IBM® Storage

IBM Hybrid Solution for Scalable Data Solutions using IBM Spectrum Scale

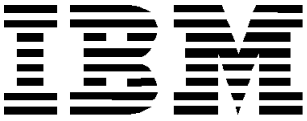

© Copyright International Business Machines Corporation 2019.

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

Contents

About this document 1
Executive summary 1
Scope 2
Prerequisites 2
Getting started: Scalable hybrid cloud solution for data agility and collaboration 3
IBM Spectrum Scale at on-premise 3
IBM Spectrum Scale on AWS 3
IBM Spectrum Scale deployment process on AWS: 3
IBM Spectrum Scale in IBM Cloud 4
IBM Spectrum Scale Active File Management (AFM) 6
Networking configuration 6
Site-to-site VPN connection between on-premise and AWS Cloud 6
Site-to-site VPN connection between on-premise and IBM Cloud 7
IBM Active File Management (AFM) configuration and use cases 8
Case study: Implementing enterprise data pipeline in hybrid cloud environments 12
Summary 13

Notices 15
Trademarks 16
Terms and conditions for product documentation 17
Applicability 17
Commercial use 17
Rights 17
Privacy policy considerations 17

© Copyright IBM Corp. 2019. All rights reserved.

About this document

This document is intended to facilitate the deployment of the scalable hybrid cloud solution for data agility and collaboration using IBM® Spectrum Scale across multiple public clouds. To complete the tasks it describes, you must understand IBM Spectrum Scale and IBM Spectrum Scale Active File Management (AFM).

The information in this document is distributed on an *as is* basis without any warranty that is either expressed or implied. Support assistance for the use of this material is limited to situations where IBM Spectrum Scale or IBM Spectrum Scale Active File Management are supported and entitled, and where the issues are specific to a blueprint implementation.

Executive summary

In today's environment, many organizations are using some form of cloud services. These cloud services can be private, public, or hybrid, and storage infrastructure is an integral part of these deployments.

Modern data agility and collaboration solutions can involve leveraging cloud-based resources. A hybrid multicloud solution is designed to run the correct tasks in the correct place across various service providers in an optimal way. However, infrastructure differences between on-premises, private, or public cloud offerings for these services can complicate the data movement, specifically in regards to a secure and cost-effective manner across various sites. The lack of metadata management and methods to intelligently move data is another problem in hybrid cloud environments.

IBM Spectrum Scale is a high-performance, highly available, clustered file system available on a variety of platforms (including the public cloud service providers). IBM Active File Management (AFM) is a scalable, high-performance, intelligent file system caching layer integrated into the IBM Spectrum Scale file system. This enables you to implement a single global name space across various sites including the public cloud offerings.

With IBM Spectrum Scale available as software-defined storage on AWS Cloud and IBM Cloud™, organizations can deploy a hybrid cloud environment. This creates an environment where data can be moved and cached between on-premise to Public Cloud offerings (AWS Cloud or IBM Cloud).

© Copyright IBM Corp. 2019. All rights reserved.

Scope

This blueprint guide provides the following information:

- A solutions architecture and related solution configuration workflows, with the following essential software components:
  - IBM Spectrum™ Scale
  - IBM Spectrum Active File Management
- Detailed technical configuration steps for building an end-to-end solution in the hybrid cloud environment

This technical report does not make the following changes:

- Provide performance analysis from a user perspective
- Replace any official manuals and documents issued by IBM

Prerequisites

This technical paper assumes basic knowledge of the following prerequisites:

- IBM Spectrum Scale™
- IBM Spectrum Scale Active File Management
- AWS Cloud
- IBM Cloud
- IP networking

Getting started: Scalable hybrid cloud solution for data agility and collaboration

This section describes the components and solution building blocks used in the hybrid cloud environment.

IBM Spectrum Scale at on-premise

IBM Spectrum Scale is a high-performance, highly available, clustered file system available on a variety of platforms. IBM Spectrum Scale is a cluster file system that provides concurrent access to a single file system or set of file systems from multiple nodes. On-premise, IBM Spectrum scale nodes can be SAN-attached, network-attached, a mixture of SAN-attached, and network-attached, or in a Shared Nothing cluster configuration.

IBM Spectrum Scale on AWS

IBM Spectrum Scale on AWS is a software-defined storage offering available through the AWS Marketplace and deployed using the Elastic Cloud Compute (Amazon EC2) instances and Amazon Elastic Block Store (Amazon EBS) volumes.

IBM Spectrum Scale deployment process on AWS:

Figure 1 shows the deployment process on AWS for IBM Spectrum Scale (such as Bring Your Own License (BYOL) and Automated deployment through the AWS).

