

APEX File Storage for AWS

Getting Started Guide

Notes, cautions, and warnings

 **NOTE:** A NOTE indicates important information that helps you make better use of your product.

 **CAUTION:** A CAUTION indicates either potential damage to hardware or loss of data and tells you how to avoid the problem.

 **WARNING:** A WARNING indicates a potential for property damage, personal injury, or death.

Chapter 1: Introduction to this guide.....	4
About this guide.....	4
APEX File Storage for AWS and PowerScale OneFS scale-out NAS overview.....	4
PowerScale OneFS on AWS guidelines.....	5
Supported configurations.....	9
Unsupported PowerScale OneFS features.....	9
Activate APEX File Storage for AWS product license.....	10
Validate the Amazon Machine Image (AMI) signature.....	10
Chapter 2: APEX File Storage for AWS storage architecture.....	12
APEX File Storage for AWS architecture overview.....	12
Chapter 3: Prerequisites and deployment checklist.....	15
AWS account ID and region.....	15
AWS subscription and permissions.....	15
Download and install the AWS Command Line Interface (AWS CLI).....	16
Chapter 4: Plan your APEX File Storage for AWS deployment	17
Visual overview of deployment steps.....	17
What you are deploying?.....	18
EC2 instance types configuration options.....	19
EBS disk volume type configuration options.....	20
Chapter 5: Complete post-AWS deployment PowerScale OneFS tasks.....	22
Appendix A: Technical Specifications Guidelines.....	23
Protocol guidelines.....	23
File system guidelines.....	26
Authentication, identity management, and access (AIMA) control guidelines.....	28
OneFS software module guidelines.....	29
Networking guidelines.....	33
Appendix B: APEX File Storage for AWS help resources.....	35
Where to get help.....	35
Support options.....	35

Introduction to this guide

This section contains the following topics:

Topics:

- [About this guide](#)
- [APEX File Storage for AWS and PowerScale OneFS scale-out NAS overview](#)
- [PowerScale OneFS on AWS guidelines](#)
- [Supported configurations](#)
- [Unsupported PowerScale OneFS features](#)
- [Activate APEX File Storage for AWS product license](#)
- [Validate the Amazon Machine Image \(AMI\) signature](#)

About this guide

This guide is intended as a supplement to the Dell PowerScale OneFS version 9.6.x.x documentation, Amazon Elastic Compute Cloud (Amazon EC2) documentation, and APEX File Storage for AWS Deployment Guide.

Amazon Web Services (AWS) administrators and PowerScale storage administrators can use this guide for getting started with a deployment of OneFS as a virtual machine on AWS. The step-by-step deployment instructions are on the following Dell Technologies InfoHub: <https://infohub.delltechnologies.com/t/cloud/>.

In order to deploy APEX File Storage for AWS, PowerScale OneFS administrators must be familiar with AWS infrastructure, including creating, scaling, replacing, and terminating the virtual cluster in AWS using the AWS CLI or AWS Management Console.

Users must be experienced in AWS virtual machine deployment, configuration procedures, and the following:

- Using the AWS Management Console and AWS CLI
- Configuring Amazon EC2 instances
- Creating AWS services, such as security groups
- Creating AWS network interfaces

APEX File Storage for AWS and PowerScale OneFS scale-out NAS overview

APEX File Storage for AWS is a software-defined cloud solution for PowerScale OneFS deployed on AWS infrastructure. The PowerScale OneFS operating system powers a virtual software platform that delivers a scalable pool of storage and global namespace.

The APEX File Storage for AWS unified software solution supports centralized administration through PowerScale OneFS cluster administration and Amazon Elastic Compute Cloud (EC2) administration.

Dell Technologies APEX File Storage for AWS administrators manage:

- A cluster that runs a distributed file system
- A maximum of six scale-out nodes that add capacity and performance
- Storage options that manage files
- Flexible data protection and high availability

As a APEX File Storage for AWS storage administrator or application owner, you can perform self-service cluster data management tasks, such as:

- Managing folders and the file hierarchy structure
- Monitoring SMB shares and NFS exports

- Managing storage pools policies
- Monitoring quotas
- Monitoring snapshots
- Viewing reports
- Managing users

You must have the following OneFS permissions to form a OneFS cluster on AWS:

- ISL_PRIV_LOGIN_SSH
- ISL_PRIV_DEVICES
- ISL_PRIV_NETWORK
- ISL_PRIV_STATISTICS

PowerScale OneFS on AWS guidelines

The following guidelines are unique to PowerScale OneFS deployed on AWS. For product specifications and limitations, see Appendix A.

Networking on AWS guidelines

Item	Description
Internal and external networks	<ul style="list-style-type: none"> • A PowerScale OneFS cluster contains an external (front-end) network over which clients can move data in and out of the cluster. The cluster also has an internal (back-end) network over which the nodes communicate with each other. The back-end network is isolated from devices that are not in the cluster. • The Amazon Virtual Private Cloud (VPC) must have sufficient IPv4 address space to host OneFS internal and external networks, and any additional clients that will be using the deployed cluster. For details on planning the network, see the Isilon OneFS External Network Connectivity Guide version 8.10.
IPv6	<ul style="list-style-type: none"> • IPv6 is not supported in the APEX File Storage for AWS v1.0 release.
Default network pool	<ul style="list-style-type: none"> • One network pool will be created by default for client connections. The network pool name is <code>groupnet0.subnet0.pool0</code>. Each node in the cluster will be assigned one IP address from this pool. • The IP address assignments for each network interface are on Amazon EC2. • The IP address assigned to a node must not be removed or reassigned. <ul style="list-style-type: none"> ◦ Note: Removing the IP address for a node from the default network pool will cause the node to be read-only. • One IP address per NIC is the primary address and cannot be deleted, changed, or reassigned. • One IP is used for <code>groupnet0.subnet0.pool0</code> on each node. <ul style="list-style-type: none"> ◦ The IP addresses used in the pool <code>groupnet0.subnet0.pool0</code> are the AWS primary addresses that cannot be moved from one node to another. ◦ The IP addresses used in the pool <code>groupnet0.subnet0.pool0</code> cannot be a dynamic pool and cannot be changed to a dynamic pool. ◦ NOTE: Mishandling of pool0 or any of the IP addresses in it can render the cluster inaccessible.
Additional network pools	<ul style="list-style-type: none"> • Once a cluster is deployed, users are allowed to create additional network pools. These new pools can use static or dynamic allocation. The remaining IP addresses can be used after the cluster deployment for creating additional pools.
IP addresses	<ul style="list-style-type: none"> • Although the cluster does not limit the number of IP addresses that can be assigned to a node, there is a limit on the maximum number of IPs depending on the instance type. • The number of IPs used in the cluster that each node serves must not exceed the maximum number of IPs allowed for the instance type. • For more information, see: Elastic Network Interfaces and https://docs.aws.amazon.com/vpc/latest/userguide/how-it-works.html#vpc-ip-addressing.

