

Generation 6 Site Preparation and Planning Guide

June 2019

Copyright © 2015-2019 Dell Inc. or its subsidiaries. All rights reserved.

Dell believes the information in this publication is accurate as of its publication date. The information is subject to change without notice.

THE INFORMATION IN THIS PUBLICATION IS PROVIDED “AS-IS.” DELL MAKES NO REPRESENTATIONS OR WARRANTIES OF ANY KIND WITH RESPECT TO THE INFORMATION IN THIS PUBLICATION, AND SPECIFICALLY DISCLAIMS IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE. USE, COPYING, AND DISTRIBUTION OF ANY DELL SOFTWARE DESCRIBED IN THIS PUBLICATION REQUIRES AN APPLICABLE SOFTWARE LICENSE.

Dell Technologies, Dell, EMC, Dell EMC and other trademarks are trademarks of Dell Inc. or its subsidiaries. Other trademarks may be the property of their respective owners. Published in the USA.

Dell EMC
Hopkinton, Massachusetts 01748-9103
1-508-435-1000 In North America 1-866-464-7381
www.DellEMC.com

CHAPTER 1

Introduction to this guide

This section contains the following topics:

- [About this guide](#)..... 4
- [Isilon scale-out NAS overview](#)..... 4
- [OneFS storage architecture](#)..... 4
- [Isilon node components](#)..... 4
- [Where to go for support](#)..... 5

About this guide

This guide describes how to prepare and plan for an Isilon hardware installation.

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

A cluster that includes an HD400 or a Generation 6 node in the data workflow includes unique characteristics of the respective node as described in this guide.

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

For detailed information about the OneFS operating system, review OneFS documentation on the Online Support site ([Online Support](#)).

Isilon scale-out NAS overview

The Isilon 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

Isilon takes 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. Segmenting and distributing data—a process known as striping—not only protects data, but also enables a user 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 slave devices do not invoke dependencies. Each node helps to control data requests, boost performance, and expand cluster capacity.

Isilon node components

As a rack-mountable appliance, a pre-Generation 6 storage node includes the following components in a 2U or 4U rack-mountable chassis with an LCD front panel:

CPUs, RAM, NVRAM, network interfaces, InfiniBand adapters, disk controllers, and storage media. An Isilon cluster is made up of three or more nodes, up to 144. The 4U chassis is always used for Generation 6. There are four nodes in one 4U chassis in Generation 6, therefore a quarter chassis makes up one node.

When you add a node to a pre-Generation 6 cluster, you increase the aggregate disk, cache, CPU, RAM, and network capacity. OneFS groups RAM into a single coherent cache so that a data request on a node benefits from data that is cached anywhere. NVRAM is grouped to write data with high throughput and to protect write operations from power failures. As the cluster expands, spindles and CPU combine to increase throughput, capacity, and input-output operations per second (IOPS). The minimum cluster for Generation 6 is four nodes and Generation 6 does not use NVRAM. Journals are stored in RAM and M.2 flash is used for a backup in case of node failure.

There are several types of nodes, all of which can be added to a cluster to balance capacity and performance with throughput or IOPS:

Node	Use Case
Gen-6 Hardware F800 and F810 (The F810 is supported with OneFS 8.1.3 only)	All flash solution
Gen-6 Hardware H-Series	<ul style="list-style-type: none"> H600, performance spinning solution H500, performance capacity H400, capacity performance
Gen-6 Hardware A-Series	<ul style="list-style-type: none"> A200, active archive A2000, deep archive
S-Series	IOPS-intensive applications
X-Series	High-concurrency and throughput-driven workflows
NL-Series	Near-primary accessibility, with near-tape value
HD-Series	Maximum capacity

The following Dell EMC Isilon nodes improve performance:

Node	Function
A-Series Performance Accelerator	Independent scaling for high performance
A-Series Backup Accelerator	High-speed and scalable backup-and-restore solution for tape drives over Fibre Channel connections

Where to go for support

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

Online support	<ul style="list-style-type: none"> Live Chat Create a Service Request
----------------	---

	For questions about accessing online support, send an email to support@emc.com .
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 are available at Dell EMC Customer Support Centers.
Isilon Community Network	The Isilon Community Network connects you to a central hub of information and experts to help you maximize your current storage solution. From this site, you can demonstrate Isilon products, ask questions, view technical videos, and get the latest Isilon product documentation.
Isilon Info Hubs	For the list of Isilon info hubs, see the Isilon Info Hubs page on the Isilon Community Network . Use these info hubs to find product documentation, troubleshooting guides, videos, blogs, and other information resources about the Isilon products and features you're interested in.

CHAPTER 2

Selecting the equipment

Isilon offers a variety of storage and accelerator nodes that you can combine to meet the storage needs.

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

Talk to an Isilon Sales Account Manager to identify the equipment that is best suited to support the workflow.

