate EBS optimized instances. For more information, see AWS documentation.</p>
Disk Per Node (DiskPerNode)	1	The number of NSD volumes to attach to each NSD server node. You can choose from 1 - 15 disks.
Disk Size (DiskSize)	500	The disk size of NSD volumes that are attached to each NSD server node, in GiBs. Supported disk sizes are 10 - 16,384 GiB.

Parameter label (name)	Default	Description
Storage Node Count (StorageNodeCount)	2	The number of EC2 instances to launch for the NSD server on the GPFS cluster. You can select from 2 - 64 instances.
Storage Instance Type (StorageInstanceType)	t2.medium	The instance type to use for the storage node. A minimum t2.medium instance type is needed. However, other types are available, and a value can be selected based on your requirements. Ensure that the chosen instance type is supported in the desired region.
StorageAMIID		<p>AMI-ID to be used for the storage node. Only Red Hat Enterprise Linux Operating system version 7.7, 7.8, 8.1 and 8.2 are supported. For more information, see Amazon Machine Images (AMI).</p> <p>Customers must specify an appropriate AMI running a supported Operating Systems. Red Hat Enterprise Linux AMIs are available from the AWS Marketplace and the AWS Community. Additionally, customers can choose to create their own custom AMI. For more information, see Chapter 4, "Creating custom AMI," on page 25.</p> <p>Important: The customer is responsible to ensure that they have the necessary support agreements in place for the Red Hat Enterprise Linux Operating Systems AMIs they specify. The support agreement is needed to facilitate resolutions with Red Hat support for any issues.</p>

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

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. A minimum t2.medium instance type is needed. However, other types are available, and a value can be selected based on your requirements. Ensure that the chosen instance type is supported in the desired region.
ComputeAMIID		<p>AMI-ID to be used for the compute node. Only Red Hat Enterprise Linux Operating system version 7.7, 7.8, 8.1 and 8.2 are supported. For more information, see Amazon Machine Images (AMI).</p> <p>Customers must specify an appropriate AMI running a supported Operating Systems. Red Hat Enterprise Linux AMIs are available from the AWS Marketplace and the AWS Community. Additionally, customers can choose to create their own custom AMI. For more information on how to build custom AMI, see Chapter 4, "Creating custom AMI," on page 25.</p> <p>Important: The customer is responsible to ensure that they have the necessary support agreements in place for the Red Hat Enterprise Linux Operating Systems AMIs they specify. The support agreement is needed to facilitate resolutions with Red Hat support for any issues.</p>

Parameter label (name)	Default	Description
Availability Zone (AvailabilityZones)	Requires input	The list of Availability Zones to use for the subnets in the VPC. Only two Availability Zones are used for this deployment, and the logical order of your selections is preserved.
VPC CIDR(VPCCIDR)	10.0.0.0 /16	The CIDR block for the VPC.
Private Subnet 1 CIDR (PrivateSubnet1CIDR)	10.0.1.0 /24	The CIDR block for the private subnet located in Availability Zone 1.
Private Subnet 2 CIDR (PrivateSubnet2CIDR)	10.0.3.0 /24	The CIDR block for the private subnet located in Availability Zone 2.
Private Subnet 3 CIDR (PrivateSubnet3CIDR)	10.0.5.0 /24	The CIDR block for the private subnet located in Availability Zone 3.
Public Subnet 1 CIDR (PublicSubnet1CIDR)	10.0.0.0 /24	The CIDR block for the public subnet located in Availability Zone 1.
Public Subnet 2 CIDR (PublicSubnet2CIDR)	10.0.2.0 /24	The CIDR block for the public subnet located in Availability Zone 2.
Public Subnet 3 CIDR (PublicSubnet3CIDR)	10.0.4.0 /24	The CIDR block for the public subnet located in Availability Zone 3.

<i>Table 16. Network Configuration (continued)</i>		
Allowed External Access CIDR (RemoteAccessCIDR)	Requires input	The CIDR block that is allowed external SSH access to the bastion hosts. For example, <i>x.x.x.x/16-28</i> . It is recommended that you set this value to a trusted CIDR block. For example, you might want to restrict access to your corporate network.

<i>Table 17. Amazon EC2 Configuration</i>		
Parameter label (name)	Default	Description
Key Pair Name (KeyPairName)	Requires input	A public or private key pair, which allows you to connect securely to your instance after it launches. When you created an AWS account, this is the key pair that you created in your preferred region.
Bastion AMI OS (BastionAMIOS)	Amazon-Linux2-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 .
Tie Breaker Instance Type	t2.medium	Instance type to use for the Tie breaker node instance.
Bastion Instance Type (BastionInstanceType)	t2.micro	The EC2 instance type for the bastion host instances.
Controller Instance Type	t2.micro	The instance type to use for the controller node instance. The list displays all the available instance types. The instance type of the controller node affects the deployment time of the IBM Spectrum Scale cluster. The larger the chosen cluster, the larger the controller instance type must be for optimal deployment time. The default value is usually preferred only for proof-of-concept deployments. For IBM Spectrum Scale clusters with larger than 30 nodes, a value with a large instance type must be specified for the controller instance type.

<i>Table 18. Personal Configuration</i>		
Parameter label (name)	Default	Description
Spectrum S3 Bucket (SpectrumS3Bucket)	Requires input	S3 bucket name to be created and used for storing state files. Ensure that the S3 bucket does not exist. You can review the rules for naming a bucket at Bucket Restriction .
Operator Email (OperatorEmail)	Requires input	The email address to which the notifications about scaling operations are sent.

<i>Table 19. License Information</i>		
Parameter label (name)	Default	Description
IBM Customer Number (IBMCustomerNumber)	Requires input	Needed to validate the customer entitlement for BYOL offering.
License Agreement Terms (LicenseAgreementTerms)	Requires input	Review the licensing terms at Licence Information , and if you agree to the terms, choose Accept .

Option 3: Deploying IBM 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 for single Availability Zone. Three private and three public networks for three Availability Zones.

Note: It is recommended for the user to have an existing VPC in three Availability Zones.

- The following are the private subnet requirements:
 1. Internet gateway (IGW) is not configured in the route table.
 2. Auto-assign public IPv4 address is disabled. The IBM 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 needed for Bastion nodes. It is recommended to use Amazon templates to create Bastion nodes. The IBM Spectrum Scale stack creation does not need public subnets.

Note: It is recommended to use the following setup:

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 subnet.
 4. Subnets that are not defined as default subnet.
- Create an S3 endpoint. An S3 endpoint is created by the CF template for VPC that is being used to create VPC.
 - VPC endpoints must be created and added to the route tab.

Note: The VPC created by the template has the value for DNS hostname set to **Yes**. The DNS hostname value is set to **No** when a new VPC is created. For more information on DNS hostnames, see <https://docs.aws.amazon.com/AmazonVPC/latest/UserGuide/vpc-dns.html#vpc-dns-hostnames>.


Table 20. File System Configurations		
Parameter label (name)	Default	Description
Block Size (BlockSize)	1 M	The file system block size. You can choose a value from 256 KiB - 16 MiB.
GPFS Mount Point(GpfsMountPoint)	/gpfs/fs1	The mount point for the IBM Spectrum Scale volume. Note: fs1 is the file system name that is created by default, and is mounted on location /gpfs/fs1. Changing the mount path to /gpfs/fs2 creates a file system with name fs2, and the file system is mounted on /gpfs/fs2.

Table 21. NSD Configurations		
Parameter label (name)	Default	Description

EBS Type (EBSType)	gp2	<p>The EBS volume type for IBM Spectrum Scale storage that is attached to each NSD server node. The following options are available:</p> <ul style="list-style-type: none"> • General Purpose SSD (gp2) • Provisioned IOPS SSD (io1) • Cold HDD (sc1) • Throughput Optimized HDD (st1) • EBS Magnetic (standard) <p>For more information about EBS volume type choices, see Amazon EBS.</p> <p>To maximize the benefit of io1 EBS volume type, it is recommended to use the appropriate EBS optimized instances. For more information, refer to the AWS documentation.</p>
Disk Per Node (DiskPerNode)	1	The number of NSD volumes to attach to each NSD server node. You can choose 1 - 15 disks.
Disk Size (DiskSize)	500	The disk size of NSD volumes that are attached to each NSD server node, in GiBs. Supported disk sizes are 500 - 16,384 GiB.

Parameter label (name)	Default	Description
Storage Node Count (StorageNodeCount)	2	The number of EC2 instances to launch for the NSD server on the GPFS cluster. You can select 2 - 64 instances.
Storage Instance Type (StorageInstanceType)	t2.medium	The instance type to use for the storage node. A minimum t2.medium instance type is needed. However, other types are available, and a value can be selected based on your requirements. Ensure that the chosen instance type is supported in the desired region.
StorageAMIID		<p>AMI-ID to be used for the storage node. Only Red Hat Enterprise Linux Operating system version 7.7, 7.8, 8.1 and 8.2 are supported. For more information, see Amazon Machine Images (AMI).</p> <p>Customers must specify an appropriate AMI running a supported Operating Systems. Red Hat Enterprise Linux AMIs are available from the AWS Marketplace and the AWS Community. Additionally, customers can choose to create their own custom AMI. For more information on how to build custom AMI, see Chapter 4, “Creating custom AMI,” on page 25.</p> <p>Important: The customer is responsible to ensure that they have the necessary support agreements in place for the Red Hat Enterprise Linux Operating Systems AMIs they specify. The support agreement is needed to facilitate resolutions with Red Hat support for any issues.</p>

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. A minimum t2.medium instance type is needed. However, other types are available, and a value can be selected based on your requirements. Ensure that the chosen instance type is supported in the desired region.
ComputeAMIID		<p>AMI-ID to be used for the compute node. Only Red Hat Enterprise Linux Operating system version 7.7, 7.8, 8.1 and 8.2 are supported. For more information, see Amazon Machine Images (AMI).</p> <p>Customers must specify an appropriate AMI running a supported Operating Systems. Red Hat Enterprise Linux AMIs are available from the AWS Marketplace and the AWS Community. Additionally, customers can choose to create their own custom AMI. For more information on how to build custom AMI, see Chapter 4, "Creating custom AMI," on page 25.</p> <p>Important: The customer is responsible to ensure that they have the necessary support agreements in place for the Red Hat Enterprise Linux Operating Systems AMIs they specify. The support agreement is needed to facilitate resolutions with Red Hat support for any issues.</p>

Parameter label (name)	Default	Description
VPC ID		ID of your existing VPC. For example, vpc-0343606e. Value can be select from the list.
Private Subnet 1 ID (PrivateSubnet1ID)		<p>Private subnet ID corresponding to Availability Zone 1 in your VPC. For example, subnet-a0246dcd.</p> <p> CAUTION: Ensure that the values for Private Subnet ID 1, Private Subnet ID 2, and Private Subnet ID 3 are:</p> <ul style="list-style-type: none"> • Either the same value that belongs to a single Availability Zone. • Or unique values that belong to a three distinct Availability Zones. <p>The same value in each Private Subnet ID field signifies that only one Availability Zone is used.</p>
Private Subnet 2 ID (PrivateSubnet2ID)		Private subnet ID corresponding to Availability Zone 2. You can choose private subnet ID corresponding to Availability Zone 1 again when only one private subnet exists in your VPC. For example, subnet-a0246dcd.