Item	Description
Event monitoring	<ul style="list-style-type: none"> • Every node in a cluster monitors maintenance events from the AWS Instance Meta Data Service (IMDS) through the external network. • If a node cannot connect to the IMDS through the external network for two minutes or more, the node will be set to read-only.
Subnets	<ul style="list-style-type: none"> • When configuring a OneFS cluster in AWS, you must allocate two ranges of IP addresses in different AWS subnets, with one for each of the back-end and front-end networks. You create two dedicated subnets for each OneFS cluster in an existing VPC. • The internal subnet must be reserved exclusively for use by a single OneFS cluster and must contain enough free IP addresses to assign one IP address for each instance in the cluster. • Nodes in a cluster are created with a network interface for external client connections. <ul style="list-style-type: none"> ◦ The external network interface is named <code>lni-name</code> <code>ext-1</code> and <code>nic-name</code> <code>en1</code>. ◦ The external subnet must have at least one free IP address for each node in the OneFS cluster. This subnet can be shared with other clients. ◦ AWS reserves the first four addresses in subnet Classless Inter-Domain Routing (CIDR). The first address is used as the default gateway address. One IP is used for <code>groupnet0.subnet0.pool0</code> on each node. You can use any remaining IPs from the external subnet CIDR range after the cluster deployment to create additional pools. <p>Administrators create subnets as follows:</p> <ul style="list-style-type: none"> • Create a dedicated subnet for the OneFS cluster internal (back-end) network interfaces. <ul style="list-style-type: none"> ◦ Create one IP address for each node's internal network interface. ◦ Do not share this subnet with other EC2 instances. • Create a dedicated subnet for the OneFS cluster external (front-end) network interfaces. <ul style="list-style-type: none"> ◦ Create one IP address for each node's external network interface. ◦ You can share the subnet with other EC2 instances. • Note: Each NIC in AWS can belong to only one subnet <code>groupnet0.subnet0</code> and all network pools must belong to <code>groupnet0.subnet0</code>.
Network failover	<ul style="list-style-type: none"> • The cluster moves the front-end dynamic IP addresses between nodes during network failover. For on-premises clusters, nodes send GARP packets immediately after the IP move, and the IP reassignment is nearly instantaneous. However, on AWS, the cluster needs to call the cloud provider API to reassign the IP address which can take approximately 20-40 seconds. • The back-end network in AWS makes use of a single network (int-a) and the infrastructure is fully managed by the cloud provider. It uses the AWS primary address of the network interface and must not be modified. • To use network failover, an additional dynamic pool must be created from the remaining addresses in the external subnet after deployment. • Network failover is slower on AWS and can take 30-40 seconds compared to a few seconds on a PowerScale OneFS on-premises cluster. • IP addresses in dynamic pools on AWS cannot be changed in the software by the running instance without also going through EC2, which requires authorization. • The AWS IAM role and policy that you provide to the cluster at deployment time allows the IAM role to unassign and assign IP addresses and describe network interfaces. • AWS cloud calls that are triggered during normal IP failover flow through the OneFS <code>isi_cloud_net</code> library.
VPC interface endpoints for network pools	<p>Administrators set up a virtual private cloud (VPC) interface endpoint, which enables calls to AWS services without having to go through the public Internet.</p> <ul style="list-style-type: none"> • OneFS clusters running in the Cloud support configuring multiple network pools. When you create a cluster, it creates a default network pool known as <code>groupnet0.subnet0.pool0</code> automatically during the initial cluster deployment. One IP address for each node from the external subnet address range will be used in this pool. These are the AWS primary address of the external network interfaces. Any remaining

Item	Description
	<p>unused addresses from the AWS subnet CIDR can be used to create additional network pools.</p> <ul style="list-style-type: none"> OneFS allows both static and dynamic allocation policies for the new pools. You can use the OneFS CLI, OneFS Platform APIs, or the OneFS WebUI to create new network pools. When an IP address is assigned to an interface on a node, the node needs to make an API call to the AWS EC2 server to associate the IP address to the network interface. OneFS does not recommend adding elastic IP addresses to the nodes to contact EC2 servers. We recommend that you create a VPC interface endpoint for nodes to connect directly to AWS EC2 services using private IP addresses, as if the EC2 service is hosted in the cluster VPC. The VPC endpoint can be created through the AWS VPC console or by using the AWS CLI. See the procedures in the APEX File Storage for AWS Deployment Guide to create an interface VPC endpoint that connects to an AWS EC2 service. Also see https://docs.aws.amazon.com/vpc/latest/privatelink/create-interface-endpoint.html
SmartConnect DNS	<ul style="list-style-type: none"> The OneFS SmartConnect DNS feature depends on the ability of the DNS server to perform delegation. You have the option of using either private DNS servers or the AWS-provided Route53. The default DNS server on AWS, known as the <i>Route53 Resolver</i>, does not support DNS delegation, although it does support forwarding rules for resolution. Therefore, forwarding rules must be set up on Route53 to use the SmartConnect DNS feature. Administrators set up Route53 Resolver endpoints, which then forward requests to the SmartConnect IP. For more information, see: https://docs.aws.amazon.com/Route53/latest/DeveloperGuide/resolver-getting-started.html.
Cluster resizing	<ul style="list-style-type: none"> Cluster resizing, by changing the number of drives in a node or by changing the size of the drives in a node, is not supported. Cluster capacity can only be changed by adding nodes or by smartfailing and deleting nodes. Adding a node that was previously removed by a Smartfail operation is not supported. The preferred alternative is to destroy and create new instances in AWS. Reformatting a node with a new configuration that changes the externally-managed IP addresses from the original configuration is not supported in the APEX File Storage for AWS v1.0 release.
External security group	<ul style="list-style-type: none"> A security group must be applied to the external interfaces in the cluster. The details of this group depend on your planned use case. For more information on creating an external security group in OneFS, see the PowerScale OneFS 9.6.0.0 Security Configuration Guide.

Disk volume subsystem and bay mapping on AWS guidelines

Item	Description
Data drives	<ul style="list-style-type: none"> All data drives are NVMe (nonvolatile memory express) types. <ul style="list-style-type: none"> Bays that are mapped by the PCI bus device function have a maximum of 25 bays (0-24). Drives do not support HDD or SSD detection interfaces. The EBS type that is used by the cluster is provided at deployment time in the instance user data. Boot drive is EBS gp3 (nvd0) EBS data drives are EBS st1 or gp3. <ul style="list-style-type: none"> All data drives in the cluster must be the same EBS type. nvd1-nvdN Drive serial number starts with "AWS", for example AWS38335C1FBB24403E3

AWS Local instance stores guidelines

Item	Description
Temporary block storage in AWS local instance stores	<p>Temporary block storage is host-local storage that is presented as virtual disks with the following characteristics:</p> <ul style="list-style-type: none">• Local instance stores contents are persistent through reboot.• Local instance stores contents are not migrated with a virtual machine (VM) from host to host.<ul style="list-style-type: none">◦ VM host migration is possible on any power-off event.◦ Standard reboots and panics do not cause or enable VM migration.• Amazon Elastic Block Store (EBS) is remote (off-rack) storage.• EBS is performance-capped and I/O to local instance stores do not count against this limit. <p>The OneFS software journal is on the first Local instance store drive in bay 0.</p> <ul style="list-style-type: none">• The remaining Local instance store drives are at the top of the bay map.• PowerScale nodes on AWS with st1 (HDD) hard drives use the remaining Local instance stores for metadata read cache.• Local instance store drives are unused if the cluster EBS type is gp3.• Local instance store drives are used for metadata that is read if the cluster EBS type is st1 (HDD) hard drive and displays as data drives.• EBS data drives are in the remaining bays. <p>The OneFS software journal is saved during orderly shutdowns as follows:</p> <ul style="list-style-type: none">• The AWS EC2 Management Console and CLI <code>Instance Stop</code> operations cause orderly shutdowns.• AWS maintenance events cause orderly shutdowns.• Subsequent boot restores the software journal from a saved copy. <p>Local instance stores are provisioned from Local NVMe types as follows:</p> <ul style="list-style-type: none">• Two for 8xlarge and 12xlarge, four for 16xlarge and 24xlarge.• <code>nvd (N+1)-nvd (N+(2 or 4))</code>.• Volume serial number starts with "vol," for example <code>vol00ea30e97ded5ff9f</code>.•

Cloud events and monitoring guidelines

A PowerScale OneFS hardware monitoring job polls the AWS instance metadata service (IMDS) for EC2 scheduled events through the front-end network every 10 seconds as follows:

EC2 Scheduled Event	Effect	Response
Instance Stop	Instance is powered off and moved at scheduled time	Node shuts down and powers itself off
Instance Retirement	Same as Instance Stop for instances with EBS backed boot drives	Node shuts down and powers itself off
Instance Reboot	Instance is rebooted through the operating system at scheduled time Local instance store data preserved	Node proactively gracefully reboots itself
System Reboot	Physical host is rebooted at scheduled time All instances rebooted Local instance store data preserved	Node shuts down and powers itself off

EC2 Scheduled Event	Effect	Response
System Maintenance	Host may lose network connectivity during scheduled time Host may lose power at scheduled time	Node shuts down and powers itself off

NOTE: Powering off protects the OneFS software journal and prepares a node to be moved. Nodes in this state must be powered back on manually through EC2.

Supported configurations

APEX File Storage for AWS has the following supported configurations:

Characteristic	Requirement
Number of Nodes in Cluster	Minimum 4, maximum 6
EC2 Instance Type of Nodes	m5dn family; minimum m5dn.8xlarge, maximum m5dn.24xlarge
EC2 Instance Types per Cluster	One - all nodes in cluster must be of same EC2 instance type
EBS Volume Types per Cluster	One - all EBS drives in cluster must be of same EBS volume type
EBS Volume Types Supported	gp3 (SSD), st1 (HDD)
EBS Volume Counts per Node	5 or 6 (st1); 5, 6, 10, 12, 15, 18, or 20 (gp3)
EBS Volumes Sizes per Cluster	One - all EBS drives in cluster must be of same raw capacity as shown in EC2
EBS Volume Sizes Supported	4TiB or 10TiB (st1); 1TiB to 16TiB (gp3)
Aggregate Raw EBS Cluster Capacity	80TiB to 360TiB (st1); 20TiB to 1PiB (gp3)
Protection	+2n

The cluster has the same user interface as a PowerScale OneFS on-premises cluster and supports the OneFS CLI and APIs.

