IBM Spectrum Scale On AWS
Version 1.2.0

IBM Spectrum Scale On AWS Guide

SC27-9283-02

IBM Spectrum Scale On AWS
Version 1.2.0

IBM Spectrum Scale On AWS Guide

SC27-9283-02

Note
Before using this information and the product it supports, read the information in "Notices" on page 51.

This edition applies to version 5 release 0 modification 3 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 xxi. When you send information to IBM, you grant IBM a nonexclusive right to use or distribute the information in any way it believes appropriate without incurring any obligation to you.

Contents

Tables

About this information

This edition applies to IBM Spectrum Scale version 5.0.3 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
IBM Spectrum Scale: Concepts, Planning, and Installation Guide	This guide provides the following information: **Product overview** • 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 • Introduction to clustered watch • 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 **Planning** • Planning for GPFS • Planning for protocols • Planning for Cloud services • Planning for AFM • Planning for AFM DR • Firewall recommendations • Considerations for GPFS applications	System administrators, analysts, installers, planners, and programmers of IBM Spectrum Scale clusters who are very experienced with the operating systems on which each IBM Spectrum Scale cluster is based

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

Information unit	Type of information	Intended users
IBM Spectrum Scale: Concepts, Planning, and Installation Guide	**Installing** • 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 • Installing clustered watch • Steps to permanently uninstall GPFS	System administrators, analysts, installers, planners, and programmers of IBM Spectrum Scale clusters who are very experienced with the operating systems on which each IBM Spectrum Scale cluster is based

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

Information unit	Type of information	Intended users
IBM Spectrum Scale: Concepts, Planning, and Installation Guide	**Upgrading** • IBM Spectrum Scale supported upgrade paths • Upgrading to IBM Spectrum Scale 5.0.x from IBM Spectrum Scale 4.2.y • Upgrading to IBM Spectrum Scale 4.2.y from IBM Spectrum Scale 4.1.x • Upgrading to IBM Spectrum Scale 4.1.1.x from GPFS V4.1.0.x • Upgrading from GPFS 3.5 • Online upgrade support for protocols and performance monitoring • Upgrading 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 • Upgrading file audit logging authentication • Upgrading watch folder callbacks • 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	System administrators, analysts, installers, planners, and programmers of IBM Spectrum Scale clusters who are very experienced with the operating systems on which each IBM Spectrum Scale cluster is based

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

Information unit	Type of information	Intended users
IBM Spectrum Scale: Concepts, Planning, and Installation Guide	• Coexistence considerations • Compatibility considerations • Considerations for IBM Spectrum Protect for Space Management • GUI user role considerations • Applying maintenance to your GPFS system • Guidance for upgrading the operating system on IBM Spectrum Scale nodes • Servicing IBM Spectrum Scale protocol nodes	

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

Information unit	Type of information	Intended users
IBM Spectrum Scale: Administration Guide	This guide provides the following information: **Configuring** • Configuring the GPFS cluster • Configuring the CES and protocol configuration • Configuring and tuning your system for GPFS • Parameters for performance tuning and optimization • Ensuring high availability of the GUI service • Configuring and tuning your system for Cloud services • Configuring file audit logging • Steps that are taken when a watch is enabled with the mmwatch command • Configuring Active File Management • Configuring AFM-based DR • Tuning for Kernel NFS backend on AFM and AFM DR **Administering** • 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 • Native NFS and GPFS • Considerations for GPFS applications • Accessing a remote GPFS file system	System administrators or programmers of IBM Spectrum Scale systems

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

Information unit	Type of information	Intended users
IBM Spectrum Scale: Administration Guide	• 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 • Administering AFM • Administering AFM DR • Highly-available write cache (HAWC) • Local read-only cache • Miscellaneous advanced administration • GUI limitations	System administrators or programmers of IBM Spectrum Scale systems

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

Information unit	Type of information	Intended users
IBM Spectrum Scale: Problem Determination Guide	This guide provides the following information: **Monitoring** • Performance monitoring • Monitoring system health through the IBM Spectrum Scale GUI • Monitoring system health by using the mmhealth command • Monitoring events through callbacks • Monitoring capacity through GUI • Monitoring AFM and AFM DR • GPFS SNMP support • Monitoring the IBM Spectrum Scale system by using call home • Monitoring remote cluster through GUI • Monitoring file audit logging • Monitoring clustered watch **Troubleshooting** • Best practices for troubleshooting • Understanding the system limitations • Collecting details of the issues • Managing deadlocks • Installation and configuration issues • Upgrade issues • Network issues • File system issues • Disk issues • Security issues • Protocol issues • Disaster recovery issues • Performance issues • GUI issues • AFM issues • AFM DR issues • Transparent cloud tiering issues • File audit logging issues • Troubleshooting watch folder • Troubleshooting mmwatch • Message queue issues • Maintenance procedures • Recovery procedures • Support for troubleshooting • References	System administrators of GPFS systems who are experienced with the subsystems used to manage disks and who are familiar with the concepts presented in the *IBM Spectrum Scale: Concepts, Planning, and Installation Guide*

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

Information unit	Type of information	Intended users
IBM Spectrum Scale: Command and Programming Reference	This guide provides the following information: **Command reference** • 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 • mmclone command • mmcloudgateway command • mmcrcluster command • mmcrfileset command • mmcrfs command	• System administrators of IBM Spectrum Scale systems • Application programmers who are experienced with IBM Spectrum Scale systems and familiar with the terminology and concepts in the XDSM standard

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

Information unit	Type of information	Intended users
IBM Spectrum Scale: Command and Programming Reference	• 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 • mmhealth command • mmimgbackup command • mmimgrestore command • mmimportfs command • mmkeyserv command • mmlinkfileset command • mmlsattr command • mmlscallback command • mmlscluster command • mmlsconfig command • mmlsdisk command • mmlsfileset command • mmlsfs command • mmlslicense command • mmlsmgr command	• System administrators of IBM Spectrum Scale systems • Application programmers who are experienced with IBM Spectrum Scale systems and familiar with the terminology and concepts in the XDSM standard

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

Information unit	Type of information	Intended users
IBM Spectrum Scale: Command and Programming Reference	• mmlsmount command • mmlsnodeclass command • mmlsnsd command • mmlspolicy command • mmlspool command • mmlsqos command • mmlsquota command • mmlssnapshot command • mmmigratefs command • mmmount command • mmmsgqueue command • mmnetverify command • mmnfs command • mmnsddiscover command • mmobj command • mmperfmon command • mmpmon command • mmprotocoltrace command • mmpsnap command • mmputacl command • mmquotaoff command • mmquotaon command • mmremotecluster command • mmremotefs command • mmrepquota command • mmrestoreconfig command • mmrestorefs command • mmrestripefile command • mmrestripefs command • mmrpldisk command • mmsdrrestore command • mmsetquota command • mmshutdown command • mmsmb command • mmsnapdir command • mmstartup command • mmtracectl command • mmumount command • mmunlinkfileset command • mmuserauth command • mmwatch command • mmwinservctl command • spectrumscale command	• System administrators of IBM Spectrum Scale systems • Application programmers who are experienced with IBM Spectrum Scale systems and familiar with the terminology and concepts in the XDSM standard

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

Information unit	Type of information	Intended users
IBM Spectrum Scale: Command and Programming Reference	**Programming reference** • IBM Spectrum Scale Data Management API for GPFS information • GPFS programming interfaces • GPFS user exits • IBM Spectrum Scale management API commands • Watch folder API	• 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
IBM Spectrum Scale: Big Data and Analytics Guide	This guide provides the following information: Hadoop Scale Storage Architecture • Elastic Storage Server (ESS) • Share Storage (SAN-based storage) • File Placement Optimizer (FPO) • Deployment model • Additional supported features about storage IBM Spectrum Scale support for Hadoop • HDFS transparency • Supported IBM Spectrum Scale storage modes • Hadoop cluster planning • Installation and configuration of HDFS transparency • Application interaction with HDFS transparency • Upgrading the HDFS Transparency cluster • Rolling upgrade for HDFS Transparency • Security • Advanced features • Hadoop distribution support • Limitations and differences from native HDFS • Problem determination IBM Spectrum Scale Hadoop performance tuning guide • Overview • Performance overview • Hadoop Performance Planning over IBM Spectrum Scale • Performance guide	• System administrators of IBM Spectrum Scale systems • Application programmers who are experienced with IBM Spectrum Scale systems and familiar with the terminology and concepts in the XDSM standard

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

Information unit	Type of information	Intended users
IBM Spectrum Scale: Big Data and Analytics Guide	Hortonworks Data Platform 3.X • Planning • Installation • Upgrading and uninstallation • Configuration • Administration • Limitations • Problem determination Open Source Apache Hadoop • Apache Hadoop 3.0.x Support BigInsights® 4.2.5 and Hortonworks Data Platform 2.6 • Planning • Installation • Upgrading software stack • Configuration • Administration • Troubleshooting • Limitations • FAQ	• 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
IBM Spectrum Scale on AWS Guide	IBM Spectrum Scale on AWS • Summary of changes • Introduction to IBM Spectrum Scale on AWS • Setting up the IBM Spectrum Scale environment in the AWS Cloud • Deploying IBM Spectrum Scale on AWS • Cleaning up the cluster and the stack • Data security and AWS Identity and Access Management • Cluster lifecycle management and debug data collection • Upgrading IBM Spectrum Scale • Active file management on AWS • Troubleshooting • Frequently Asked Questions	• System administrators of IBM Spectrum Scale systems • Application programmers who are experienced with IBM Spectrum Scale systems and familiar with the terminology and concepts in the XDSM standard

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

Information unit	Type of information	Intended users
IBM Spectrum Scale Erasure Code Edition Guide	IBM Spectrum Scale Erasure Code Edition • Introduction to IBM Spectrum Scale Erasure Code Edition • Planning for IBM Spectrum Scale Erasure Code Edition • Installing IBM Spectrum Scale Erasure Code Edition • Incorporating IBM Spectrum Scale Erasure Code Edition in an Elastic Storage Server (ESS) cluster • Upgrading IBM Spectrum Scale Erasure Code Edition • Administering IBM Spectrum Scale Erasure Code Edition • Troubleshooting • IBM Spectrum Scale RAID Administration [1] **Note:** [1] For PDF or EPUB format of IBM Spectrum Scale RAID Administration documentation, see Elastic Storage Server for Power documentation on IBM Knowledge Center.	• System administrators of IBM Spectrum Scale systems • Application programmers who are experienced with IBM Spectrum Scale systems and familiar with the terminology and concepts in the XDSM standard

Prerequisite and related information

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

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

Conventions used in this information

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

Note: Users of IBM Spectrum Scale for Windows must be aware that on Windows, UNIX-style file names need to be converted appropriately. For example, the GPFS cluster configuration data is stored in the `/var/mmfs/gen/mmsdrfs` file. On Windows, the UNIX namespace starts under the `%SystemDrive%\cygwin64` directory, so the GPFS cluster configuration data is stored in the `C:\cygwin64\var\mmfs\gen\mmsdrfs` file.