Parameter label (name)	Default	Description
Private Subnet 3 ID (PrivateSubnet3ID)		Private subnet ID corresponding to Availability Zone 3. You can choose private subnet ID corresponding to Availability Zone 1 again when only one private subnet exists in your VPC. For example, subnet-a0246dcd.
Allowed External Access CIDR (RemoteAccessCIDR)	Requires input	The CIDR block that is allowed external SSH access to the Bastion hosts. For example, <i>x.x.x.x/16 - 28</i> . It is recommended that you set this value to a trusted CIDR block. For example, you might want to restrict access to your corporate network.

Note: For a single-AZ deployment, add the same PrivateSubnetID in all three places. For multi-AZ deployment, select a unique ID in each place.

Parameter label (name)	Default	Description
Key Pair Name (KeyPairName)	Requires input	A public or private key pair, which is used to connect securely to your instance after it launches. When you created an AWS account, this is the key pair that you created in your preferred region.
Tie Breaker Instance Type	t2.medium	Instance type to use for the Tie breaker node instance.
Bastion Security Group ID		Bastion host security group ID to enable SSH connections. For example, sg-5f16e910.
Controller Instance Type	t2.micro	The instance type to use for the controller node instance. The list displays all the available instance types. The instance type of the controller node affects the deployment time of the IBM Spectrum Scale cluster. The larger the chosen cluster, the larger the controller instance type must be for optimal deployment time. The default value is usually preferred only for proof-of-concept deployments. For IBM Spectrum Scale clusters with larger than 30 nodes, a value with a large instance type must be specified for the controller instance type.

Parameter label (name)	Default	Description
Spectrum S3 Bucket (SpectrumS3Bucket)	Requires input	S3 bucket name to be created and used for storing state files. Ensure that the S3 bucket does not exist. You can review the rules for naming a bucket at Bucket Restriction .
Operator Email (OperatorEmail)	Requires input	The email address to which the notifications about scaling operations are sent.

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 .

Table 27. License Information (continued)

IBM Customer Number		IBM Customer number to validate entitlement for BYOL offering.
---------------------	--	--

Chapter 4. Creating custom AMI

Custom AMI enables user to use their own AMI as a compute node or storage node in the cluster deployment.

Ensure that the custom AMI runs an operating system that is supported for IBM Spectrum Scale on AWS BYOL 1.3.1. For more information, see [“IBM Spectrum Scale instance types and operating systems” on page 3](#).

Note: During custom AMI creation process, the root volume must be specified to be 100 GB or lesser to ensure compatibility with IBM Spectrum Scale on AWS 1.3.1 deployment.

Run the `precheck-customami` script on the EC2 instance that is used as the source instance for the custom AMI. This script ensures that all IBM Spectrum Scale prerequisites are met. You can download the `precheck-customami` script from [ibm-spectrum-scale-cloud-install](#). For more information, see [Creating an instance store-backed Linux AMI](#).

Note:

After the creation of the custom AMI, you must pass the custom AMI ID as either a Compute or a Storage AMI while you deploy the CloudFormation stack.

Chapter 5. Cluster lifecycle management

IBM Spectrum Scale on AWS provides convenient utilities for the expansion and contraction of cluster, enabling GUI, and deleting stack.

The **mmcloudworkflows** utility manages the IBM Spectrum Scale cluster workflow. The utility must be run from the controller node, and has the following functions:

- Add compute and storage nodes to the IBM Spectrum Scale cluster.
- Remove the compute nodes from the IBM Spectrum Scale cluster.
- Enabling the IBM Spectrum Scale Management GUI.
- Removing the entire IBM Spectrum Scale cluster.

mmcloudworkflows utility

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

Synopsis

```
mmcloudworkflows cluster expand [-h]
                                --node_type {compute,storage}
                                --num-instance NUM_INSTANCES
                                {--stack_name STACK_NAME }
```

or

```
mmcloudworkflows cluster contract [-h]
                                   --ip_address IP_ADDRESS, [IP_ADDRESS...]
                                   {--stack_name STACK_NAME }
```

or

```
mmcloudworkflows cluster info [-h]
                               --stack_name STACK_NAME
```

or

```
mmcloudworkflows cluster destroy [-h][--force]
                                   {--stack_name STACK_NAME }
```

or

```
mmcloudworkflows gui_service start {--stack_name STACK_NAME }
```

Availability

Available with IBM Spectrum Scale on AWS BYOL 1.3.1.

Description

Use the **mmcloudworkflows** utility to manage the IBM Spectrum Scale cluster lifecycle and workflow on AWS.

Parameters

cluster expand

Adds the compute or storage nodes to the IBM Spectrum Scale cluster. The nodes that are added have the same node type and configuration as the nodes that were added during the initial cloud deployment.

--node-type {compute, server}

Specifies the IBM Spectrum Scale node type.

--num-instances NUM_INSTANCES

Specifies the number of nodes to be added.

--stack-name STACK NAME

Specifies the CloudFormation stack name.

cluster contract

Removes the compute nodes from the IBM Spectrum Scale cluster.

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

Specifies the IP addresses of the nodes to be removed. The IP address that is provided must be the one listed in the **mm1scluster** command.

--stack-name STACK NAME

Specifies the CloudFormation stack name.

Note: IBM Spectrum Scale storage nodes cannot be removed by using the **mmcloudworkflows cluster contract** command. Only IBM Spectrum Scale compute nodes can be removed from the cluster by using this command.

cluster info

Display the nodes in the existing cluster.

--stack-name STACK NAME

Specifies the CloudFormation stack name.

cluster destroy

Deletes the entire cluster.

--force

Skips user intervention.

--stack-name STACK NAME

Specifies the CloudFormation stack name.

gui_service start

Starts the IBM Spectrum Scale GUI service, and indicates the IP address of the node where the service can be accessed.

--stack-name STACK NAME

Specifies the CloudFormation stack name.

Exit status

0

Successful completion.

nonzero

A failure has occurred.

Security

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

Examples

1. To add compute nodes, run the following command:

```
mmcloudworkflows cluster expand --node_type compute --num-instances 1 --stack_name scale-workflow-demo
```

The system displays output similar to this:

```
2020-06-17 07:38:00,100 - INFO - Logging in to file:
/var/adm/ras/ibm_cloud_workflow_logs/mm_cloud_workflow_add_nodes.log_2020-Jun-17-07-38-00
2020-06-17 07:38:00,101 - INFO - 1. Performing prerequisite check
2020-06-17 07:38:01,962 - INFO - 2. Preparing cloud-install, install-infra configuration
2020-06-17 07:38:06,457 - INFO - 3. Downloading terraform state file from s3 bucket
2020-06-17 07:38:06,639 - INFO - 4. Downloading aws_scale_instances.inputs.tfvars.json file from s3 bucket
2020-06-17 07:38:06,759 - INFO - 5. Downloading scale_clusterdefinition.json file from s3 bucket
2020-06-17 07:38:06,956 - INFO - 6.1. Existing compute instance ids: ['i-0fdff5fd63cbae4b2c',
'i-0d3f82b396e048a59', 'i-041db38e0cf9e15c2']
2020-06-17 07:38:06,957 - INFO - 6.2. Existing storage instance ids: ['i-0401f3df84f38be2d',
'i-08bc78c2b70b11cb3']
2020-06-17 07:38:06,958 - INFO - 6.3. Existing compute instance ips: ['10.0.1.25', '10.0.1.64', '10.0.1.75']
2020-06-17 07:38:06,958 - INFO - 6.4. Existing storage instance ips: ['10.0.1.49', '10.0.1.193']
2020-06-17 07:39:16,536 - INFO - 6.5. Compute instance ids obtained after expansion of resources: ['i-0fdff5fd63cbae4b2c', 'i-0d3f82b396e048a59',
'i-041db38e0cf9e15c2', 'i-007db34046660a79d']
2020-06-17 07:39:16,537 - INFO - 6.6. Compute instance ips obtained after expansion of resources: ['10.0.1.25', '10.0.1.64', '10.0.1.75',
'10.0.1.36']
2020-06-17 07:39:16,537 - INFO - 7.1. Waiting to obtain 'ok' status for instance id(s): ['i-007db34046660a79d']
2020-06-17 07:41:32,172 - INFO - 7.2. Newly spun compute instance ips are: ['10.0.1.36']
2020-06-17 07:41:32,172 - INFO - 8. Proceeding to add newly spun compute instances to cluster (this could take more than a minute)
2020-06-17 07:44:08,047 - INFO - Adding new compute instance(s) (['10.0.1.36']) to cluster completed successfully.
```

2. To remove nodes, run the following command:

```
mmcloudworkflows cluster contract --ip-addresses 10.0.1.75 --stack_name scale-workflow-demo
```