Unsupported PowerScale OneFS features

APEX File Storage for AWS includes a list of PowerScale OneFS features that are not currently supported in APEX File Storage for AWS. Use of these unsupported OneFS features could cause issues with your APEX File Storage for AWS cluster.

The following PowerScale OneFS features are not supported in APEX File Storage for AWS:

- IPv6 is disabled by default in APEX File Storage for AWS
 - IPv4 addresses are supported on the OneFS front-end network
- The following client protocols are disabled by default: HDFS, HTTP, FTP, SCP, and SFTP
- NDMP-based backup functionality is disabled by default
- Large file support
- SmartPools is supported but tiering is not supported
- STIG Hardening
- SmartLock (Compliance and Enterprise modes)

The following features are not applicable to a PowerScale OneFS virtual cloud infrastructure:

- Data at rest encryption (DARE)
- Instant Secure Erase (ISE)
- L3 cache
- LAGG
- RMDA and RoCE requirements that require specialized networking endpoints
- Secure Boot

- Secure Remote Services (SRS) - gateway connections
- SupportAssist - gateway connections (direct connections are supported)

Activate APEX File Storage for AWS product license

You must buy your product license from Dell or its partners. Then use the standard activation process through Dell Software Licensing Central (SLC).

About this task

The license for APEX File Storage for AWS contains the ONEFS feature. All other features are auto-licensed. There are no licensing alerts for any other features except for ONEFS capacity and duration.

For information about generating and applying a license, see the *PowerScale OneFS 9.6.x.x Web Administration Guide* section about licensing APEX File Storage for AWS.

Validate the Amazon Machine Image (AMI) signature

After you have fulfilled the prerequisites and provided Dell with your AWS account ID, your Dell Technologies APEX File Storage for AWS File Services representative will provide AWS with a preconfigured Amazon Machine Image (AMI) template for your EC2 instance.

The [Amazon Machine Image \(AMI\)](#) template packages the bits you need for your deployment.

Locate your OneFS AMI image in the AWS Management Console. For step-by-step instructions, see the section entitled, "Find the OneFS AMI ID" in the APEX File Storage for AWS **Deployment Guide**: <https://infohub.delltechnologies.com/t/cloud/>.

Validate the AMI signature

You must verify a signed manifest using OpenSSL. A signed manifest is a file that contains information about the files in a package and a digital signature to ensure the integrity and authenticity of the package. The procedure includes extracting the certificate and public key, converting the signature file to binary, and verifying the signature using OpenSSL.

Verifying a signed manifest using OpenSSL ensures the integrity and authenticity of a package. By using the following steps, you can verify the signed manifest and ensure that the package is valid.

Prerequisites

Before verifying the signed manifest, ensure that the following prerequisites are met:

- OpenSSL is installed on the system.
- The signed manifest (.rsig) file and certificate (.rcerts) files are available.

Procedure

Follow the steps below to verify the signed manifest:

1. Extract the "cer.pem" from the certificate file (.rcerts) using the command: `openssl x509 -in <certificate file> -out cert.pem`. For example:

```
openssl x509 -in onefs-9.6.0.0-391768d-manifest.txt.rcerts -out cert.pem
```

2. Extract the public key from the certificate file (cer.pem) using the command:

```
openssl x509 -pubkey -noout -in cert.pem > pubkey.pem
```

3. Convert the signature file (.rsig) to binary using the command: `openssl base64 -d -in <signature file> -out <signature file>-binary`. For example:

```
openssl base64 -d -in onefs-9.6.0.0-391768d-manifest.txt.rsig -out onefs-9.6.0.0-391768d-manifest.txt.rsig-binary
```

4. Verify the signature using the public key and the binary signature file using the command: `openssl dgst -sha256 -verify pubkey.pem -signature <binary signature file> <signed manifest file>`. For example:

```
openssl dgst -sha256 -verify pubkey.pem -signature onefs-9.6.0.0-391768d-  
manifest.txt.rsig-binary onefs-9.6.0.0-391768d-manifest.txt
```

If the signature is valid, the command will output `Verified OK`.

APEX File Storage for AWS storage architecture

This section contains the following topics:

Topics:

- [APEX File Storage for AWS architecture overview](#)

APEX File Storage for AWS architecture overview

APEX File Storage for AWS allows you to create and deploy PowerScale OneFS clusters through the resources available with Amazon Elastic Compute Cloud (Amazon EC2) on the Amazon Web Services (AWS) Cloud infrastructure.

APEX File Storage for AWS is a software-defined and customer-managed scale-out storage solution running on AWS cloud infrastructure. It brings the Dell Technologies PowerScale OneFS distributed file system into the public cloud to provide users with the same management experience as an on-premises PowerScale cluster. You can run OneFS on multiple EC2 instances backed by EBS volumes, and then form a OneFS cluster using the EC2 instances virtual nodes.

In general, cluster performance is correlated to your cluster size. The more cluster nodes (up to a maximum of six), the higher the throughput and IOPS. Therefore, before creating your cluster in AWS, consider the capacity of storage and the number of nodes that you need for your business workflow.

The following diagram shows the architecture of APEX File Storage for AWS.

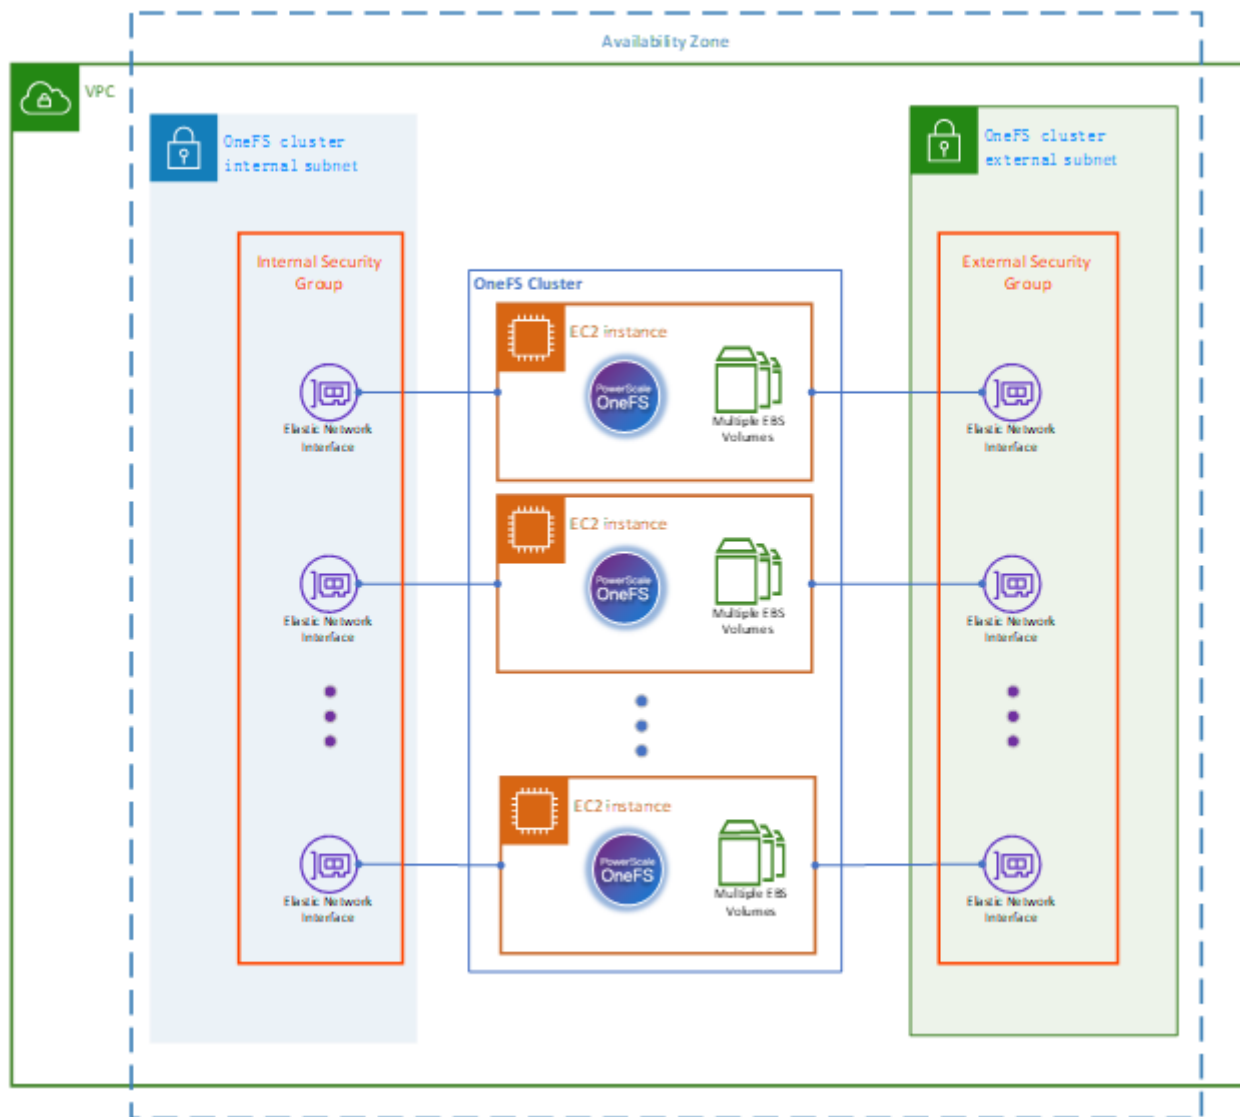


Figure 1. APEX File Storage for AWS Architecture

Resource	Description
Availability zone	APEX File Storage for AWS is designed to run in a single availability zone to obtain the best performance.
VPC	APEX File Storage for AWS requires an AWS VPC to provide network connectivity.
OneFS cluster internal subnet	Cluster nodes communicate with each other through the internal subnet. The internal subnet must be isolated from instances that are not in the cluster . Therefore, a dedicated subnet is required for the internal network interfaces of cluster nodes and that do not share the internal subnets with other EC2 instances.
OneFS cluster external subnet	Cluster nodes communicate with clients via through the external subnet by using different protocols, such as NFS, SMB, and S3.
OneFS cluster internal network interfaces	Cluster nodes network interfaces are located in the internal subnets.

Resource	Description
OneFS cluster external network interfaces	Cluster nodes network interfaces are located in the external subnets.
OneFS cluster internal security group	Security group applies to the cluster internal network interfaces, which allows traffic between the internal network interfaces of cluster nodes.
OneFS cluster external security group	Security group applies to the cluster external network interfaces, which allows specific ingress traffic from clients.
EC2 instance nodes	Cluster nodes which run the OneFS file system backed by EBS volumes.

Prerequisites and deployment checklist

Use the following checklist to record your deployment prerequisites for later use.

AWS account ID		
AWS region		
AWS availability zone		
AWS VPC ID		
OneFS cluster front-end network (external subnet)	AWS subnet ID	
	IPv4 CIDR	
	IP range for OneFS cluster	
OneFS cluster back-end network (internal subnet)	AWS subnet ID	
	IPv4 CIDR	
	IP range for OneFS cluster	

NOTE: It is recommended that you plan six IPs in each subnet for the OneFS cluster because the maximum number of cluster nodes for APEX File Storage for AWS is six.

Topics:

- [AWS account ID and region](#)
- [AWS subscription and permissions](#)
- [Download and install the AWS Command Line Interface \(AWS CLI\)](#)

AWS account ID and region

APEX File Storage for AWS is designed to run in a single availability zone. You must choose the availability zone and note what region it is in.

Provide the following information to Dell Technologies so that Dell can grant you access to the OneFS AMI image:

- Your AWS account ID
- The region in which OneFS will run

AWS subscription and permissions

An AWS subscription is required to deploy APEX File Storage for AWS.

To set up an AWS account, go to: <https://aws.amazon.com/getting-started/>.

APEX File Storage for AWS is publicly listed in AWS Marketplace, providing an overview of the product. Reach out to your Dell account teams to receive approval and purchase a license. After Dell reviews and approves your request, Dell shares the APEX File Storage for AWS AMI with your AWS customer account and adds your account to the allowed list. Once your AWS account is added to the allowed list, you can deploy the APEX File Storage for AWS AMI from your AWS Management Console.

You must have AWS permissions to perform the following cluster administration tasks using the AWS Management Console or AWS CLI.

Tasks	Supplemental AWS documentation
Create IAM policy and IAM role	<ul style="list-style-type: none"> • AWS Identity and Access Management (IAM) • Creating an IAM user in your AWS account • Using IAM roles
Create EC2 M5 instance and instance profile	<ul style="list-style-type: none"> • Creating Amazon EC2 M5 instances
Create EBS volumes	<ul style="list-style-type: none"> • Creating Amazon EBS volume types
Create subnets on an existing Amazon Virtual Private Cloud (VPC)	<ul style="list-style-type: none"> • Amazon virtual private cloud (VPC)
Create network interfaces	<ul style="list-style-type: none"> • Creating elastic network interfaces
Create security group	<ul style="list-style-type: none"> • Creating Amazon EC2 security groups
Create placement group	<ul style="list-style-type: none"> • Creating placement groups

The following links provide additional information about Amazon EC2:

- [Amazon Elastic Computer Cloud \(EC2\) documentation](#)
- [Working with the AWS Management Console](#)
- [Amazon Machine Image \(AMI\)](#)
- [Best practices for Amazon EC2](#)

Required AWS permissions

sts:GetCallerIdentity	ec2:ModifyInstanceAttribute	ec2:AuthorizeSecurityGroupIngress
ec2:DescribeVpcs	ec2:DescribeAccountAttributes	ec2:AuthorizeSecurityGroupEgress
ec2:DescribeVpcAttribute	ec2:DescribePlacementGroups	ec2:RevokeSecurityGroupIngress
ec2:DescribeVolumes	ec2>DeletePlacementGroup	ec2:RevokeSecurityGroupEgress
ec2:DescribeTags	ec2>CreatePlacementGroup	ec2>DeleteSecurityGroup
ec2:CreateTags	ec2:DescribeNetworkInterfaces	ec2>CreateSecurityGroup
ec2:DescribeRouteTables	ec2:AttachNetworkInterface	ec2:DescribeSecurityGroups
ec2:DescribeInstanceTypes	ec2>DeleteNetworkInterface	ec2>DeleteSubnet
ec2:DescribeInstances	ec2>CreateNetworkInterface	ec2>CreateSubnet
ec2:DescribeInstanceAttribute		ec2:DescribeSubnets
ec2:TerminateInstances		
ec2:RunInstances		

Download and install the AWS Command Line Interface (AWS CLI)

The AWS Command Line Interface v2 (AWS CLI v2) is a unified tool to manage your AWS services. You must have the latest AWS CLI installed and configured correctly to access your AWS subscription using the AWS CLI.

For information about downloading and installing the AWS CLI, see: <https://aws.amazon.com/cli/>. Select the operating system for the system that you are using to access AWS and to install the AWS CLI.

Plan your APEX File Storage for AWS deployment

This section contains the following topics:

Topics:

- [Visual overview of deployment steps](#)
- [What you are deploying?](#)

Visual overview of deployment steps

This section provides a visual overview of the steps that are required to deploy the AWS infrastructure and form a OneFS cluster.

Use the procedures in the **APEX File Storage for AWS Deployment Guide** located [here](#) for detailed deployment instructions.