Table 2. Conventions

Convention	Usage
bold	**Bold** words or characters represent system elements that you must use literally, such as commands, flags, values, and selected menu options. Depending on the context, **bold** typeface sometimes represents path names, directories, or file names.

Table 2. Conventions (continued)

Convention	Usage
bold underlined	**bold underlined** keywords are defaults. These take effect if you do not specify a different keyword.
`constant width`	Examples and information that the system displays appear in `constant-width` typeface. Depending on the context, `constant-width` typeface sometimes represents path names, directories, or file names.
italic	*Italic* words or characters represent variable values that you must supply. *Italics* are also used for information unit titles, for the first use of a glossary term, and for general emphasis in text.
<*key*>	Angle brackets (less-than and greater-than) enclose the name of a key on the keyboard. For example, **<Enter>** refers to the key on your terminal or workstation that is labeled with the word *Enter*.

\	In command examples, a backslash indicates that the command or coding example continues on the next line. For example:
	`mkcondition -r IBM.FileSystem -e "PercentTotUsed > 90" \`
	`-E "PercentTotUsed < 85" -m p "FileSystem space used"`

{*item*}	Braces enclose a list from which you must choose an item in format and syntax descriptions.
[*item*]	Brackets enclose optional items in format and syntax descriptions.
<**Ctrl-***x*>	The notation <**Ctrl-***x*> indicates a control character sequence. For example, <**Ctrl-c**> means that you hold down the control key while pressing <**c**>.
item...	Ellipses indicate that you can repeat the preceding item one or more times.
|	In *synopsis* statements, vertical lines separate a list of choices. In other words, a vertical line means *Or*. In the left margin of the document, vertical lines indicate technical changes to the information.

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

How to send your comments

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

mhvrcfs@us.ibm.com

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

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

gpfs@us.ibm.com

Summary of changes

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

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 name space across multiple instances with good performance, without requiring an in-depth knowledge of IBM Spectrum Scale.

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

Note: In IBM Spectrum Scale, the term `Node` is typically used to refer to any running instance of an operating system. The nodes deployed in this Amazon Machine Images (AMIs) are all Amazon Elastic Compute Cloud (Amazon EC2) instances, so the term instance is used in the place of node.

AWS Services

The following section gives a brief introduction to the various AWS services that are available.

AWS CloudFormation

AWS CloudFormation allows you to create and manage a collection of related AWS resources, and provision and update them in an orderly and predictable way. You use a template to describe all the AWS resources that are needed, while the AWS CloudFormation creates and configures the resources, and figures out dependencies. For more information, see AWS CloudFormation Documentation.

Amazon VPC

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

Amazon EC2

The Amazon EC2 service enables you to launch virtual machine instances with a variety of

operating systems. You can choose from the existing AMIs or import your own virtual machine images. For more information, see Amazon Elastic Compute Cloud Documentation.

Amazon Auto Scaling
The Amazon Auto Scaling service helps maintain high availability and manage capacity by automatically increasing or decreasing the EC2 instance fleet. You can use auto scaling to run your fleet at optimal utilization by increasing instance capacity during demand spikes and decreasing capacity during down times. For more information, see AWS Auto Scaling Documentation.

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

Amazon S3
Amazon Simple Storage Service (Amazon S3) is a storage for the Internet. You can use Amazon S3 to store and retrieve any amount of data at any time, from anywhere on the web. You can accomplish these tasks using the simple and intuitive web interface of the AWS Management Console. For more information, see Amazon Simple Storage Service Documentation.

Amazon CloudWatch
Amazon CloudWatch is a monitoring service for AWS Cloud resources and the applications you run on AWS. You can use CloudWatch to collect and track metrics, collect and monitor log files, set alarms, and automatically react to changes in your AWS resources. For more information, see Amazon CloudWatch Documentation.

Amazon IAM
AWS Identity and Access Management (IAM) enables you to securely control access to 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.

Regions and availability zones

IBM Spectrum Scale deployed on the AWS cloud uses regions to place or replicate resources and data in multiple availability zones. If the `multiple availability zone` template is selected, this architecture deploys the IBM Spectrum Scale cluster nodes across two or more availability zones within an AWS region by default. Within each region there can be two or more availability zones.

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

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

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

- Either 25 Gbps or 10 Gbps of bandwidth for a multi-flow traffic 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 avalability zone` template configuration for an IBM Spectrum Scale cluster with replication, each element of the data and the metadata is replicated in a separate availability zone to avoid the loss of data when hardware failures occur in one of the availability zones. By default, instances are distributed evenly between the availability zones. Therefore, each availability zone has one private subnet. However, these zones remain a single or logical cluster.

- If you choose to deploy a `single availability zone` configuration, all the server and compute instances are deployed into a single availability zone, and IBM Spectrum Scale creates only one copy of each data element.
- If you choose to deploy a `multi availability zone` configuration, all the server and compute instances are deployed into multiple availability zones, and IBM Spectrum Scale creates multiple copies of each data element.

IBM Spectrum Scale instance types and operating system

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

It is recommended that you benchmark the environment to make sure you achieve the level of performance you need before starting a production deployment. The deployment launches an EC2 instance running RHEL 7.6. For more information on EC2 instances, see Amazon EC2 Instance Types.

IBM Spectrum Scale usage restrictions

This version does not support the following features of IBM Spectrum Scale:

- Protocol support, including the use of Cluster Export Services (CES) nodes and protocol access such as Network File System (NFS), Object, and Server Message Block (SMB)
- Transparent Cloud Tiering (TCT)
- Compression
- Encryption
- Data Management API (DMAPI) support, including Hierarchical Storage Management (HSM) to tape
- Hadoop Distributed File System (HDFS) connector support
- Call home
- Multi-cluster support: Exporting an IBM Spectrum Scale file system from one IBM Spectrum Scale cluster to another IBM Spectrum Scale
- IBM Spectrum Scale Graphical User Interface
- User name space management and quota management
- Snapshots and clones

Additional limitations include the following:

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

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

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

Note: You can choose to create a new VPC for the IBM Spectrum Scale deployment or use your existing VPC on AWS. The template that deploys the setup into an existing VPC skips the first four components of this list:

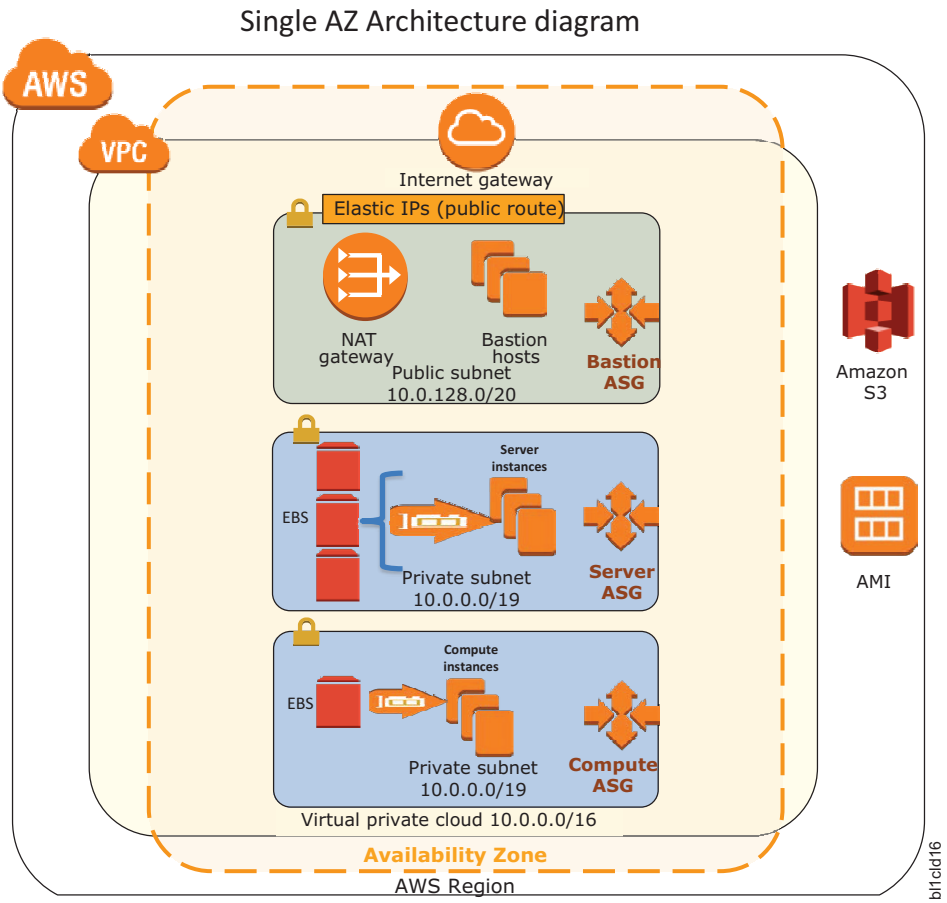

Figure 1. Single availability zone architecture diagram

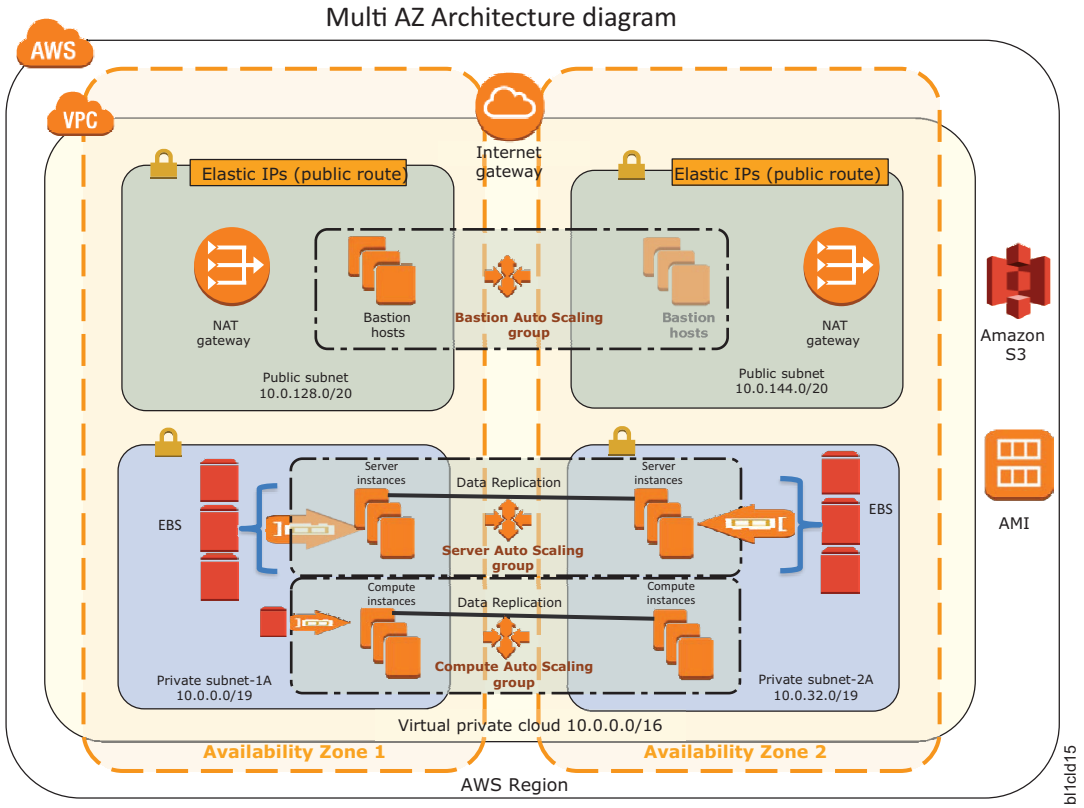

Figure 2. Multi availability zone architecture diagram

- A VPC that spans two availability zones and includes two public and two private subnets, for security and high availability. The number of private subnets is configurable. For more information, see "Optimal setup considerations" on page 7.
- An internet gateway to allow access to the internet.
- The managed NAT gateways in the public subnets allow outbound internet access for resources in the private subnets. For more information, see "Optimal setup considerations" on page 7.
- A bastion host in a public subnet provides Secure Shell (SSH) access to the IBM Spectrum Scale cluster. The bastion host instance is managed by an auto scaling group of one. This ensures that there is always at least one host available.
- An AWS Identity and Access Management instance role with fine-grained permissions for access to AWS services necessary for the deployment process.
- Appropriate security groups for each instance or function to restrict access to only necessary protocols and ports. The AWS offering opens only ports for SSH and the IBM Spectrum Scale daemon.

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

- A100-GB Amazon Elastic Block Store (Amazon EBS) volume for the root device.
- By default, one 500-GB EBS volume is attached for use as an NSD storage per NSD server . You can change the number and size of the NSD storage EBS volumes that is attached per NSD server during the deployment.

Each IBM Spectrum Scale compute instance has:

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

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

Optimal setup considerations

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

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

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

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

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

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

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

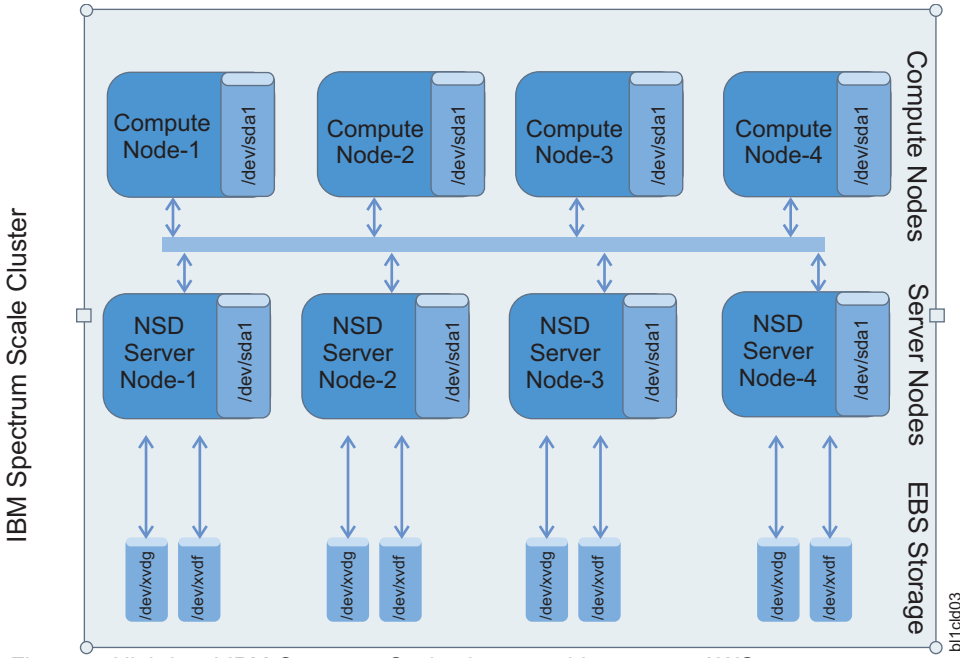

Figure 3. High-level IBM Spectrum Scale cluster architecture on AWS

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

Chapter 3. Deploying IBM Spectrum Scale on AWS

Follow these given steps to deploy IBM Spectrum Scale on AWS:

1. Prepare your AWS account:
   a. Create an AWS account at AWS, if you do not have one.
   b. Use the region selector in the navigation bar to choose the AWS region where you want to deploy IBM Spectrum Scale on AWS.
   c. Create a key pair in your preferred region. For information on how to create a key pair, see Amazon EC2 Key Pairs
   d. If necessary, request a service limit increase for the EC2 instance types that you intend to deploy. To do this, in the AWS Support Center, choose **Create Case** > **Service Limit Increase** > **EC2 instances**, and then complete the fields in the **Limit Increase** form.
2. Launch the IBM Spectrum Scale AWS stack:
   a. Choose one of the following options to launch the AWS CloudFormation template into your AWS account.
      - Option 1: Deploying IBM Spectrum Scale on a new Amazon VPC with a single availability zone
      - Option 2: Deploying IBM Spectrum Scale on a new Amazon VPC with multiple availability zones
      - Option 3: Deploying IBM Spectrum Scale on an existing Amazon VPC.

      For a single availability zone deployment configuration, IBM Spectrum Scale does not replicate any of the file system data. It is recommended that this option only be used for cases in which a higher probability of data loss is acceptable in order to potentially improve performance and save on data movement costs. For example, in a scratch file system scenario.

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

      Note: This is where the network infrastructure for IBM Spectrum Scale is built. The template is launched in the US East (Ohio) region by default. You can change the region.
   b. On the **Select Template** page, keep the default setting for the **Template URL**, and then choose **Next**.
   c. On the **Specify Details** page, change the stack name if needed. Review the parameters for the template.

      Note: Provide values for the parameters that require input. For all other parameters, review the default settings and customize them as necessary.
   d. In the **Options** page, you can specify tags or key-value pairs for resources in your stack, and set **Advanced options**.

      Note: For more information on how to specify tags, see AWS CloudFormation Resource Tags Type. For setting stack options in the **Advanced options** section, see Setting AWS CloudFormation Stack Options.
   e. In the **Review** page, review and confirm the template settings. Under **Capabilities**, select the check box to acknowledge that the template creates IAM resources.

f. Click **Create** to deploy the stack.

g. Monitor the status of the stack. When the status displays `CREATE_COMPLETE`, the IBM Spectrum Scale environment is ready.

h. View the resources that were created for the stack in the URLs displayed in the **Outputs** tab.

Note: This IBM Spectrum Scale AWS stack deployment is automated by the nested AWS CloudFormation templates. The main template builds the network-related resources first, using the VPC template, and then launches separate stacks for the bastion host and IBM Spectrum Scale cluster. Deleting the stack created by the main template deletes the entire IBM Spectrum Scale deployment stack. However, you still need to delete CloudWatch alerts manually.

3. Connect to the IBM Spectrum Scale cluster.

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

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

   a. Open the Amazon EC2 console at https://console.aws.amazon.com/ec2/.

   b. In the navigation pane, under **Instances**, check for the public DNS (IPv4) value for the instance named `LinuxBastion`.

   c. Use your AWS private key to connect to the bastion host using SSH.

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

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

      Note: For more information about using an SSH agent to forward your private key on connection, see Using SSH agent forwarding.

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

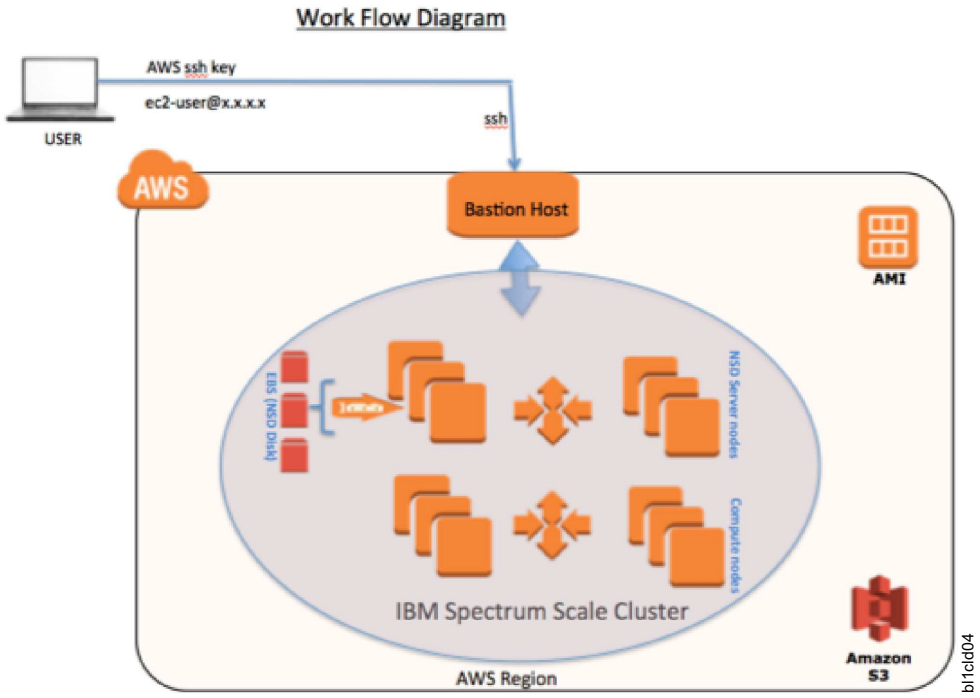

Figure 4. High-level IBM Spectrum Scale cluster architecture to connect from host

4. Test the deployment using IBM Spectrum Scale commands

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

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

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

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

GPFS cluster information
========================
  GPFS cluster name:         ip-10-0-1-110.ap-south-1.compute.internal
  GPFS cluster id:           12901386493707864068
  GPFS UID domain:           ip-10-0-1-110.ap-south-1.compute.internal
  Remote shell command:      /usr/bin/ssh
  Remote file copy command:  /usr/bin/scp
  Repository type:           CCR

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

   For more information on the **mmlscluster** 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 an output similar to the following:

```
[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 an output similar to the following:

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

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

               =============                         ==================== ===================
(total)            1053818880                         1041760256 ( 99%)         17584 ( 0%)

Inode Information
-----------------
Number of used inodes:            4038
Number of free inodes:          497722
Number of allocated inodes:     501760
Maximum number of inodes:      1025024
```

For more information on the **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.

Table 3. File System Configurations

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

Table 4. NSD Configurations

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

Table 5. Server Node Configurations

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

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

Table 6. Compute Node Configurations

Parameter label (name)	Default	Description
Compute Node Count (ComputeNodeCount)	2	The number of IBM Spectrum Scale compute node instances. You can select 1-64 instances.
Compute Instance Type (ComputeInstanceType)	t2.medium	The instance type to use for the compute node instances.