The system displays output similar to this:

```
2020-06-17 08:15:56,304 - INFO - Logging in to file:
/var/adm/ras/ibm_cloud_workflow_logs/mm_cloud_workflow_rm_nodes.log_2020-Jun-17-08-15-56
2020-06-17 08:15:56,305 - INFO - 1. Performing prerequisite check
2020-06-17 08:15:58,106 - INFO - 2. Preparing cloud-install, install-infra configuration
2020-06-17 08:16:02,640 - INFO - 3. Downloading terraform state file from s3 bucket
2020-06-17 08:16:02,839 - INFO - 4. Downloading aws_scale_instances.inputs.auto.tfvars.json file from s3 bucket
2020-06-17 08:16:03,001 - INFO - 5. Downloading scale_clusterdefinition.json file from s3 bucket
2020-06-17 08:16:03,179 - INFO - 6.1. Existing compute instance ids: ['i-0fdff5fd63cbae4b2c',
'i-0d3f82b396e048a59', 'i-041db38e0cf9e15c2', 'i-007db34046660a79d']
2020-06-17 08:16:03,180 - INFO - 6.2. Existing storage instance ids: ['i-0401f3df84f38be2d',
'i-08bc78c2b70b11cb3']
2020-06-17 08:16:03,181 - INFO - 6.3. Existing compute instance ips: ['10.0.1.25', '10.0.1.64', '10.0.1.75',
'10.0.1.36']
2020-06-17 08:16:03,181 - INFO - 6.4. Existing storage instance ips: ['10.0.1.49', '10.0.1.193']
2020-06-17 08:16:03,189 - INFO - Identified local node ipv4 address: 10.0.1.49
2020-06-17 08:16:03,189 - INFO - 7.1. Compute instance ips targeted for removal: ['10.0.1.75']
2020-06-17 08:16:03,189 - INFO - 7.2. Storage instance ips targeted for removal: []
2020-06-17 08:16:03,191 - INFO - 8.1. Identified compute node(s) associated terraform module:
[module.compute_instances.aws_instance.main_with_0_data[2]]
2020-06-17 08:16:03,202 - INFO - 9. Proceeding to remove node(s) from cluster (this could take few minutes depending on data size)
2020-06-17 08:16:17,445 - INFO - 10. Proceeding to destroy resources associated with opted node(s)
2020-06-17 08:16:25,310 - INFO - Performing Terraform destroy
-target=module.compute_instances.aws_instance.main_with_0_data[2]
2020-06-17 08:17:33,145 - INFO - Removing ['10.0.1.75'] ip(s) from cluster completed successfully.
```

3. To delete the entire cluster, run the following command:

```
mmcloudworkflows cluster destroy --stack_name scale-workflow-demo
```

The system displays output similar to this:

```
2020-06-17 08:27:48,571 - INFO - Logging in to file:
/var/adm/ras/ibm_cloud_workflow_logs/mm_cloud_workflow_tearardown.log_2020-Jun-17-08-27-48
=====
| ! Danger Zone ! |
=====
| This workflow, will result in teardown of IBM Spectrum Scale |
| cluster and resources. However, it will not destroy VPC, Bastion |
| Host resources and the s3 bucket. |
| Notes: |
| 1. Ensure to STOP all your applications before proceeding further. |
| 2. All IBM Scale Scale instances must be in either 'running', |
| 'pending', 'stopping' or 'stopped' state. |
=====
Do you want to continue teardown [y/N]: y
2020-06-17 08:27:53,109 - INFO - Proceeding for tear down ..
2020-06-17 08:27:53,109 - INFO - 1. Performing prerequisite check
2020-06-17 08:27:54,916 - INFO - 2. Downloading terraform state file from s3 bucket
2020-06-17 08:27:55,150 - INFO - 3. Obtaining IBM SpectrumScale resources
2020-06-17 08:27:56,268 - INFO - 4. Performing IAM role deletion
2020-06-17 08:27:58,856 - INFO - 5. Performing security group deletion
2020-06-17 08:28:03,038 - INFO - 6. Performing instance termination
2020-06-17 08:28:05,034 - INFO - Operation (IBM Spectrum Scale cluster teardown) completed successfully.
2020-06-17 08:28:05,034 - INFO - Make sure VPC, Bastion resources are cleaned via cloudformation. [root@ip-10-0-1-49 ec2-user]# Connection to
10.0.1.49 closed by remote host.
```

4. To start the GUI node, run the following command:

```
mmcloudworkflows gui_service start --stack_name s0
```

The system displays output similar to this:

```
2020-06-23 10:02:51,800 - INFO - Logging in to file:
/var/adm/ras/ibm_cloud_workflow_logs/mm_cloud_workflow_gui_start.log_2020-Jun-23-10-02-51
2020-06-23 10:02:51,801 - INFO - 1. Performing prerequisite check
```

```
2020-06-23 10:02:53,059 - INFO - 2. Preparing install-infra configuration
2020-06-23 10:02:54,481 - INFO - 3. Downloading scale_clusterdefinition.json file from s3 bucket
2020-06-23 10:02:55,119 - INFO - 4. Identified IBM Spectrum Scale management GUI ip: 10.0.1.118
2020-06-23 10:02:55,119 - INFO - 5. Starting IBM Spectrum Scale management GUI service
2020-06-23 10:03:18,871 - INFO - Successfully started IBM Spectrum Scale management GUI on IP address: 10.0.1.118
```

Location

/usr/lpp/mmfs/bin

Chapter 6. Accessing IBM Spectrum Scale GUI in AWS

In order to access to IBM Spectrum Scale Graphical User Interface (GUI), a user needs to follow these steps: .

1. Allow TCP package forwarding from the Bastion host.
 - a. Comment out `AllowTcpForwarding no` in the `/etc/ssh/sshd_config` file on the Bastion host.
 - b. Restart the ssh service.
2. Run the following command to start the GUI service on the node:

```
# mmcloudworkflows gui_service start --stack_name scale-deploy
```

The system gives an output similar to the following:

```
2020-07-15 09:29:14,497 - INFO - Logging in to file: /var/adm/ras/ibm_cloud_workflow_logs/mm_cloud_workflow_gui_start.log_2020-Jul-15_09-29-14
2020-07-15 09:29:14,498 - INFO - 1. Performing prerequisite check
2020-07-15 09:29:15,463 - INFO - 2. Preparing install-infra configuration
2020-07-15 09:29:16,841 - INFO - 3. Downloading scale_clusterdefinition.json file from s3 bucket
2020-07-15 09:29:17,059 - INFO - 4. Identified IBM Spectrum Scale management GUI ip: 10.0.1.12
2020-07-15 09:29:17,059 - INFO - 5. Starting IBM Spectrum Scale management GUI service
2020-07-15 09:29:55,845 - INFO - Successfully started IBM Spectrum Scale management GUI on IP address: 10.0.1.12
[root@ip-10-0-1-12 ~]#
```

3. Identify the Security Group which owns the GUI node.
 - a. Go to **Services** > **EC2**, and select the EC2 instance which has the IP address of GUI node.
 - b. Select the corresponding security group.
4. Add Rule to Security Group.
 - a. Click **Edit inbound rules**.
 - b. Add HTTPS (443) Type to allow the GUI traffic from the Bastion host to the GUI node.
5. Run the following command from the local machine:

```
- eval `ssh-agent`
- ssh-add -k <path_of_region_specific_key>
- ssh -A -L 22443:<GUI_node_IP>:443 -N ec2-user@<bastion_host_IP>
```

6. Open the browser on the local machine, and run `https://localhost:22443`.

For more information, see *Introduction to IBM Spectrum Scale GUI* in the [IBM Spectrum Scale Knowledge Center](#).

Chapter 7. Active file management on AWS

Active File Management (AFM) is a scalable, high-performance, intelligent file system caching layer integrated into the IBM Spectrum Scale file system. AFM enables the sharing of data across clusters, even if the network is unreliable or has high latency. These attributes make AFM an ideal choice for hybrid cloud environments, where data has to be transferred between the on-premise and cloud environments across the internet.

A hybrid cloud solution is designed to deploy and run a customer's workload in the most optimal environment. This could be either on-premise or on cloud and requires the data to be available at the appropriate environment.

However, infrastructure differences between on-premise and cloud environments can often complicate the data movement in a secure and cost-effective manner. The lack of metadata management and methods to intelligently move data is another hurdle in realizing a hybrid cloud deployment. AFM addresses these challenges by integrating with the IBM Spectrum Scale file system and providing a way to make data available between the on-premise and cloud environments completely transparent to applications that require it.

AFM has the following properties:

- Enables data mobility and sharing of data across various clusters.
- Allows asynchronous-data cross-cluster caching utility.
- Configures on-premise cluster as home, which acts as the primary storage.
- Configures AWS public cloud end points as cache clusters.
- Uses NFSv3 protocol for communication between home and cache sites.

AFM on IBM Spectrum Scale can be used for the following:

- Enable the customers to implement a single global name space across various sites including on-premise and the AWS cloud.
- Enables the seamless movement and caching of data between the on-premise and cloud environments.

For more information on AFM on IBM Spectrum Scale, see [Active File Management \(AFM\) quick reference](#).

Preparing the environment for AFM

The section details the steps to prepare a hybrid cloud environment for AFM:

Follow these steps to set up the hybrid cloud:

1. Set up the IBM Spectrum Scale on-premises environment.

An IBM Spectrum Scale cluster must be first set up in the customer's data center. After an on-premises IBM Spectrum Scale cluster is installed and configured, ensure that NFS services are installed and configured through either Cluster Export Services (CES) or Kernel NFS (kNFS). For more information, see [NFS protocol quick reference](#).

2. Set up the IBM Spectrum Scale cluster on AWS.

An IBM Spectrum Scale cluster must be set up and configured in the AWS cloud. The IBM Spectrum Scale cluster on AWS can be provisioned through the [Amazon Web Services Marketplace](#), and requires a [Bring Your Own License \(BYOL\)](#) that can be purchased from [Passport Advantage](#) or from other Business Partners. The IBM Spectrum Scale on AWS offering provides a fully automated deployment of IBM Spectrum Scale by using the [AWS CloudFormation](#) templates. For more information on deployment, see [Chapter 3, "Deploying IBM Spectrum Scale on AWS," on page 9](#).

3. Set up the site-to-site VPN connection between the home cluster and AFM cache cluster.

Hybrid cloud environments require special network configuration to establish connectivity between the on-premises and public cloud resources, and securely move data between them. The data that

flows between the two environments typically does so through the public internet, posing significant security and privacy risks. To ensure that the in-flight data is protected, a secure connection that uses IPsec based Virtual Private Network (VPN) between the on-premises and public cloud networks must be set up. This ensures that all the data that is passed between the on-premises and AWS cloud cluster is encrypted and protected. The overall performance of the data transfer between the on-premises and AWS clusters depends on the network link connection and the configuration of the VPN. You must ensure that the site-to-site VPN setup does not introduce long latencies, and meets the performance requirements.

The following diagram illustrates a typical example of a site-to-site VPN connection between an on-premises and AWS IBM Spectrum Scale instances.