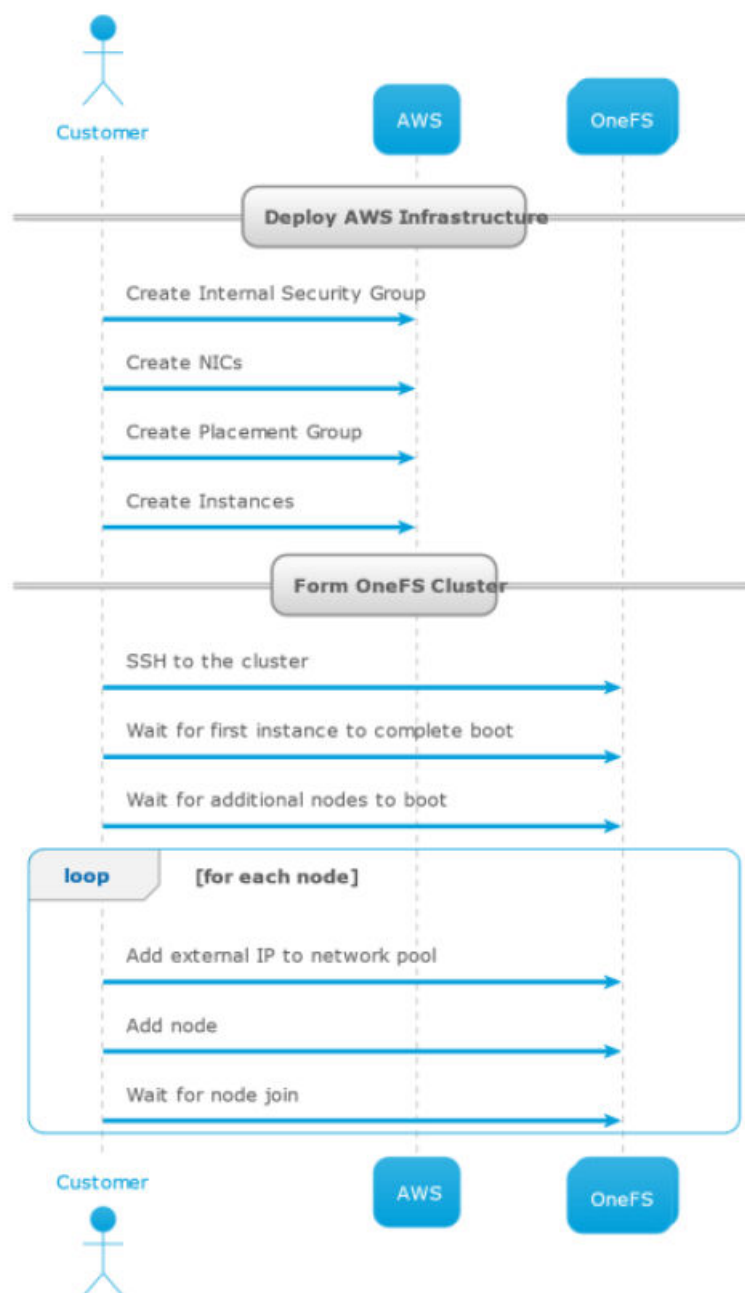


Figure 2. Overview of deployment procedures

What you are deploying?

Throughout your deployment, you are using the AWS Management Console and the AWS CLI to instantiate the required AWS cloud resources on Amazon EC2. Adding nodes, replacing failed nodes and drives requires working with both OneFS and AWS.

An administrator creates the following Amazon EC2 instances, policies, roles, groups, nodes (VMs), disks, network interfaces, and clusters using the [AWS Management Console](#), the [AWS CLI](#), and the OneFS CLI:

- **IAM policies, roles, and instance profiles**
 - OneFS nodes require an instance profile to be attached.
 - The `ec2:AssignPrivateIpAddresses` permission must be assigned to the cluster network interfaces at runtime so that they can interact with the AWS APIs.

- Creating the IAM policy, role, and instance profile for a OneFS cluster is one-time work for the same AWS account. These resources are reusable when deploying additional OneFS clusters.
- **Spread placement groups**
 - A node VM spread placement group is required to ensure that nodes are placed on distinct hardware to ensure high availability.
 - A placement group for the cluster is defined at the time of deployment.
 - Virtual machine (VM) host isolation is obtained by placing each VM in the cluster in a different spread group. Each spread group maps to a different rack in the data center.
 - EC2 allows a maximum of seven running instances per availability zone.
 - Each node VM runs on a distinct real-world rack.
 - Each rack has its own network and power.
 - Each cluster has a maximum of six nodes.
 - See [Amazon EC2 spread placement groups](#) for more information.
 - EBS volumes are provided by separate hardware.
- **Security groups**
 - Create a security group for the OneFS cluster front-end network (external security group) to allow specific ingress traffic from clients.
 - Create a security group for the OneFS cluster back-end network (internal security group) to allow all traffic between cluster nodes internal network interfaces only.
- **Network interfaces**
 - Create network interfaces for the OneFS cluster front-end network (external interfaces).
 - Create network interfaces for the OneFS cluster back-end network (internal interfaces).
- **EC2 instance user data**
 - OneFS requires user data in a JSON format file that is provided in the bits for your package. The JSON file provides new instances running OneFS with the information that is needed to create a PowerScale cluster.
- **EC2 instance types: m5dn family**
 - Sizes: 8xlarge through 24xlarge

EC2 instance types configuration options

APEX File Storage for AWS supports the following EC2 instance types.

- m5dn.8xlarge, m5dn.12xlarge, m5dn.16xlarge, m5dn.24xlarge
- m5d.24xlarge*
- i3en.12xlarge*

All of the above instance types work with st1 and gp3 type EBS volumes.

*These instance types are allowed for Proof of Concept (POC). Talk to your account team to start a POC.

The following table shows the regional availability of the supported ec2 instance types:

Region Name	m5dn	m5d.24xlarge*	i3en.12xlarge*
NAM			
USA East (N.Virginia)	Y	Y	Y
USA East (Ohio)	Y	Y	Y
USA West (N. California)	N	Y	Y
USA West (Oregon)	Y	Y	Y
Canada (Central)	N	Y	Y
SAM			
South America (Sao Paulo)	N	Y	Y
EMEA			
Europe (Frankfurt)	Y	Y	Y
Europe (Ireland)	Y	Y	Y

Region Name	m5dn	m5d.24xlarge*	i3en.12xlarge*
Europe (London)	N	Y	Y
Europe (Milan)	N	Y	Y
Europe (Paris)	N	Y	Y
Europe (Spain)	N	Y	Y
Europe (Stockholm)	N	Y	Y
Europe (Zurich)	N	Y	Y
Israel (Tel Aviv)	N	Y	Y
Middle East (Bahrain)	N	Y	Y
Middle East (UAE)	N	Y	Y
Africa (Cape Town)	N	Y	Y
Asia			
Hongkong	N	Y	Y
Hyderabad	N	Y	Y
Jakarta	N	Y	Y
Melbourne	N	Y	Y
Mumbai	N	Y	Y
Osaka	N	Y	Y
Seoul	N	Y	Y
Singapore	Y	Y	Y
Sydney	N	Y	Y
Tokyo	Y	Y	Y

For additional information, see:

- [Amazon EC2 Instance Types](#)
- [Amazon EC2 M5 Instances](#)

EBS disk volume type configuration options

Depending on your workload, you have multiple EBS disk volume types available that allows you to optimize storage performance and cost for a broad range of applications.

These volume types are divided into two broad categories as described below:

- SSD-backed storage for transactional workloads
 - **gp3**: General-purpose SSD
- HDD-backed storage for throughput intensive workloads
 - **st1**: Streaming-optimized hard drive

Volume Type	Durability	Baseline IOPS/Volume	Baseline Throughput/Volume	Max IOPS/Volume	Max Throughput/Volume	Max IOPS/Instance	Max Throughput/Instance
gp3	99.8-99.9%	3,000	125 MiBps	16,000	1,000 MB/s	260,000	10,000 MB/s
st1	99.8-99.9%	500	40 MB/s per TB**	500 MB/s	500 MB/s		10,000 MB/s

** bursting to 250 MB/s per TB

Source: <https://aws.amazon.com/ebs/volume-types/>

Complete post-AWS deployment PowerScale OneFS tasks

After you have completed the deployment of APEX File Storage for AWS on your Amazon EC2 instance using either the AWS Management Console or AWS CLI and the **Deployment Guide** located [here](#), complete the cluster formation and configuration tasks.

You can access the Command Line Interface of the first node in the cluster in any of the following ways:

- Set up a VPN to the Virtual Private Cloud (VPC)
- Use the serial console of the first cluster node
- SSH from a client in the VPC

See the following AWS white paper: [Amazon Virtual Private Cloud Connectivity Options](#) for details on your VPC connectivity options.

Configure NFS exports, SMB shares, and SyncIQ transfers

See the PowerScale OneFS documentation for details on the following configuration tasks:

Tasks	Documentation
Create and mount NFS exports	<ul style="list-style-type: none"> • By default NFS relies on a trusted network. If the network is not trusted, NFS on OneFS should be configured securely. For more details, see the "Create an NFS export" section in the <i>PowerScale OneFS 9.6.x.x CLI Administration Guide</i>.
Create and map SMB shares	<ul style="list-style-type: none"> • The default configuration of SMB on OneFS does not encrypt in-flight data. For more information, see the "Use compensating controls to protect files that are sent in cleartext" section in the PowerScale OneFS 9.6.0.0 Security Configuration Guide. • See the "SMB security" section in the <i>PowerScale OneFS 9.6.x.x CLI Administration Guide</i>. • Also see the "SMB best practices" section in the PowerScale OneFS 9.6.0.0 Security Configuration Guide for recommendations regarding enabling signing or encryption of SMB when operating over a nontrusted network.
Migrate encrypted data using SyncIQ	<ul style="list-style-type: none"> • Use the integrated capabilities of SyncIQ to encrypt the data during transfers between PowerScale clusters and protect the in-flight data during intercluster replications. SyncIQ policies support end-to-end encryption for cross-cluster communications. • For more information, see the section entitled "Data Encryption with SyncIQ" in the <i>PowerScale OneFS 9.6.x.x CLI Administration Guide</i>.