- [Safety and EMI Compliance](#)..... 8
- [Heavy equipment and tipping warning](#)..... 8
- [Grounding guidelines](#)..... 9
- [Shock and Vibration](#)..... 9
- [Storage node specifications](#)..... 9
- [Racks and rails](#)..... 30
- [Switches and cables](#)..... 46
- [Leaf-Spine topology](#)..... 53
- [Network topology](#)..... 54
- [Assisting with installation](#)..... 55

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**.

Heavy equipment and tipping warning

Review all safety warnings before moving or installing heavy equipment.

DANGER

- **CAUTION - HEAVY.** Do NOT move or install heavy equipment alone. Installation of this equipment requires at least two people and the use of a server lift or appropriate lift-assist equipment for the installation.
- Heavy equipment should be installed at the lowest available rack or cabinet slot.
- If you are handling a unit that you can extend away from the rack, never extend more than one unit away from the cabinet or rack. If more than one unit is extended from the cabinet or rack at a time, the cabinet or rack may tip.
- Do NOT install heavy equipment in a cabinet or rack without using a stabilizer bracket or affixing the cabinet/rack to the sub-floor.
- Read all installation instructions and take note of all warnings and guidance before attempting to install equipment into any cabinet or rack.
- Dell EMC shall not be liable for any damage, injury, or loss that results from a failure to follow installation instructions and other documentation that is provided for the equipment.
- Before installation, ensure that the target cabinets and lab floors are rated to support the total weight of equipment that is installed and being installed.

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. [Floor load bearing requirements](#) and [Safety and EMI Compliance](#) 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	<ul style="list-style-type: none"> • 1.15 Grms • 1 hour Frequency range 1–200 Hz

Storage node specifications

You can review the specifications of EMC Isilon nodes to learn node weight and dimensions, attributes, options, and industry certifications.

Generation 6 node specifications

The Generation 6 chassis is a high storage density solution for large scale out systems. Generation 6 includes the F-series (F800), H-series (H5600 (OneFS 8.2), H400, H500, H600), and the A-series (A200, and A2000) platforms.

Note: Generation 6 clusters must be running OneFS 8.1.0.1 or later to use Self-Encrypting Drives (SEDs).

Dimensions and weight

Height	Width	Depth	Weight
Generation 6 deep (A2000)	17.60in (45cm)	40.4in (102.6cm) Note: Depth is from the front NEMA rail to rear 2.5in SSD cover ejector: 40.4in.	285 lbs (129 kg)
Generation 6 standard	17.60in (45cm)	35.8in (90.9cm) Note: Standard depth is from the front NEMA rail to rear 2.5in SSD cover ejector: 35.8in.	250 lbs (113 kg)

Generation 6 supports both 2.5in, and 3.5 in drives in the same enclosure.


Generation 6 F-Series node specifications

The Generation 6 F-Series is an all-flash solution that includes the F810 and F800 node models.

Note: Currently, the F810 nodes are supported for OneFS 8.1.3 only. To take advantage of F810 in your environment, contact your Dell EMC Account team.

F810 node attributes and options

Attribute	Capacity		
	3.84 TB SSD	7.68 TB SSD	15.36 TB SSD
Chassis capacity	230 TB	460 TB	924 TB
Number of nodes	4	4	4
Self-Encrypting Drives (SED, SSD) option	No	No	Yes
SSD drives (2.5 in) per chassis	Per node = 15, per chassis (4 nodes), the value is 60	Per node = 15, per chassis (4 nodes), the value is 60	Per node = 15, per chassis (4 nodes), the value is 60
OneFS Version Required	8.1.3 only		
System Memory	256 GB		

Attribute	Capacity		
	3.84 TB SSD	7.68 TB SSD	15.36 TB SSD
Front-end Networking	2 x 10GbE (SFP+ or twin-ax copper) or 2 x 40GbE (QSFP+)		
Network interfaces	Network interfaces support IEEE 802.3 standards for 1 Gbp/s, 10Gbp/s, 40 Gbp/s, and 100Mbps network connectivity		
Drive Controller	SATA-3, 6 Gbp/s, SAS-3 12 Gb/s		
CPU Type	Intel® Xeon® Processor E5-2967A v4 (40M Cache, 2.60 GHz)		
Infrastructure Networking	2 x 40GbE (QSFP+)  Note: InfiniBand connections are not supported for the F810 node back-end networking.		
Typical Thermal Rating	4,440 BTU per hour		

F800 node attributes and options



Attributes and Options	Capacity				
	1.6 TB SSD	3.2 TB SSD	3.84 TB SSD	7.68 TB SSD	15.36 TB SSD
Chassis capacity	96 TB	192 TB	230 TB	460 TB	924 TB
Number of nodes per chassis	4	4	4	4	4
Self-Encrypting Drives (SED, SSD) option	Yes	Yes	No	No	Yes
SSD drives (2.5 in) per chassis	Per node = 15, per chassis (4 nodes), the value is 60	Per node = 15, per chassis (4 nodes), the value is 60	Per node = 15, per chassis (4 nodes), the value is 60	Per node = 15, per chassis (4 nodes), the value is 60	Per node = 15, per chassis (4 nodes), the value is 60
OneFS Version Required	8.1 or later except for self-encrypting drive options which require Isilon OneFS 8.1.0.1 or later				
System Memory	256 GB per node				
Front-end Networking (per node)	2 x 10GE (SFP+) or 2 x 40GbE (QSFP+)				
Network interfaces	Network interfaces support IEEE 802.3 standards for 1 Gbp/s, 10Gbp/s, 40 Gbp/s, and 100Mbps network connectivity				
Drive Controller	SATA-3, 6 Gbp/s, SAS-3 12 Gb/s				
CPU Type	Intel® Xeon® Processor E5-2967A v4 (40M Cache, 2.60 GHz)				
Infrastructure Networking	2 InfiniBand connections with quad data rate (QDR) link or 2 X 40GbE (QSFP+)				
Typical Thermal Rating	4,440 BTU per hour				