Note: This is only an example deployment. Factors like on-premises network topology, network infrastructure and its configuration, other security configurations, choice of VPN solution and others determine exactly how the site-to-site VPN is installed and configured. It is the user's responsibility to ensure that the site-to-site VPN is set up and configured optimally and securely.

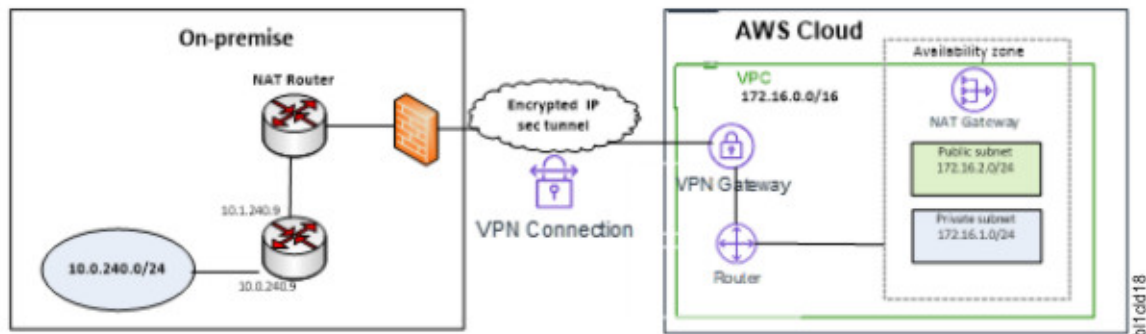


Figure 6. An example on-premises to AWS hybrid cloud networking architecture

The following components form a part of a site-to-site VPN on-premises:

Router

Used for configuring the IPsec tunnel configuration-premises.

NAT router

Used for generating the Network Address Translation (NAT) of the public IP address to the private subnet used by the storage and compute devices.

Networking subnet

Used for configuring the compute and storage devices on-premises.

The following components form a part of a site-to-site VPN on the AWS Cloud:

VPC

Used to launch the resources into a virtual network.

VPN Gateway

Used to create a virtual private gateway and attach it to the VPC from which you want to create the site-to-site VPN connection.

Router

Used for routing the traffic from the AWS private subnet to the on-premises private subnet.

Private subnet

Used for configuring the compute and storage devices in AWS.



Attention:

There are numerous methods to establish a site-to-site VPN between the on-premises environment and AWS environments. The VPN connection depends on both the network infrastructure and configuration, and the components employed on the AWS cloud. For examples on how to install and configure a site-to-site VPN, see [IBM Solutions for Hybrid Cloud Networking Configuration](#).

4. Set up the authentication-premises and authorization infrastructure to enable multi-user restricted access to the data on AWS.

In a hybrid cloud environment, a common authentication and authorization infrastructure are needed across the AFM home and cache sites to ensure that the user access to files is protected and controlled. AFM requires that the User IDs (UID) and Group IDs (GID) be managed the same way across the AFM cache and AFM home clusters.

In addition to caching metadata and data, AFM also caches extended attributes, and access control lists (ACLs) to ensure that protected user access is enforced. AFM does not perform ID mapping between the AFM cache and AFM home clusters. You can configure either the Active Directory (AD) or the Lightweight Directory Access Protocol (LDAP) by using an external server for the ID-mapping information. The same set of authentication mechanisms needs to be configured on both the locations – AFM cache and AFM home cluster - for effective UID mapping.

The setup and configuration of authentication infrastructure on the AWS cloud depends on the authentication infrastructure that is configured in the on-premises environment. It is important that the AWS cloud has the same UID and GID mapping that exists on-premises. AWS provides directory services such as [AWS Managed Microsoft AD](#) that also has LDAP connectors that allow AWS to communicate through AD or LDAP protocols. For more information, see the *Requirements for UID and GID on the cache and home clusters* section in the [IBM Knowledge Center](#)

5. Time-sync the AFM home and cache clusters.

AFM relies on the file modification time to enable data movement between the AFM home and cache sites. Therefore, it is critical that both the AFM home and cache clusters be time-synced through tools like NTPD, Cronty. For more information, see [Setting the Time for Your Linux Instance and Keeping Time With Amazon Time Sync Service](#).

Note: The on-premises and AWS IBM Spectrum Scale clusters can reside in and be configured to use different time zones. However, ensure that both clusters are time-synced to their respective time zones.

AFM cache modes

The following section describes the different cache modes.

Four types of cache modes are available:

- [“Read only \(RO\) cache mode” on page 35](#)
- [Local Update \(LU\) cache mode](#)
- [“Single Writer \(SW\) cache mode” on page 37](#)
- [“Independent Writer \(IW\) cache mode” on page 38](#)

Read only (RO) cache mode

The RO cache mode has the following characteristics:

- The files are created and reside in the home cluster.
- The data in the cache cluster is read-only.
- The data is copied from the home cluster to the cache cluster when files are accessed or during a prefetch operation.

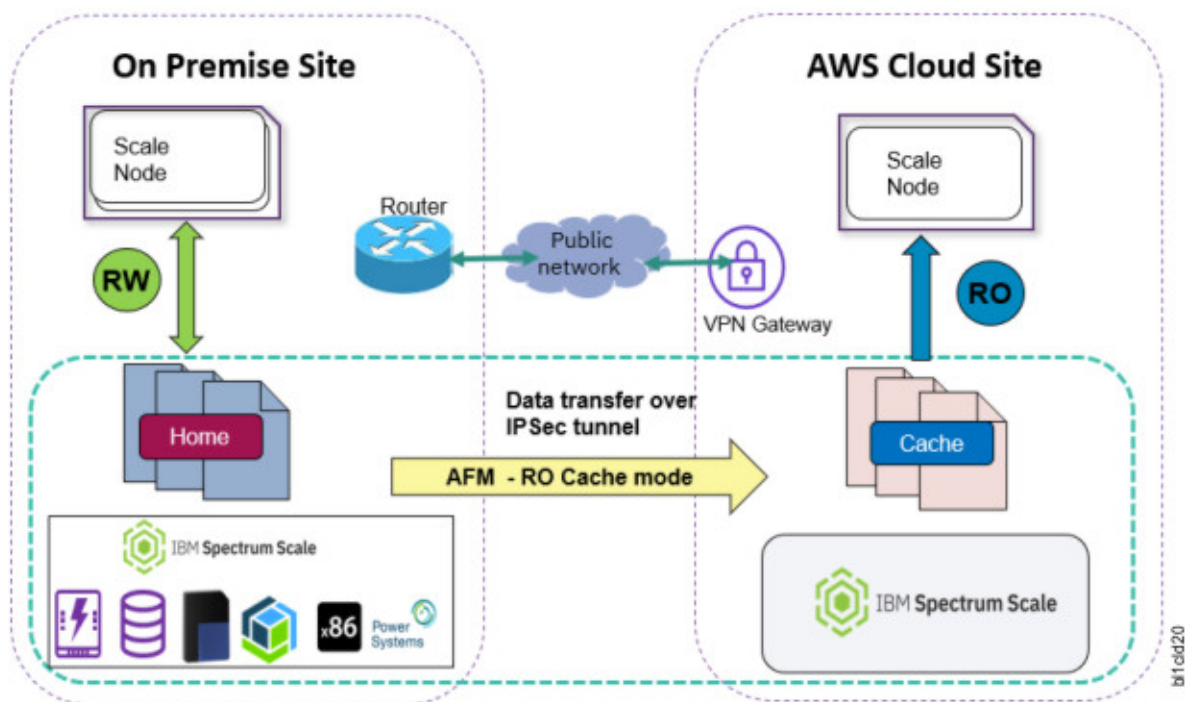


Figure 7. Read only cache mode

Local Update (LU) cache mode

The LU cache mode has the following characteristics:

- The data in the cache cluster can be modified or new files can be created.
- The files that are created or modified in the cache cluster are considered local updates.
- Local updates that are made on the cache cluster are never pushed back to the home cluster.
- If changes are done on a file in the cache cluster, any subsequent changes that are made to that file in the home cluster no longer reflect in the cache cluster.
- If changes are done to a directory in the cache cluster, any subsequent changes that are made to that directory in the home cluster no longer reflect in the cache cluster. Changes can include renaming the directory, removing a directory, etc. However, any changes that are made to the files within the modified directory in the home cluster continue to reflect in the cache cluster. The changes reflect in the cache cluster if the files are not modified before in the cache cluster.

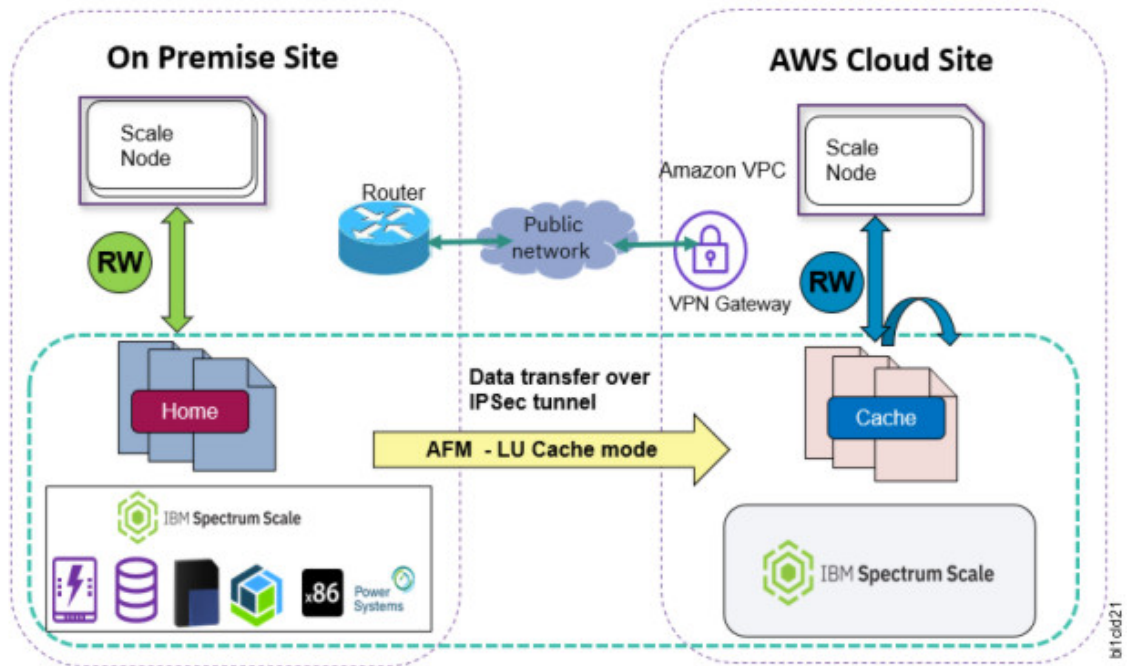


Figure 8. Local update cache mode

Single Writer (SW) cache mode

The SW cache mode has the following characteristics:

- Upon creation, if the home cluster contains data, the single-writer cache copies all the data from home cluster into the cache cluster.
- After this copy is complete, the single writer cache no longer checks the home cache for changes. This is because the single writer cache assumes that it always contains the latest version of a file and therefore never checks the home cluster for updates.
- All files that are created or changed in a single-writer cache are pushed to home cluster.