Table 7. Network Configuration

Parameter label (name)	Default	Description
Availability Zone (AvailabilityZones)	requires input	Availability zone to use for the subnet in the VPC.
VPC CIDR(VPCCIDR	10.0.0.0/16	The CIDR block for the VPC.
Private Subnet CIDR (PrivateSubnetCIDR)	10.0.1.0/24	The CIDR block for the private subnet 1A located in availability zone 1.
Public Subnet CIDR (PublicSubnetCIDR)	10.0.3.0/24	The CIDR block for the public DMZ subnet 1 located in availability zone 1.
Allowed External Access CIDR (RemoteAccessCIDR)	requires input	The CIDR block that is allowed external SSH access to the bastion hosts, e.g., *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 8. Amazon EC2 Configuration

Parameter label (name)	Default	Description
Key Pair Name (KeyPairName)	requires input	A public/private key pair, which allows you to connect securely to your instance after it launches. When you created an AWS account, this is the key pair you created in your preferred region.
Bastion AMI OS (BastionAMIOS)	Amazon-Linux- HVM	The Linux distribution for the AMI to be used for the bastion host instances. If you choose CentOS, make sure that you have a subscription to the CentOS AMI in AWS Marketplace..
Bastion Instance Type (BastionInstanceType)	t2.micro	TheEC2 instance type for the bastion host instances.

Table 9. Personal Configuration

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

Table 10. License Information

Parameter label (name)	Default	Description
IBM Customer Number (IBMCustomerNumber)	requires input	Required to validate the customer entitlement for BYOL offering.
License Agreement Terms (LicenseAgreementTerms)	requires input	Review the licensing terms at Licence Information, and if you agree to the terms, choose **Accept**.

Option 2: Deploying IBM Spectrum Scale on a new Amazon VPC with multiple availability zones

The following section details the parameters for a new VPC deployment with multiple availability zones.

Table 11. File System Configurations

Parameter label (name)	Default	Description

Table 11. File System Configurations (continued)

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

Table 12. NSD Configurations

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

Table 13. Server Node Configurations

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

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

Table 14. Compute Node Configurations

Parameter label (name)	Default	Description
Compute Node Count (ComputeNodeCount)	2	The number of IBM Spectrum Scale compute node instances. You can select 1-64 instances.
Compute Instance Type (ComputeInstanceType)	t2.medium	The instance type to use for the compute node instances.

Table 15. Network Configuration

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

Table 15. Network Configuration (continued)

Public Subnet 2 CIDR (PublicSubnet2CIDR)	10.0.2.0/24	The CIDR block for the public subnet located in availability zone 2.
Allowed External Access CIDR (RemoteAccessCIDR)	requires input	The CIDR block that is allowed external SSH access to the bastion hosts, e.g., *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 16. Amazon EC2 Configuration

Parameter label (name)	Default	Description
Key Pair Name (KeyPairName)	requires input	A public/private key pair, which allows you to connect securely to your instance after it launches. When you created an AWS account, this is the key pair you created in your preferred region.
Bastion AMI OS (BastionAMIOS)	Amazon-Linux- HVM	The Linux distribution for the AMI to be used for the bastion host instances. If you choose CentOS, make sure that you have a subscription to the CentOS AMI in AWS Marketplace..
Bastion Instance Type (BastionInstanceType)	t2.micro	TheEC2 instance type for the bastion host instances.

Table 17. Personal Configuration

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

Table 18. License Information

Parameter label (name)	Default	Description
IBM Customer Number (IBMCustomerNumber)	requires input	Required to validate the customer entitlement for BYOL offering.
License Agreement Terms (LicenseAgreementTerms)	requires input	Review the licensing terms at Licence Information, and if you agree to the terms, choose **Accept**.

Option 3: Deploying IBM Spectrum Scale on an existing Amazon VPC

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

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

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

  Note: It is recommended to use the following:

  1. Dedicated private subnet for the IBM Spectrum Scale cluster
  2. Subnet masks that are large enough to support maximum number of nodes
  3. Same subnet mask for all private subnets
  4. Subnets that are not defined as default subnet
- Create an S3 endpoint. An S3 endpoint is created by the CF template for VPC that is being used to create VPC.
- VPC endpoints must be created and added to the route tab.

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

Table 19. File System Configurations

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

Table 20. NSD Configurations

Parameter label (name)	Default	Description
EBS Type (EBSType)	gp2	The EBS volume type for IBM Spectrum Scale storage attached to each NSD server node. Options are: General Purpose SSD (gp2), Provisioned IOPS SSD (io1), Cold HDD (sc1), Throughput Optimized HDD (st1), and EBS Magnetic (standard). For more information about EBS volume type choices, see Amazon EBS.
Disk Per Node (DiskPerNode)	1	The number of NSD volumes to attach to each NSD server node. You can choose 1-15 disks.
Disk Size (DiskSize)	500	The disk size of NSD volume(s) attached to each NSD server node, in GiBs. Supported disk sizes are 10-16,384 GiB.

Table 21. Server Node Configurations

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

Table 22. Compute Node Configurations

Parameter label (name)	Default	Description
Compute Node Count (ComputeNodeCount)	2	The number of IBM Spectrum Scale compute node instances. You can select 1-64 instances.
Compute Instance Type (ComputeInstanceType)	t2.medium	The instance type to use for the compute node instances.

Table 23. Network Configuration

Parameter label (name)	Default	Description
Availability Zone (AvailabilityZones)	requires input	The list of Availability Zones to use for the subnets in the VPC. Only two Availability Zones are used for this deployment, and the logical order of your selections is preserved.
VPC CIDR(VPCCIDR	10.0.0.0/16	The CIDR block for the VPC.
Private Subnet 1 CIDR (PrivateSubnet1CIDR)	10.0.1.0/24	The CIDR block for the private subnet located in Availability Zone 1.
Private Subnet 2 CIDR (PrivateSubnet2CIDR)	10.0.3.0/24	The CIDR block for the private subnet located in Availability Zone 2.
Public Subnet 1 CIDR (PublicSubnet1CIDR)	10.0.0.0/24	The CIDR block for the public subnet located in Availability Zone 1.
Public Subnet 2 CIDR (PublicSubnet2CIDR)	10.0.2.0/24	The CIDR block for the public subnet located in Availability Zone 2.
Allowed External Access CIDR (RemoteAccessCIDR)	requires input	The CIDR block that is allowed external SSH access to the bastion hosts, e.g., *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 24. Amazon EC2 Configuration

Parameter label (name)	Default	Description
Key Pair Name (KeyPairName)	requires input	A public/private key pair, which allows you to connect securely to your instance after it launches. When you created an AWS account, this is the key pair you created in your preferred region.
Bastion AMI OS (BastionAMIOS)	Amazon-Linux- HVM	The Linux distribution for the AMI to be used for the bastion host instances. If you choose CentOS, make sure that you have a subscription to the CentOS AMI in AWS Marketplace..
Bastion Instance Type (BastionInstanceType)	t2.micro	TheEC2 instance type for the bastion host instances.

Table 25. Personal Configuration

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

Table 26. License Information

Parameter label (name)	Default	Description
License Agreement Terms (LicenseAgreementTerms)	requires input	Review the licensing terms at Licence Information, and if you agree to the terms, choose **Accept**.

Chapter 4. Cleaning up the cluster and the stack

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

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

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

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

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

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

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

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

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

Chapter 5. Data security and AWS Identity and Access Management

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

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

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

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

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

AWS Identity and Access Management (IAM)

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

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

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

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

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

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

ComputeSecurityGroup

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

Chapter 6. Cluster lifecycle management and debug data collection

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

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

mmaws utility

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

Synopsis

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

or

```
mmaws remove_nodes{--ip-addresses IP_ADDRESSES [IP_ADDRESSES...] | --hostnames HOSTNAMES [HOSTNAMES...]}
```

or

```
mmaws list_instances
```

or

```
mmaws start_nodes {--ip-addresses  IP_ADDRESSES [IP_ADDRESSES...] |
                   --instance-ids INSTANCE_IDS [INSTANCE_IDS...]}
```

or

```
mmaws stop_nodes {--all |  --compute  |
                  --ip-addresses  IP_ADDRESSES |
                  --hostnames HOSTNAMES [HOSTNAMES...] |
                  --nodeclass NODECLASS [NODECLASS...]}
```

or

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

or

```
mmaws collect_debug_data
```

Availability

Available with IBM Spectrum Scale on AWS.

Description

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

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

Parameters

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

--node-type {compute,server}
Specifies the IBM Spectrum Scale node type.

--num-instances *NUM_INSTANCES*
Specifies the number of nodes to be added.

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

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

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

list_instances
Lists the instances in the VPC.

start_nodes
Starts the specified nodes in an IBM Spectrum Scale cluster.

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

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

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

stop_nodes
Stops the specified nodes in an IBM Spectrum Scale cluster:

--all
Stops all the nodes in the cluster.

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

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

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

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

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

create_lambda_functions
Creates a lambda function.

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

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

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

collect_debug_data
Collects debug data.

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

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

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

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

Exit status

0 Successful completion.

nonzero
A failure has occurred.

Security

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

Examples

1. To add compute nodes:

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

The system displays output similar to this:

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

2. To add server nodes:

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

The system displays output similar to this:

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

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

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

The system displays output similar to this:

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

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

4. To list the nodes:

```
# mmaws list_nodes
```

The system displays output similar to this:

```
Compute Node/Instances List:
|---------------------|---------------|---------------|------------------|------------|
|    Instance Id      |   Private IP  | Instance Type | AvailabilityZone |   State    |
|---------------------|---------------|---------------|------------------|------------|
| i-0601f80f96812e70f |    10.0.1.237 |     t2. micro |       us-east-2a |    running |
| i-0b25af79f2f1fd4ff |     10.0.3.96 |     t2. micro |       us-east-2b |    running |
|---------------------|---------------|---------------|------------------|------------|

NSD Server Node/Instances List:
|---------------------|---------------|---------------|------------------|------------|
|    Instance Id      |   Private IP  | Instance Type | AvailabilityZone |   State    |
|---------------------|---------------|---------------|------------------|------------|
| i-0693826b9e5648e58 |    10.0.3.175 |      t2.micro |       us-east-2b |    running |
| i-0979e758a545da2ed |    10.0.1.226 |      t2.micro |       us-east-2a |    running |
|---------------------|---------------|---------------|------------------|------------|
```

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

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

The system displays output similar to this:

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

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

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

The system displays output similar to this:

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

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

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

The system displays output similar to this:

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

8. To collect debug data:

```
# mmaws collect_debug_data
```

The system displays output similar to this:

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

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

Location

`/usr/lpp/mmfs/bin`

Lambda function to start nodes

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

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

Consider the following before creating a Lambda function:

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

```
# mmaws create_lambda_functions --function-path /usr/lpp/mmfs/bin/ibm_aws_workflows/ibm_aws_lambda_funcs/mm_aws_lambda_start_nodes.py --iam-role Scale_lambda_exec_role
2019-04-04 13:22:06,977 - mm_aws_create_lambda_funcs - INFO - Logging in to file: /var/adm/ras/aws_scale_logs/mm_aws_create_lambda_funcs.log_2019-Apr-04_13-22-06
2019-04-04 13:22:06,977 - mm_aws_create_lambda_funcs - INFO - A. Performing prerequisite check
2019-04-04 13:22:06,978 - mm_aws_create_lambda_funcs - INFO - B. Performing current stack resource identification
2019-04-04 13:22:07,723 - mm_aws_create_lambda_funcs - INFO - Provided IAM role (Scale_lambda_exec_role) contains required "AWSLambdaBasicExecutionRole" policy.
2019-04-04 13:22:07,723 - mm_aws_create_lambda_funcs - INFO - C. Performing Lambda function vs. filepath mapping
2019-04-04 13:22:07,723 - mm_aws_create_lambda_funcs - INFO - Identified Lambda function(s) vs. filepath: {'start_nodes': '/usr/lpp/mmfs/bin/ibm_aws_workflows/ibm_aws_lambda_funcs/mm_aw
2019-04-04 13:22:07,723 - mm_aws_create_lambda_funcs - INFO - D. Performing Lambda function creation
2019-04-04 13:22:08,234 - mm_aws_create_lambda_funcs - INFO - Lambda function (start_nodes) created successfully.
```

Note: `Scale_lambda_exec_role` is the name of the IAM role. The user should provide the name of the IAM role which is created with the `AWSLambdaBasicExecutionRole` policy.

