PowerScale Node Site Preparation and Planning Guide

May 2021

DELL EMC

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.

Contents

1

Introduction to this guide

This section contains the following topics:

Topics:

About this guide

This guide describes how to prepare and plan for an F900, F600 and F200 node hardware installation.

Before implementing a cluster in the data workflow, it is important to identify equipment and software requirements. Be sure to confirm that the facility is ready to support the cluster.

The information in this guide provides a framework for the research that a System Administrator or Facility Manager must conduct before powering on a node.

For detailed information about the OneFS operating system, review OneFS documentation on the Online Support site.

Scale-out NAS overview

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

The unified software platform provides centralized web-based and command-line administration to manage the following features:

- 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

OneFS storage architecture

F900, F600 and F200 nodes take a scale-out approach to storage by creating a cluster of nodes that runs a distributed file system. OneFS combines the three layers of storage architecture—file system, volume manager, and data protection—into a scale-out NAS cluster.

Each node adds resources to the cluster. Because each node contains globally coherent RAM, as a cluster becomes larger, it becomes faster. Meanwhile, the file system expands dynamically and redistributes content, which eliminates the work of partitioning disks and creating volumes.

Nodes work as peers to spread data across the cluster. Striping—the process of segmenting and distributing data—protects data and enables users connecting to any node to take advantage of the entire cluster's performance.

OneFS uses distributed software to scale data across commodity hardware. Master devices do not control the cluster, and secondary devices do not invoke dependencies. Each node helps to control data requests, boost performance, and expand cluster capacity.

Node components

F900, F600, and F200 nodes can be joined to existing clusters.

The *Generation 6 Site Preparation and Planning Guide* provides details on Generation 6, and earlier node types.

The F600 and F200 nodes are 1U models that require a minimum cluster size of three nodes. The F900 nodes are 2U models that require a minimum cluster size of four nodes. Clusters can be expanded to a maximum of 252 nodes in single node increments. The F900, F600, and F200 nodes are all-flash solutions, with software inline data compression, and data deduplication.

2

Selecting the equipment

The requirements for mixed-node clusters section in the *OneFS Supportability and Compatibility Guide* provides information on installing more than one type of node in a cluster.

Talk to a Sales Account Manager to identify the equipment that is best suited to support your workflow.

Topics:

Safety and EMI Compliance

This IT Equipment is compliant with the electromagnetic compatibility (EMC) and product safety regulations/standards required by the countries in which the product is sold. Compliance is based on FCC part 15, CISPR22/CISPR24 and EN55022/EN55024 standards, including applicable international variations.

EMC compliant Class A products are marketed for use in business, industrial, and commercial environments. Product Safety compliance is based on IEC 60950-1 and EN 60951-1 standards, including applicable national deviations.

This IT Equipment is in compliance with EU RoHS Directive 2011/65/EU.

The devices that are used in this product are approved under a unique regulatory model identifier. The regulatory model that is affixed to each device rating label may differ from any marketing or product family name in this data sheet.

For more information, go to Online Support. In the **Product and Support Tools** area, click **Safety & EMI Compliance Information**.

Grounding guidelines

To eliminate shock hazards and facilitate the operation of circuit-protective devices, ensure that the rack is grounded.

- The rack must have an earth-ground connection as required by applicable codes. Connections such as a grounding rod or building steel provide an earth ground.
- The electrical conduits must be made of rigid metallic material that is securely connected or bonded to panels and electrical boxes, which provides continuous grounding.
- The ground must have the correct low impedance to prevent buildup of voltage on equipment or exposed surfaces. Low-impedance grounding and lightning protection are recommended.
- The electrical system must meet local and national code requirements. Local codes might be more stringent than national codes. Site floor load-bearing requirements and Safety and EMI Compliance sections in this manual provide more information.

Shock and vibration

Products have been tested to withstand the shock and random vibration levels. The levels apply to all three axes and should be measured with an accelerometer on the equipment enclosures within the cabinet and shall not exceed:

Platform condition	Response measurement level
Non operational shock	10 G's, 7 ms duration
Operational shock	3 G's, 11 ms duration
Non operational random vibration	0.40 Grms, 5–500 Hz, 30 minutes
Operational random vibration	0.21 Grms, 5–500 Hz, 10 minutes

Systems that are mounted on an approved package have completed transportation testing to withstand the following shock and vibrations in the vertical direction only and shall not exceed:

Packaged system condition	Response measurement level
Transportation shock	10 G's, 12ms duration
Transportation random vibration	• 1.15 Grms • 1 hour Frequency range 1–200 Hz

Storage node specifications

To verify weight and dimensions, attributes, options, and industry certifications, review the F900, F600, and F200 node specifications in this section.

F900, F600 and F200 node specifications

In your environment, OneFS 9.2.0.0 and later is required for F900 nodes, and OneFS 9.0.0.0 and later is required for F600 and F200 nodes

See the *F900 Node Installation Guide*, or the *F200 and F600 Node Installation Guide*, or contact your account team for assistance with node installation and configuration.

NOTE: InfiniBand connections are supported for F900, F600, and F200 node back-end networking with OneFS 9.1.0.0 and later.

Table 1. F900 node attributes and options

Attribute	Capacity			
	1.92 TB SSD x24	3.84 TB SSD x24	7.68 TB SSD	15.36 TB SSD
Node capacity (raw)	46.08 TB	92.16 TB	184.32 TB	368.64 TB
Self-Encrypting Drives (SED, SSD) option	Yes	Yes	Yes	Yes
NVMe SSD Drives (2.5")	Per node = 24	Per node = 24	Per node = 24	Per node = 24
OneFS Version Required	OneFS 9.2.0.0 and later			
ECC Memory (per node)	736 GB (23x32GB Dual Rank DDR4 RDIMMs)			
Front-end networking	2 x 10 GbE (QSFP+), 2 x 25 GbE (QSFP+), 2 x 40 Gbe (SFP+), 2 x 100 GbE (SFP+)			
Network interfaces	Network interfaces support IEEE 802.3 standards for 1 Gbp/s, 10Gbp/s, 40 Gbp/s, 100Mbp/s and InfiniBand network connectivity			

Table 1. F900 node attributes and options (continued)

Attribute	Capacity			
	1.92 TB SSD x24	**3.84 TB SSD x24**	**7.68 TB SSD**	**15.36 TB SSD**
Drive Controller	24 x 2.5" NVMe SSD drives Samsung 1733 1.92, 3.84, 7.68, 15.36TB			
CPU type (per node)	Dual socket Intel® Processor			
Infrastructure networking (per node)	2 x 25 GbE (QSFP+), 2 x 100 GbE (SFP+), or Dual Port QSFP+ InfiniBand			
Typical Thermal Rating	3753.4 BTU per hour			