Note: The AFM SW mode moves any modified data from the cache cluster to the home cluster. Data movement from AWS cluster to the on-premises cluster incurs higher charges as compared to moving data the other way around.

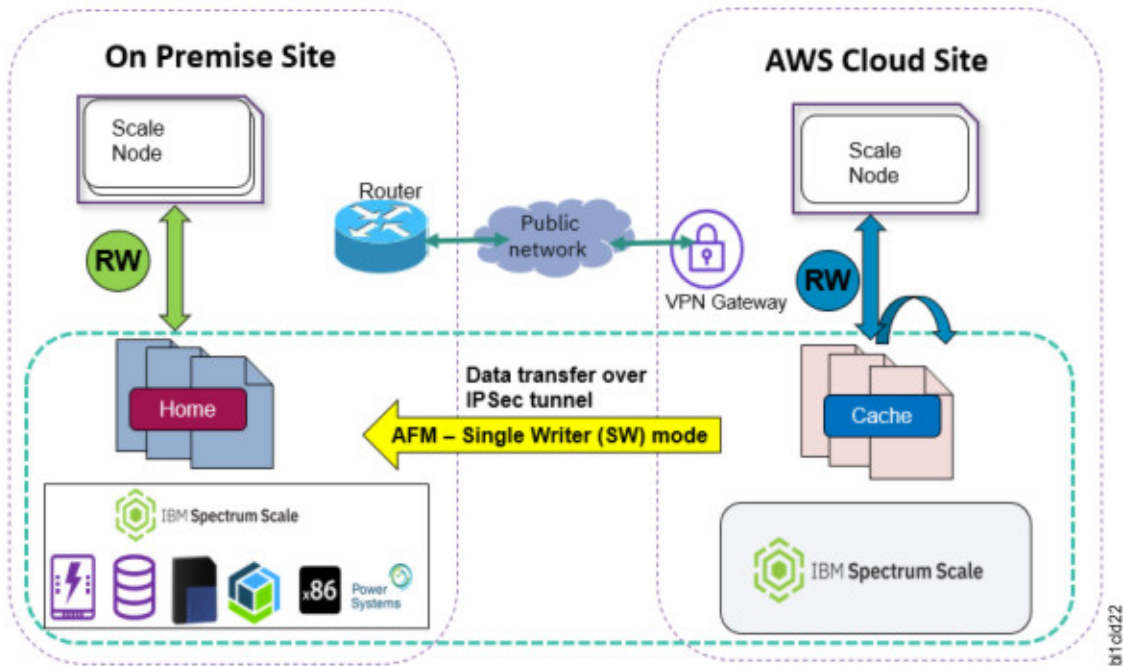


Figure 9. Single writer cache mode

Independent Writer (IW) cache mode

The IW cache mode has the following characteristics:

- Data is read and written on both the home cluster and the cache cluster.
- Changes that are made in the home cluster are synchronized with the cache cluster and vice versa.

Note: The AFM IW mode moves any modified data from the cache cluster to the home cluster. Data movement from AWS cluster to the on-premises cluster incurs higher charges as compared to moving data the other way around.

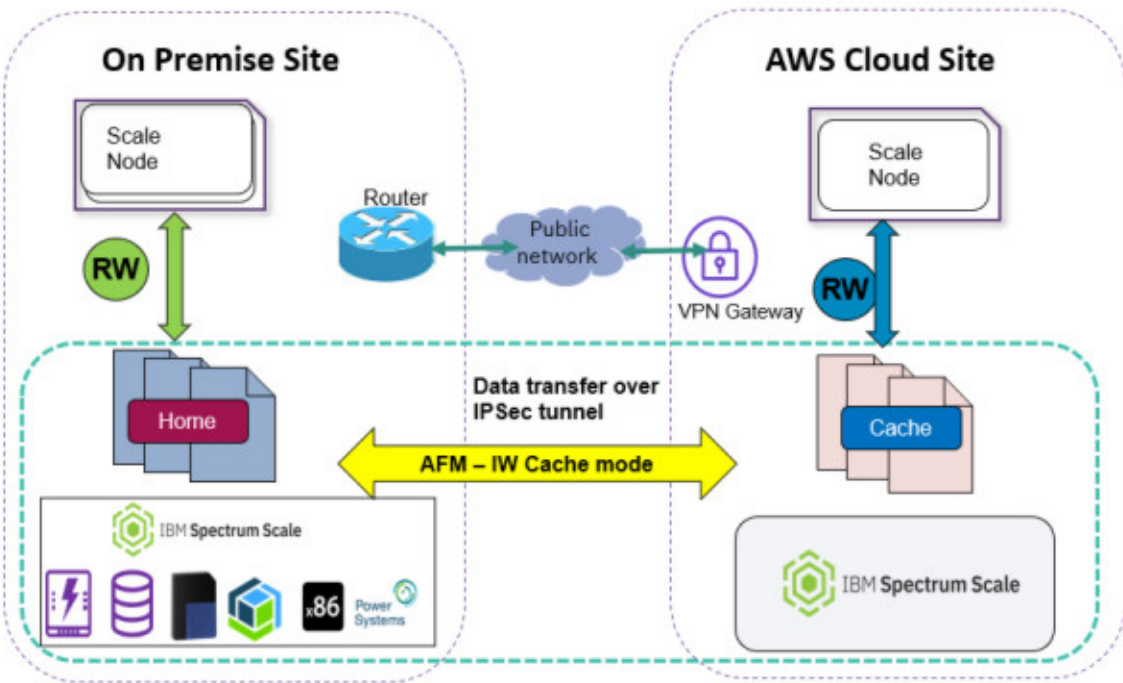


Figure 10. Independent writer cache mode

Deploying AFM on AWS

IBM Spectrum Scale Active File Management (AFM) is used for data movement and caching between the on-premises and public cloud environments. Public cloud end points are only supported as cache sites.

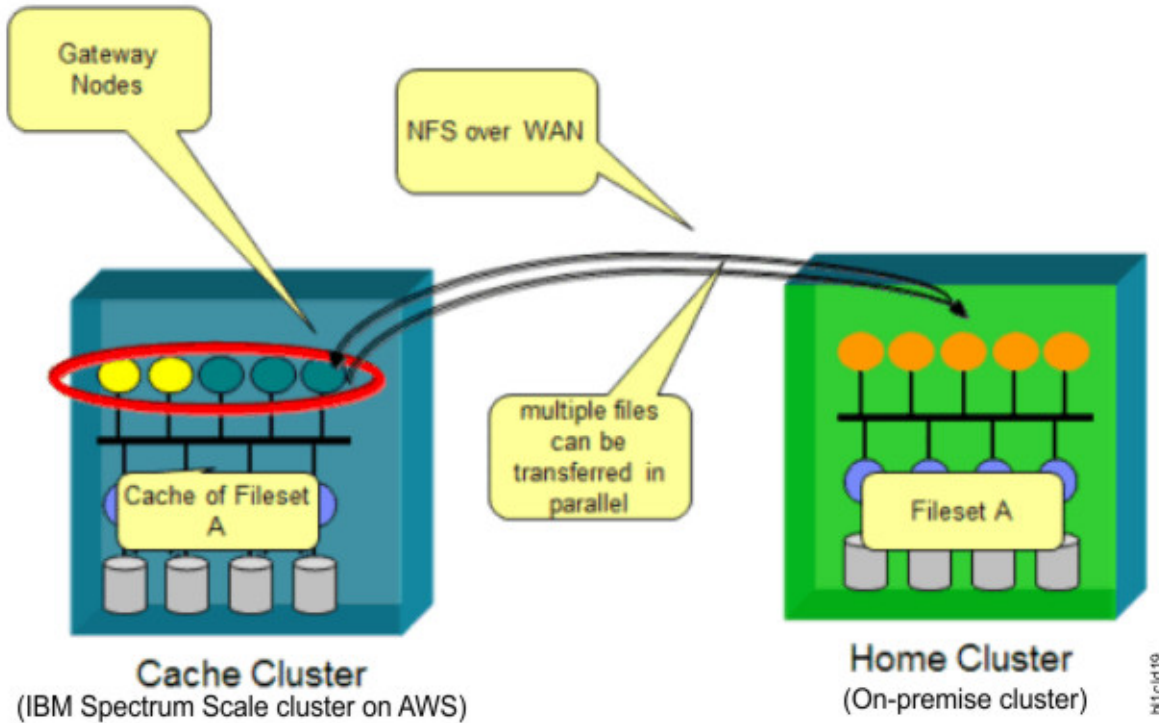


Figure 11. Example of an on-premises to AWS hybrid cloud networking architecture with AFM

Follow these steps to deploy AFM:

1. Set up the on-premises cluster.

The on-premises IBM Spectrum Scale cluster acts as the AFM home cluster. An AFM home cluster is defined as an IBM Spectrum Scale cluster that can make the NFS v3 exports available to other IBM Spectrum Scale clusters. Follow these steps to set up the on-premises cluster:

- a. Run the **mmnfs** command to define the NFS exports and provide access to the AFM gateway nodes:

```
mmnfs export add /ibm/gpfs0/onprem-aws -client "172.16.1.237  
(Access_Type=RW,Squash=root_squash)"
```

The command gives the following output:

```
mmnfs: The NFS export was created successfully
```

For more information, see the *IBM Spectrum Scale: Command and Programming Reference* guide in the [IBM Spectrum Scale Knowledge Center](#).

- b. Run the **mmafmconfig** command to enable support of extended attributes or sparse files on the AFM home exports:

```
mmafmconfig enable /ibm/gpfs0/onprem-aws
```

For more information, see the *IBM Spectrum Scale: Command and Programming Reference* guide in the [IBM Spectrum Scale Knowledge Center](#).