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

```
mmaws create_lambda_func --iam-role Scale_lambda_exec_role
```

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

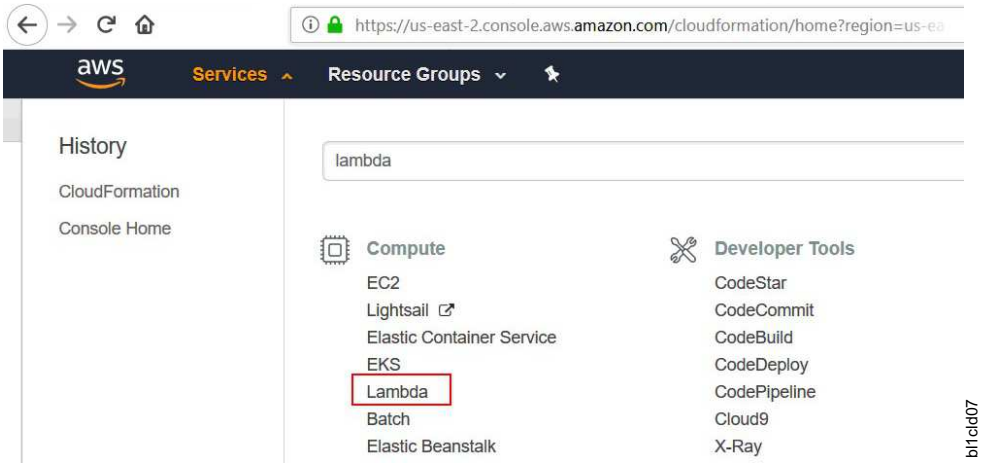

Figure 5. Lambda console

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

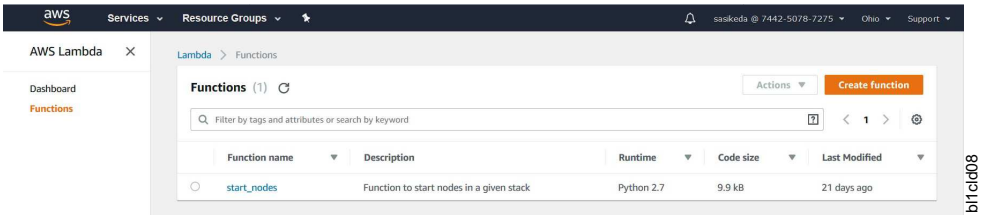

Figure 6. Viewing Lambda function details

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

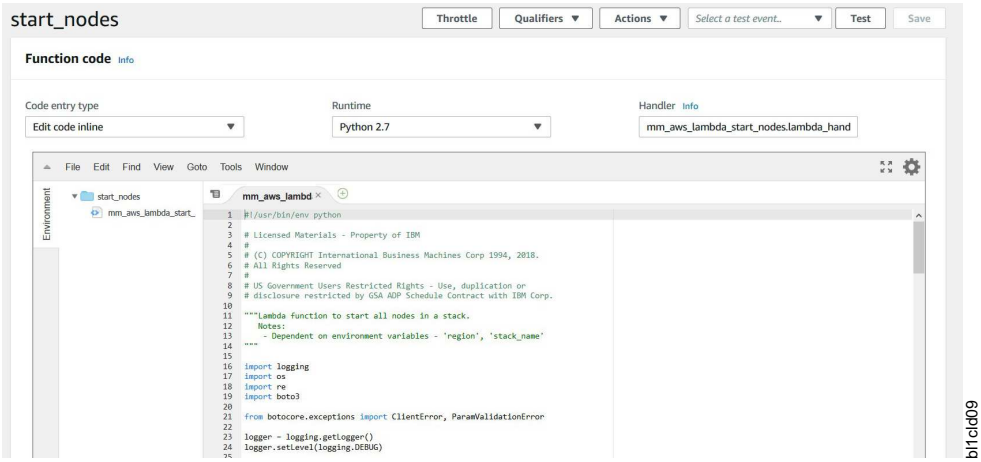

Figure 7. Function code information

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

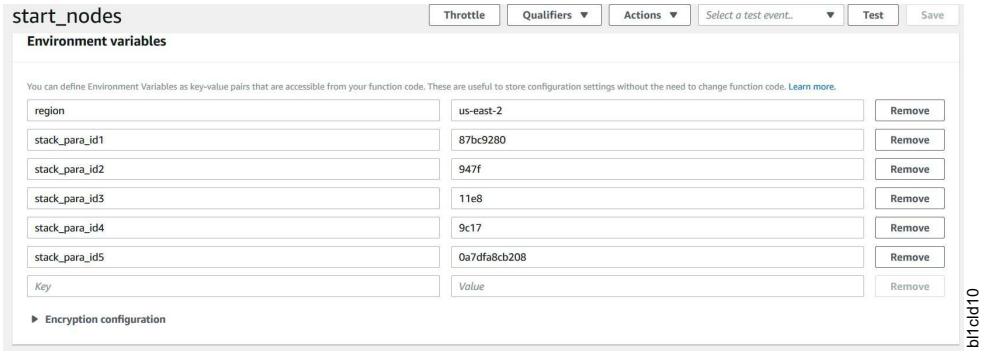

Figure 8. Environment variables

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

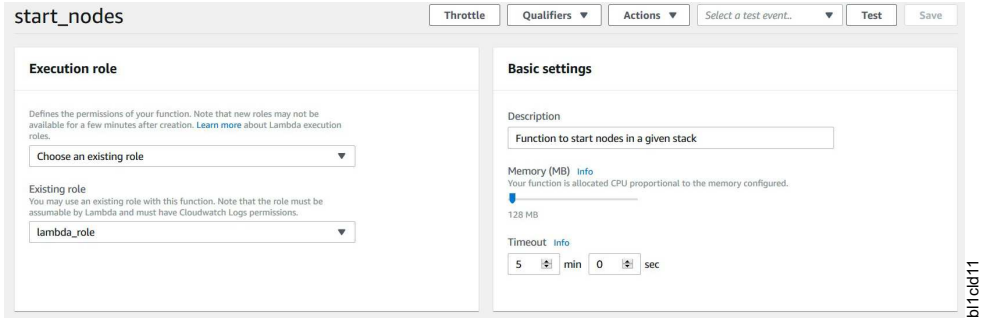

Figure 9. Environment variables

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

Follow these steps to execute the Lambda function:

1. Click on **Test**.

Figure 10. Execute Lambda function

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

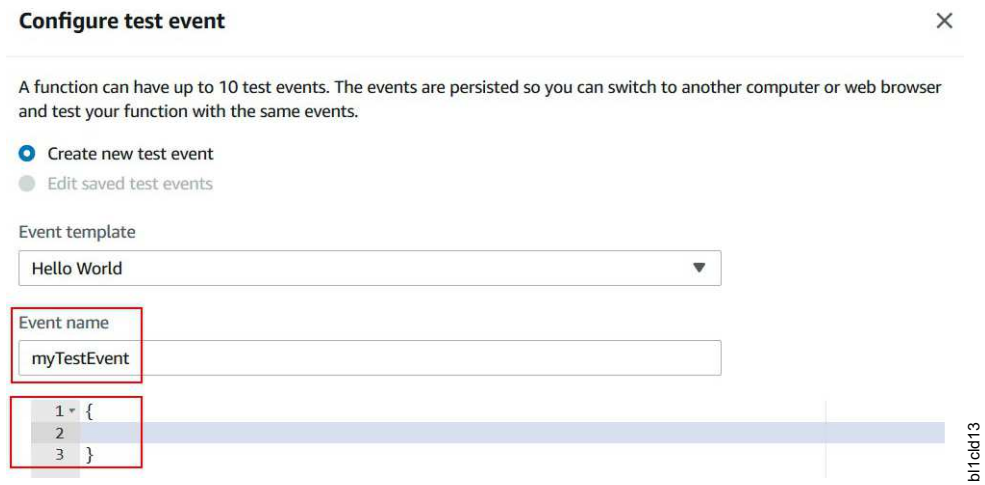

Figure 11. Execute Lambda function

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

Figure 12. Select test event

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

Chapter 7. Upgrading IBM Spectrum Scale

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

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

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

   The command has the following options:

   -h Displays the options for this script.

   -s Specifies the s3bucket that contains IBM Spectrum Scale Data Management bundle.

   -l Specifies the local path of the directory or folder that contains IBM Spectrum Scale Data Management bundle.

   -b Specifies the build file name to be upgraded.

   -V Provides the version information to be upgraded to.

   -v Displays the verbose output.

   -f Specifies the force option to silently accept the IBM Spectrum Scale electronic license agreement.

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

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

   The following code displays an example of the command to upgrade from IBM Spectrum Scale version 5.0.3.0 to IBM Spectrum Scale version 5.0.3.1:

   ```
   [ec2-user@ip-10-0-1-187]$/usr/lpp/mmfs/bin/ibm_aws_workflows/aws_upgrade_scale_1 -s s3_bucket/
   5.0.3.1_PTF -V 5.0.3.1 -b Spectrum_Scale_Protocols_Data_Management-5.0.3.1-x86_64-Linux-install -v
   ```

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

Chapter 8. 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 sections 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-premise environment.

   A IBM Spectrum Scale cluster must be first setup in the customer's datacenter. Once an on-premise 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 on NFS services, see NFS protocol quick reference.
2. Set up the IBM Spectrum Scale cluster on AWS.