Power consumption values

Table 1 Power consumption values

F810 and F800	Value (Watts)
Watts per chassis at 25° Celsius	1300 typical power consumption at 240v
Watts per chassis	1800 maximum power consumption
Maximum power N+1 (per node pair)	900 maximum power at 240v
Power per node (idle)	181
Power per node (typical)	325

Table 2 Cluster attributes

Attribute	F810 cluster	F800 cluster
Number of chassis	1 to 36	1 to 36
Number of nodes	4 to 144	4 to 144
Capacity (raw)  Note: Usable capacity is slightly lower than the raw capacity reflected in this table.	230 TB to 33 PB	96 TB to 33 PB
Effective Cluster capacity  Note: Based on an 80% storage utilization rate. For Isilon F810, effective capacity is also based on a data compression ratio of up to 3:1. The actual data compression ratio varies by dataset.	77 TB to 26.5 PB	184 TB up to 79.6 PB
Memory	184 TB up to 79.6 PB	77 TB to 26.5 PB
Rack Units	4 to 144	4 to 144
Scalability	Scales from 4 to 144 nodes in a single cluster with up to 33 PB capacity (raw) and up to 79.6 PB of effective storage capacity.	Scales from 4 to 144 nodes in a single cluster with up to 33 PB capacity (raw).

Generation 6 H-Series node specifications

The Generation 6 H-Series includes the H5600, H400, H500, and H600 product configurations.

Table 3 H5600 node attributes and options

Attribute	Capacity
Chassis capacity	800 TB

Table 3 H5600 node attributes and options (continued)

Attribute	Capacity
Number of nodes	4
Cache (per Node) Solid state drives (1x3.2 TB)	1 or 2
Self-Encrypting Drives (SED, SSD) option	Yes
SAS drives (2.5 in 512n) per chassis	Per node = 30, per chassis (4 nodes), the value is 120
SAS drives (3.5 in) with long sleds	Per node = 20, per chassis (4 nodes), the value is 80
OneFS Version Required	8.2 and later
System Memory	256 GB
Front-end Networking	2 x 10GE (SFP+) or 2 x 40GbE (QSFP+)
Network interfaces	Network interfaces support IEEE 802.3 standards for 1Gbps, 10Gbps, 40Gbps and 100Mbps network connectivity
Drive Controller	10 TB Non-SED SATA 5 GB/s per 500TB
CPU Type	Intel® Xeon® Processor E5-2680 v4
Infrastructure Networking	40GbE front-end/back-end connections with quad data rate (QDR) link or 2 x 40GbE (QSFP+)
Typical Thermal Rating	6619 BTU/hr

H400 node attributes and options

Attribute	Capacity		
	2 TB SSD	4 TB SSD	8 TB SSD
Chassis capacity	120 TB	240 TB	480 TB
Number of nodes	4	4	4
Self-Encrypting Drives (SED, SSD) option	Yes	Yes	Yes
SSD drives (2.5 in) per chassis	Per node = 15, per chassis (4 nodes), the value is 60	Per node = 15, per chassis (4 nodes), the value is 60	Per node = 15, per chassis (4 nodes), the value is 60
OneFS Version Required	8.1 and later		
System Memory	64 GB		

Attribute	Capacity		
	2 TB SSD	4 TB SSD	8 TB SSD
Front-end Networking	2 x 10GE (SFP+)		
Network interfaces	Network interfaces support IEEE 802.3 standards for 1Gbps, 10Gbps, and 100Mbps network connectivity		
Drive Controller	SATA-3 6 Gb/s, SAS-3 12 Gb/s		
CPU Type	Intel® Xeon® Processor D-1527		
Infrastructure Networking	2 InfiniBand connections with quad data rate (QDR) link or 2 X 10 GbE (SFP+)		
Typical Thermal Rating	3800 BTU/hr		

H500 node attributes and options

Attribute	Capacity		
	2 TB SATA	4 TB SATA	8 TB SATA
Chassis capacity	120 TB	240 TB	480 TB
Number of nodes	4	4	4
Cache (per Node) Solid state drives (1.6 TB, or 3.2 TB)	1 or 2	1 or 2	1 or 2
Self-Encrypting Drives (SED, SSD) option	Yes	Yes	Yes
HDD drives (3.5" 4kn SATA) per chassis)	Per node = 15, per chassis (4 nodes), the value is 60	Per node = 15, per chassis (4 nodes), the value is 60	Per node = 15, per chassis (4 nodes), the value is 60
OneFS Version Required	8.1 and later		
System Memory	128 GB		
Front-end Networking	2 x 10GE (SFP+) or 2 x 40GbE (QSFP+)		
Network interfaces	Network interfaces support IEEE 802.3 standards for 1Gbps, 10Gbps, 40Gbps and 100Mbps network connectivity		
Drive Controller	SATA-3 6 Gb/s, SAS-3 12 Gb/s		
CPU Type	Intel® Xeon® Processor E5-2630 v4		
Infrastructure Networking	2 InfiniBand connections with quad data rate (QDR) link or 2 X 40GbE (QSFP+)		
Typical Thermal Rating	4540 BTU/hr		

H600 node attributes and options

Attribute	Capacity	
	600 GB SAS	1.2 TB SAS
Chassis capacity	72 TB	144 TB
Number of nodes	4	4
Cache (per Node) Solid state drives (1.6 TB, or 3.2 TB)	1 or 2	1 or 2
Self-Encrypting Drives (SED, SSD) option	Yes	Yes
SAS drives (2.5 in 512n) per chassis	Per node = 30, per chassis (4 nodes), the value is 120	Per node = 30, per chassis (4 nodes), the value is 120
OneFS Version Required	8.1 and later	
System Memory	128 GB	
Front-end Networking	2 x 10GE (SFP+) or 2 x 40GbE (QSFP+)	
Network interfaces	Network interfaces support IEEE 802.3 standards for 1Gbps, 10Gbps, 40Gbps and 100Mbps network connectivity	
Drive Controller	SATA-3 6 Gb/s, SAS-3 12 Gb/s	
CPU Type	Intel® Xeon® Processor E5-2680 v4	
Infrastructure Networking	2 InfiniBand connections with quad data rate (QDR) link or 2 X 40GbE (QSFP+)	
Typical Thermal Rating	5840 BTU/hr	

Power consumption values

Generation 6 product configuration	Typical power consumption at 240v	Maximum power at 240v
H400	1120 Watts per chassis at 25°Celsius	1560 Watts per chassis
H500	1330 Watts per chassis at 25°Celsius	1870 Watts per chassis
H600	1700 Watts per chassis at 25°Celsius	1990 Watts per chassis

Cluster attributes

Cluster attribute	H400	H500	H600
Number of Chassis	1 to 36	1 to 36	1 to 36
Number of nodes	4 to 144	4 to 144	4 to 144

Cluster attribute	H400	H500	H600
Cluster Capacity	120 TB to 17.2 PB	120 TB to 17.2 PB	72 TB to 144 TB
Rack units	4 to 144	4 to 144	4 to 144

Generation 6 A-Series node specifications

The Generation 6 A-Series includes the A200, and A2000 product configurations.

A200 node attributes and options

Attribute	Capacity		
	2 TB HDD	4 TB HDD	8 TB HDD
Chassis capacity	120 TB	240 TB	480 TB
Number of nodes	4	4	4
Self-Encrypting Drives (SED, SSD) option	Yes	Yes	No
SSD drives (2.5 in) per chassis	Per node = 15, per chassis (4 nodes), the value is 60	Per node = 15, per chassis (4 nodes), the value is 60	Per node = 15, per chassis (4 nodes), the value is 60
OneFS Version Required	8.1 and later		
System Memory	16 GB		
Cache (per node) Solid state drives (400 GB)	1 or 2		
Front-end Networking	2 x 10GE (SFP+)		
Network interfaces	Network interfaces support IEEE 802.3 standards for 10 Gbp/s, 1 Gbp/s, and 100 Mbp/s network connectivity		
Drive Controller	SATA-3, 6 Gbp/s, SAS-3 12 Gb/s		
CPU Type	Intel® Pentium® Processor D1508		
Infrastructure Networking	2 InfiniBand connections with quad data rate (QDR) links or 2 X 10GbE (SFP+)		
Typical Thermal Rating	3600 BTU per hour		

A2000 node attributes and options

Attribute	Capacity
	10 TB HDD
Chassis capacity	800 TB
Number of nodes	4
HDD drive (3.5" 4kn SATA)	Per node = 20, per chassis (4 nodes), the value is 80

Attribute	Capacity
	10 TB HDD
Self-encrypting drive (SED HDD) option	Yes
OneFS Version Required	8.1 and later
System Memory	16 GB
Front-end Networking	2 x 10GE (SFP+)
Network interfaces	Network interfaces support IEEE 802.3 standards for 10Gbps, 1Gbps, and 100Mbps network connectivity
Drive Controller	SATA-3 6 Gb/s, SAS-3 12 Gb/s
CPU Type	Intel® Pentium® Processor D1508
Infrastructure Networking	2 InfiniBand connections supporting QDR links or 2 X 10 GbE (SFP+)
Typical Thermal Rating	3800 BTU/hr

Power consumption values

Generation 6 product configuration	Typical power consumption @ 240v	Maximum power consumption @ 240v
A200	1060 Watts per chassis at 25°Celsius	1460 Watts per chassis
A2000	1120 Watts per chassis at 25°Celsius	1520 Watts per chassis

Cluster attributes

Cluster attributes	A200	A2000
Number of Chassis	1 to 36	1 to 36
Number of nodes	4 to 144	4 to 144
Cluster Capacity	120 TB to 17.3 PB	800 TB to 28.8 PB
Rack units	4 to 144	4 to 144

S210 node specifications

The Isilon S210 node is a 2U storage option in the Isilon S-Series product line. To install an Isilon S210 node, the cluster must be running OneFS 7.1.1 or later.

Dimensions and weight

Height	Width	Depth	Weight
3.48 in (8.8 cm)	18.87 in (48 cm)	30.5 in (77.47 cm)	55.0 lbs (25.0 kg)

Note: The S210 is approximately 27 in. without the front panel cable bend (approximately 3 in.) resulting in the 30.5 in. depth.

Node attributes and options

Attribute	300 GB HDDs	600 GB HDDs	900 GB HDDs	1.2 TB HDDs
Capacity (HDD/SSD)	6.6–7.2 TB / 0–1.6 TB	13.2–14.4 TB / 0–1.6 TB	16.2–21.6 TB / 0–4.8 TB	21.6–28.8 TB/0–4.8TB
Hard Drives (2.5" 10k RPM)	22–24	22–24	18–24	18–24
Solid State Drives (200GB)	0–2	0–2	0–6	0–6
System Memory	32, 64, 128, or 256 GB			
Front-end Networking	2 copper 1000 Base-T (GE) and 2 x 10GE (SFP+ or twin-ax copper)			
Drive Controller	SAS-2, 6 Gb/s			
CPU Type	Dual, 6-core Intel® Xeon® processor			
Infrastructure Networking	2 InfiniBand connections supporting DDR and QDR links			
Non-volatile RAM (NVRAM)	2 GB			
Typical Power Consumption @100v and @240v	425W			
Typical Thermal Rating	1,500 BTU/hr			
Maximum Thermal Rating	2350 BTU, 680 Watts			

Cluster attributes

Number of nodes	Capacity	Memory	Rack Units
3–144	16.2 TB to 4.15 PB	96 GB to 36.8 TB	6–288

X210 node specifications

The Isilon X210 node is a 2U storage option in the Isilon X-Series product line. To install an Isilon X210 node, the cluster must be running OneFS 7.2.1 or later.

Dimensions and weight

Height	Width	Depth	Weight
3.48 in (8.8 cm)	18.87 in (47.9 cm)	28.5 in (72.4 cm)	61 lbs (27.7 kg)

Node attributes and options

Attribute	1 TB HDDs	2 TB HDDs	3 TB HDDs	4 TB HDDs
Capacity	6–12 TB	12–24 TB	18–36 TB	24–48 TB
Hard Drives (3.5" 7200 RPM)	6–12	6–12	6–12	6–12
Self-Encrypting Drives (SED HDD) option (7200 RPM)	No	No	Yes	Yes
Solid State Drives (SSDs) (200 GB, 400 GB, or 800 GB)	Up to 6	Up to 6	Up to 6	Up to 6
Self-Encrypting Drives (SED SSD) option (800 GB)	No	No	Yes (0, 2, or 4)	Yes (0, 2, or 4)
System ECC Memory	24 GB or 48 GB			
Front-end Networking	2 x GbE and 2 x 10 GbE (SFP+ or twin-ax copper)			
Network Interfaces	Isilon network interfaces support IEEE 802.3 standards for 10Gbps, 1Gbps, and 100MBps network connectivity			
Drive Controller	SATA-3, 6 Gb/s			
CPU Type	Single Intel® Xeon® Processor E5-2407v2 @ 2.4 GHz, 4 core			
Infrastructure Networking	2 InfiniBand connections supporting DDR and QDR links			
Non-volatile RAM (NVRAM)	2 GB			
Typical Power Consumption @100v	400 Watts			
Typical Power Consumption @240v	400 Watts			

Attribute	1 TB HDDs	2 TB HDDs	3 TB HDDs	4 TB HDDs
Typical Thermal Rating	1,370 BTU/hr			

Cluster attributes

Number of nodes	Capacity	Memory	Rack Units
3–144	18 TB to 6.9 PB	72 GB to 6.9 TB	6–288

X410 node specifications

The Isilon X410 node is a 4U storage option in the Isilon X-Series product line. To install an Isilon X410 node, the cluster must be running OneFS 7.1.1 or later.

Dimensions and weight

Height	Width	Depth	Weight
6.96 in (17.7 cm)	18.90 in (48 cm)	28.5 in (72.4 cm)	120 lbs (54.5 kg)

Node attributes and options

Attribute	1 TB HDDs	2 TB HDDs	3 TB HDDs	4 TB HDDs
Capacity	30–36 TB	60–72 TB	90–108 TB	120–144 TB
Hard Drives (3.5" SATA 7200 RPM)	30–36	30–36	30–36	30–36
Solid State Drives	0–6	0–6	0–6	0–6
Self-Encrypting Drives (SED SSD) option (800 GB)	No	No	Yes (0, 2, 4, or 6)	Yes (0, 2, 4, or 6)
System ECC Memory	64 GB, 128 GB, or 256 GB			
Front-end Networking	Use either of the following: <ul style="list-style-type: none"> 2 x 1 Gigabit Ethernet and 2 x 10GE (SFP+ or twin-ax copper) 2 x 1 Gigabit Ethernet and 2 x 40GbE (QSFP+) 			
Network interfaces	Isilon network interfaces support IEEE 802.3 standards for 10Gbps, 1Gbps, and 100Mbps network connectivity			
Drive Controller	SATA-3, 6 GB/s			
CPU Type	Dual, 8-core Intel® Xeon® processor			
Infrastructure Networking	2 InfiniBand connections supporting DDR and QDR links			
Non-volatile RAM (NVRAM)	2 GB			

Attribute	1 TB HDDs	2 TB HDDs	3 TB HDDs	4 TB HDDs
Typical Power Consumption @100v	700 Watts			
Typical Power Consumption @240v	700 Watts			
Typical Thermal Rating	2,400 BTU/hr			

Cluster attributes

Number of nodes	Capacity	Memory	Rack Units
3–144	108 TB to 20.7 PB	192 GB to 36.8 TB	12–576

NL410 node specifications

The Isilon NL410 node is a 4U storage option in the Isilon NL-Series product line. To install an Isilon NL410 node, the cluster must be running OneFS 7.2.1 or later.

Dimensions and weight

Height	Width	Depth	Weight
6.96 in (17.7 cm)	18.9 in (48 cm)	31.25 in (79.4 cm)	114 lbs (51.7 kg)

Node attributes and options

Attribute	1 TB HDDs	2 TB HDDs	3 TB HDDs	4 TB HDDs	6 TB HDDs
Capacity	35–36 TB	70–72 TB	105–108 TB	140 TB –144 TB	210 TB
Hard Drives (3.5" SATA 7200 RPM)	35 or 36	35 or 36	35 or 36	35 or 36	35
Self-Encrypting Drives (SEDs) option (7200 RPM)	No	No	Yes FIPS 140-2 level 2 validated SEDs	Yes	No
Solid State Drives (SSDs) (200 GB or 800 GB)	0 or 1 (optional)	0 or 1 (optional)	0 or 1 (optional)	0 or 1 (optional)	1 (required)
Self-Encrypting Drives (SEDs SSD) option (800 GB)	No	No	No	Yes	No

Attribute	1 TB HDDs	2 TB HDDs	3 TB HDDs	4 TB HDDs	6 TB HDDs
System ECC Memory	24 GB or 48 GB				
Front-end Networking	2 x GbE and 2 x 10 GbE (SFP+ or twin-ax copper)				
Network Interfaces	Isilon network interfaces support IEEE 802.3 standards for 10Gbps, 1Gbps, and 100Mbps network connectivity				
Drive Controller	SATA-3, 6 Gb/s				
CPU Type	Intel® Xeon® Processor E5-2407v2				
Infrastructure Networking	2 InfiniBand connections supporting QDR links				
Non-volatile RAM (NVRAM)	2 GB				
Typical Power Consumption @100v	800W				
Typical Power Consumption @240v	720W				
Typical Thermal Rating	2,500 BTU/hr				

Cluster attributes

Number of nodes	Capacity	Memory	Rack Units
3–144	108 TB to 30.2 PB	36 GB to 6.9 TB	12–576

S200 node specifications

The Isilon S200 node is a 2U storage option in the Isilon S-Series product line. To install an Isilon S200 node, the cluster must be running OneFS 6.5.5 or later.

Dimensions and weight

Height	Width	Depth	Weight
3.48 in (8.8 cm)	18.87 in (48 cm)	28.5 in (72.4 cm)	55.0 lbs (25.0 kg)

Node attributes and options

Attribute	300 GB HDDs	600 GB HDDs	900 GB HDDs	1.2 TB HDDs
Capacity (HDD/SSD)	5.4–7.2 TB / 0–4.8 TB	10.8–14.4 TB / 0–4.8 TB	16.2–21.6 TB / 0–4.8 TB	21.6–28.8 TB / 0–4.8 TB

Attribute	300 GB HDDs	600 GB HDDs	900 GB HDDs	1.2 TB HDDs
Hard Drives (2.5" 10k RPM)	18–24	18–24	18–24	18–24
Solid State Drives (200GB, 400GB, or 800GB)	0–6	0–6	0–6	0–6
System Memory	24 GB, 48 GB, or 96 GB			
Front-end Networking	4 copper 1000 Base-T (GigE) or 4 x GigE (copper) or 2 x GigE and 2 x 10GigE (SFP+ or twin-ax copper)	4 copper 1000 Base-T (GigE) or 4 x GigE (copper) or 2 x GigE and 2 x 10GigE (SFP+ or twin-ax copper)	2 x GigE and 2 x 10GigE (SFP+)	2 x GigE and 2 x 10GigE (SFP+)
Drive Controller	SAS-2, 6 Gb/s			
CPU Type	Dual, 4-core Intel® Xeon® Processor			
Infrastructure Networking	2 InfiniBand connections supporting DDR links			
Non-volatile RAM (NVRAM)	512 MB			
Typical Power Consumption @100v and @240v	450W			
Typical Thermal Rating	1,550 BTU/hr			

Cluster attributes

Number of nodes	Capacity	Memory	Rack Units
3–144	16.2 TB to 4.15 PB	72 GB to 13.8 TB	6–288

X200 node specifications

The Isilon X200 node is a 2U storage option in the Isilon X-Series product line.

Dimensions and weight

Height	Width	Depth	Weight
3.48 in (8.8 cm)	18.87 in (48 cm)	28.5 in (72.4 cm)	61 lbs (27.7 kg)

Node attributes and options

Attribute	1 TB HDDs	2 TB HDDs	3 TB HDDs
Capacity	7.2–12 TB	13.2–24 TB	19.2–36 TB

Attribute	1 TB HDDs	2 TB HDDs	3 TB HDDs
Hard Drives (3.5" 7200 RPM)	6–12	6–12	6–12
Solid State Drives (200GB or 800GB)	Up to 6	Up to 6	Up to 6
Self-Encrypting Drives (SED HDD) option (7200 RPM)	No	No	Yes
Self-Encrypting Drives (SED SSD) option (800GB)	No	No	Yes (0, 2, or 4)
System Memory	6, 12, 24, or 48 GB		
Front-end Networking	4 x GigE or 2 x GigE and 2 x 10GigE (SFP+ or twin-ax copper)		
Drive Controller	SATA-3, 6 Gb/s		
CPU Type	Intel® Xeon® Processor		
Infrastructure Networking	2 InfiniBand connections with double data rate (DDR) links.		
Non-volatile RAM (NVRAM)	512 MB		
Typical Power Consumption @100v and @240v	400W		
Typical Thermal Rating	1,370 BTU/hr		

Cluster attributes

Number of nodes	Capacity	Memory	Rack Units
3–144	24 TB to 5.2 PB	18 GB to 6.9 TB	6–288

X400 node specifications


The Isilon X400 node is a 4U storage option in the Isilon X-Series product line.

Dimensions and weight

Height	Width	Depth	Weight
6.96 in (17.7 cm)	18.90 in (48 cm)	31.25 in (79.4 cm)	127 lbs (57.7 kg)

Node attributes and options

Attribute	1 TB HDDs	2 TB HDDs	3 TB HDDs	4 TB HDDs
Capacity	33.6–36 TB	67.2–72 TB	97.6–108 TB	122.4–144 TB
Hard Drives (3.5" SATA)	32–36	33–36	32–36	30–36

Attribute	1 TB HDDs	2 TB HDDs	3 TB HDDs	4 TB HDDs
Solid State Drives	0, 2, or 4	0 or 3	0 or 4	0 or 6
Self-Encrypting Drives (SEDs) option (7200 RPM)	No	No	Yes  Note: FIPS 140-2 level 2 validated SEDs with unique AES-256 bit strength keys that are assigned to each drive.	No
System Memory	24, 48, 96, or 192 GB			
Front-end Networking	<ul style="list-style-type: none"> 4 x GigE or 2 x GigE and 2 x 10 GigE (SFP+ or twin-ax copper) 4 TB HDDs: 4 x GigE or 2 x GigE and 2 x 10GigE (SFP+) 			
Network Interface	Isilon network interfaces support IEEE 802.3 standards for 10Gps, 1Gps, and 100Mbps network connectivity			
Drive Controller	SATA-3, 6 GB/s			
CPU Type	Dual, 6-core Intel® Xeon® Processor			
Infrastructure Networking	2 InfiniBand connections with double data rate (DDR) links.			
Non-volatile RAM (NVRAM)	512 MB			
Typical Power Consumption @100v and @240v	725W			
Typical Thermal Rating	2,500 BTU/hr			

Cluster attributes

Number of nodes	Capacity	Memory	Rack Units
3–144	108 TB to 20.7 PB	72 GB to 27.6 TB	12–576


NL400 node specifications

The Isilon NL400 node is a 4U storage option in the Isilon NL-Series product line.

Dimensions and weight

Height	Width	Depth	Weight
6.96 in (17.7 cm)	18.90 in (48 cm)	31.25 in (79.4 cm)	127 lbs (57.7 kg)

Node attributes and options

Attribute	1 TB HDDs	2 TB HDDs	3 TB HDDs	4 TB HDDs
Capacity	36 TB	72 TB	108 TB	144 TB
Hard Drives (3.5" SATA)	36	36	36	36
Self-Encrypting Drives (SEDs) option (7200 RPM)	No	No	Yes  Note: FIPS 140-2 level 2 validated SEDs with unique AES-256-bit strength keys assigned to each drive.	No
System Memory	12, 24, or 48 GB			
Front-end Networking	4 x GigE or 2 x GigE and 2 x 10 GigE (SFP+ or twin-ax copper)			
Network Interface	Isilon network interfaces support IEEE 802.3 standards for 10Gps, 1Gps, and 100Mbps network connectivity			
Drive Controller	SATA-3, 6 GB/s			
CPU Type	Intel® Xeon® Processor			
Infrastructure Networking	2 InfiniBand connections with double data rate (DDR).			
Non-volatile RAM (NVRAM)	512 MB			
Typical Power Consumption @100v and @240v	725W			
Typical Thermal Rating	2,500 BTU/hr			