2. Set up the AWS cluster.

The IBM Spectrum Scale cluster on AWS acts as the AFM cache cluster. The AFM cache cluster hosts the filesets that are configured to participate in a relationship with the NFS exports in the AFM home cluster. Follow these steps to set up the AWS cluster:

- a. Run the following command to verify that the NFS client is installed on the nodes that are going to server as the AFM gateway:

```
rpm -qi nfs-utils
```

If the NFS client package is not installed, run the following command to install the NFS client on the nodes that are going to serve as the AFM gateway.

```
yum install nfs-utils
```

Note: A gateway node can communicate with the home cluster to transfer data. Refer the [“Configuration best practices”](#) on page 41 section before you select a node to act as the AFM gateway node.

- b. Run the **mmchnode** command to define the AFM gateway nodes on the nodes selected in [Step 2a](#).

```
mmchnode --gateway -N ip-172-16-1-237.eu-central-1.compute.internal
```

The command gives the following output:

```
mmchnode: Processing node ip-172-16-1-237.eu-central-1.compute.internal
mmchnode: Propagating the cluster configuration data to all
affected nodes. This is an asynchronous process.
```

For more information, see the *IBM Spectrum Scale: Command and Programming Reference* guide in the [IBM Spectrum Scale Knowledge Center](#).

- c. Run the **mmcrfileset** command to create an AFM fileset with a relationship to the NFS exports server from the on-premises or home cluster.

Note: AFM cache mode must be specified based on the needs of the customer use case. For more information, see [“AFM cache modes”](#) on page 35.

```
mmcrfileset fs1 onprem-aws-cache -p afmTarget=10.0.240.43:/ibm/gpfs0/onprem-aws -p
afmmode=read-only --inode-limit 999999 --inode-space new
```

The command gives the following output:

```
Fileset onprem-aws-cache created with id 21 root inode 5767171.
```

For more information, see the *IBM Spectrum Scale: Command and Programming Reference* guide in the [IBM Spectrum Scale Knowledge Center](#).

- d. Run the **mmlinkfileset** command to link the fileset to the junction folder:

Note: A junction folder is a special directory entry, like a POSIX hard link, that connects a name in the directory of the parent fileset to the root directory of a child fileset.

```
mmlinkfileset fs1 onprem-aws-cache -J /gpfs/fs1/onprem-aws-cache
```

The command gives following output:

```
Fileset onprem-aws-cache linked at /gpfs/fs1/onprem-aws-cache
```

For more information, see the *IBM Spectrum Scale: Command and Programming Reference* guide in the [IBM Spectrum Scale Knowledge Center](#).

You can prefetch the file metadata and data from the AFM home cluster before an application requests the contents. Prefetching files before an application starts can reduce the network delay when an application requests a file. Prefetch can be used to pro-actively manage WAN traffic patterns by moving files over the WAN during a period of low WAN usage. For more information, see the *Prefetch* section in the [IBM Spectrum Scale Knowledge Center](#).

For more information, see the *Setting up the cache cluster* and *Creating an AFM relationship by using the NFS protocol* sections in the [IBM Spectrum Scale Knowledge Center](#).

Configuration best practices

The following section describes the best practices for configuring AFM on AWS.

- Ensure that the network link between the home and cache clusters has enough bandwidth to support the volume of data that is being copied between them.
- Ensure that the site-to-site VPN servers on the on-premise cluster and the AWS cloud cluster have sufficiently large resources. This is because all the data between the home and cache clusters must traverse through the VPN servers. If the VPN servers do not have the requisite amount of CPU and memory availability, the VPN servers might introduce latencies that result in AFM disconnects.
- Set the following configuration parameters on the AWS cluster to a sufficiently high value to compensate for the possible large latencies that are found in the network between the on-premise and AWS clusters:
 - `afmAsyncOpWaitTimeout`
 - `afmRevalOpWaitTimeout`
 - `afmSyncOpWaitTimeout`

Since all the data between the home and cache clusters traverses through the VPN servers, it is inevitable that larger latencies are introduced. These parameters must be set to sufficiently large values to account for larger latency and to ensure that AFM does not disconnect. They are typically set to 1800. These configuration parameters can be modified using the `mmchconfig` command. For more information on the `mmchconfig` command, see *IBM Spectrum Scale: Command and Programming Reference* guide in the [IBM Spectrum Scale Knowledge Center](#).

- Ensure that the AFM gateway nodes have enough resources for optimal performance. Therefore, it is important that the AWS instances with the requisite amount of CPU and memory availability are chosen during the initial cluster provisioning.
- Ensure that enough AFM gateways are provisioned on the cache cluster to provide the desired HA redundancy and performance characteristics.
- Ensure that a node is not assigned both the `NSD server` node role and the `AFM gateway` node role simultaneously to avoid resource contention.
- AFM gateways must ideally not be assigned to a node that is running any customer workload.

Limitations of AFM on AWS

The following section describes the limitations of AFM on AWS.

- The AWS IBM Spectrum Scale cluster can only act as an AFM cache.
- No support for AFM over GPFS protocol. Only AFM over NFS v3 is supported.
- No support for AFM DR.

Chapter 8. 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 controller node.

Note: All these steps must be run on the controller node.

2. Create new directory `mkdir -p /root/pem_key`.
3. Run the following command to copy the existing `id_rsa` key to the newly created directory:

```
cp /root/.ssh/id_rsa /root/pem_key/id_rsa
```

4. Run the following command to change the format of the PEM inside the newly created directory so that it does not modify the existing `id_rsa`:

```
ssh-keygen -p -N "" -m pem -f /root/pem_key/id_rsa
```

5. Run the following script to extract IBM Spectrum Scale Data Management Bundle:

```
< Path_of_bundle> -silent
```

6. Go to the IBM Spectrum Scale installation toolkit directory.

Note: The IBM Spectrum Scale installation toolkit is used to upgrade the cluster to the wanted IBM Spectrum Scale version. For more information, see *Upgrading IBM Spectrum Scale components with the installation toolkit* section in the *IBM Spectrum Scale: Concepts, Planning, and Installation Guide*.

7. Specify the newly created modified `id_rsa` key during the following installation toolkit setup command:

```
cd /usr/lpp/mmfs/{scale_version}/installer/; ./spectrumscale setup -s  
<controller_node_ip> -i /root/pem_key/id_rsa
```

8. Run the **config populate** command:

```
./spectrumscale config populate -N <node>
```

Note: The node is the IP address or hostname of any storage node.

9. Run the **setup** command again:

```
[root@test-vm1 installer]# ./spectrumscale setup -s <controller_node_ip> -i /root/pem_key/  
id_rsa
```

Note: The **setup** command is run as the **config populate** command overrides the `id_rsa` key configuration if the `-i` option is not specified.

10. Run the following command to export the flag for SNC:

```
export SS_SNC_ROLLING_UPGRADE=true
```

11. Run the following upgrade commands:

```
./spectrumscale upgrade --precheck  
./spectrumscale upgrade run
```

12. Run the following command to create a new directory on the controller node that can be used to store all IBM Spectrum Scale packages for further use after IBM Spectrum Scale is upgraded:

```
mkdir -p /opt/IBM/<scale version>/gpfs_cloud_packages_<scale version>
```

13. Go to the `gpfs_rpms` package directory of the upgraded IBM Spectrum Scale, and copy the following packages to the directory created in step 12.

```
cd /usr/lpp/mmfs/{{ scale_version }}/gpfs_rpms
cp gpfs.adv* /opt/IBM/<scale version>/gpfs_cloud_packages_<scale version>/
cp gpfs.base* /opt/IBM/<scale version>/gpfs_cloud_packages_<scale version>/
cp gpfs.crypto* /opt/IBM/<scale version>/gpfs_cloud_packages_<scale version>/
cp gpfs.docs* /opt/IBM/<scale version>/gpfs_cloud_packages_<scale version>/
cp gpfs.gpl* /opt/IBM/<scale version>/gpfs_cloud_packages_<scale version>/
cp gpfs.gskit* /opt/IBM/<scale version>/gpfs_cloud_packages_<scale version>/
cp gpfs.gui* /opt/IBM/<scale version>/gpfs_cloud_packages_<scale version>/
cp gpfs.gpl* /opt/IBM/<scale version>/gpfs_cloud_packages_<scale version>/
cp gpfs.java* /opt/IBM/<scale version>/gpfs_cloud_packages_<scale version>/
cp gpfs.license.dm* /opt/IBM/<scale version>/gpfs_cloud_packages_<scale version>/
cp gpfs.msg.en_US* /opt/IBM/<scale version>/gpfs_cloud_packages_<scale version>/
```

14. Go to the `zimon_rpms` package directory of the upgraded IBM Spectrum Scale, and copy the following packages to the directory created in step 12:

```
cd /usr/lpp/mmfs/{{ scale_version }}/zimon_rpms/rhel7/
cp gpfs.pmcollector* /opt/IBM/<scale version>/gpfs_cloud_packages_<scale version>/
cp gpfs.pmsensors* /opt/IBM/<scale version>/gpfs_cloud_packages_<scale version>/

cd /usr/lpp/mmfs/{{ scale_version }}/zimon_rpms/rhel8/
cp gpfs.pmcollector* /opt/IBM/<scale version>/gpfs_cloud_packages_<scale version>/
cp gpfs.pmsensors* /opt/IBM/<scale version>/gpfs_cloud_packages_<scale version>/
```

Note: After you upgrade the IBM Spectrum Scale cluster, any `mmcloudworkflows` commands must be run only after the following environment variables are set:

```
export SPECTRUM_SCALE_PACKAGE_LOC= <Location of the IBM Spectrum Scale packages created in
step 12>
For example: export SPECTRUM_SCALE_PACKAGE_LOC=/opt/IBM/<scale version>/
gpfs_cloud_packages_<scale version>/
```

Chapter 9. Cleaning up the cluster and the stack

Cleaning up the cluster involves a two-step process.

Follow these steps to clean up the stack:

1. Run the **mmcloudworkflows cluster destroy** command from the controller node.

Note: You must be a root user to run the **mmcloudworkflow** administrative command.

Important: Running the **mmcloudworkflows cluster destroy** command enables the customer to remove and deprovision all the nodes in the IBM Spectrum Scale cluster. Running the **mmcloudworkflows cluster destroy** command also removes all the data that is stored in the storage nodes. The data that existed in the IBM Spectrum Scale cluster cannot be retrieved after it is deleted.