   An IBM Spectrum Scale cluster must be setup 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 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-premise and public cloud resources, and securely move data between them. The data that flows between the two environments typically does so via the public internet, posing significant security and privacy risks. To ensure that the in-flight data is protected, a secure connection using IPSec based Virtual Private Network (VPN) between the on-premise and public cloud networks must be set up. This ensures that all the data passing between the on-premise and AWS cloud cluster is encrypted and protected. The overall performance of the data transfer between the on-premise and AWS clusters is dependent 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-premise and AWS IBM Spectrum Scale instances.

   Note: This is only an example deployment. Factors like on-premise network topology, network infrastructure and its configuration, other security configurations as well as choice of VPN solution etc. 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 setup and configured optimally and securely.

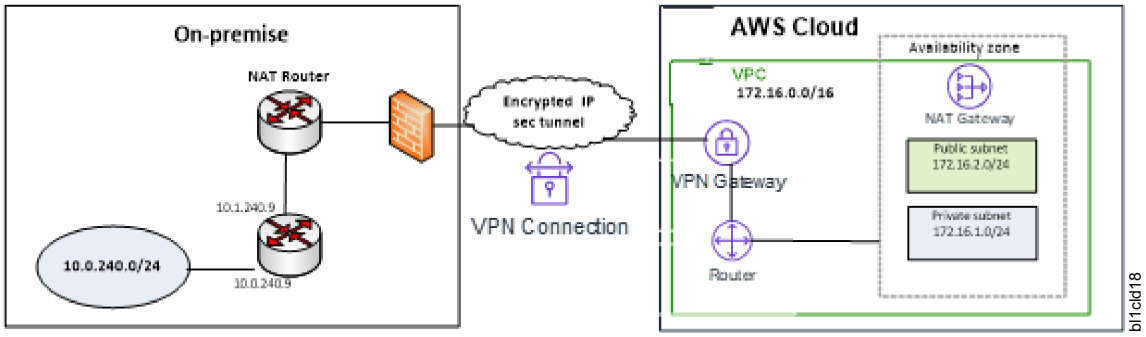

Figure 13. An example on-premise to AWS hybrid cloud networking architecture

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

Router
: Used for configuring the IPsec tunnel configuration.

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-premise.

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-premise 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-premise and AWS environments. The VPN connection depends on both the network infrastructure and configuration on-premise as well as 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 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 is required across the AFM home and cache sites to ensure that the user access to files are 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 meta data 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) using an external server for the ID mapping information. The same set of authentication mechanisms need 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 has been configured in the on-premise environment. It is important that the AWS cloud has the same UID and GID mapping that exists on-premise. AWS provides directory services such as AWS Managed Microsoft AD that also has LDAP connectors that allow AWS to communicate via AD or LDAP protocols. For more information on UID and GID mapping, see the *Requirements for UID and GID on the cache and home clusters* section in the *IBM Spectrum Scale: Concepts, Planning, and Installation Guide* in 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, Chrony etc. For more information on setting up time syncs, see Setting the Time for Your Linux Instance and Keeping Time With Amazon Time Sync Service.

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

AFM cache modes

The following section describes the different cache modes

There are four types of cache modes:

- "Read only (RO) cache mode"
- Local Update (LU) cache mode
- "Single Writer (SW) cache mode" on page 39
- "Independent Writer (IW) cache mode" on page 40

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.

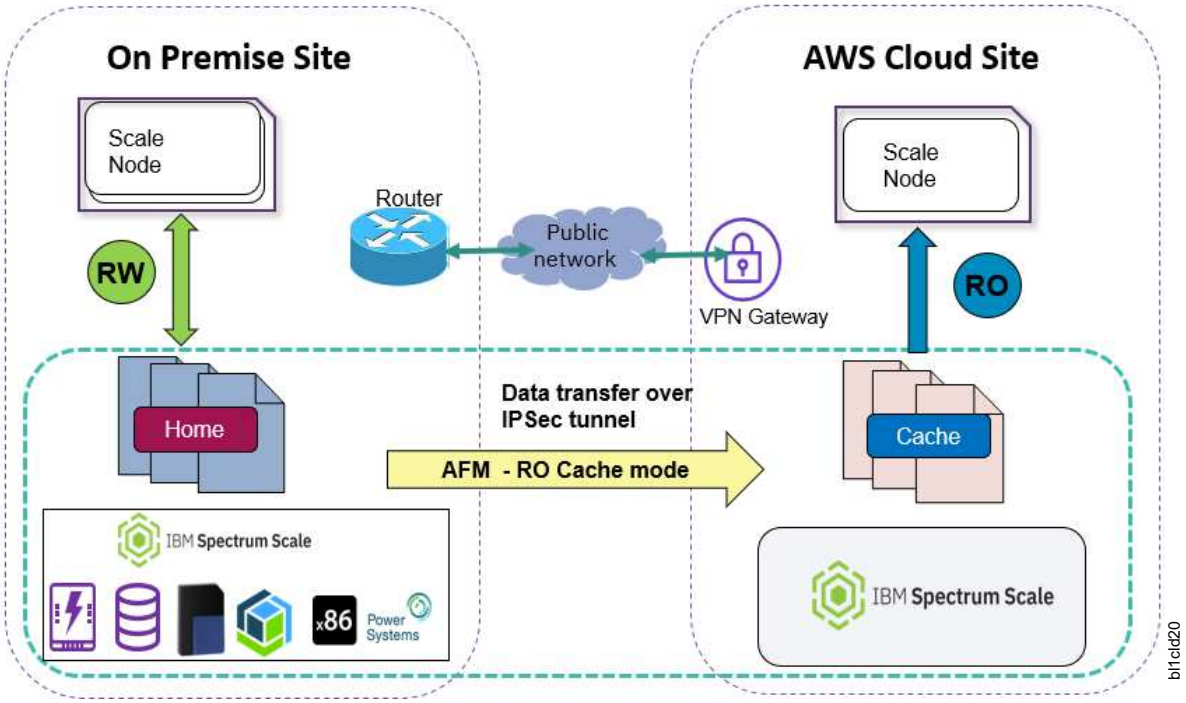

Figure 14. Read only cache mode

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 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 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 made to that directory in the home cluster will no longer reflect in the cache cluster. Changes can include renaming the directory, removing a directory, etc. However, any changes made to the files within the modified directory in the home cluster continue to reflect in the cache cluster, as long as the files have not been modified before in the cache cluster.

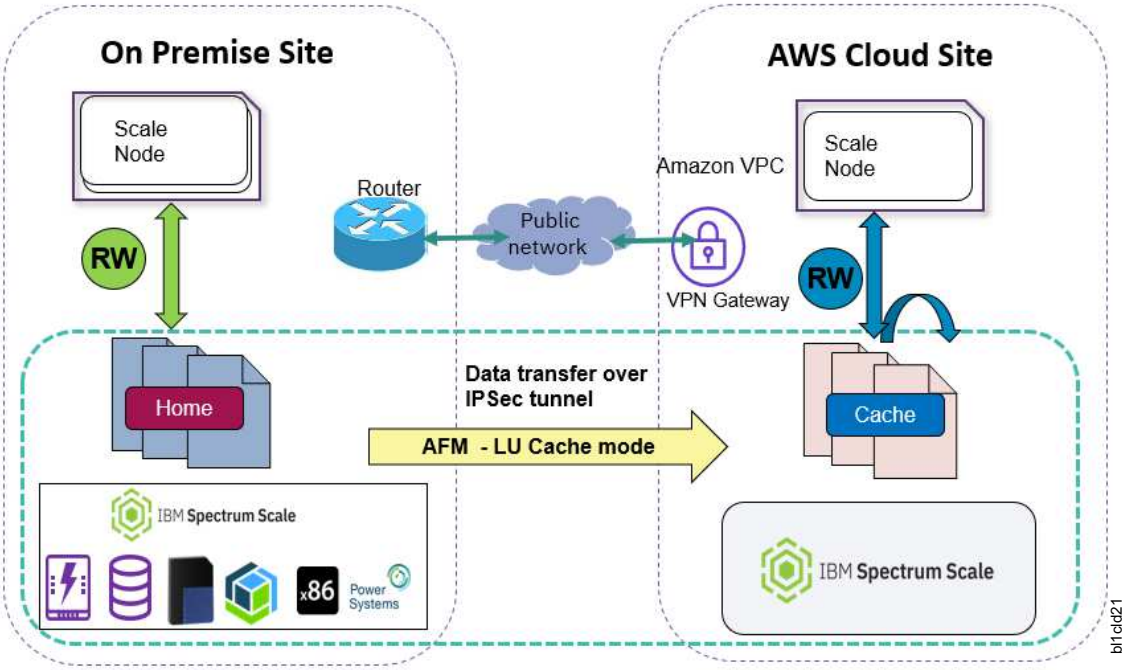

Figure 15. Local update cache mode

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.
- Once 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 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. Date movement from AWS cluster to the on-premise cluster incurs higher charges as compared to moving data the other way round.

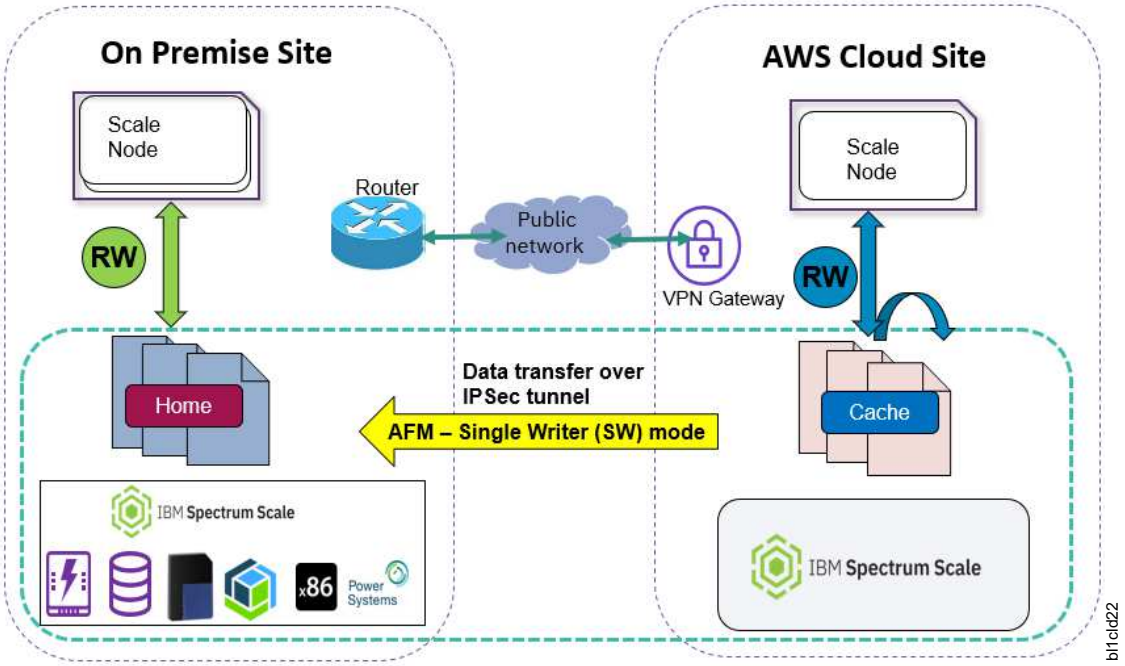

Figure 16. Single writer cache mode

Independent Writer (IW) cache mode

The IW cache mode has the following characteristics:

- Data is read and written on both the home cluster as well as the cache cluster.
- Changes 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. Date movement from AWS cluster to the on-premise cluster incurs higher charges as compared to moving data the other way round.

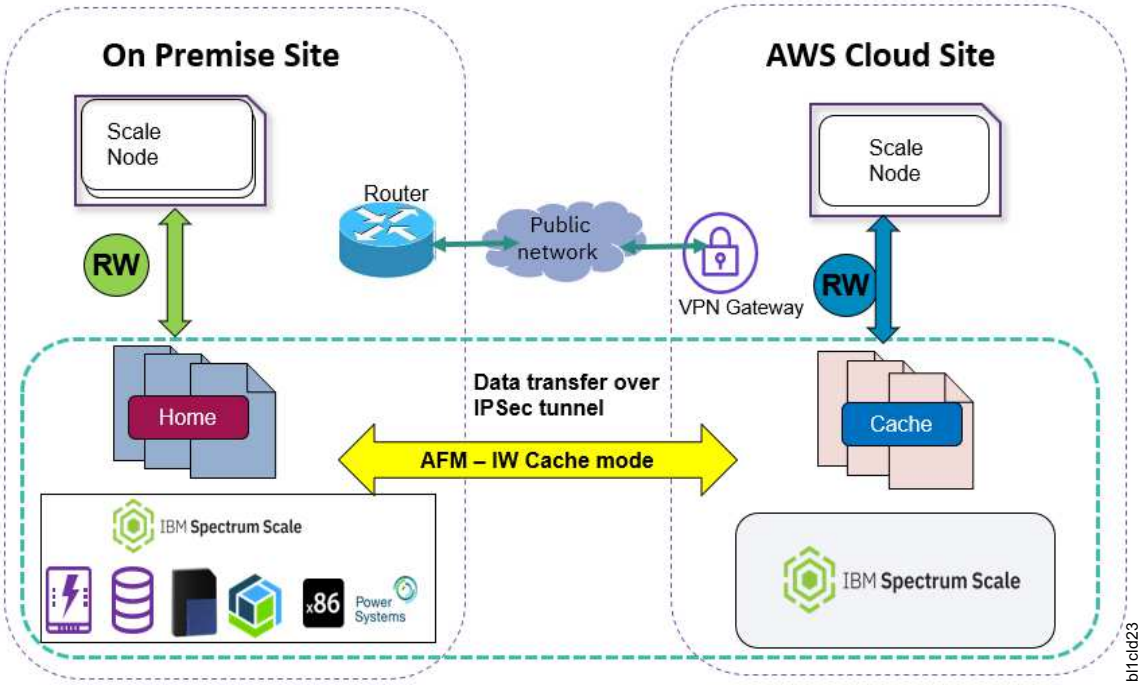

Figure 17. Independent writer cache mode

Deploying AFM on AWS

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

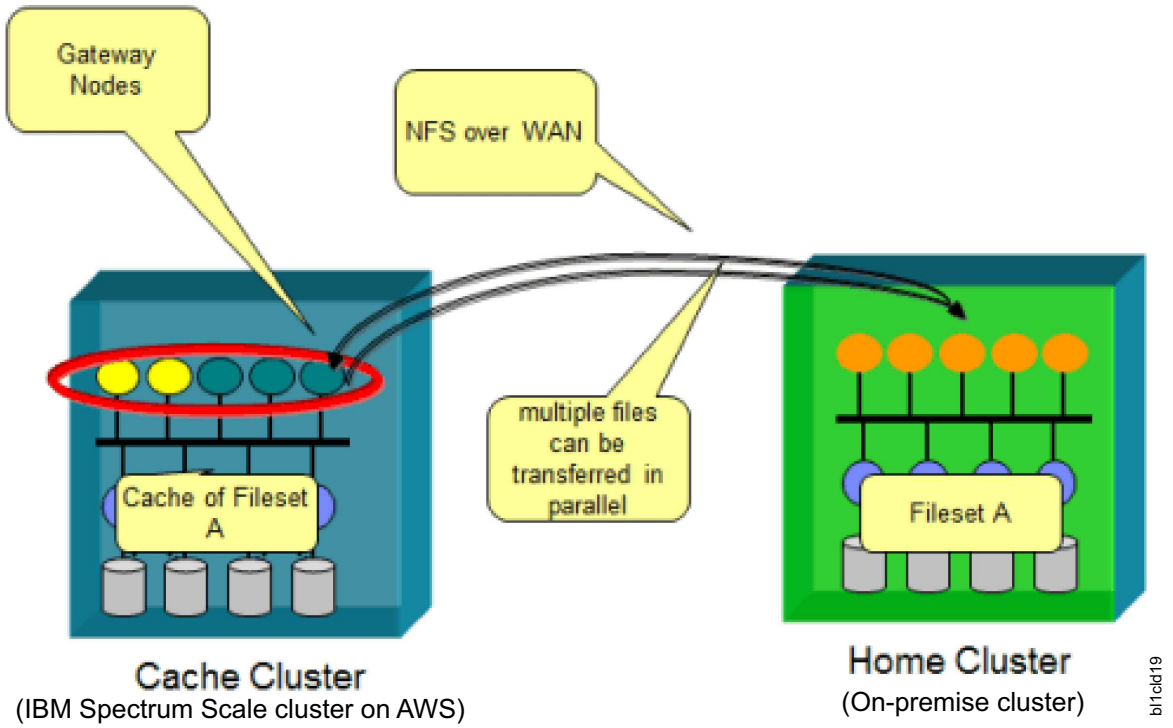

Figure 18. Example of an on-premise to AWS hybrid cloud networking architecture with AFM

Follow these steps to deploy AFM:

1. Set up the on-premise cluster.

   The on-premise IBM Spectrum Scale cluster acts as the AFM home cluster. An AFM home cluster is defined as a IBM Spectrum Scale cluster that has the ability to make the NFS v3 exports available to other IBM Spectrum Scale clusters. Follow these steps to set up the on-premise 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 an output similar to the following:

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

      For more information on the **mmnfs** command, 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 on the **mmafmconfig** command, 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 has been installed on the nodes that are going to server as the AFM gateway:

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

      If the NFS client package has not been 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 43 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 an output similar to the following:

      ```
      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 on the **mmchconfig** command, 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-premise or home cluster.

      Note: AFM cache mode should be specified based on the needs of the customer use case. For more information on cache modes, see "AFM cache modes" on page 37.

      ```
      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 an output similar to the following:

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

      For more information on the **mmcrfileset** command, 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 an output similar to the following:

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

For more information on the **mmlinkfileset** command, 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 on AFM prefect, see the *Prefetch* section in the *IBM Spectrum Scale: Concepts, Planning, and Installation Guide* in the IBM Spectrum Scale Knowledge Center.

For more information on setting up an AFM cache cluster, see the *Setting up the cache cluster* and *Creating an AFM relationship by using the NFS protocol* sections in the *IBM Spectrum Scale: Administration Guide* 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 9. Troubleshooting

The following issues have been encountered in AWS cloud:

CREATE_FAILED error with timeout message is encountered on launching AMI

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

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

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

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

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

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

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

Size limitation error on deploying the AWS CloudFormation templates

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

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

Stack creation failure message encountered

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

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

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

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

Chapter 10. Frequently Asked Questions

This page details some frequently asked questions.

Functional Support Matrices

Table 27. IBM Spectrum Scale BYOL AWS Marketplace Functional Support Matrix

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

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

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

How to resolve issues encountered while running IBM Spectrum Scale?

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

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

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

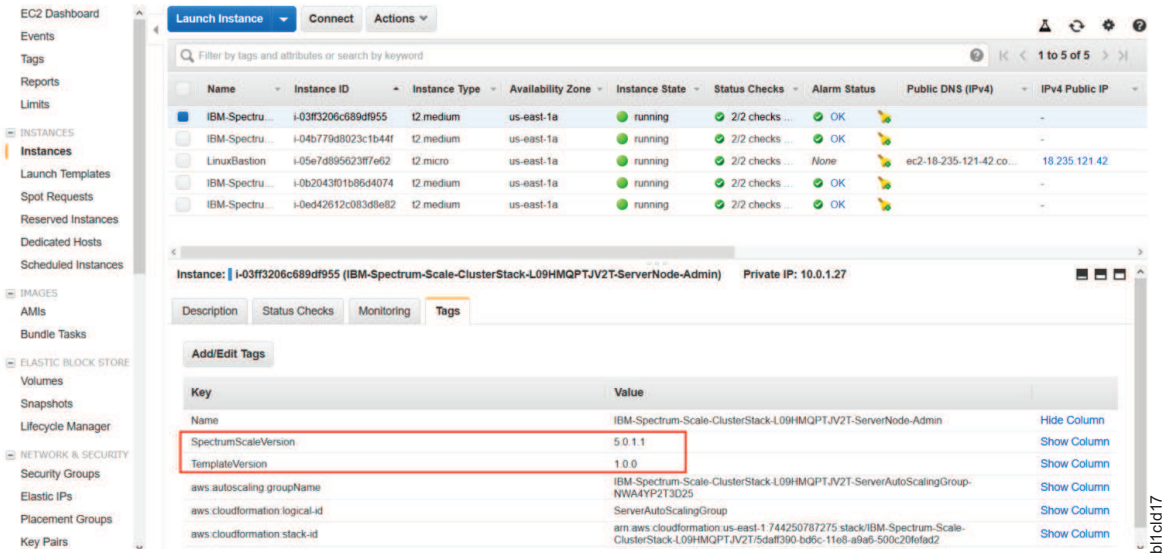

Figure 19. IBM Spectrum Scale stack outputs

Accessibility features for IBM Spectrum Scale

Accessibility features help users who have a disability, such as restricted mobility or limited vision, to use information technology products successfully.

Accessibility features

The following list includes the major accessibility features in IBM Spectrum Scale:

- Keyboard-only operation
- Interfaces that are commonly used by screen readers
- Keys that are discernible by touch but do not activate just by touching them
- Industry-standard devices for ports and connectors
- The attachment of alternative input and output devices

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

Keyboard navigation

This product uses standard Microsoft Windows navigation keys.

IBM and accessibility

See the IBM Human Ability and Accessibility Center (www.ibm.com/able) for more information about the commitment that IBM has to accessibility.

Notices

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

Information concerning non-IBM products was obtained from the suppliers of those products, their published announcements or other publicly available sources. IBM has not tested those products and cannot confirm the accuracy of performance, compatibility or any other claims related to non-IBM products. Questions on the capabilities of non-IBM products should be addressed to the suppliers of those products.