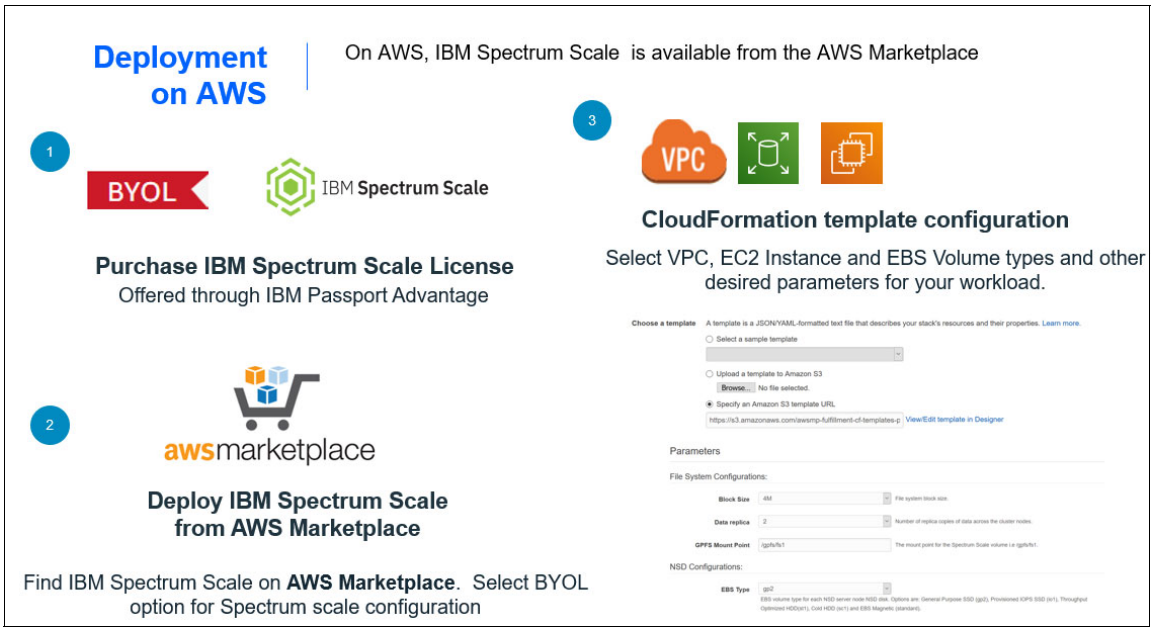

Figure 1 IBM Spectrum Scale ordering and deployment on AWS

The deployment process includes the following options:

- Bring Your Own License (BYOL): IBM Spectrum Scale follows the BYOL model for deployment in the AWS Cloud. Customers can purchase the IBM Spectrum Scale license using IBM Passport Advantage® or from business partners.

- Automated Deployment using the AWS Marketplace: IBM Spectrum Scale is available in the AWS Marketplace, and deployment is automated using the Cloud Formation template.

For detailed deployment instructions, use the following guide on AWS Marketplace:

https://aws.amazon.com/marketplace/pp/B07DRLMG2W

IBM Spectrum Scale in IBM Cloud

IBM Spectrum Scale in IBM Cloud is a software-defined storage offering and it can be built using the standard Bare Metal servers and Block Storage available in the cloud.

Figure 2 shows the IBM Spectrum Scale deployment process in IBM Cloud (using BYOL or renting Bare Metal servers).

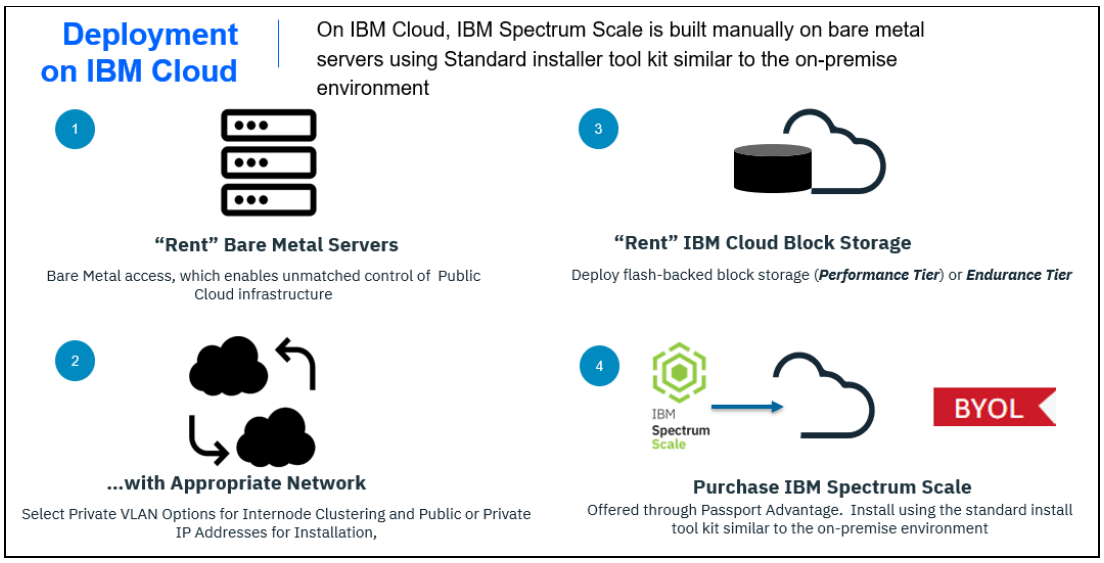

Figure 2 IBM Spectrum Scale ordering and deployment on IBM cloud

The deployment process includes the following options:

- Bring Your Own License (BYOL): IBM Spectrum Scale follows the BYOL model for deployment in the IBM Cloud. Customers can purchase the IBM Spectrum Scale license through Passport Advantage or from business partners.
- Rent the Bare Metal Servers using IBM Cloud console: Select the appropriate hardware model and OS version as per the IBM Spectrum Scale requirements. See the following link for more information about validating the IBM Spectrum Scale requirements:

  https://www.ibm.com/support/knowledgecenter/STXKQY/gpfsclustersfaq.html#linux

These systems are used as the NSD nodes, and the minimum requirement is for three servers, noted in the following list:

- Networking configuration: Select the Portable private VLAN network for the IBM Spectrum Scale networking purpose
- Storage Configuration: IBM Spectrum Scale NSD servers can be configured in two ways:
  - External SAN Storage model that uses iSCSI Block Storage from the IBM Cloud. IBM Cloud provides two types of disks:
    - Performance
    - Endurance

Flash drives are recommended for performance work loads. See Figure 3.

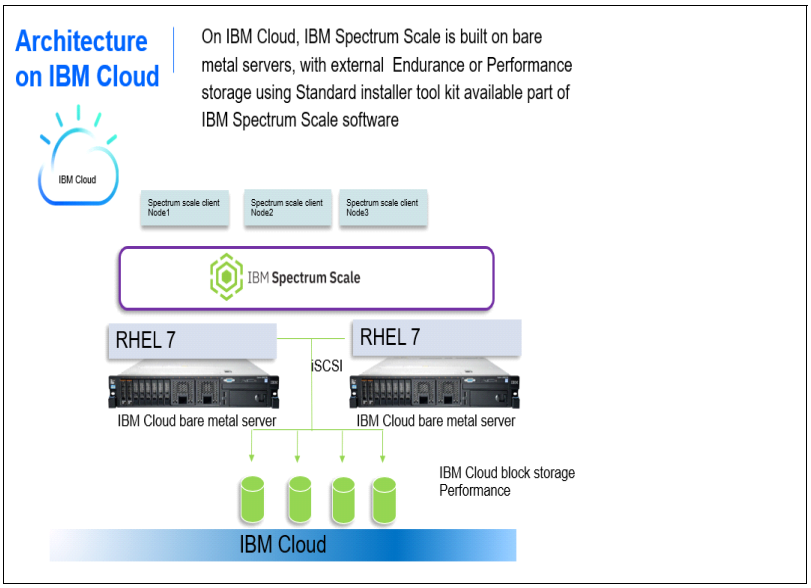

Figure 3 IBM Spectrum Architecture on IBM Cloud using external SAN attached drives

- Internal drives of the Bare Metal servers: In this configuration, the Shared Nothing model is used and IBM Spectrum Scale file-system level replication is used for high availability. See Figure 4.

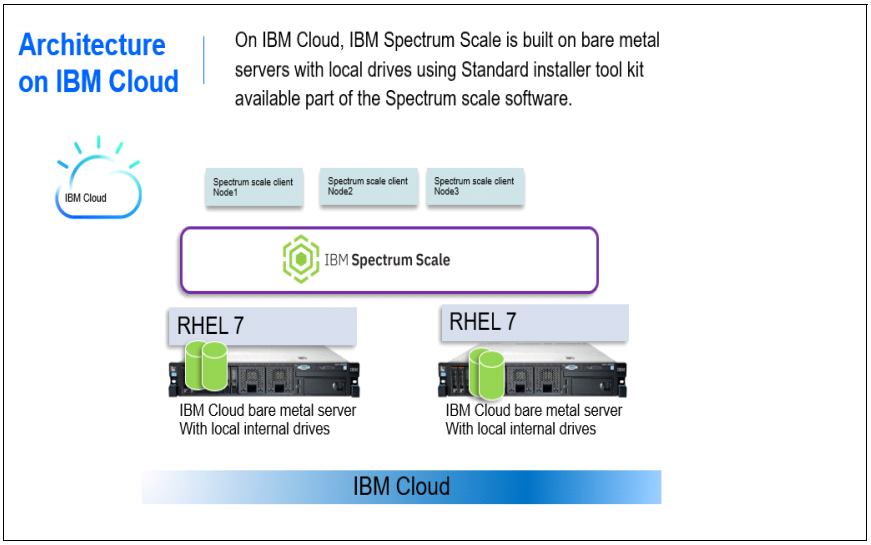

Figure 4 IBM Spectrum architecture on IBM cloud using internal drives

- IBM Spectrum Scale deployment: This model is user-deployed, and customers build their own cluster. The customer is responsible for building the cluster using the standard tool kit available as part of the IBM Spectrum Scale software (such as for the on-premise environment).

  For detailed deployment instructions and guidance, use the following link to download the IBM Redbooks® publication *IBM Spectrum Scale Deployment on IBM Softlayer:*

  `http://www.redbooks.ibm.com/redpapers/pdfs/redp5410.pdf`

IBM Spectrum Scale Active File Management (AFM)

IBM Active File Management (AFM) is a scalable, high-performance, intelligent file system caching layer integrated into the IBM Spectrum Scale file system. This enables implementation of a single global name space across various sites, including the public cloud offerings. Consider the following additional opportunities:

- Enables data mobility and sharing of data across various clusters
- An asynchronous data, cross-cluster caching utility
- On-premise is configured as *home* and acts as a primary storage
- Public cloud end points (AWS & IBM Cloud) are Cache only
- NFS protocol is used for communication between home (on-premise) and Cache sites (AWS and IBM Cloud)

Networking configuration

Hybrid cloud is an environment that combines both private and public cloud resources by allowing data and applications to be shared. This requires network configuration and establishing connectivity between on-premise private cloud and public cloud resources. The most commonly used solution is the public internet for data transfer.

However, this has some privacy and security concerns about data flowing over the public internet. Users can make a secure connection using an IPSec-based Virtual Private Network (VPN) between private and public cloud networks. This encrypts the data and provides a point-to-point tunnel between private and public networking devices. The performance is dependent on the network link speed of the connection.

Site-to-site VPN connection between on-premise and AWS Cloud

The following diagram illustrates the site-to-site VPN connectivity between the on-premise private cloud and the AWS Cloud. An IBM LAB environment was used to test this scenario for the IBM Proof of Concept validation for this solution blueprint (shown in Figure 5).

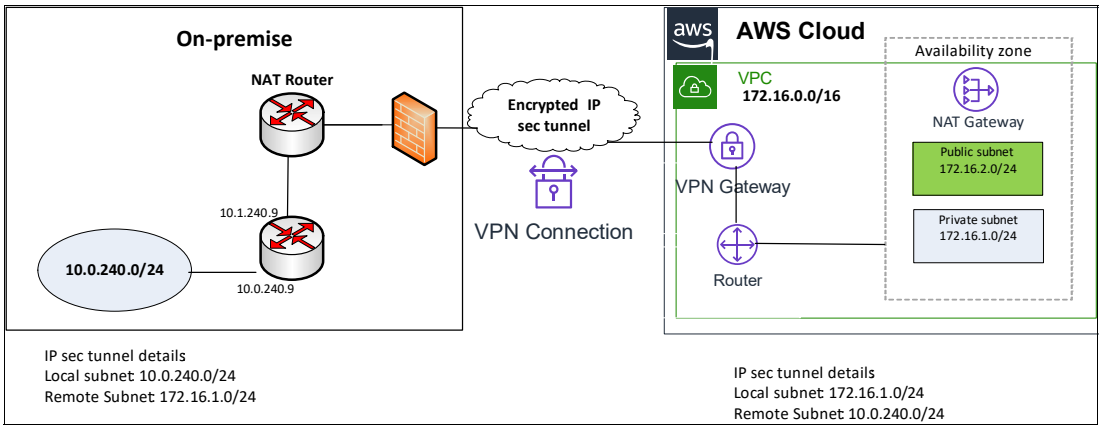

Figure 5 On-premise to AWS Cloud networking architecture

The on-premise network components are noted in the following list:

- Router: A VyOS-based software defined networking router is used for configuring the IPsec tunnel
- NAT router: Is used for NATing the public IP address to the private subnet of the storage and compute devices

- Networking subnet: `10.0.240.0/24` is used for configuring the compute and storage devices at the on-premise site

The AWS Cloud network components are noted in the following list:

- VPC: Enables you to launch the resources into a virtual network
- VPN Gateway: A *virtual private gateway* is the VPN concentrator on the AWS side of the Site-to-Site VPN connection. You 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 subnets.
- Private subnet: Used for configuring the compute and storage devices in AWS.

See the following solution blueprint for configuration steps and commands:

`http://www.redbooks.ibm.com/Redbooks.nsf/RedbookAbstracts/redp5542.html?Open`

The Site-to-Site configurations for AWS and On-premise are noted in the following lists:

- AWS configuration:
  - Create Customer gateway: This step allows you to define the customer on-premise end configuration at the AWS
  - Create Virtual private gateway: This step allows you to create the virtual private gateway and attach it to the VPC
  - Configure site-to-site VPN configuration: Select your VPC and customer gateway created earlier for configuring the VPN
- On-premise configuration: For LAB validation purposes, a VyOS based software defined router for VPN configurations is used in this example.
  - Configure the Interfaces: Define the private and public network interface and TCP IP networking parameters
  - Configure routing: Define the default gateway for private networks and IPsec tunnels
  - Enable NAT traversing: Allows establishing the IPSec tunnel using the NAT IP address
  - Allow selected remote networks: Define the remote subnets in the public cloud which can be allowed through the IPSec tunnel
  - Configure IPsec Parameters: Define the IKE and other required configuration parameters for authentication
  - Configure the tunnels: Define the local and remote subnets of the tunnel for networking communication

See the following solution blueprint for configuration steps and commands:

`http://www.redbooks.ibm.com/Redbooks.nsf/RedbookAbstracts/redp5542.html?Open`

Note: For a larger enterprise customer, the on-premises network configuration, infrastructure, and security configurations can vary per organization policies. The choice of connectivity depends on the customers' environments and their preferences.

Site-to-site VPN connection between on-premise and IBM Cloud

The following diagram illustrates the site-to-site VPN connectivity between the on-premise private cloud and IBM Cloud used for the solution validation in the LAB. See Figure 6 on page 8.

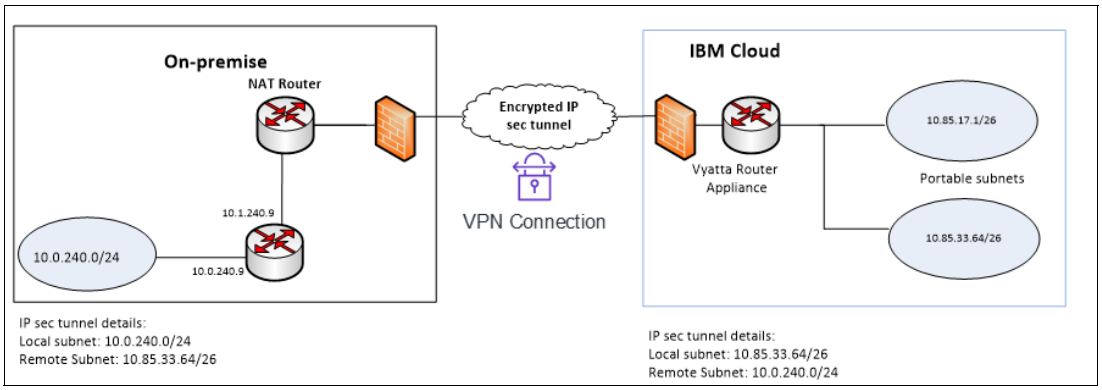

Figure 6 On-premise to IBM Cloud cloud networking architecture

The IBM Cloud components are noted in the following list:

- Vyatta Gateway Appliance: Required for managing the networks and allowing the routing of multiple VLAN's and VPN traffic
- Portable subnets: Used for network communication across compute and storage devices. Configure this portable subnet as routable for communication with other subnets and remote traffic across the VPN tunnel
  - Vyatta Gateway Router configuration information:
    - Configure the interfaces: Define the private and public network interface and TCP IP networking parameters
    - Configure routing: Define the default gateway for private networks and IPsec tunnels
    - Enable NAT traversing: Necessary to establish the IPSec tunnel using the NAT IP address
    - Allow selected remote networks: Define the remote subnets in the public cloud that are allowed through the IPSec tunnel
    - Configure IPsec parameters: Define the IKE and other required configuration parameters for authentication
    - Configure the tunnels: Define the local and remote subnets of the tunnel for networking communication

See the following solution blueprint for configuration steps and commands:

http://www.redbooks.ibm.com/Redbooks.nsf/RedbookAbstracts/redp5542.html?Open

IBM Active File Management (AFM) configuration and use cases

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

The following steps are the high-level tasks for configuring Active File Management:

1. Home Site (On-premise) configuration:
   a. Configure the VPN IPSec tunnel as described in the earlier section.
   b. Define the fileset (NFS exports) and configure access to the gateway node(s) that are configured in the Public Cloud. See Figure 7 on page 9.

```
[root@scalenode1 ~]# mmnfs export add /ibm/gpfs0/onprem-aws --client "172.16.1.237(Access_Type=RW)"
mmnfs: The NFS export was created successfully
[root@scalenode1 ~]#
```

Figure 7 Fileset and node access configuration

c. Enable AFM configuration on the home exports using **mmafmconfig**. See Figure 8.

```
[root@scalenode1 ~]# mmafmconfig enable /ibm/gpfs0/onprem-aws
[root@scalenode1 ~]#
```

Figure 8 Using the mmafmconfig command

2. Cache Site (AWS or IBM Cloud):

   a. Set up the VPN tunnel on the Public Cloud as described in “Site-to-site VPN connection between on-premise and IBM Cloud” on page 7.

   b. Install NFS client on the compute node(s). Specifically, the node(s) used as the gateway node(s) for the AFM configuration. See Figure 9.

```
root@ip-172-16-1-237 ~]# yum install nfs-utils
oaded plugins: amazon-id, rhui-lb, search-disabled-repos
```

Figure 9 Installing NFS client

   c. Define the node with NFS client (same node as configured in Step b as the gateway node) at the IBM Spectrum Scale cluster using **mmchnode**, as shown in Figure 10.

```
[root@ip-172-16-1-237 ~]# mmchnode --gateway -N ip-172-16-1-237.eu-central-1.compute.internal
Fri May 24 07:30:36 UTC 2019: mmchnode: Processing node ip-172-16-1-237.eu-central-1.compute.interna
l
mmchnode: Propagating the cluster configuration data to all
  affected nodes.  This is an asynchronous process.
[root@ip-172-16-1-237 ~]#
```

Figure 10 Using the mmchnode command

   d. On the Cache site, create a fileset with an AFM relationship to the filesets exported from the Home (on-premise) site. See Figure 11.

```
[root@ip-172-16-1-237 ~]# mmcrfileset fs1 onprem-aws-cache -p afmTarget=10.0.240.43:/ibm/gpfs0/onpre
m-aws -p afmmode=read-only --inode-space new --inode-limit 999999
Fileset onprem-aws-cache created with id 21 root inode 5767171.
[root@ip-172-16-1-237 ~]#
```

Figure 11 Creating filesets with AFM relationship

   e. Link the fileset to the Junction folder. See Figure 12.

```
[root@ip-172-16-1-237 ~]# mmlinkfileset fs1 onprem-aws-cache -J /gpfs/fs1/onprem-aws-cache
Fileset onprem-aws-cache linked at /gpfs/fs1/onprem-aws-cache
[root@ip-172-16-1-237 ~]#
```

Figure 12 Linking fileset and folder

Authentication and authorization

In the hybrid cloud AFM environment, a common authentication and authorization infrastructure is required across the entire home and cache home filesets. AFM requires the same global namespace authorization to keep user ID/group ID (UID/GID) the same across all AFM filesets. AFM caches all file data, extended attributes and access control lists (ACLs) from the writing fileset (either home or cache) to all other associated AFM filesets.

AFM does not perform ID mapping between the home site and the cache site. This is because the remote AFM cache resides in a separate IBM Spectrum Scale cluster system configured in the Public cloud environments. You can configure either Active Directory (AD) or Lightweight Directory Access Protocol (LDAP) using an external server for the ID mapping information. The same set of authentication mechanism needs to be configured at both locations: Home (On-premise) and Cache Sites (AWS or IBM Cloud) for effective User ID mapping purposes.

AFM modes and use cases

This section describes use cases with various AFM cache modes configured in the public cloud environment:

- Read-only cache mode
- Local-update (LU)
- Independent-writer

Read-only cache mode

Read-only Cache mode functions as noted in the following list and shown in Figure 13:

- In AFM, Read Only (RO) mode files are created and reside in the on-premise Home.
- File stubs (inodes) are created in and metadata is Cached upon AFM configuration at the Public cloud locations: AWS and IBM Cloud.
- Data is ready only at the Cache Sites: AWS and IBM Cloud.
- Data is being copied from Home (on-premise) to Cache (cloud) upon file access or pre-fetch operation.

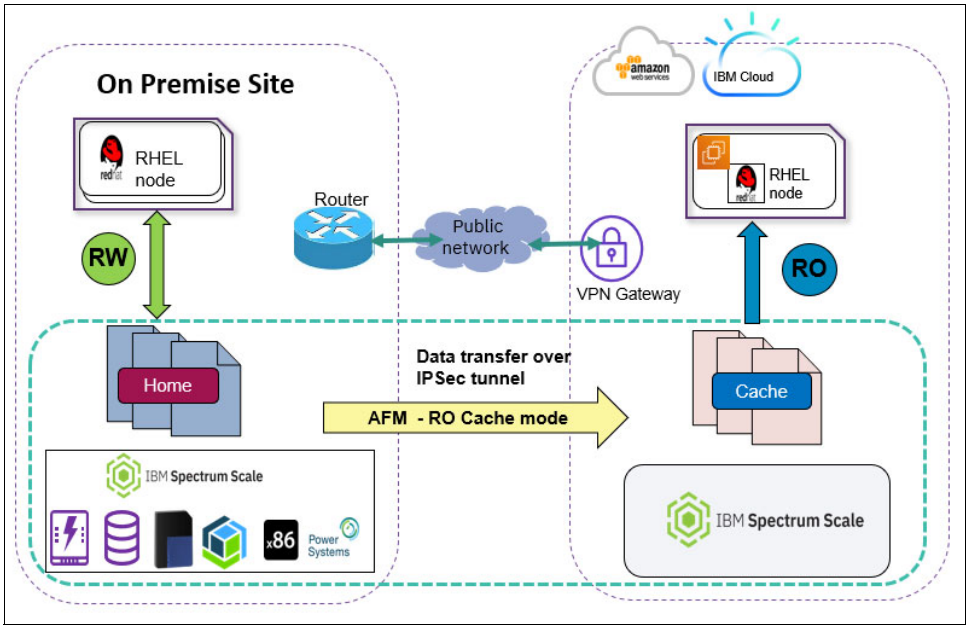

Figure 13 AFM Read-Only (RO) mode architecture

Use cases:

- Collaboration and data sharing: In this use case, AFM enables you to share the on-premise cloud data sets with other colleagues in read-only mode. This is useful for universities, research institutes, and corporate organizations when data sharing is required for collaboration and cloud-bursting purposes.
- Data Streaming from cloud end points: AFM Read-only mode enables streaming the selected data sets from public cloud end points configured in different regions.

Local-update (LU)

Local-update functions as noted in the following list and shown in Figure 14.

- Cached data (data from home) is read-only. The difference between cached data and read-only mode is that you can create and modify files in the cache fileset.
- Files that are created or modified in the cache sites configured in public cloud locations (AWS or IBM Cloud) are considered local updates.
- Local-updates at cache sites (AWS or IBM Cloud) are never pushed back to the home site configured at the on-premise location.

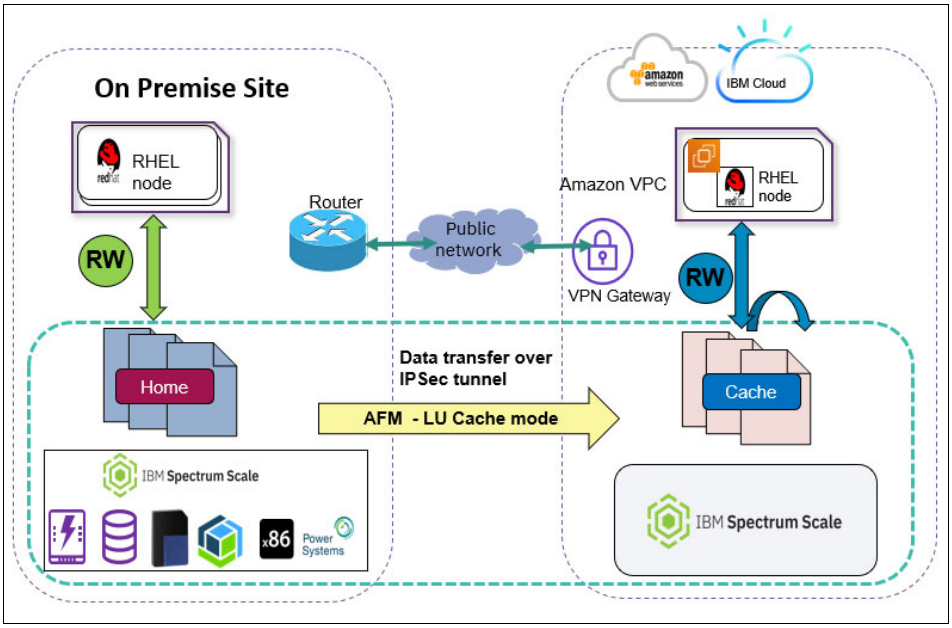

Figure 14 AFM Local Update (LU) mode architecture

Use cases:

- Data re-use: Local-update (LU) mode can be used for running the test and development environment in public cloud locations. Data can be cached from the on-premise home site and will be used for development and testing in the public cloud location. All the changes made in the public cloud locations are local and not pushed back to the home sites
- Data simulation and Analytics: LU mode allows you to run the various batch jobs for simulations and analytics purpose in the public clouds. AFM fetches only the required data blocks not the complete data set. It also enables you to perform the required simulations on the data sets fetched into the public cloud cache end points.

Public cloud providers (AWS and IBM Cloud) charge *egress* fees to move data out of the public clouds to the on-premise location. AFM Local-update mode allows you to make changes in the public cloud locations, and local updates that are made are never pushed back to the on-premise location. Therefore, no *egress* fee is incurred, because data is never pushed out of the public cloud locations.

Independent-writer

Independent-writer functions as noted in the following list and shown in Figure 15:

- Data is read and write at both home (on-premise) and cache sites (AWS or IBM Cloud).
- Changes at cache site are synchronized with home and vice versa.

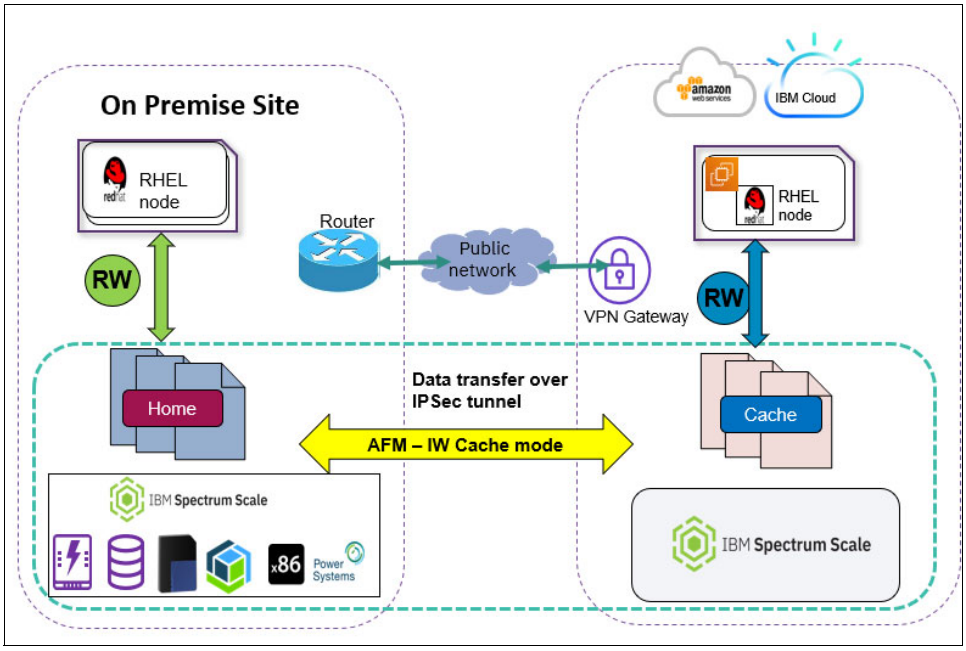

Figure 15 AFM Independent Writer (IW) mode architecture

Use cases:

- This mode is recommended for workloads where minimal changes are expected in the cache sites (AWS or IBM Cloud) and changed data can be transmitted back to the home site on-premise locations.
- Huge data changes in the public cloud results in moving larger data sets back home, and incur higher egress costs from the cloud providers. There are no 'egress cost guard rails' and the Architect, operator and dev ops team should be educated on the potential of incurring huge costs.

Case study: Implementing enterprise data pipeline in hybrid cloud environments

Enterprise Analytics and AI environments typically involve multiple AI and Analytics applications, but share a common asset in the form of the enterprise data. These environments are typically designed and implemented as workflows representing stages from data ingest to final insights. These implementations are generally referred to as *enterprise data pipelines*. Data flows through multiple stages of these data pipelines in which various applications touch it for analysis/processing.

Figure 16 is a generalized representation of an enterprise data pipeline.

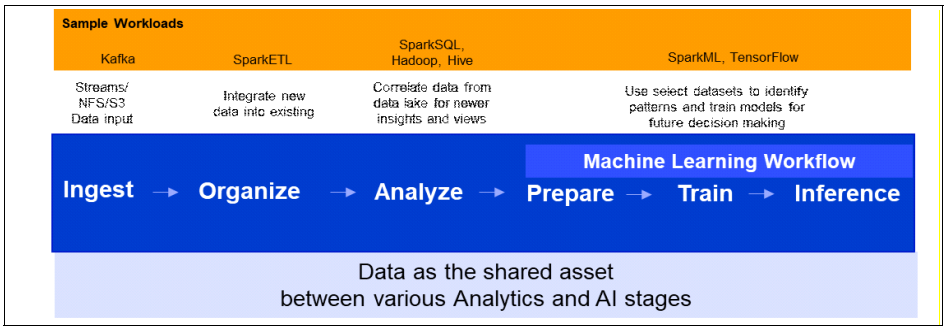

Figure 16 AI workflow model

Now with the wide adoption of hybrid cloud within enterprises, it is often a requirement to implement such an enterprise data pipeline in hybrid cloud setups, where some of the stages are implemented on-premise and some are implemented on public cloud.

For example, an enterprise has an enterprise data lake on-premise in which **Ingest** → **Organize** →**Analyze** stages are happening. And now there is a requirement to enhance their Hadoop Analytics capabilities by augmenting it with AI. Therefore, they require some GPU-enabled compute to run model training on selected data sets from the data lake. This machine learning workload can be run on GPU-enabled compute on AWS and results (trained model) can be shared back to the Hadoop data lake.

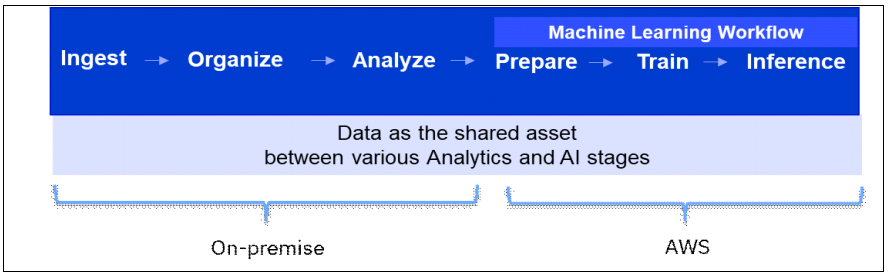

Figure 17 Hybrid cloud workflow

The hybrid cloud data solution using IBM Spectrum Scale is the perfect option to implement this type of hybrid cloud data pipeline implementation. Leveraging the AFM capabilities discussed in this paper. IBM Spectrum Scale can be used as a storage layer for implementing various stages in an enterprise data pipeline. It provides a common data management layer between various applications in the workflow.

AFM capabilities easily enables you to extend this implementation in hybrid cloud environments, as shown in Figure 17. This helps avoid creating multiple copies of the same data and helps reduce total cost of the solution variable operation costs for hybrid cloud setups, and instead enables on-demand caching of the required data set for machine learning workload.

Summary

IBM Spectrum Scale allows you to build an effective scalable data solution in a hybrid cloud environment. The IBM Active File Management (AFM) technology feature enables you to move the data in a secure and optimized way across various cloud locations.

This ability is required for different use cases, such as data sharing or collaboration, devops, data simulation, and analytics. AFM has the capability because it uses caching technology for data sharing, and fetches only the required data blocks as per the workload requirement in the public cloud end points.

This technology feature allows you to build scalable data solutions in a cost-effective manner on public clouds. Organizations using AFM caching technology save on cloud storage costs because the caching technology only requires limited storage.

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

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.

The performance data and client examples cited are presented for illustrative purposes only. Actual performance results may vary depending on specific configurations and operating conditions.

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.

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.

© Copyright IBM Corp. 2019. All rights reserved.

Trademarks

IBM, the IBM logo, and ibm.com are trademarks or registered trademarks of International Business Machines Corporation, 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 http://www.ibm.com/legal/copytrade.shtml

The following terms are trademarks or registered trademarks of International Business Machines Corporation, and might also be trademarks or registered trademarks in other countries.

Redbooks (logo) ®
IBM®
IBM Cloud™
IBM Spectrum™
IBM Spectrum Scale™
Passport Advantage®
Redbooks®

The following terms are trademarks of other companies:

Other company, product, or service names may be trademarks or service marks of others.

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.

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.

Privacy policy considerations

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 in 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.

© Copyright IBM Corporation

July 2019

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

Please recycle

ISBN 0738457876

REDP-5549-00