For example, to remove the entire cluster, run the following command:

```
mmcloudworkflows cluster destroy --stack_name scale-workflow-demo
```

The system displays the following output:

```
2020-06-17 08:27:48,571 - INFO - Logging in to file:
/var/adm/ras/ibm_cloud_workflow_logs/mm_cloud_workflow_takedown.log_2020-Jun-17_08-27-48

=====
|           ! Danger Zone !           |
=====
| This workflow, will result in takedown of IBM Spectrum Scale |
| cluster and resources. However, it will not destroy VPC, Bastion |
| Host resources and the s3 bucket.                               |
| Notes:                                                         |
| 1. Ensure to STOP all your applications before proceeding further. |
| 2. All IBM Scale Scale instances must be in either 'running',   |
|    'pending', 'stopping' or 'stopped' state.                   |
=====

Do you want to continue takedown [y/N]: y
```

For more information, see the [“mmcloudworkflows utility” on page 27](#) command.

On successful completion, the entire IBM Spectrum Scale cluster is deleted. You can verify that the cluster is deleted from the AWS EC2 console.

2. Delete the AWS CloudFormation Stack from the AWS Console.
 - a. Open the CloudFormation console, and select the root stack to be deleted.
 - b. Go to **Actions > Delete Stack**.

Chapter 10. 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.

It is the customer's responsibility to perform periodic operating system updates, including security patch applications, for all the IBM Spectrum Scale storage and compute instances. It is advisable to subscribe to the IBM Spectrum Scale Security bulletins through the IBM support portal for related updates. For information on managing subscription, refer to [IBM Security Vulnerability Management](#).

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 and on existing VPC.

Access to the IBM Spectrum Scale cluster is only allowed through a Bastion host. It is a special purpose server instance that is designed to be the primary access point from the Internet and acts as a proxy to the other IBM Spectrum Scale instances. If the deployment is in a new VPC, the Bastion host is launched automatically. In case of deployment in existing VPC, it is a prerequisite to configure a Bastion host.

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

The EC2-user user account of the instances in a cluster can be accessed only by using the SSH key that is specified during the deployment process. AWS does not store these SSH keys, so the loss of the SSH key can lead to loss of access to these instances. Updating operating system and security patches are the user's responsibility, and must 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 that are created and assigned to the individual instances as part of this solution are restricted as much as possible. However, these security groups do have access to the various functions needed by IBM Spectrum Scale. It is recommended to review the security groups to further restrict access as needed after the cluster is up and running. The initial setup 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 nodes and storage nodes.

It is the customer's responsibility to perform instances (storage, compute) operating system updates (including security patches) periodically. It is advisable to subscribe to the IBM Spectrum Scale Security bulletins through the IBM support portal. For more information, see [IBM Security Vulnerability Management](#).

ControllerSecurityGroup

This group is for IBM Spectrum Scale controller node. It allows SSH access for BastionSecurityGroup, and enables communication between the controller node and the IBM Spectrum Scale compute and storage nodes.

Chapter 11. Diagnosing and cleaning-up deployment failures

The following section details the steps to diagnose the deployment failures for AWS in IBM Spectrum Scale.

Diagnosing deployment failures

Follow these steps to diagnose deployment failures for AWS in IBM Spectrum Scale:

1. Examine the **Stack Info** tab to see a summary of the failure that is found in the Status and Status Reason fields. Note the name of the task that displays a CREATE_FAILED state.
2. Go to the **Events** tab, and look for the task that displayed a CREATE_FAILED state in Step 1.
3. Examine the corresponding Status Reason field for the failed event to find the hyperlink to the CloudWatch.
4. Examine the CloudWatchLogs to determine the error.

Cleaning-up deployment failures

If an IBM Spectrum Scale cluster deployment fails, further clean-up might be needed. Follow these steps to ensure that all provisioned resources are deleted:

1. Access the **AWS Console**, and go to **EC2 service > Instances**, and search for instances with any of the following names, and terminate the instances:
 - <stack-name>-compute
 - <stack-name>-storage
 - <stack-name>-tiebreaker-desc
2. In the **AWS Console**, go to the **EC2 service > Network and Security > Security Groups**, and search for the following security groups, and delete them:
 - Storage-Sec-group
 - Compute-Sec-group
3. Navigate to **IAM Role service > Roles**, and search for the following IAM Roles, and delete them:
 - <stack-name><region-name>-Scaleworkflows
 - <stack-name>-Cluster-<ARN>
4. Access the S3 bucket that was specified in the CloudFormation template, and download all the files for further debug. After download is complete, delete all files to empty the S3 bucket.
5. Proceed to delete the stack from the **AWS Console CloudFormation service** page.

Chapter 12. Collecting debug data

If a failure occurs during deployment or lifecycle management, or an operational error is detected in the IBM Spectrum Scale cluster, debug logs are collected to facilitate problem resolution.

Collecting debug data for IBM Spectrum Scale deployment failure

IBM Spectrum Scale deployment process stores all the important information that is related to resource allocation and the CloudFormation stack in the customer-specified S3 bucket. It usually consists of the following information:

- Template input values file
- Resource allocation status file
- IBM Spectrum Scale cluster blueprint
- Deployment status file
- Deployment log file

Collecting debug data for IBM Spectrum Scale operational failure

During an IBM Spectrum Scale operational failure, all information from the S3 bucket that is listed in the previous section must be collected. To collect all IBM Spectrum Scale debug information, log on to any storage node, and run the **gpfs.snap** command. For more information, see the **gpfs.snap command** section in *IBM Spectrum Scale: Command and Programming Reference*.

Collecting debug data for a cluster lifecycle management failure

During a cluster lifecycle management failure, the following information must be collected:

- All information from the S3 bucket listed in the previous section
- Logs from the controller node, where the **mmcloudworkflows** command was run:
 - The log files from `/var/adm/ras/ibm_cloud_workflow_logs`.
 - The `ansible.log` from `/var/logs` directory, if it exists.
 - The `ibm_specsacle_ansible.log` from the `/var/logs` directory, if it exists

Chapter 13. 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. You can see the details about the stack in the CloudWatch logs. To access the logs, go to **CloudWatch > Logs > Log groups > <stack_name> > ScaleDeploymentLogGroup**.

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 [AWS Quick 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 14. Frequently Asked Questions

This page details some frequently asked questions.

Functional Support Matrices

IBM Spectrum Scale BYOL AWS Marketplace version	IBM Spectrum Scale version
BYOL 1.0.0	IBM Spectrum Scale 5.0.1.0
BYOL 1.1.0	IBM Spectrum Scale 5.0.2.1
BYOL 1.2.0	IBM Spectrum Scale 5.0.3.0
BYOL 1.3.1	IBM Spectrum Scale 5.0.5.3

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.

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 \(www.ibm.com/support/knowledgecenter\)](http://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 \(www.ibm.com/able\)](http://www.ibm.com/able) for more information about the commitment that IBM has to accessibility.

Notices

This information was developed for products and services offered in the US. This material might be available from IBM in other languages. However, you may be required to own a copy of the product or product version in that language in order to access it.

IBM may not offer the products, services, or features discussed in this document in other countries. Consult your local IBM representative for information on the products and services currently available in your area. Any reference to an IBM product, program, or service is not intended to state or imply that only that IBM product, program, or service may be used. Any functionally equivalent product, program, or service that does not infringe any IBM intellectual property right may be used instead. However, it is the user's responsibility to evaluate and verify the operation of any non-IBM product, program, or service.

IBM may have patents or pending patent applications covering subject matter described in this document. The furnishing of this document does not grant you any license to these patents. You can send license inquiries, in writing, to:

IBM Director of Licensing IBM Corporation North Castle Drive, MD-NC119 Armonk, NY 10504-1785 US

For license inquiries regarding double-byte character set (DBCS) information, contact the IBM Intellectual Property Department in your country or send inquiries, in writing, to:

Intellectual Property Licensing Legal and Intellectual Property Law IBM Japan Ltd. 19-21, Nihonbashi-Hakozakicho, Chuo-ku Tokyo 103-8510, Japan

INTERNATIONAL BUSINESS MACHINES CORPORATION PROVIDES THIS PUBLICATION "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF NON-INFRINGEMENT, MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. Some jurisdictions do not allow disclaimer of express or implied warranties in certain transactions, therefore, this statement may not apply to you.

This information could include technical inaccuracies or typographical errors. Changes are periodically made to the information herein; these changes will be incorporated in new editions of the publication. IBM may make improvements and/or changes in the product(s) and/or the program(s) described in this publication at any time without notice.

Any references in this information to non-IBM websites are provided for convenience only and do not in any manner serve as an endorsement of those websites. The materials at those websites are not part of the materials for this IBM product and use of those websites is at your own risk.

IBM may use or distribute any of the information you provide in any way it believes appropriate without incurring any obligation to you.

Licensees of this program who wish to have information about it for the purpose of enabling: (i) the exchange of information between independently created programs and other programs (including this one) and (ii) the mutual use of the information which has been exchanged, should contact:

IBM Director of Licensing IBM Corporation North Castle Drive, MD-NC119 Armonk, NY 10504-1785 US

Such information may be available, subject to appropriate terms and conditions, including in some cases, payment of a fee.

The licensed program described in this document and all licensed material available for it are provided by IBM under terms of the IBM Customer Agreement, IBM International Program License Agreement or any equivalent agreement between us.

The performance data discussed herein is presented as derived under specific operating conditions. Actual results may vary.

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and

cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

Statements regarding IBM's future direction or intent are subject to change or withdrawal without notice, and represent goals and objectives only.

All IBM prices shown are IBM's suggested retail prices, are current and are subject to change without notice. Dealer prices may vary.

This information is for planning purposes only. The information herein is subject to change before the products described become available.

This information contains examples of data and reports used in daily business operations. To illustrate them as completely as possible, the examples include the names of individuals, companies, brands, and products. All of these names are fictitious and any similarity to actual people or business enterprises is entirely coincidental.

COPYRIGHT LICENSE:

This information contains sample application programs in source language, which illustrate programming techniques on various operating platforms. You may copy, modify, and distribute these sample programs in any form without payment to IBM, for the purposes of developing, using, marketing or distributing application programs conforming to the application programming interface for the operating platform for which the sample programs are written. These examples have not been thoroughly tested under all conditions. IBM, therefore, cannot guarantee or imply reliability, serviceability, or function of these programs. The sample programs are provided "AS IS", without warranty of any kind. IBM shall not be liable for any damages arising out of your use of the sample programs.