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

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

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

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

COPYRIGHT LICENSE:

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

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

© (your company name) (year).
Portions of this code are derived from IBM Corp.
Sample Programs. © Copyright IBM Corp. _enter the year or years_.

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

Trademarks

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

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

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

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

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

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

Terms and conditions for product documentation

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

Applicability

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

Personal use

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

Commercial use

You may reproduce, distribute and display these publications solely within your enterprise provided that all proprietary notices are preserved. You may not make derivative works of these publications, or reproduce, distribute or display these publications or any portion thereof outside your enterprise, without the express consent of IBM.

Rights

Except as expressly granted in this permission, no other permissions, licenses or rights are granted, either express or implied, to the publications or any information, data, software or other intellectual property contained therein.

IBM reserves the right to withdraw the permissions granted herein whenever, in its discretion, the use of the publications is detrimental to its interest or, as determined by IBM, the above instructions are not being properly followed.

You may not download, export or re-export this information except in full compliance with all applicable laws and regulations, including all United States export laws and regulations.

IBM MAKES NO GUARANTEE ABOUT THE CONTENT OF THESE PUBLICATIONS. THE PUBLICATIONS ARE PROVIDED "AS-IS" AND WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESSED OR IMPLIED, INCLUDING BUT NOT LIMITED TO IMPLIED WARRANTIES OF MERCHANTABILITY, NON-INFRINGEMENT, AND FITNESS FOR A PARTICULAR PURPOSE.

IBM Online Privacy Statement

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

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

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

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

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

Glossary

This glossary provides terms and definitions for IBM Spectrum Scale.

The following cross-references are used in this glossary:

- *See* refers you from a nonpreferred term to the preferred term or from an abbreviation to the spelled-out form.
- *See also* refers you to a related or contrasting term.

For other terms and definitions, see the IBM Terminology website (www.ibm.com/software/globalization/terminology) (opens in new window).

B

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

C

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

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

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

cluster manager
The node that monitors node status using disk leases, detects failures, drives recovery, and selects file system managers. The cluster manager must be a quorum node. The selection of the cluster manager node favors the quorum-manager node with the lowest node number among the nodes that are operating at that particular time.

Note: The cluster manager role is not moved to another node when a node with a lower node number becomes active.

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

D

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

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

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

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

disk leasing
A method for controlling access to storage devices from multiple host systems. Any host that wants to access a storage device configured to use disk leasing registers for a lease; in the event of a perceived failure, a host system can deny access,

preventing I/O operations with the storage device until the preempted system has reregistered.

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

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

E

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

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

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

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

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

F

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

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

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

FEK See *file encryption key*.

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

fileset snapshot
A snapshot of an independent fileset plus all dependent filesets.

file clone
A writable snapshot of an individual file.

file encryption key (FEK)
A key used to encrypt sectors of an individual file. See also *encryption key*.

file-management policy
A set of rules defined in a policy file that GPFS uses to manage file migration and file deletion. See also *policy*.

file-placement policy
A set of rules defined in a policy file that GPFS uses to manage the initial placement of a newly created file. See also *policy*.

file system descriptor
A data structure containing key information about a file system. This information includes the disks assigned to the file system (*stripe group*), the current state of the file system, and pointers to key files such as quota files and log files.

file system descriptor quorum
The number of disks needed in order to write the file system descriptor correctly.

file system manager
The provider of services for all the nodes using a single file system. A file system manager processes changes to the state or description of the file system, controls the regions of disks that are allocated to each node, and controls token management and quota management.

fixed-block architecture disk device (FBA disk device)
A disk device that stores data in blocks of fixed size. These blocks are addressed by block number relative to the beginning of the file. See also *extended count key data device*.

fragment
The space allocated for an amount of data too small to require a full block. A fragment consists of one or more subblocks.

G

global snapshot
A snapshot of an entire GPFS file system.

GPFS cluster
A cluster of nodes defined as being available for use by GPFS file systems.

GPFS portability layer
The interface module that each installation must build for its specific hardware platform and Linux distribution.

GPFS recovery log
A file that contains a record of metadata activity, and exists for each node of a cluster. In the event of a node failure, the recovery log for the failed node is replayed, restoring the file system to a consistent state and allowing other nodes to continue working.

I

ill-placed file
A file assigned to one storage pool, but having some or all of its data in a different storage pool.

ill-replicated file
A file with contents that are not correctly replicated according to the desired setting for that file. This situation occurs in the interval between a change in the file's replication settings or suspending one of its disks, and the restripe of the file.

independent fileset
A fileset that has its own inode space.

indirect block
A block containing pointers to other blocks.

inode The internal structure that describes the individual files in the file system. There is one inode for each file.

inode space
A collection of inode number ranges reserved for an independent fileset, which enables more efficient per-fileset functions.

ISKLM
IBM Security Key Lifecycle Manager. For GPFS encryption, the ISKLM is used as an RKM server to store MEKs.

J

journaled file system (JFS)
A technology designed for high-throughput server environments, which are important for running intranet and other high-performance e-business file servers.

junction
A special directory entry that connects a name in a directory of one fileset to the root directory of another fileset.

K

kernel The part of an operating system that contains programs for such tasks as input/output, management and control of hardware, and the scheduling of user tasks.

M

master encryption key (MEK)
A key used to encrypt other keys. See also *encryption key*.

MEK See *master encryption key*.

metadata
Data structures that contain information that is needed to access file data. Metadata includes inodes, indirect blocks, and directories. Metadata is not accessible to user applications.

metanode
The one node per open file that is responsible for maintaining file metadata integrity. In most cases, the node that has had the file open for the longest period of continuous time is the metanode.

mirroring
The process of writing the same data to multiple disks at the same time. The

mirroring of data protects it against data loss within the database or within the recovery log.

Microsoft Management Console (MMC)
A Windows tool that can be used to do basic configuration tasks on an SMB server. These tasks include administrative tasks such as listing or closing the connected users and open files, and creating and manipulating SMB shares.

multi-tailed
A disk connected to multiple nodes.

N

namespace
Space reserved by a file system to contain the names of its objects.

Network File System (NFS)
A protocol, developed by Sun Microsystems, Incorporated, that allows any host in a network to gain access to another host or netgroup and their file directories.

Network Shared Disk (NSD)
A component for cluster-wide disk naming and access.

NSD volume ID
A unique 16 digit hex number that is used to identify and access all NSDs.

node An individual operating-system image within a cluster. Depending on the way in which the computer system is partitioned, it may contain one or more nodes.

node descriptor
A definition that indicates how GPFS uses a node. Possible functions include: manager node, client node, quorum node, and nonquorum node.

node number
A number that is generated and maintained by GPFS as the cluster is created, and as nodes are added to or deleted from the cluster.

node quorum
The minimum number of nodes that must be running in order for the daemon to start.

node quorum with tiebreaker disks
A form of quorum that allows GPFS to run with as little as one quorum node available, as long as there is access to a majority of the quorum disks.

non-quorum node
A node in a cluster that is not counted for the purposes of quorum determination.

P

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

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

pool A group of resources with similar characteristics and attributes.

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

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

private IP address
A IP address used to communicate on a private network.

public IP address
A IP address used to communicate on a public network.

Q

quorum node
A node in the cluster that is counted to determine whether a quorum exists.

quota The amount of disk space and number of inodes assigned as upper limits for a specified user, group of users, or fileset.

quota management
The allocation of disk blocks to the other nodes writing to the file system, and comparison of the allocated space to quota limits at regular intervals.

R

Redundant Array of Independent Disks (RAID)
A collection of two or more disk physical drives that present to the host an image of one or more logical disk drives. In the event of a single physical device failure, the data can be read or regenerated from the other disk drives in the array due to data redundancy.

recovery
The process of restoring access to file system data when a failure has occurred. Recovery can involve reconstructing data or providing alternative routing through a different server.

remote key management server (RKM server)
A server that is used to store master encryption keys.

replication
The process of maintaining a defined set of data in more than one location. Replication involves copying designated changes for one location (a source) to another (a target), and synchronizing the data in both locations.

RKM server
See *remote key management server*.

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

S

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

Scale Out Backup and Restore (SOBAR)
A specialized mechanism for data protection against disaster only for GPFS file systems that are managed by IBM Spectrum Protect Hierarchical Storage Management (HSM).

secondary GPFS cluster configuration server
In a GPFS cluster, the node chosen to maintain the GPFS cluster configuration data in the event that the primary GPFS cluster configuration server fails or becomes unavailable.

Secure Hash Algorithm digest (SHA digest)
A character string used to identify a GPFS security key.

session failure
The loss of all resources of a data management session due to the failure of the daemon on the session node.

session node
The node on which a data management session was created.

Small Computer System Interface (SCSI)
An ANSI-standard electronic interface that allows personal computers to communicate with peripheral hardware, such as disk drives, tape drives, CD-ROM drives, printers, and scanners faster and more flexibly than previous interfaces.

snapshot
An exact copy of changed data in the active files and directories of a file system or fileset at a single point in time. See also *fileset snapshot*, *global snapshot*.

source node
The node on which a data management event is generated.

stand-alone client
The node in a one-node cluster.

storage area network (SAN)
A dedicated storage network tailored to a specific environment, combining servers, storage products, networking products, software, and services.

storage pool
A grouping of storage space consisting of volumes, logical unit numbers (LUNs), or addresses that share a common set of administrative characteristics.

stripe group
The set of disks comprising the storage assigned to a file system.

striping
A storage process in which information is split into blocks (a fixed amount of data) and the blocks are written to (or read from) a series of disks in parallel.

subblock
The smallest unit of data accessible in an I/O operation, equal to one thirty-second of a data block.

system storage pool
A storage pool containing file system control structures, reserved files, directories, symbolic links, special devices, as well as the metadata associated with regular files, including indirect blocks and extended attributes The **system storage pool** can also contain user data.

T

token management
A system for controlling file access in which each application performing a read or write operation is granted some form of access to a specific block of file data. Token management provides data consistency and controls conflicts. Token management has two components: the token management server, and the token management function.

token management function
A component of token management that requests tokens from the token management server. The token management function is located on each cluster node.

token management server
A component of token management that controls tokens relating to the operation of the file system. The token management server is located at the file system manager node.

transparent cloud tiering (TCT)
A separately installable add-on feature of IBM Spectrum Scale that provides a native cloud storage tier. It allows data center administrators to free up on-premise storage capacity, by moving out cooler data to the cloud storage, thereby reducing capital and operational expenditures. .

twin-tailed
A disk connected to two nodes.

U

user storage pool
A storage pool containing the blocks of data that make up user files.

V

VFS See *virtual file system*.

virtual file system (VFS)
A remote file system that has been mounted so that it is accessible to the local user.

virtual node (vnode)
The structure that contains information about a file system object in a virtual file system (VFS).

Index

L

M

N

O

R

S

T

U

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

Printed in USA

SC27-9283-02