Technical Specifications Guidelines

This chapter contains the following topics:

Topics:

- [Protocol guidelines](#)
- [File system guidelines](#)
- [Authentication, identity management, and access \(AIMA\) control guidelines](#)
- [OneFS software module guidelines](#)
- [Networking guidelines](#)

Protocol guidelines

This section presents guidelines for configuring protocols for OneFS.

For assistance, contact your PowerScale account representative or PowerScale Technical Support.


Table 1. OneFS protocol specifications

Item	OneFS 9.6.0.0 on AWS	Description
NFS exports per cluster	12,000	The recommended limit for NFS exports per cluster. Exceeding this limit might negatively affect the cluster performance and client connections. Evaluate the workflow and workloads for your cluster to determine the value that works best for your environment.
NFS max read size	1 MB	The limit for NFS read size, or rsize, for NFS3 and NFS4. When you mount NFS exports from a cluster, a larger read size for remote procedure calls can improve throughput. The default read size in OneFS is 128 KB. An NFS client uses the largest supported size by default. As a result, avoid setting the size on your clients. Setting the value too low on a client overrides the default value and can undermine performance.
NFS max write size	1 MB	The limit for NFS write size, or wsize, for NFS3 and NFS4. When you mount NFS exports from a cluster, a larger write size for remote procedure calls can improve throughput. The default write size in OneFS is 512 KB. An NFS client uses the largest supported size by default. As a result, avoid setting the size on your clients. Setting the value too low on a client overrides the default value and can undermine performance.
NFS3 connections per node	300 active connections	The recommended limit for active NFS3 connections. The maximum has not been established; however, the number of available TCP sockets can limit the number of NFS connections. The number of connections that a node can process depends on the ratio of active-to-idle connections and on the resources that are available to process the sessions. Monitoring the

Table 1. OneFS protocol specifications (continued)

Item	OneFS 9.6.0.0 on AWS	Description
		number of NFS connections to each node helps prevent overloading a node with connections.
NFS4 connections per node	300 active connections	The recommended limit for active and passive NFS4 connections. The maximum has not been established; however, the number of available TCP sockets can limit the number of NFS connections. The number of connections that a node can process depends on the ratio of active-to-idle connections and on the resources that are available to process the sessions. Monitoring the number of NFS connections to each node helps prevent overloading a node with connections.
Concurrent PAPI processes per node	The number of PAPI processes are 60 PAPI processes per node for clusters with 3-6 nodes	The limit for the process pool for the PAPI daemon. This limit scales automatically based on the size of the cluster. This limit affects the number of PAPI requests that can be processed concurrently.
RAN attribute key length	200 B	The limit of the key length for the OneFS extended user attribute (x-isi-ifs-attr- <code><name></code>).
RAN attribute value length	1 KB	The limit of the value length for the OneFS extended user attribute (x-isi-ifs-attr- <code><name></code>).
Maximum RAN concurrent connections per node	50 (default) 300 (custom)	The limit of RAN concurrent connections per node using default parameters. You can obtain higher scalability for RAN by using nondefault configuration parameters. The maximum limit depends on many parameters and can be specific to a clusters workflow. Contact your Dell EMC PowerScale account team or PowerScale Technical Support for help with configuring the nondefault parameters. For more information, see the PowerScale knowledge base article 304701, How to update RAN scalability parameters (restricted).
RAN URI length	8 KB	The limit for the URI length that is used for the RAN HTTP operation.
RAN user attributes	126	The limit for extended user attributes that OneFS supports.
S3 object key length	1024 bytes	The maximum object key length used to identify objects uniquely within a bucket can be 1024 bytes.
S3 maximum number of objects per bucket	1,000,000	This is the limit of objects per bucket. This affects only the number of direct children of a prefix, not the total number of objects that can be stored within a root bucket. Exceeding this limit might negatively affect the cluster performance and client connections. Evaluate the workflow and workloads for your cluster to determine the value that works best for your environment.
S3 buckets per cluster	12,000 total buckets	Total number of S3 buckets that can be created on the cluster. There is also a limit of 1000 buckets per user.

Table 1. OneFS protocol specifications (continued)

Item	OneFS 9.6.0.0 on AWS	Description
S3 metadata size	Key length: 200 bytes. Value length: 1024 bytes.	Objects may have arbitrary keys that consist of maximum of 200 bytes of UTF-8 encoded, case-sensitive alphanumeric characters, period ('.'), and underscore ('_') characters. Values of the attributes are arbitrary binary data of not more than 1 KB. Although objects on OneFS can support up to 128 extended attributes with a total size of 8 KB, S3 file upload operations support a lower limit as we are limited by a maximum HTTP header size of 8 KB.
S3 connections per node	150	The limit for concurrent S3 connections per node.
S3 maximum object size	4,398 TB (4 TiB)	The maximum size for a file for all PowerScale clusters. Files larger than 1 TB can negatively affect job engine performance.
S3 multi-part upload: part size	5 MB to 5 GB	This limit is the same as that for Amazon S3.
SMB share names	80 characters	SMB share names of length limited to 80 characters are supported. Unicode characters are supported except control characters (0x00-0x1F). The following characters are illegal in a share name: " \ / [] : < > + = ; , * ?
SMB shares per cluster	24,000	This is the recommended limit for SMB shares per cluster.
SMB 2 request size	1 MB	OneFS supports the large 1 MB maximum transmission unit (MTU) that the SMB2.1 introduced. The MTU is the size of the largest data unit that the SMB protocol can transmit and receive. The large MTU can improve the overall throughput of SMB transmissions.
SMB 2 and SMB 3 connections per node	1,000 active connections 9,000 idle connections	The number of active SMB 2 or SMB 3 connections that a node can process depends on the type of node. The more CPUs and RAM that a node has, the more active connections the node can process. The kernel imposes memory constraints on the OneFS protocol daemons, such as the input-output daemon (lwio), and these constraints limit the number of SMB 2 or SMB 3 connections. To ensure that a node does not become overloaded with connections, you should monitor the number of SMB connections to each node.  NOTE: SMB 3 features require increased memory and CPU processing. Enabling continuous availability or encryption on a share reduces these limits.
SSH connections per node	200	The recommended limit for SSH connections per node. The maximum number of SSH connections per node has not been established.

File system guidelines

This section presents guidelines for configuring the OneFS file system.

For assistance, contact your PowerScale account representative or PowerScale Technical Support.

Table 2. OneFS file system specifications

Item	OneFS 9.6.0.0 on AWS	Description
Block size	8 KB	The maximum block size limit. This limit cannot be changed.
Cluster name length	40 characters	The maximum length for the cluster name.
Cluster size	6 nodes	The maximum number of nodes that a cluster can have.
Custom access pattern templates	5	The limit for custom file-system-tunable templates. This limit is in addition to the default templates of "random," "streaming," and "default."
Directories per directory	100,000	The recommended limit for the number of directories in a directory. Exceeding this limit might negatively affect the cluster performance and client connections. Evaluate the workflow and workloads for your cluster to determine the value that works best for your environment.
Directory depth	509	Maximum recommended depth of a directory tree is 509.
Protection policies	+2n	The following FEC options are supported: +2nOneFS protection is defined at the node pool level. A cluster with multiple node pools can have multiple protection schemes simultaneously. The recommended protection level depends on the size of the node pool and node types. For information about disk pools, node pools, and tiers, see the white paper Storage Tiering with Dell PowerScale SmartPools .
File clones per file	32,766	The maximum number of references for a single block in a shadow store. When the limit for file clones per file is exceeded, a new shadow store is created.
File name length	Up to 1024 unicode characters in namelength domains. 1024 bytes in regular directories.	In namelength domains, OneFS can support upto 1024 unicode characters. In regular directories, OneFS supports a maximum filename length of 1024 bytes. Most Unicode character encodings, such as UTF-8, specify that a character can have multiple bytes. UTF-8 can have up to 4 bytes per characters. The characters in some languages, such as Japanese, are likely to have multiple bytes per character. OneFS supports UTF-8 by default.
Standard file size	4,398 TB (4 TiB)	The maximum size for a file for all PowerScale clusters. Files larger than 1 TB can negatively affect job engine performance.
File system size	1 PB	The maximum capacity for the file system. The capacity size does not include overhead for the

Table 2. OneFS file system specifications (continued)

Item	OneFS 9.6.0.0 on AWS	Description
		OneFS operating system, the file system, or data protection.
Files per directory	1,000,000	The recommended limit for files per directory. Exceeding this limit might negatively affect the cluster performance and client connections. Evaluate the workflow and workloads for your cluster to determine the value that works best for your environment. To improve performance when managing large numbers of files in a single directory, use nodes that have solid-state drives (SSDs).
Hard links per file	1,000	The default and maximum number of hard links per file. You can set the maximum number of hard links per file with the <code>efs.ifm.max_links</code> system control. Setting the number higher than the default value can slow snapshot operations and file deletions. For more information, see the EMC Isilon knowledge base article 447064, OneFS: Sysctl: efs.ifm.max_links .
Inodes per cluster	Billions	OneFS dynamically allocates new inodes from free file system blocks. The maximum number of possible inodes depends on the number and density of nodes in the cluster, as expressed by the following formulas: <ul style="list-style-type: none"> 4Kn drives: $((\text{number of nodes in the cluster}) * (\text{node raw TB}) * 1000^4 * .99) / (8192 * (\text{number of inode mirrors}))$ 512n drives: $((\text{number of nodes in the cluster}) * (\text{node raw TB}) * 1000^4 * .73) / (512 * (\text{number of inode mirrors}))$ See the guideline for files per directory. The limit for files per directory can limit the number of files that the system can store.
Logical Node Numbers (LNNs)	6	The limit for logical node numbers.
Node pools per cluster	1	The recommended and maximum limits for node pools per cluster. The number of node pools that can be created is limited by the number of nodes in the cluster.
Open files per node	315,000	The maximum number of open files per node depends on the maximum number of vnodes that are available on the node. The amount of available vnodes depends on how much RAM the node has. The maximum number of open files per node is 90% of the maximum number of vnodes on that node, as expressed in the following formula: <code>kern.maxfiles = kern.maxvnodes * .9</code> The OneFS protocol daemons, such as the input-output daemon (<code>lwio</code>), might impose additional constraints on the number of files that a node can have open. The protocol daemons typically impose such constraints because the kernel places limits on per-process memory consumption.

Table 2. OneFS file system specifications (continued)

Item	OneFS 9.6.0.0 on AWS	Description
Path length	4096 bytes	<p>The maximum length for a pathname. The length is the maximum length of a directory path that can be passed into a system call; it does not represent the absolute depth of nested directories. Shorter path and file names require fewer lookup operations. As a best practice, keep your path and file names as short as possible, especially in workflows that include many lookups. OneFS features like NDMP and SyncIQ may not work as expected on paths longer than the maximum limit.</p> <p>NOTE: For symbolic links, the path length of the target is restricted to 1024 bytes if the symlink source is in a restricted domain.</p>
Device IDs	65,535	Device IDs are unique identifiers for nodes. Device IDs are not reused when nodes are replaced. To reach the limit of Device IDs in a three-node cluster, you must replace nodes 65,532 times.
User attribute keys	16	The limit of attribute keys that can be created within any file system object. The user attribute term refers to custom file system metadata that the FreeBSD extattr API creates. These extended attribute data types can be acted on by SmartPools, for example, by choosing the <code>File Attribute</code> file pool policy filter. Extended attributes exist as "name=value" pairs within a file system object.
User attribute key size	24 bytes	The limit size for the user attribute key.
User attribute value size	128 bytes	The limit size for the user attribute value.
User attribute total size	1 KB	The limit for the size of the custom metadata that is associated with the file system object.

Authentication, identity management, and access (AIMA) control guidelines

This section presents guidelines for configuring directory services and OneFS access zones.

For assistance, contact your PowerScale account representative or Dell Technologies Support.

Table 3. OneFS AIMA specifications

Item	OneFS 9.6.0.0 on AWS	Description
Access zones	15	The recommended limit for access zones. Exceeding this limit might negatively affect the cluster performance and client connections. Evaluate the workflow and workloads for your cluster to determine the value that works best for your environment. The maximum limit has not been established.
ACEs per ACL	1,000	The limit for Access Control Entries (ACEs) per Access Control List (ACL). ACEs are stored and evaluated linearly. Large numbers of ACEs

Table 3. OneFS AIMA specifications (continued)

Item	OneFS 9.6.0.0 on AWS	Description
		per ACLs increase the number of authorization checks that must be performed, which might negatively affect system performance.
Kerberos token size	64 KB	The size limit for the Kerberos token.
LDAP domains	15	The recommended limit for Lightweight Directory Access Protocol (LDAP) domains. This guideline represents unique LDAP domains. See the entry for access zones.
Local groups (per cluster)	7,500	The recommended limit for local groups per cluster. Exceeding this limit might negatively affect the cluster performance and client connections. Evaluate the workflow and workloads for your cluster to determine the value that works best for your environment.
Local users (per cluster)	7,500	The recommended limit for local users per cluster. Exceeding this limit might negatively affect the cluster performance and client connections. Evaluate the workflow and workloads for your cluster to determine the value that works best for your environment.
Microsoft Active Directory domains	15	The recommended limit for Active Directory domains. See the entry for access zones.
NIS domains	15	The recommended limit for Network Information Service (NIS) domains. The guideline represents unique NIS domains. See the entry for access zones. Although you can specify multiple NIS domains in an access zone, NFS users benefit only from the NIS configuration that is defined in the system access zone.
RBAC roles	200	The recommended limit for role-based access control (RBAC) roles. Exceeding this limit might negatively affect the cluster performance and client connections. Evaluate the workflow and workloads for your cluster to determine the value that works best for your environment. The maximum limit has not been established.
User mapper rules	1,000	The recommended limit for user mapper rules. Exceeding this limit might negatively affect the cluster performance and client connections. Evaluate the workflow and workloads for your cluster to determine the value that works best for your environment. The maximum limit has not been established.

OneFS software module guidelines

This section presents guidelines for configuring OneFS software modules.

For assistance, contact your PowerScale account representative or Dell Technologies Support.

Table 4. OneFS software module specifications


Item	OneFS 9.6.0.0 on AWS	Description
Anti-virus: file size	ICAP: 2 GB CAVA: 4 TiB	The recommended and maximum allowed file size limit for anti-virus scans. Evaluate the workflow and workloads for your cluster to determine the value that works best for your environment.
Anti-virus: scan report entries	10000	The maximum number of anti-virus scan and threat reports that can be fetched from a configuration at a given time. Reports beyond the limit can be fetched by configuring the offset parameter for the configuration.
Audit: CEE servers	6 servers per cluster	OneFS must ping all the Common Event Enabler (CEE) servers within a single heartbeat window. The number of servers that can be contacted and that can respond during the window is estimated to be 252. The network topology and cluster bandwidth might require a lower limit.
Audit: Events forwarded to CEE	4500 events per second	The sustained number of audit events, per second, that can be forwarded to a CEE server. This limit might be higher in some circumstances, depending on the workload, the type of node, and the CEE server configuration.
Audit: log expiration	User configurable	Audit logs can be autodeleted from the system by specifying a retention period. Minimum retention period that can be specified is 1 day. Logs can also be deleted manually by specifying a delete-before date.  NOTE: Logs are not deleted until all the contained events have been forwarded to a CEE server.
Audit: log file size	1 GB	The size limit for audit log files. When a log file reaches the maximum size, the log file is closed and a new log file is created. Old log files can be deleted from the cluster using manual or auto-delete methods.
Audit: maximum size of an audit event	65535 bytes	This is the maximum supported size for an audit event. If the size of an audit event is greater than 65,535, that log event is discarded and the file access operation that caused the event fails.
CloudPools: account name	768 characters	The maximum length for a CloudPool account name.
CloudPools: account username	Service provider sets this limit	The maximum length for a CloudPool account . This limit is set by the service provider. Check with your cloud provider for more information.
CloudPools: account password	255 characters	The maximum length for a CloudPool account password.
CloudPools: pool name	768 characters	The maximum length for a CloudPool name
CloudPools: vendor name	2048 characters	The maximum length for a CloudPool vendor name.
CloudPools: description	4096 characters	The maximum length for a CloudPool description.
CloudPools: accounts to tier to	80 accounts 30 active accounts	The maximum number of accounts that a CloudPool account can tier to. The number of

Table 4. OneFS software module specifications (continued)

Item	OneFS 9.6.0.0 on AWS	Description
		accounts that can be active is limited by the maximum number of file pool policies.
CloudPools: containers in cloud	The service provider sets this limit	The maximum number of containers in the cloud. This limit is set by the service provider. Check with your cloud provider for more information.
CloudPools: cloud container size	The service provider sets this limit	The size of the cloud container. The service provider sets this limit. Check with your cloud provider for more information.
CloudPools: storage size per CloudPool account	The service provider sets this limit	The storage size for a CloudPool account. The service provider sets this limit. Check with your cloud provider for more information.
CloudPools: file size tiered to cloud	4.398 TB (4 TiB)	The size of files that can be archived to the cloud and retrieved from the cloud. The service provider sets this limit. Check with your cloud provider for more information.
CloudPools: proxy limits	Proxy name: 1024 characters Proxy hostname: 1024 characters Proxy username: 1024 characters Proxy password: 256 characters	The maximum lengths for a CloudPool proxy name, hostname, username, and password.
Datamover Jobs	300	The limit for concurrently executed datamover jobs. Datamover can only replicate one same dataset to at most ~126 target locations on the same cluster, else the job can not complete successfully.
Datamover Policies	300	The recommended limit for datamover policies. Evaluate the workflow and workloads for your cluster to determine the value that works best for your environment. The maximum limit has not been established.
Datamover Accounts	300	The recommended limit datamover accounts. Evaluate the workflow and workloads for your cluster to determine the value that works best for your environment. The maximum limit has not been established.
Datamover Datasets	300	The recommended limit for datamover datasets. Evaluate the workflow and workloads for your cluster to determine the value that works best for your environment. The maximum limit has not been established.
File pool policies: AND and OR conditions	3 ORs and 5 ANDs	A file pool policy can have 3 OR disjunctions, and each term joined by the ORs can contain at most 5 ANDs. For example: (A and B and C and D and E) or (F and G and H and I and J) or (K and L and M and N and O).
File pool policies: number of file pool policies per cluster	80	The recommended limit for file policies per cluster. Exceeding this limit might negatively affect the cluster performance and client connections. Evaluate the workflow and workloads for your cluster to determine the value that works best for your environment.

Table 4. OneFS software module specifications (continued)

Item	OneFS 9.6.0.0 on AWS	Description
Job Engine: concurrent jobs	3	The number of concurrent jobs that the job engine can run. However, the job Exclusion Sets (restripe or marking) determine which jobs can be run simultaneously. Concurrent job execution is also governed by job priority and overall cluster health. For more information, see the OneFS Job Engine White Paper .
SmartDedupe: block size	8 KB	SmartDedupe works on file system blocks that are 8 KB.
SmartDedupe: maximum paths per job	10	The recommended limit for paths per job for SmartDedupe. Exceeding this limit might negatively affect the cluster performance and client connections. Evaluate the workflow and workloads for your cluster to determine the value that works best for your environment.
SmartDedupe: minimum file size	32 KB	The minimum file size that SmartDedupe can process. SmartDedupe will not deduplicate files that are smaller than 32 KB.
SmartDedupe: shadow stores	32,000	Each shadow store can have 32,000 pointers. This limit is imposed by the kernel. The OneFS shadow store is a metadata structure that references physical blocks to decrease the physical storage that is required to store data, which maximizes storage efficiency.
SmartPools: Tiers	1	The recommended limit for SmartPools tiers.
SmartQuotas: directory depth	509	The maximum limit for directory depths for SmartQuotas. Directory depths deeper than 275 directories might negatively affect system performance.
SmartQuotas: number of quotas per cluster	150,000	The recommended limit for quotas per cluster. The maximum number of quotas per cluster has not been established. Exceeding this recommended limit might negatively affect the cluster performance and client connections. Listing of quotas in the WebUI is expected to take time. For assistance, contact your PowerScale account representative or Dell Technologies Technical Support.
SnapshotIQ: directory depth	509	The maximum limit for directory depths for SnapshotIQ. Directory depths deeper than 275 directories might negatively affect system performance.
SnapshotIQ: number of snapshots	6,000	The limit for snapshots per cluster
SnapshotIQ: Number of writable snapshots	Default: 30 Maximum supported: 2048 with limitations*	The limit for writable snapshots per cluster. Limitations: Do not delete all writable snapshots at the same time which can lead to filling up of Job Engine queue.
SyncIQ: defined policies	1,000	The recommended limit for defined SyncIQ policies. The maximum limit of defined policies has not been established. If the number of policies exceeds the recommended limit, you should keep in mind the following effects:

Table 4. OneFS software module specifications (continued)

Item	OneFS 9.6.0.0 on AWS	Description
		<ul style="list-style-type: none"> • SyncIQ is bound by the limit on the number of concurrently running policies. If many policies are running on schedules, the queue to run the jobs might become so large that OneFS can never complete all the jobs in the queue. • Each policy represents a set of snapshots on the source and the destination clusters. More snapshots mean that more jobs must run to delete the snapshots, and the increase in the number of jobs can negatively affect the cluster performance.
SyncIQ: running policies	15	The recommended limit of running SyncIQ policies. For clusters with 3 or fewer nodes, the limit depends on the number of CPU cores per node. There can be one worker per CPU core, with each worker running 4 policies. The recommended limit for smaller clusters is: 4 * number of CPU cores per cluster. Exceeding this limit might negatively affect the cluster performance and client connections. Evaluate the workflow and workloads for your cluster to determine the value that works best for your environment.
SyncIQ: workers per node (policy setting)	3	The recommended limit for workers per node. Exceeding this limit might negatively affect the cluster performance and client connections. Evaluate the workflow and workloads for your cluster to determine the value that works best for your environment.
SyncIQ: workers per policy	40	The recommended limit for workers per policy. Exceeding this limit might negatively affect the cluster performance and client connections. Evaluate the workflow and workloads for your cluster to determine the value that works best for your environment.

Networking guidelines

This section presents guidelines for OneFS networking configurations.

Table 5. OneFS networking specifications

Item	OneFS 9.6.0.0 on AWS	Description
Default routes per node	1	The limit for default routes per node. OneFS does not support default routes per interface.
DNS configurations per cluster	1 per groupnet	The recommended limit for DNS configurations per cluster. In OneFS, you can specify multiple DNS resolver configurations with a limit of one DNS resolver configuration per groupnet. You can have as many groupnets as there are access zones.
DNS name servers per configuration	3	The limit for DNS name servers per configuration.
Groupnets	1 per cluster	The limit for groupnets per access zone. Groupnets are optional and should be used only if the access zone requires an alternate DNS

Table 5. OneFS networking specifications (continued)

Item	OneFS 9.6.0.0 on AWS	Description
		server. The number of access zones should not exceed 50.
DNS search suffixes per configuration	6	The limit for DNS search suffixes per configuration.
Network pools per cluster	5	The recommended limit for network pools per cluster. The maximum limit has not been established. The number of network pools should be kept under 100 pools across all subnets and groupnets in the cluster
SmartConnect DNS zone names	5	The limit for SmartConnect DNS zone names per cluster. See the "Network pools per cluster" entry for more information.
SmartConnect DNS zone name aliases	5	The recommended limit for SmartConnect DNS zone name aliases. The maximum limit has not been established. The number of DNS zone name aliases should be kept under 100 in the cluster.
Subnets per cluster	1	The limit for subnets per cluster.

APEX File Storage for AWS help resources

Topics:

- [Where to get help](#)

Where to get help

The Dell Technologies Support site (<https://www.dell.com/support>) contains important information about products and services including drivers, installation packages, product documentation, knowledge base articles, and advisories.

A valid support contract and account might be required to access all the available information about a specific Dell Technologies product or service.

Support options

This section contains resources for getting answers to questions about PowerScale OneFS products.

Dell Technologies Support	<ul style="list-style-type: none">• https://www.dell.com/support/incidents-online/en-us/contactus/product/isilon-onefs
Telephone support	<ul style="list-style-type: none">• United States: 1-800-SVC-4EMC (1-800-782-4362)• Canada: 1-800-543-4782• Worldwide: 1-312-725-5401• Local phone numbers for a specific country or region are available at https://www.dell.com/support/incidents-online/contactus/product/isilon-onefs.
PowerScale OneFS Documentation Info Hubs	<ul style="list-style-type: none">• https://www.dell.com/support/kbdoc/en-us/000152189/powerscale-onefs-info-hubs
APEX File Storage for AWS Deployment Guide on PowerScale Cloud InfoHub	<ul style="list-style-type: none">• https://infohub.delltechnologies.com/t/cloud/