Each copy or any portion of these sample programs or any derivative work must include a copyright notice as follows:

© (your company name) (year).

Portions of this code are derived from IBM Corp.

Sample Programs. © Copyright IBM Corp. _enter the year or years_.

If you are viewing this information softcopy, the photographs and color illustrations may not appear.

Trademarks

IBM, the IBM logo, and ibm.com are trademarks or registered trademarks of International Business Machines Corp., registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the Web at [Copyright and trademark information at www.ibm.com/legal/copytrade.shtml](http://www.ibm.com/legal/copytrade.shtml).

Intel is a trademark of Intel Corporation or its subsidiaries in the United States and other countries.

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

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 \(www.ibm.com/software/globalization/terminology\)](http://www.ibm.com/software/globalization/terminology) (opens in new window).

B

block utilization

The measurement of the percentage of used subblocks per allocated blocks.

C

cluster

A loosely coupled collection of independent systems (nodes) organized into a network for the purpose of sharing resources and communicating with each other. See also *GPFS cluster*.

cluster configuration data

The configuration data that is stored on the cluster configuration servers.

Cluster Export Services (CES) nodes

A subset of nodes configured within a cluster to provide a solution for exporting GPFS file systems by using the Network File System (NFS), Server Message Block (SMB), and 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.

clustered watch folder

Provides a scalable and fault-tolerant method for file system activity within an IBM Spectrum Scale file system. A clustered watch folder can watch file system activity on a fileset, inode space, or an entire file system. Events are streamed to an external Kafka sink cluster in an easy-to-parse JSON format. For more information, see the *mmwatch command* in the *IBM Spectrum Scale: Command and Programming Reference*.

control data structures

Data structures needed to manage file data and metadata cached in memory. Control data structures include hash tables and link pointers for finding cached data; lock states and tokens to implement distributed locking; and various flags and sequence numbers to keep track of updates to the cached data.

D

Data Management Application Program Interface (DMAPI)

The interface defined by the Open Group's XDSM standard as described in the publication *System Management: Data Storage Management (XDSM) API Common Application Environment (CAE) Specification C429*, The Open Group ISBN 1-85912-190-X.

deadman switch timer

A kernel timer that works on a node that has lost its disk lease and has outstanding I/O requests. This timer ensures that the node cannot complete the outstanding I/O requests (which would risk causing file system corruption), by causing a panic in the kernel.

dependent fileset

A fileset that shares the inode space of an existing independent fileset.

disk descriptor

A definition of the type of data that the disk contains and the failure group to which this disk belongs. See also *failure group*.

disk leasing

A method for controlling access to storage devices from multiple host systems. Any host that wants to access a storage device configured to use disk leasing registers for a lease; in the event of a perceived failure, a host system can deny access, preventing I/O operations with the storage device until the preempted system has reregistered.

disposition

The session to which a data management event is delivered. An individual disposition is set for each type of event from each file system.

domain

A logical grouping of resources in a network for the purpose of common management and administration.

E**ECKD**

See *extended count key data (ECKD)*.

ECKD device

See *extended count key data device (ECKD device)*.

encryption key

A mathematical value that allows components to verify that they are in communication with the expected server. Encryption keys are based on a public or private key pair that is created during the installation process. See also *file encryption key, master encryption key*.

extended count key data (ECKD)

An extension of the count-key-data (CKD) architecture. It includes additional commands that can be used to improve performance.

extended count key data device (ECKD device)

A disk storage device that has a data transfer rate faster than some processors can utilize and that is connected to the processor through use of a speed matching buffer. A specialized channel program is needed to communicate with such a device. See also *fixed-block architecture disk device*.

F**failback**

Cluster recovery from failover following repair. See also *failover*.

failover

(1) The assumption of file system duties by another node when a node fails. (2) The process of transferring all control of the ESS to a single cluster in the ESS when the other clusters in the ESS fails. See also *cluster*. (3) The routing of all transactions to a second controller when the first controller fails. See also *cluster*.

failure group

A collection of disks that share common access paths or adapter connections, and could all become unavailable through a single hardware failure.

FEK

See *file encryption key*.

fileset

A hierarchical grouping of files managed as a unit for balancing workload across a cluster. See also *dependent fileset*, *independent fileset*.

fileset snapshot

A snapshot of an independent fileset plus all dependent filesets.

file audit logging

Provides the ability to monitor user activity of IBM Spectrum Scale file systems and store events related to the user activity in a security-enhanced fileset. Events are stored in an easy-to-parse JSON format. For more information, see the *mmaudit* command in the *IBM Spectrum Scale: Command and Programming Reference*.

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

journalized 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.

Non-Volatile Memory Express (NVMe)

An interface specification that allows host software to communicate with non-volatile memory storage media.

P

policy

A list of file-placement, service-class, and encryption rules that define characteristics and placement of files. Several policies can be defined within the configuration, but only one policy set is active at one time.

policy rule

A programming statement within a policy that defines a specific action to be performed.

pool

A group of resources with similar characteristics and attributes.

portability

The ability of a programming language to compile successfully on different operating systems without requiring changes to the source code.

primary GPFS cluster configuration server

In a GPFS cluster, the node chosen to maintain the GPFS cluster configuration data.

private IP address

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 consists of copying designated changes for one location (a source) to another (a target) and synchronizing the data in both locations.

RKM server

See *remote key management server*.

rule

A list of conditions and actions that are triggered when certain conditions are met. Conditions include attributes about an object (file name, type or extension, dates, owner, and groups), the requesting client, and the container name associated with the object.

S**SAN-attached**

Disks that are physically attached to all nodes in the cluster using Serial Storage Architecture (SSA) connections or using Fibre Channel switches.

Scale Out Backup and Restore (SOBAR)

A specialized mechanism for data protection against disaster only for GPFS file systems that are managed by IBM 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).

W

watch folder API

Provides a programming interface where a custom C program can be written that incorporates the ability to monitor inode spaces, filesets, or directories for specific user activity-related events within IBM Spectrum Scale file systems. For more information, a sample program is provided in the following directory on IBM Spectrum Scale nodes: `/usr/lpp/mmfs/samples/util` called `tswf` that can be modified according to the user's needs.

Index

A

- accessibility features for IBM Spectrum Scale [57](#)
- active file management
 - AWS [33](#)
 - cache modes [35](#)
 - deployment [39](#)
 - prepare the environment [33](#)
- 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 [1–3](#), [5](#), [7](#), [9](#), [13](#), [16](#), [20](#), [25](#), [27](#), [31](#), [33](#), [35](#), [39](#), [41](#), [43](#), [45](#), [47](#), [49](#), [51](#), [53](#), [55](#)
- Amazon web services)IBM Spectrum Scale [41](#)
- availability zone [20](#)
- Availability Zone
 - new Amazon VPC
 - multiple Availability Zone [16](#)
- Availability Zones [2](#)
- AWS
 - AFM
 - best practices [41](#)
 - AFM configuration best practices [41](#)
 - AFM limitations [41](#)
 - debug data
 - collect [25](#), [51](#)
 - deployment failure
 - diagnose [31](#), [49](#)
- AWS Cloud
 - AFM
 - cache modes [35](#)
 - configuration [41](#)
 - deployment [39](#)
 - limitations [41](#)
 - preparation [33](#)
 - availability zones [2](#), [7](#)
 - Bastion host [47](#)
 - Data security [2](#), [47](#)
 - deployment
 - deployment options [13](#)
 - Deployment [9](#)
 - EC2-user [47](#)
 - Frequently Asked Questions [55](#)
 - IBM Spectrum Scaleupgrade [43](#)
 - Instance types [3](#)
 - Operating system [3](#)
 - regions [2](#)
 - Setup
 - Optimal setup [7](#)
 - Troubleshooting [53](#)
 - usage restrictions [3](#)
- AWS CloudFormation [1](#)

- AWS CloudFormation templates deployment [53](#)
- AWS Lambda function [1](#)

B

- Bastion host [27](#), [47](#)

C

- cache modes
 - independent writer [35](#)
 - IW [35](#)
 - local update [35](#)
 - LU [35](#)
 - read only [35](#)
 - RO [35](#)
 - single writer [35](#)
 - SW [35](#)
- collect debug data
 - AWS [25](#), [51](#)

D

- Data security [27](#), [47](#)
- Deployment
 - existing VPC [20](#)
 - multiple Availability Zone [16](#)
 - new VPC
 - single Availability Zone [13](#)
- deployment options
 - existing Amazon VPC [9](#), [13](#)
 - new Amazon VPC
 - multiple availability zone [9](#), [13](#)
 - single availability zone [9](#), [13](#)
- diagnose deployment failure
 - AWS [31](#), [49](#)

E

- EC2-user [27](#), [47](#)
- existing Amazon VPC [20](#)

F

- Frequently Asked Questions [55](#)

G

- GPFS
 - mmaws utility [27](#), [45](#)

I

- IBM Spectrum Scale information units [vii](#)

IBM Spectrum Scale on AWS [1–3](#), [5](#), [7](#), [9](#), [13](#), [16](#), [20](#), [25](#), [27](#),
[31](#), [33](#), [35](#), [39](#), [41](#), [43](#), [45](#), [47](#), [49](#), [51](#), [53](#), [55](#)
Identity and Access Management
IAM [47](#)
Instance types [3](#)

M

mmcloudworkflows [27](#), [45](#)

N

new Amazon VPC
multiple Availability Zone [16](#)
single Availability Zone [13](#)

O

Operating system [3](#)

R

regions [2](#)

S

Service.RequestLimitExceeded error [53](#)
setup
optimal setup [7](#)
Setup [5](#)
size limitation error [53](#)
Stack creation failure [53](#)
Stack creation failure message [53](#)

T

Troubleshooting
AWS CloudFormation templates deployment [53](#)
AWS CloudIBM Spectrum Scale [53](#)
CREATE_FAILED error [53](#)
no timeout message [53](#)
Service.RequestLimitExceeded error [53](#)
size limitation error [53](#)
Stack creation failure [53](#)
Stack creation failure message [53](#)

U

usage restrictions [3](#)
utility
mmaws [27](#), [45](#)



Part Number:
Product Number: 5641-DM1
5641-DM3
5641-DM5
5641-DA1
5641-DA3
5641-DA5
5737-F34
5737-I39
5765-DME
5765-DAE

SC27-9283-03



(1P) P/N: