

PowerScale OneFS

Technical Specifications Guide

9.3.0.0

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.....	4
About this guide.....	4
Scale-out NAS overview.....	4
Where to get help.....	4
Additional options for getting help.....	5
Chapter 2: Technical Specifications Guidelines.....	6
Protocol guidelines.....	6
File system guidelines.....	10
Authentication, identity management, and access (AIMA) control guidelines.....	13
OneFS software module guidelines.....	14
Networking guidelines.....	18
Hardware guidelines.....	19

Introduction

This chapter provides information about this guide and contains the following topics:

Topics:

- [About this guide](#)
- [Scale-out NAS overview](#)
- [Where to get help](#)

About this guide

This document presents guidelines and recommendations for configuring OneFS. Configuration guidelines are provided for protocols, file system features, software and hardware components, and network settings.

Your suggestions help us to improve the accuracy, organization, and overall quality of the documentation. Send your feedback to <http://bit.ly/isilon-docfeedback>. If you cannot provide feedback through the URL, send an email message to docfeedback@dell.com.

Scale-out NAS overview

The scale-out NAS storage platform combines modular hardware with unified software to harness unstructured data. The OneFS operating system powers the platform to deliver a scalable pool of storage with a global namespace.

The unified software platform supports centralized administration through OneFS and through Dell Technologies APEX File Storage Services (File Services). OneFS administrators and Dell Technologies APEX File Storage Services administrators manage:

- A cluster that runs a distributed file system
- Scale-out nodes that add capacity and performance
- Storage options that manage files and tiering
- Flexible data protection and high availability
- Software modules that control costs and optimize resources.

If you are a File Services storage administrator or application owner, you request services through your Dell Technologies APEX File Storage Services service provider. As a File Services 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, NFS exports, and HDFS access
- Managing storage pools policies
- Monitoring quotas
- Monitoring snapshots
- Viewing reports
- Managing users

See the PowerScale APEX File Storage Services Administration Guide for details.

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.

Additional options for getting help

This section contains resources for getting answers to questions about PowerScale 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: 1-508-497-7901• Local phone numbers for a specific country or region are available at https://www.dell.com/support/incidents-online/en-us/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
Dell Community Board for self-help	<ul style="list-style-type: none">• https://www.dell.com/community

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](#)
- [Hardware 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 9.3.0.0 and later	Description
FTP connections per node	200	The recommended limit for FTP connections per node. This number is the tested limit. If the number of FTP connections to a node exceeds the recommended limit, FTP performance might be negatively affected. The limit for FTP connections per node assumes anonymous access that requires no authentication.
HDFS block size	64 MB–512 MB	The recommended range for HDFS block sizes. For best results, the block size should not be smaller than 4 KB or larger than 1 GB. The specific value varies by workflow. Smaller block sizes require more tasks; however, you want a large enough number of tasks to take advantage of all the slots on the cluster.
HDFS root directory	1 per access zone	The number of HDFS root directories per access zone that OneFS supports. The limitation for access zones and authentication providers is the same for HDFS and other protocols.
Files and directories per HDFS fsimage	30,000,000	HDFS supports a dataset of 30,000,000 objects (files or directories) for the generation of an fsimage in each zone.
Encryption zone keys for HDFS	999	Transparent Data Encryption for HDFS protocol stores encrypted data in a directory tree that is called the encryption zone. Each encryption zone is defined by a KMS key. Each OneFS cluster supports up to 999 keys. The same key can be used in multiple zones, so this does not limit the creation or management of encryption zones themselves.
HTTP connections per node	500	The limit for HTTP connections per node. OneFS runs version 2 of the Apache HTTP

Table 1. OneFS protocol specifications (continued)

Item	OneFS 9.3.0.0 and later	Description
		Server, which includes the Apache Multi-Processing Module (MPM) that implements a hybrid multiprocess, multithreaded server. The Apache MPM configuration limits the number of simultaneous connections that OneFS services. OneFS queues connections after the connection limit is reached and processes them as resources become available. 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.
NDMP block size	512 KB	The size limit for tape blocks. If you back up tape blocks that are larger than the size limit, OneFS backs up 256 KB blocks.
NDMP connections per node	64	The limit for the number of NDMP connections that are allowed per node.
NFS exports per cluster	40,000	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 <code>rsize</code> , 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 <code>wsize</code> , 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	1,000 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 number of NFS connections to each node helps prevent overloading a node with connections.
NFS4 connections per node	1,000 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

Table 1. OneFS protocol specifications (continued)

Item	OneFS 9.3.0.0 and later	Description
		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.
NFS over RDMA connections per node	32 connections	The recommended maximum limit for NFS over RDMA connections per node.
Concurrent PAPI processes per node	For 8.2.2 and later, the number of PAPI processes per node increases by 20 for the following: <ul style="list-style-type: none"> 60 PAPI processes per node for clusters with 3-20 nodes 62-98 PAPI processes per node for clusters with 21-39 nodes 100 PAPI processes per node for clusters with more than 40 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 (<code>x-isi-ifs-attr- <name></code>).
RAN attribute value length	1 KB	The limit of the value length for the OneFS extended user attribute (<code>x-isi-ifs-attr- <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	40,000 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 9.3.0.0 and later	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	500	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 expanded object size	17.6 TB (16 TiB)	The maximum size for a file that can be supported with specific PowerScale hardware configurations.
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	80,000	This is the recommended limit for SMB shares per cluster.
SMB 1 connections per node	1,000	The number of SMB 1 connections that a node can process depends on the type of node and whether the connections are active or idle. 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 1 connections.
SMB 1 request size	64 KB	The SMB1 protocol determines the request size limit.
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	3,000 active connections 27,000 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

Table 1. OneFS protocol specifications (continued)

Item	OneFS 9.3.0.0 and later	Description
		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. i 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.
SWIFT connections per node	150	The SWIFT limit depends on the memory pattern of concurrent GET/PUT connections and is capped at 512 MB of memory. Because memory use varies by process, lower connection numbers are possible.
SWIFT objects per container	10,000	The number of objects that can be listed in a container GET or HEAD request.

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 9.3.0.0	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	252 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.
FEC protection	+4n	The following FEC options are supported: +1n, +2n, +2d:1n, +3d:1n, +3d:1n1d, +3n, +4d:1n, +4d:2n, +4n. OneFS 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,

Table 2. OneFS file system specifications (continued)

Item	OneFS 9.3.0.0	Description
		node pools, and tiers, see the white paper Storage Tiering with Dell EMC Isilon SmartPools .
Mirrored protection	8x (maximum)	Mirroring options from 2x to 8x are supported. The recommended value depends on the node pool size.
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.
Expanded file size	17.6 TB (16 TiB)	The maximum size for a file that can be supported with specific PowerScale hardware configurations.
File system size	66 PB	The maximum capacity for the file system. The capacity size does not include overhead for the 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:

Table 2. OneFS file system specifications (continued)

Item	OneFS 9.3.0.0	Description
		<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: $((\text{number of nodes in the cluster}) * (\text{node raw TB}) * 1000^4 * .73) / (512 * (\text{number of inode mirrors}))$ <p>See the guideline for files per directory. The limit for files per directory can limit the number of files that the system can store.</p>
Logical Node Numbers (LNNs)	252	The limit for logical node numbers.
Node pools per cluster	20	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 (lwio), 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.
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 File Attribute file pool policy filter.

Table 2. OneFS file system specifications (continued)

Item	OneFS 9.3.0.0	Description
		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 PowerScale Technical Support.

Table 3. OneFS AIMA specifications

Item	OneFS 9.3.0.0	Description
Access zones	50	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 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	50	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)	25,000	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)	25,000	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	50	The recommended limit for Active Directory domains. See the entry for access zones.
NIS domains	50	The recommended limit for Network Information Service (NIS) domains. The guideline represents

Table 3. OneFS AIMA specifications (continued)

Item	OneFS 9.3.0.0	Description
		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 PowerScale Technical Support.

Table 4. OneFS software module specifications

Item	OneFS 9.3.0.0	Description
Anti-virus: file size	ICAP: 2 GB CAVA: 16 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	1 (max) audit server per node 252 (max) audit 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.

Table 4. OneFS software module specifications (continued)


Item	OneFS 9.3.0.0	Description
		 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	256 accounts 30 active accounts	The maximum number of accounts that a CloudPool account can tier to. The number of 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.
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

Table 4. OneFS software module specifications (continued)

Item	OneFS 9.3.0.0	Description
		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	256	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.
InsightIQ 4.1.1 and later: monitored clusters	8 clusters	The number of clusters that InsightIQ 4.1.1 and later can monitor. If you are running a different version of InsightIQ, see the release notes for the version of InsightIQ that you are running.
InsightIQ 4.1.1 and later: monitored nodes	80 nodes for one cluster; 150 nodes for multiple clusters	The number of nodes for one cluster or multiple clusters that InsightIQ 4.1.1 and later can monitor. If you are running a different version of InsightIQ, see the release notes for the version of InsightIQ that you are running.
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 .
MFTv3: concurrent downloads	1	Maximum number of concurrent downloads through EMC Secure Remote Services (ESRS).
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	5	The recommended limit for SmartPools tiers. 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.
SmartQuotas: directory depth	509	The maximum limit for directory depths for SmartQuotas. Directory depths deeper than 275 directories might negatively affect system performance.

Table 4. OneFS software module specifications (continued)

Item	OneFS 9.3.0.0	Description
SmartQuotas: number of quotas per cluster	500,000	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	20,000	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: <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	50 - for clusters with 4 or more nodes OR 4 * number of CPU cores per cluster - for clusters with 3 or fewer nodes	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.

Table 4. OneFS software module specifications (continued)

Item	OneFS 9.3.0.0	Description
		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 9.3.0.0	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 access zone	The limit for groupnets per access zone. Groupnets are optional and should be used only if the access zone requires an alternate DNS 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	100	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	100	The limit for SmartConnect DNS zone names per cluster. See the "Network pools per cluster" entry for more information.
SmartConnect DNS zone name aliases	100	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	100	The limit for subnets per cluster.
VLANs per cluster	100	The limit for VLANs per cluster.

Hardware guidelines

This section presents guidelines for OneFS hardware configurations.

Table 6. OneFS hardware specifications

Item	OneFS 9.3.0.0	Description
Backup accelerator: tape device paths	4 (1 path per FC port that is connected)	The limit for device paths per backup accelerator node.
InfiniBand cable length	Varies by node type	Nodes with InfiniBand use QDR IB cables and support cable lengths of up to 100 meters.

Leaf-Spine clusters

Leaf-Spine networking on OneFS version 8.2, and later clusters requires Dell Z9100-ON switches:

- 12 Z9100-ON Leaf switches on the front-end (external) network
- 5 Z9100-ON Spine switches on the back-end (internal) network