Table 2. F600 node attributes and options

Attribute	Capacity			
	1.92 TB SSD	**3.84 TB SSD**	**7.68 TB SSD**	**15.36 TB SSD**
Node capacity (raw)	15.36 TB	30.72 TB	61.44 TB	122.88 TB
Self-Encrypting Drives (SED, SSD) option	Yes	Yes	Yes	Yes
NVMe SSD Drives (2.5")	Per node = 8	Per node = 8	Per node = 8	Per node = 8
OneFS Version Required	OneFS 9.0.0.0 and later			
ECC Memory (per node)	128 GB, 192 GB, 384 GB			
Front-end networking	2 x 10 GbE (SFP+) or 2 x 25 GbE (QSFP+) or 2 x 100 GbE (SFP+)			
Network interfaces	Network interfaces support IEEE 802.3 standards for 1 Gbp/s, 10Gbp/s, 40 Gbp/s, 100Mbp/s and InfiniBand network connectivity			
Drive Controller	8 x 2.5" NVMe SSD drives			
CPU type (per node)	Dual socket Intel® Processor			
Infrastructure networking (per node)	2 x 25 GbE (QSFP+), 2 x 100 GbE (SFP+), or Dual Port QSFP+ InfiniBand			
Typical Thermal Rating	1593.5 BTU per hour			

Table 3. F200 node attributes and options

Attribute	Capacity		
	960 GB SSD	**1.92 TB SSD**	**3.84 TB SSD**
Node capacity (Raw)	3.84 TB	7.68 TB	15.36 TB
Self-Encrypting Drives (SED, SSD) option	No	No	No
SSD drives (2.5 in) (per node)	Per node = 4	Per node = 4	Per node = 4
OneFS Version Required	OneFS 9.0.0.0 and later		
ECC Memory (per node)	48 GB or 96 GB		
Front-end Networking (per node)	2 x 10 GbE (SFP+) or 2 x 25 GbE (QSFP+)		
Network interfaces	Network interfaces support IEEE 802.3 standards for 1 Gbp/s, 10Gbp/s, 40 Gbp/s, 100Mbp/s and InfiniBand network connectivity		
Drive Controller	SAS-3 12 Gb/s		

Table 3. F200 node attributes and options (continued)

Attribute	Capacity		
	960 GB SSD	1.92 TB SSD	3.84 TB SSD
CPU Type	Single socket Intel® Processor		
Infrastructure Networking (per node)	2 x 10 GbE, 2 x 25 GbE (QSFP+), or Dual Port QSFP+ InfiniBand		
Typical Thermal Rating	815.5 BTU per hour		

Power consumption

Table 4. Power consumption values

Attribute	F900 value (Watts)	F600 value (Watts)	F200 value (Watts)
Typical power consumption at 25° Celsius	750 at 240v	467 at 240v	239 at 240v
Maximum power N+1 (per node)	718.6 at 240v	718.6 at 240v	394.6 at 240v

Table 5. Cluster attributes

Attribute	F900 cluster	F600 cluster	F200 cluster
Number of nodes	3-252	3–252	3–252
Cluster capacity (raw) **NOTE:** Usable capacity is slightly lower than the raw capacity reflected in this table.		3.8 PB to 15 PB	968 TB to 3.8 PB
Memory		77 TB to 26.5 PB	184 TB up to 79.6 PB
Rack Units	3-252	3–252	3–252

Racks and rails

You can secure F900, F600, and F200 nodes to standard storage racks with a sliding rail system.

Rail kits are in all node packaging and are compatible with racks with the following types of holes:

- 3/8 inch square holes
- 9/32-inch round holes
- 10-32, 12-24, M5X.8, or M6X1 prethreaded holes

Rail kit mounting brackets adjust in length from 24 to 36 inches to accommodate different rack depths. When you select a rack, ensure that the rack supports the minimum and maximum rail kit sizes for the nodes.

Titan S and Titan HD rail kits can support NEMA spacing 24 or 29 inches.

Table 6. Rack solutions per node type

Node type	Rack type
• F900 • F600 • F200	The following racks with adequate cable, door (if exists), and PDU (if exists) clearance: • Titan S • Titan SS • Titan D • Titan HD

Table 7. Rack dimensions

Rack type	Dimensions (front bezel to rear door)
Titan	24in W x 39in D x 40U H
Titan D	24in W x 44in D x 40U H
Titan S	600 mm W x 1147 mm D x 42U H
Titan SS	600 mm W x 997 mm D x 42U H
Titan HD	28in W x 48in D x 42U H
NOTE: The depth that is listed in the table is from the front bezel to the rear door. More depth (approximately 50 mm to 60 mm or 2 inches to 2.5 inches) is required from the front bezel to the rear door.	

Rail kit components for 2U systems

The sliding rail assemblies are used to secure the node in the cabinet, and extended from the cabinet so that the system cover can be removed to access the internal FRUs. The sliding rail assembly (2U) is used for installation of the F900 nodes.

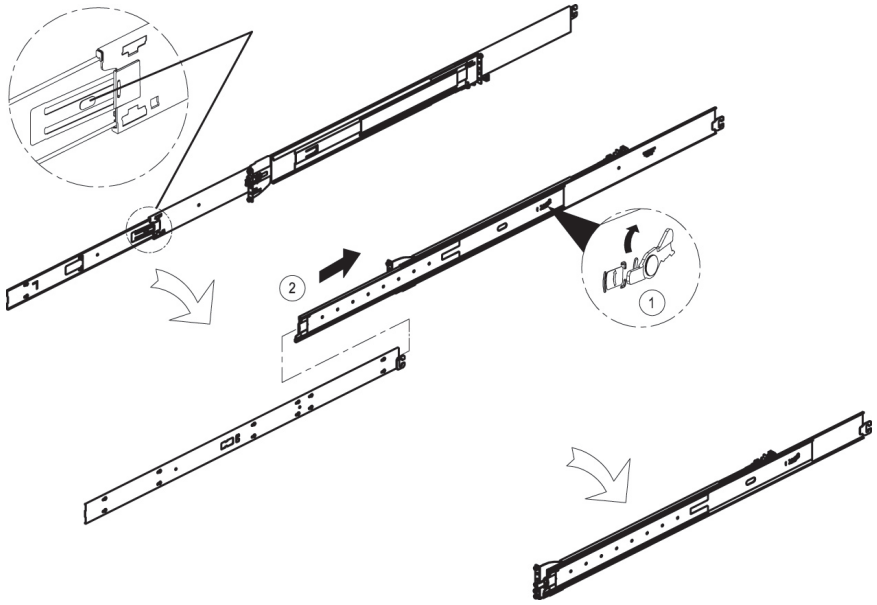

Figure 1. Sliding rail assembly - 2U systems

Rail kit components for 1U systems

The sliding rail assemblies are used to secure the node in the cabinet, and extended from the cabinet so that the system cover can be removed to access the internal FRUs. The A10 sliding rail assembly (1U) is used for installation of the F200 and F600 nodes.

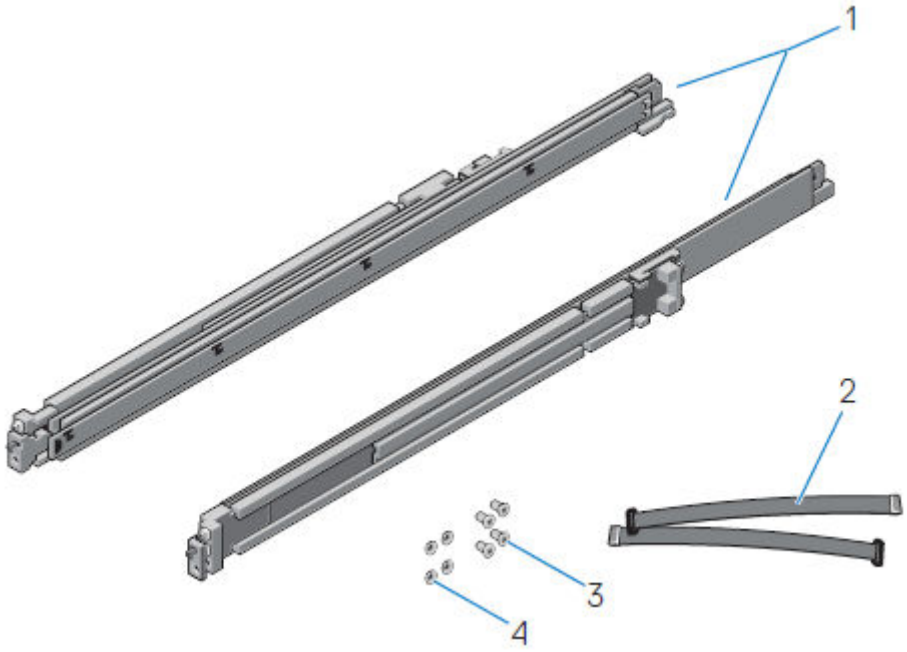

Figure 2. A10 sliding rail assembly - 1U systems

1. A10 sliding rail (2)
2. velcro strap (2)
3. screw (4)
4. washer (4)

Connect and route cords and cables

Steps

1. Connect the power cables and I/O cables as described in documentation for your system.
2. If the system uses a cable management arm (CMA), install it as described in the document that is shipped with the CMA.
3. If the system does not use a CMA, use the two velcro straps to route and secure cords and cables at the rear of the system:
   a. Locate the CMA bracket slots on the rear end of both the rails.
   b. Bundle the cables gently, pulling them clear of the system connectors to the left and right sides.

      NOTE: Ensure that there is enough space for the cables to move when you slide the system out of the rack.

   c. Thread the straps through the CMA bracket slots on each side of the system to hold the cable bundles.

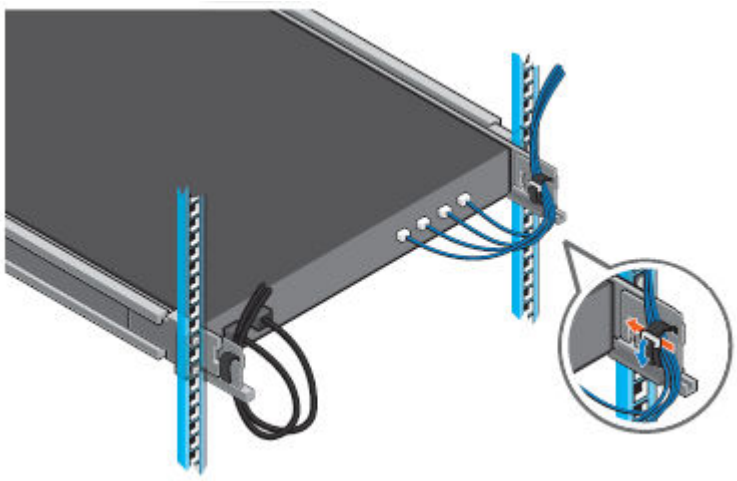

Figure 3. CMA bracket slots

Node ports

About this task

The back-end ports are the private network connections to the nodes. Port 1 from all nodes connects to one switch, and port 2 from all the nodes connects to a second switch. Both back-end switches are provided.

The front-end ports are for the client network connections.

NOTE: In the F900 and F600 nodes, the rNDC does not provide network connectivity. In the F200, the rNDC can provide 10 GbE or 25 GbE connections for front-end networking.

Network topology

External networks connect the cluster to the outside world.

NOTE: Do not use the 1 GbE ports on nodes for data services such as AV scanning, replication, data access, and so on. Use of the 1 GbE ports often results in errors and issues.

Subnets can be used in external networks to manage connections more efficiently. Specify the external network subnets depending on the topology of the network.

In a basic network topology in which each node communicates to clients on the same subnet, only one external subnet is required.

More complex topologies require several different external network subnets. For example, nodes that connect to one external IP subnet, nodes that connect to a second IP subnet, and nodes that do not connect externally. Configure the default external IP subnet by using IPv4.

External networks provide communication outside the cluster. OneFS supports network subnets, IP address pools, and network provisioning rules to facilitate the configuration of external networks.

The internal network supports communication among the nodes that form a cluster and is intentionally separate from the external, front-end network. The back-end, internal network is InfiniBand-based for pre-Generation 6, Generation 6 is based on 10 GbE or 40 GbE, and InfiniBand is optional for older clusters. The F900, F600 and F200 node specifications section in this guide provides details on front-end and back-end networking options.

NOTE: InfiniBand is supported for F900, F600, and F200 nodes on OneFS 9.1.0.0 and later.

To configure the cluster, set up an initial network. Optionally, set up an alternate interface as a failover network.

CAUTION: Information that is exchanged on the back-end network is not encrypted. Connecting third-party nodes to the back-end switch creates a security risk.

Leaf-Spine topology

OneFS 9.0.0.0 and later releases support Leaf-Spine network topology for internal networks that communicate with the nodes that form clusters up to 252 nodes. For large clusters that are intended to grow significantly over time, the Leaf-Spine topology is recommended.

NOTE: To connect 252 nodes on a OneFS 8.2.2 Leaf-Spine cluster, you must update the cluster with the latest OneFS 8.2.2 rollup patch. For more information, see the Current OneFS patches guide.

Architecture

In a Leaf-Spine topology, Z9100-ON switches are arranged in a two-level hierarchy. The bottom level switches with the nodes connected are called Leaf switches. Leaf switches are connected to the top level switches called Spine switches. OneFS requires two Leaf-Spine networks for intra-cluster traffic.

The following table lists the main Leaf-Spine components in a cluster.

Table 8. Leaf-Spine network components

Component	Description	Connection considerations
Spine	Z9100-ON 32-port switch	Back-end network with 100GbE (uplink) connects to the Leaf switch
Leaf	Z9100-ON 32-port switch	• 22 performance nodes or 88 archive nodes with breakout cables for downlinks from the Leaf switch • Up to 10 x 100GbE uplinks from the Leaf switch
Performance - 40 GbE back-end node	F600, F800, H5600, H500, and H600 Isilon nodes	• Performance nodes support a 40GbE connection to the Leaf • F600 nodes support a 40GbE or 100GbE connection to the Leaf switch
Archive - 10 GbE back-end node	F200, A200, A2000, and H400 Isilon nodes	• Archive nodes support a 10GbE connection to the Leaf. Use a breakout cable to connect up to four archive nodes to a single Leaf port • F200 nodes support a 10GbE or 25GbE connection to the Leaf switch
Breakout cable	40 GbE cable that breaks out to four 10 GbE cables	Connects the 40GbE Leaf switch to four 10GbE archive nodes
Uplink	Leaf to Spine connection	There must be the same number of uplinks on every Leaf switch, and that number should be the number of uplinks required by the Leaf switch with the most downlinks
Downlink	Leaf to node connection	

Assisting with installation

Contact an Account Manager for help with planning the best workflow for your environment.

Installation and implementation details

Three to four weeks before the installation date, the Professional Services team helps to gather the information necessary to configure the cluster.

The project team helps to complete the *Configuration Guide* worksheet, which documents technical details that are needed for the installation and implementation of the cluster.

Be prepared to discuss the following information with the project team:

- Data workflow, Including the role of the cluster in that workflow: including, production, test, or disaster recovery.
- The OneFS version to install on the cluster.
- Network connectivity details, including IP ranges, for the client and networks.
- The DNS configuration details, including name servers and search lists.
- The directory services such as Active Directory, LDAP, NIS, or local user groups.
- File sharing protocols such as SMB and NFS, and advanced file sharing options such as FTP and HTTP.
- The data protection levels, anti-virus solutions, and NDMP backup.
- Cluster alert solutions such as SupportIQ and SNMP monitoring.

Switches and cables

Select network switches and cables that are compatible with F900, F600, and F200 nodes that support the network topology.

A complete list of qualified switches and cables is in the *OneFS Supportability and Compatibility Guide*.

Front-end switches must meet the following minimum specifications:

- OneFS 9.2.0.0 and later includes front-end support for:
- F900: 10 GbE, 25 GbE, 40 GbE 100 GbE
- OneFS 9.0.0.0 and later includes front-end support for:
  - F600: 10 GbE, 25 GbE, 100 GbE
  - F200: 10 GbE, 25 GbE
- Nonblocking fabric switch
- Minimum of 1 MB per port of packet buffer memory
- Support for jumbo frames
- Separate switches are required for the front-end and back-end interfaces.

InfiniBand is supported for the back-end (internal) traffic on F900, F600 and F200 nodes for existing clusters only.

CAUTION: Information that is exchanged on the back-end network is not encrypted. Connecting third-party nodes to the back-end switch creates a security risk.

Cable management

To protect the cable connections, organize cables for proper airflow around the cluster, and to ensure fault-free maintenance of the nodes.

Protect cables

NOTE: 1GbE management interface on Generation 6 hardware is designed to handle SSH traffic only.

Damage to the cables can affect the cluster performance. Consider the following to protect cables and cluster integrity:

- Never bend cables beyond the recommended bend radius. The recommended bend radius for any cable is at least 10–12 times the diameter of the cable. For example, if a cable is 1.6 inches, round up to 2 inches and multiply by 10 for an acceptable bend radius. Cables differ, so follow the recommendations of the cable manufacturer.
- As illustrated in the following figure, the most important design attribute for bend radius consideration is the minimum mated cable clearance (Mmcc). Mmcc is the distance from the bulkhead of the chassis through the mated connectors/strain relief including the depth of the associated 90 degree bend. Multimode fiber has many modes of light (fiber optic) traveling through the core. As each of these modes moves closer to the edge of the core, light and the signal are more likely to be reduced, especially if the cable is bent. In a traditional multimode cable, as the bend radius is decreased, the amount of light that leaks out of the core increases, and the signal decreases.

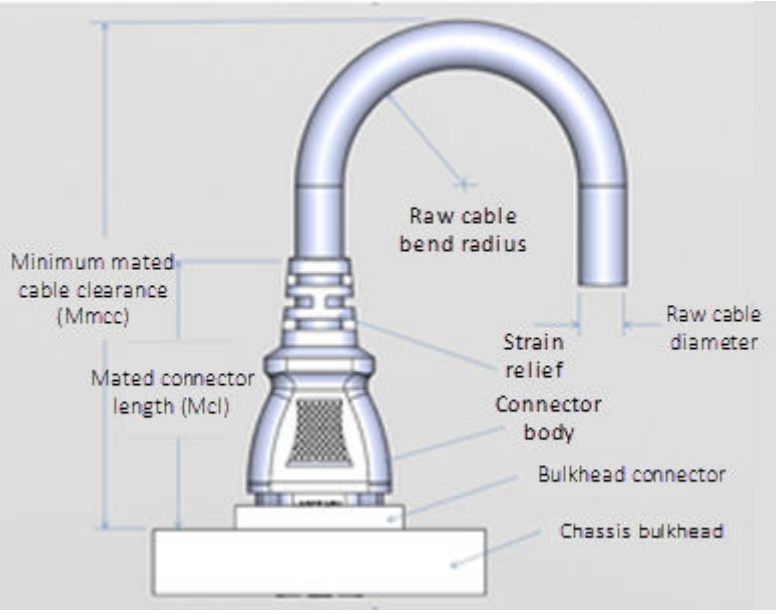

Figure 4. Cable design

- Keep cables away from sharp edges or metal corners.
- Never bundle network cables with power cables. If network and power cables are not bundled separately, electromagnetic interference (EMI) can affect the data stream.
- When bundling cables, do not pinch or constrict the cables.
- Avoid using zip ties to bundle cables, instead use velcro hook-and-loop ties that do not have hard edges, and can be removed without cutting. Fastening cables with velcro ties also reduces the impact of gravity on the bend radius.

  NOTE: Gravity decreases the bend radius and results in the loss of light (fiber optic), signal power, and quality.
- For overhead cable supports:
  - Ensure that the supports are anchored adequately to withstand the significant weight of bundled cables. Anchor cables to the overhead supports, then again to the rack to add a second point of support.
  - Do not let cables sag through gaps in the supports. Gravity can stretch and damage cables over time. You can anchor cables to the rack with velcro ties at the mid-point of the cables to protect your cable bundles from sagging.
  - Place drop points in the supports that allow cables to reach racks without bending or pulling.
- If the cable is running from overhead supports or from underneath a raised floor, be sure to include vertical distances when calculating necessary cable lengths.

Ensure airflow

Bundled cables can obstruct the movement of conditioned air around the cluster.

- Secure cables away from fans.
- To keep conditioned air from escaping through cable holes, employ flooring seals or grommets.

Prepare for maintenance

To accommodate future work on the cluster, design the cable infrastructure. Think ahead to required tasks on the cluster, such as locating specific pathways or connections, isolating a network fault, or adding and removing nodes and switches.

- Label both ends of every cable to denote the node or switch to which it should connect.

- Leave a service loop of cable behind nodes. Service technicians should be able to slide a node out of the rack without pulling on power or network connections. In the case of Generation 6 nodes, you should be able to slide any of the four nodes out of the chassis without disconnecting any cables from the other three nodes.

  WARNING: If adequate service loops are not included during installation, downtime might be required to add service loops later.
- Allow for future expansion without the need for tearing down portions of the cluster.

Optical cabling to a switch or server

Optical cables connect the small form-factor pluggable (SFP+ and QSFP+) modules on the storage processors (SPs) to an external Fibre Channel, 40GbE or 10GbE environment. It is strongly recommended that you use OM3, and OM4 50 µm cables for all optical connections.

Table 9. Optical cables

Cable type	Operating speed	Length
50 µm	1.0625 Gb	2 m (6.6 ft) minimum to 500 m (1,650 ft) maximum
	2.125 Gb	2 m (6.6 ft) minimum to 300 m (990 ft) maximum
	4 Gb	2 m (6.6 ft) minimum to 150 m (492 ft) maximum
	8 Gb	OM3: 1 m (3.3 ft) minimum to 150 m (492 ft) maximum
		OM2: 1 m (3.3 ft) minimum to 50 m (165 ft) maximum
	10 Gb	OM3: 1 m (3.3 ft) minimum to 300 m (990 ft) maximum
		OM2: 1 m (3.3 ft) minimum to 82 m (270 ft) maximum
62.5 µm	1.0625 Gb	2 m (6.6 ft) minimum to 300 m (985 ft) maximum
	2.125 Gb	2 m (6.6 ft) minimum to 150 m (492 ft) maximum
	4 Gb	2 m (6.6 ft) minimum to 70 m (231 ft) maximum
Note: Dual LC for 10GbE cables have a bend radius of 3 cm (1.2 in) minimum. You can obtain MPO connector ends for optical 40GbE cables.		

The maximum length that is listed in the preceding table for the 50 µm or 62.5 µm optical cables includes two connections or splices between the source and the destination.

NOTE: It is not recommended to mix 62.5 µm and 50 µm optical cables in the same link. In certain situations, you can add a 50 µm adapter cable to the end of an already installed 62.5 µm cable plant. Contact the service representative for details.

Supported switches

Switches ship with the proper rails or tray to install the switch in the rack.

The following internal network switches ship with rails to install the switch. The switch rails are adjustable to fit NEMA front rail to rear rail spacing ranging from 22 in to 34 in.

Table 10. Z9100-ON Ethernet switch

Switch	Maximum number of ports	Network
Z9100-ON	128-port	32x100 GbE, 32x40 GbE, 128x10 GbE (with breakout cables)
The Z9100-ON is a fixed 1U Ethernet switch which can accommodate high port density (lower and upper RUs) and multiple interface types (32 ports of 100 GbE or 40 GbE in QSFP28 or 128 ports of 25 GbE or 10 GbE with breakout) for maximum flexibility.		

NOTE: In OneFS 8.2.0 and later, the Z9100-ON switch is required for Leaf-Spine networking of large clusters.

Table 11. Z9264F-ON Ethernet switch

Switch	Maximum number of ports	Network
Z9264F-ON	128-port	64x100 GbE, 64x40 GbE, 256x10 GbE (with breakout cables)
The Z9264F-ON is a fixed 2U Ethernet switch which provides industry-leading density of either 64 ports of 100GbE or 40 GbE in QSFP28 or 128 ports of 25 GbE or 10 GbE by breakout. Breakout cables are only used in the odd-numbered ports and using one in odd-numbered port disables the corresponding even-numbered port.		

Table 12. S4148F-ON Ethernet switch

Switch	Maximum number of ports	Network
S4148F-ON	48-port	2x40 GbE 48x10 GbE
The S4148F-ON is the next generation family of 10 GbE (48 ports) top-of-rack, aggregation-switch, or router products that aggregates 10 GbE server or storage devices and provides multi speed uplinks for maximum flexibility and simple management.		

Table 13. S4112F-ON Ethernet switch

Switch	Maximum number of ports	Network
S4112F-ON	12-port	3x100 GbE (with breakout, connect 12x10 GbE nodes using the 3x100 GbE ports) 12 x10GbE
The S4112F-ON supports 10/100GbE with 12 fixed SFP+ ports to implement 10 GbE and three fixed QSFP28 ports to implement 4x10 or 4x25 using breakout. A total of 24 10 GbE connections including the three fixed QSFP28 ports using 4x10 breakout cables.		

Table 14. Infiniband switches

Switch	Ports	Network
Mellanox Neptune MSX6790	36-port	QDR InfiniBand
Mellanox SX6018	18-port	
Mellanox MIS5022Q-1BFR	8-port	
Mellanox MSX6506-NR	108-port	QDR chassis switch
Mellanox MSX6512-NR	216-port	QDR chassis switch

Installing rails for an F900 rack

Install all rails for nodes and switches.

Install and secure rails in the following order:

1. Install all rails for F900 nodes.
2. Install all rails for front and back-end switches.

Do not install any nodes or switches until all rails are installed in the rack.

Install rails for F900 nodes

Install all F900 rails and secure them to the rack.

About this task

NOTE: Rails are labeled as Left and Right and must be installed on the correct side of the rack. Left and Right are based on your point of reference as you look at the front side of the rack.

Steps

1. Take out the inner member and slide the immediate member back.
   a. Press and take out the inner member.
   b. Press down according to the arrow's direction, and slide the intermediate member back.

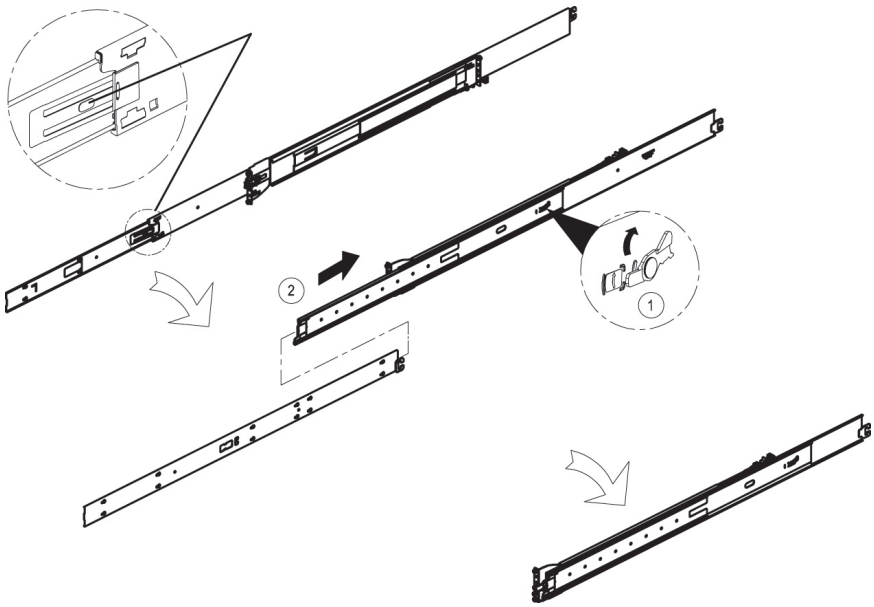

2. Install the inner member onto the chassis and secure it with the screw.
   Pay attention to the installation direction.

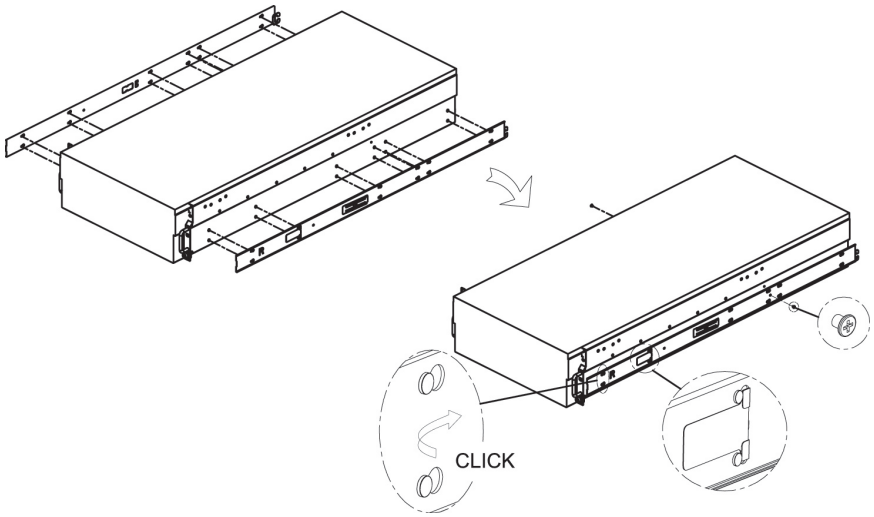

3. Secure the outer member and bracket onto the rack with the screws for both the left and right sides.

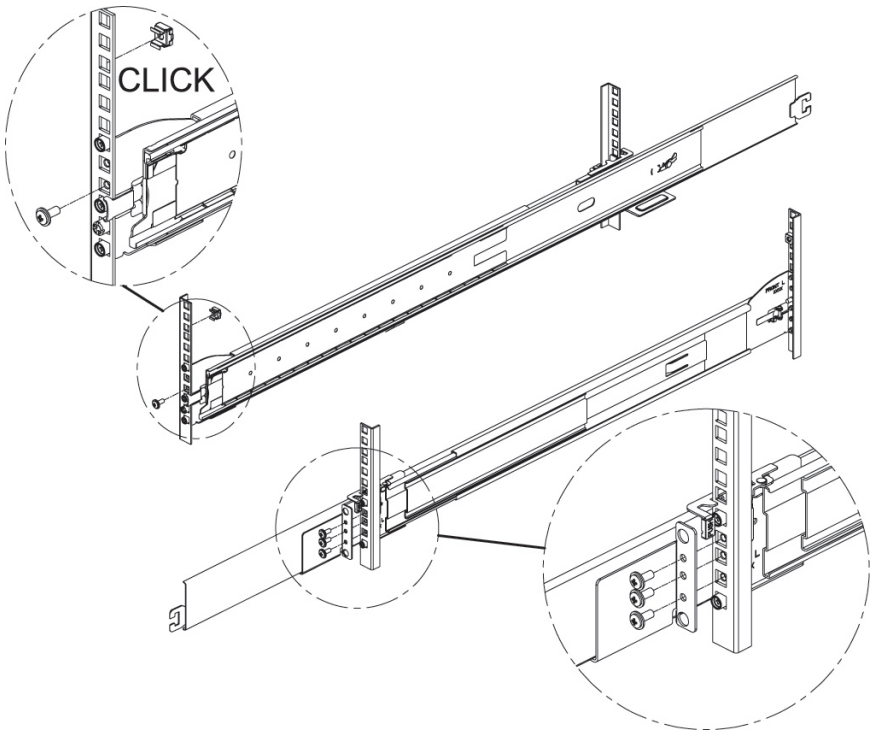

4. Secure the outer rail and L-brackets use the "L" and "R" labeling when viewing from the rear of the rack with three screws for both the left and right sides.

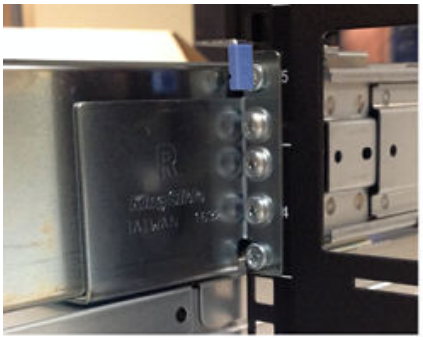

5. Repeat all steps to install rails for all F900 nodes.

Installing the switch

About this task

The switches ship with the proper rails or tray to install the switch in the rack.

NOTE: If the installation instructions in this section do not apply to the switch you are using, follow the procedures provided by your switch manufacturer.

CAUTION: If the switch you are installing features power connectors on the front of the switch, it is important to leave space between appliances to run power cables to the back of your rack. There is no 0U cable management option available at this time.

Steps

1. Remove rails and hardware from packaging.
2. Verify that all components are included.
3. Locate the inner and outer rails and secure the inner rail to the outer rail.
4. Attach the rails assembly to the rack using the eight screws as illustrated in the following figure.

   NOTE: The rail assembly is adjustable for NEMA, front to rear spacing extends from 22 in to 34 in.

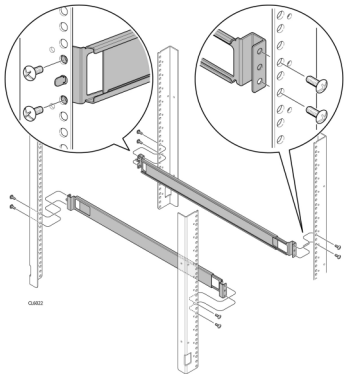

Figure 5. Install the inner rail to the outer rail

5. Attach the switch rails to the switch by placing the larger side of the mounting holes on the inner rail over the shoulder studs on the switch. Press the rail even against the switch.

   NOTE: The orientation of the rail tabs for the front NEMA rail are located on the power supply side of the switch.

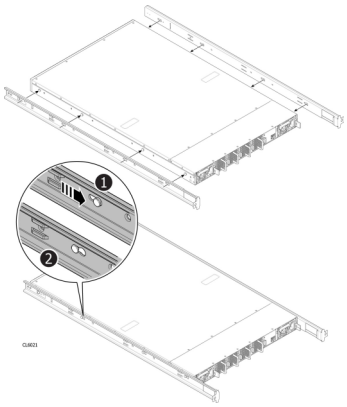

Figure 6. Install the inner rails

6. Slide the inner rail towards the rear of the switch slide into the smaller side of each of the mounting holes on the inner rail. Ensure the inner rail is firmly in place.
7. Secure the switch to the rail, securing the bezel clip and switch to the rack using the two screws as illustrated in the following figure.

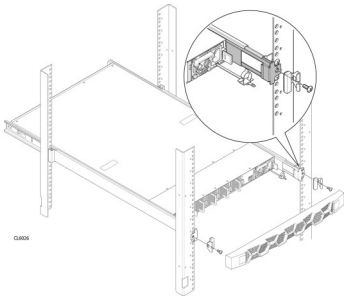

Figure 7. Secure the switch to the rail

8. Snap the bezel in place.

3

Adding functionality to the cluster

Advanced cluster features can be obtained through OneFS software modules.

To enable a OneFS module after the cluster is installed, activate a license by entering a license key in OneFS.

More information about features that are offered through optional software modules, and licensing is available in the *OneFS web Administration Guide* or the *OneFSCLI Administration Guide*, or contact a sales representative.

Topics:

Data management modules

Isilon offers software modules that add advanced data management features to your cluster.

You can install advanced data management modules to optimize storage performance.

SmartQuotas

The SmartQuotas module is a quota-management tool that monitors and enforces administrator-defined storage limits.

Through the use of accounting and enforcement quota limits, reporting capabilities, and automated notifications, you can manage and monitor storage utilization, monitor disk storage, and issue alerts when storage limits are exceeded.

A storage quota defines the boundaries of storage capacity that are allowed for a group, a user, or a directory on a cluster. The SmartQuotas module can provision, monitor, and report disk-storage usage and can send automated notifications when storage limits are approached or exceeded. SmartQuotas also provides flexible reporting options that can help you analyze data usage.

SmartDedupe

The SmartDedupe software module enables you to save storage space on your cluster by reducing redundant data. Deduplication maximizes the efficiency of your cluster by decreasing the amount of storage required to store multiple files with similar blocks.

SmartDedupe deduplicates data by scanning an Isilon cluster for identical data blocks. Each block is 8 KB. If SmartDedupe finds duplicate blocks, SmartDedupe moves a single copy of the blocks to a hidden file called a shadow store. SmartDedupe then deletes the duplicate blocks from the original files and replaces the blocks with pointers to the shadow store.

Deduplication is applied at the directory level, targeting all files and directories underneath one or more root directories. You can first assess a directory for deduplication and determine the estimated amount of space you can expect to save. You can then decide whether to deduplicate the directory. After you begin deduplicating a directory, you can monitor how much space is saved by deduplication in real time.

You can deduplicate data only if you activate a SmartDedupe license on a cluster. However, you can assess deduplication savings without activating a SmartDedupe license.

4

Preparing your facility

To ensure an optimal data center, and the long-term health of the cluster, prepare and maintain the environment as described in this section.

Topics:

Environmental requirements

Prepare the site to meet the required parameters.

Environmental item	Description
	10°C to 35°C (50°F to 95°F) with no direct sunlight on the equipment. **NOTE:** There are four independent cooling zones so that a cooling failure in one node, does not impact other nodes.
	5% to 95% percent relative humidity with 33°C (91°F) maximum dew point. Atmosphere must be non-condensing at all times. **NOTE:** The cluster can be qualified to operate outside these limits. Product-specific documentation for system specifications provides more information.
	F600 and F200, 21.9 kg (48.28 lbs) F900, 26.3 kg (57.98 lb)
	0 meters to 2439 meters (0 to 8,000 ft) above sea level operating altitude.

The cluster might be qualified to operate outside of these limits. Refer to the product-specific documentation for system specifications.

Power requirements

Depending on the cabinet configuration and input AC power source, single or three-phase listed in the Single-phase power connection requirements and Three-phase power connection requirements. The cabinet requires between two and six independent power sources. To determine the site requirements, use the published technical specifications and device rating labels to provide the current draw of the devices in each rack. The total current draw for each rack can then be calculated by using the Power Calculator.

Table 15. Single-phase power connection requirements

Specification	North American 3 wire connection (2 L and 1 G)[a]	International and Australian 3 wire connection (1 L, 1 N, and 1 G)
Input nominal voltage	200–240 V ac +/- 10% L - L nom	220–240 V ac +/- 10% L - L nom
Frequency	50–60 Hz	50–60 Hz
Circuit breakers	30 A	32 A

Table 15. Single-phase power connection requirements (continued)

Specification	North American 3 wire connection (2 L and 1 G)[a]	International and Australian 3 wire connection (1 L, 1 N, and 1 G)
Power zones	Two	Two
Power requirements at site (minimum to maximum)	• Single-phase: six 30A drops, two per zone • Three-phase Delta: two 50A drops, one per zone • Three-phase Wye: two 32A drops, one per zone **NOTE:** The options for the single-phase PDU interface connector are listed in the following table, Single-phase AC power input connector options.	

a. L = line phase, N = neutral, G = ground

Table 16. Single-phase AC power input connector options

Single-phase rack connector options	Customer AC source interface receptacle	Site
 NEMA L6-30P	 NEMA L6-30R	North America and Japan
 Russellstoll 3750DP	 Russellstoll 9C33U0	North America and Japan
 IEC-309 332P6	 IEC-309 332C6	International
 CLIPSAL 56PA332	 CLIPSAL 56CSC332	Australia

Table 17. Three-phase AC power connection requirements

Specification	North American (Delta) 4 wire connection (3 L and 1 G)[a]	International and Australian (Wye) 5 wire connection (3 L, 1 N, and 1 G)
Input nominal voltage	200–240 V ac +/- 10% L - L nom	220–240 V ac +/- 10% L - N nom
Frequency	50–60 Hz	50–60 Hz
Circuit breakers	50 A	32 A
Power zones	Two	Two
Power requirements at site (minimum to maximum)	North America (Delta): • One to two 50 A, three-phase drops per zone • Each rack requires a minimum of two drops to a maximum of four drops. The system configuration and the power requirement for that configuration determine the number of drops. International (Wye): • One 32 A, three-phase drop per zone • Each Wye rack requires two 32 A drops.	

Table 17. Three-phase AC power connection requirements (continued)

Specification	North American (Delta) 4 wire connection (3 L and 1 G)[a]	International and Australian (Wye) 5 wire connection (3 L, 1 N, and 1 G)
	NOTE: The interface connector options for the Delta and Wye three-phase PDUs are listed in the following tables, and .	

a. L = line phase, N = neutral, G = ground

Table 18. Three-phase Delta-type AC power input connector options

Three-phase Delta rack connector options	Customer AC source interface receptacle	Site
Russellstoll 9P54U2 	Russellstoll 9C54U2 	North America and International
Hubbell CS-8365C 	Hubbell CS-8364C 	North America

Table 19. Three-phase Wye-type AC power input connector options

Three-phase Wye rack connector options	Customer AC source interface receptacle	Site
GARO P432-6 	GARO S432-6 	International

Node power values

Table 20. F900, F600 and F200 node power values

	F900	F600	F200
Input power	859 watts 2931 btu p/hr	438 watts 1494.5 btu p/hr	221 watts 754.1 btu p/hr
Input current	3.9 amps	2 amps	1 amps
Sound power	8.1 bels	7.7 bels	6.8 bels
Air temperature rise	15.5 C 27.9 F	16.4 C 29.5 F	131 C 23.6 F
Max potential power	1297.1 watts 4425.9 btu p/hr	632.3 watts 2157.5 btu p/hr	385.1 watts 1314 btu p/hr
Airflow	99.5 CFM 47 l/s	48 CFM 22.7 l/s	30.5 CFM 14.4 l/s
Weight	33.1 Kg 73 lbs	21.9 Kg 48.3 lbs	21.9 Kg 48.3 lbs
Power supply capacity	1100 watts 3753.4 btu p/hr	750 watts 2559.1 btu p/hr	n/a

Equipment power requirements

To support the recommended power parameters of equipment, prepare the site.

Plan to set up redundant power for each rack that contains F200 and F600 nodes. Supply the power with a minimum of two separate circuits on the electrical system. If one of the circuits fails, one or more remaining circuits could handle the full power load of the rack.

- Power each power distribution panel (PDP) within the rack with a separate power circuit.
- Power the two IEC 60320 C14 power input connectors in the node, by separate PDPs within the rack.

When calculating the power requirements for circuits that supply power to the rack, consider the power requirements for network switches and for nodes.

Each circuit should be rated appropriately for the node types and input voltage. Refer to product specifications for power requirements specific to each node type.

CAUTION: If a node loses power, the NVRAM battery sustains the cluster journal on the NVRAM card for five days. If you do not restore power to the node after five days, it is possible that you will lose data.

Radio Frequency Interference (RFI) requirements

Electromagnetic fields that include radio frequencies can interfere with the operation of electronic equipment.

The hardware is certified to withstand radio frequency interference in accordance with standard EN61000-4-3. In data centers that employ intentional radiators, such as cell phone repeaters, the maximum ambient RF field strength should not exceed 3 Volts/meter.

Take field measurements at multiple points close to the equipment. Consult with an expert before you install any emitting device in the data center. If you suspect high levels of RFI, contract an environmental consultant to evaluate RFI field strength and address mitigation efforts.

The ambient RFI field strength is inversely proportional to the distance and power level of the emitting device. To determine if the cell phone repeater or other intentional radiator device is at a safe distance from the equipment, use the following table as a guide.

Table 21. Minimum recommended distance from RF emitting device

Repeater power level*	Recommended minimum distance
1 Watt	3 meters
2 Watt	4 meters
5 Watt	6 meters
7 Watt	7 meters
10 Watt	8 meters
12 Watt	9 meters
15 Watt	10 meters

* Effective Radiated Power, ERP

Hardware acclimation

Systems and components must acclimate to the operating environment before power is applied to them. Once unpackaged, the system must reside in the operating environment for up to 16 hours to thermally stabilize and prevent condensation.

If the last 24 hours of the TRANSIT/STORAGE environment was:		...and the OPERATING environment is:	...then let the system or component acclimate in the new environment this many hours:
Temperature	Relative Humidity	Temperature	Acclimation Time

If the last 24 hours of the TRANSIT/STORAGE environment was:		...and the OPERATING environment is:	...then let the system or component acclimate in the new environment this many hours:
Nominal 68–72°F (20–22°C)	Nominal 40–55% RH	Nominal 68–72°F (20–22°C) 40–55% RH	1 hour or less
Cold <68°F (20°C)	Dry <30% RH	<86°F (30°C)	4 hours
	Damp >30% RH		
Hot >72°F (22°C)	Dry <30% RH		
	Humid 30–45% RH		
Hot >72°F (22°C)	Humid 45–60% RH	<86°F (30°C)	8 hours
Hot >72°F (22°C)	Humid >60% RH	<86°F (30°C)	16 hours
Unknown			

IMPORTANT:

- If there are signs of condensation after the recommended acclimation time has passed, allow an additional eight (8) hours to stabilize.
- System components must not experience changes in temperature and humidity that are likely to cause condensation to form on or in that system or component. Do not exceed the shipping and storage temperature gradient of 45°F/hr (25°C/hr).
- To facilitate environmental stabilization, open both front and rear cabinet doors.

Air quality requirements

F200 and F600 nodes are designed to be consistent with the air quality requirements and thermal guidelines of the American Society of Heating, Refrigeration and Air Conditioning Engineers (ASHRAE).

For specifics, see the ASHRAE *Environmental Standard Handbook* and the most current revision of *Thermal Guidelines for Data Processing Environments, Second Edition, ASHRAE 2009b*.

Most products are best suited for Class 1 datacom environments, which consist of tightly controlled environmental parameters including temperature, dew point, relative humidity and air quality. These facilities house mission-critical equipment and are typically fault-tolerant, including the air conditioners.

The data center should maintain a cleanliness level as identified in ISO 14664-1, class 8 for particulate dust and pollution control. The air entering the data center should be filtered with a MERV 11 filter or better. The air within the data center should be continuously filtered with a MERV 8 or better filtration system. Take measures to prevent conductive particles such as zinc whiskers from entering the facility.

The allowable relative humidity level is 20% to 80% non condensing. However, the recommended operating environment range is 40% to 55%. Lower temperatures and humidity minimize the risk of hardware corrosion and degradation, especially in data centers with gaseous contamination such as high sulfur content. Minimize humidity fluctuations within the data center. Prevent outside air contaminants and humidity from entering the facility by positively pressurizing the data center and installing air curtains on entryways.

For facilities below 40% relative humidity, use grounding straps when contacting the equipment to avoid the risk of Electrostatic discharge (ESD), which can harm electronic equipment.

As part of an ongoing monitoring process for the corrosiveness of the environment, place copper and silver coupons (per ISA 71.04-1985, Section 6.1 Reactivity) in airstreams representative of those in the data center. The monthly reactivity rate of the coupons should be less than 300 Angstroms. If the monitored reactivity rate exceeds 300 Angstroms, analyze the coupon for material species, and put a corrective mitigation process in place.

Site floor load-bearing requirements

Install the cabinet in raised or non-raised floor environments capable of supporting at least 1,495kg (3300 lbs) per cabinet. The system may weigh less, but requires extra floor support margin to accommodate equipment upgrades and/or reconfiguration.

In a raised floor environment:

- 24 x 24 inches or (60 x 60 cm) heavy-duty, concrete filled steel floor tiles are recommended.
- Use only floor tiles and stringers that are rated to withstand:
  - Concentrated loads of two casters or leveling feet, each weighing up to 1,800 lb (818 kg).
  - Minimum static ultimate load of 3,300 lb (1,495 kg).
  - Rolling loads of 1,800 (818 kg) per floor tile. On floor tiles that do not meet the 1,800 lb rolling load rating, use coverings such a plywood to protect floors during system roll.
- Position that is adjacent cabinets with no more than two casters or leveling feet on a single floor tile.
- Cutouts in 24 x 24 in tiles must be no more that 8 inches (20.3 cm) wide by 6 inches (15.3 cm) deep, and centered on the tiles, 9 inches (22.9 cm) from the front and rear and 8 inches (20.3 cm) from the sides. Since cutouts weaken the tile, you can minimize deflection by adding pedestal mounts adjacent to the cutout. The number and placement of additional pedestal mounts relative to a cutout must be in accordance with the floor tile manufacturer's recommendations.

When positioning the cabinet, take care to avoid moving a caster into a floor tile cutout.

Ensure that the combined weight of any other objects in the data center does not compromise the structural integrity of the raised floor and/or the subfloor (non-raised floor).

It is recommended that a certified data center design consultant inspect the site to ensure that the floor can support the system and surrounding weight.

NOTE: The actual cabinet weight depends on the specific product configuration. Calculate the total by using the Power Calculator.

Shipping and storage requirements

CAUTION: Systems and components must not experience changes in temperature and humidity that are likely to cause condensation to form on or in that system or component. Do not exceed the shipping and storage temperature gradient of 45°F/hr (25°C/hr).

Table 22. Shipping and storage requirements

Requirement	Description
Ambient temperature	-40° F to 149°F (-40°C to 65°C)
Temperature gradient	45°F/hr (25°C/hr)
Relative humidity	10% to 90% noncondensing
Elevation	-50 to 35,000 ft (-16 to 10,600 m)
Storage time (unpowered) Recommendation	Do not exceed 6 consecutive months of unpowered storage.

Fire suppressant disclaimer

Fire prevention equipment in the computer room should always be installed as an added safety measure. A fire suppression system is the responsibility of the customer. When selecting appropriate fire suppression equipment and agents for the data center, choose carefully. An insurance underwriter, local fire marshal, and local building inspector are all parties that you should consult during the selection of a fire suppression system that provides the correct level of coverage and protection.

Equipment is designed and manufactured to internal and external standards that require certain environments for reliable operation. We do not make compatibility claims of any kind nor do we provide recommendations on fire suppression systems. It is not recommended to position storage equipment directly in the path of high pressure gas discharge streams or loud fire sirens so as to minimize the forces and vibration adverse to system integrity.

NOTE: The previous information is provided on an “as is” basis and provides no representations, warranties, guarantees or obligations on the part of our company. This information does not modify the scope of any warranty set forth in the terms and conditions of the basic purchasing agreement between the customer and the manufacturer.

Getting Help

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

Dell Technologies support	• Support tab on the Dell home page: https://www.dell.com/support/incidents-online. Once you identify your product, the **Contact Us** gives you the option of email, chat, or telephone support.
Telephone support	• 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 Dell EMC Customer Support Centers.
PowerScale OneFS Documentation Info Hubs	• OneFS Info Hubs: https://www.dell.com/support/article/sln318794