Cluster attributes

Number of nodes	Capacity	Memory	Rack Units
3–144	108 TB to 20.7 PB	72 GB to 27.6 TB	12–576

HD400 node specifications

The Isilon HD400 node is a 4U high density, deep archive storage solution.

Dimensions and weight

Height	Width	Depth	Weight
6.96 in (17.7 cm)	18.90 in (48 cm)	35.00 in (88.9 cm)	220 lbs (100 kg)

- The HD400 node is 35 inches deep, which is longer than other Isilon nodes.
- You can install the HD400 node in standard racks, but you are not able to close the rear doors on most standard cabinets.
- The 35 inch depth of the HD400 node does not include additional space for cable management arms.

Node attributes and options

Attribute	6 TB HDDs
Capacity	354 TB
Hard Drives (3.5" SATA, 7200 RPM)	59
Solid State drives (800GB)	1
Self-Encrypting Drives (SEDs) option	No
OneFS Version Required	7.2 and later
System Memory	24 GB or 48 GB
Front-end Networking	2 x Gigabit Ethernet and 2 x 10GE (SFP+ or twin-ax copper)
Network interfaces	Isilon network interfaces support IEEE 802.3 standards for 10Gbps, 1Gbps, and 100Mbps network connectivity
Drive Controller	SATA-3, 6 Gb/s
CPU Type	Intel® Xeon® Processor E5-2407 v2 (10M Cache, 2.40 GHz)
Infrastructure Networking	2 InfiniBand connections with quad data rate (QDR) link
Non-volatile RAM (NVRAM)	2 GB
Typical Power Consumption @100v	Not applicable
Typical Power Consumption @240v	1100 Watts

Attribute	6 TB HDDs
Typical Thermal Rating	2900 BTU/hr to 3750 BTU/hr

Cluster attributes

Number of nodes	Capacity	Memory	Rack Units
3–144	1.06 PB to 50.9 PB	72 GB to 6.9 TB	12–576

Accelerator nodes

With Isilon accelerators, backup and performance, you can scale the performance of the cluster without increasing capacity.

Backup accelerator

The backup accelerator provides a simple and flexible solution to address these challenges and scale backup performance to fit within the backup windows.

The backup accelerator integrates quickly and easily into the Isilon scale-out NAS cluster and can be installed in about a minute, without system downtime.

The backup accelerator is also designed to integrate easily with current tape infrastructures, as well as with leading data backup and recovery software technologies and processes. Each backup accelerator node supports multi-paths and can drive multiple uncompressed LTO-5 and LTO-6 tape drives while supporting both two-way and three-way NDMP to meet specific data backup requirements.

Performance accelerator

The performance accelerator is a next-generation performance accelerator that is designed for demanding applications and workloads that require maximum performance.

With performance accelerator nodes, latency can be reduced and concurrent reads that are increased with throughput for a cached dataset while gaining the ability to scale performance independent of capacity. Isilon A100 performance accelerator nodes also allow accelerated cluster operations, including disk and node rebuilds, file striping, and file-based replication.

With 256 GB of L1 cache per node, performance accelerator nodes can hold a huge, recently accessed dataset in memory. This dataset can be read with extremely high performance, low latency, and concurrent throughput.

Each performance accelerator node delivers an aggregate throughput of up to 1100 MB/s per node and allows you to scale performance without adding capacity, while supporting the most demanding high performance work flows.

A100 accelerator specifications

The Isilon A100 performance accelerator reduces latency and increases throughput for clusters supporting demanding applications and workloads. The backup accelerator scales performance for tape backup, disaster recovery, compliance, and archiving.

Dimensions and weight

Height	Width	Depth	Weight
1.72 in (4.4 cm)	18.87 in (47.9 cm)	28.6 in (72.6 cm)	30 lbs (13.6 kg)

Node attributes

Attribute	A100 backup accelerator	A100 performance accelerator
Front-end networking	<ul style="list-style-type: none"> 8 Gb Fibre Channel x Four (4) 1 Gb Ethernet X Four (4) (Copper) 	<ul style="list-style-type: none"> 10 Gb Ethernet X Two (2) (Fibre Channel or Copper) 1 Gb Ethernet X Four (4) (Copper)
Memory	64 GB	256 GB
CPU type	Intel Sandy Bridge Hex-Core 2 GHz X Two (2)	
Infrastructure networking	Two (2) QDR InfiniBand connections	
AC power input	100–240 VAC, 60–50 Hz, auto-ranging	
Power supply	Dual, 750W, redundant, hot-swappable	
Typical power consumption @100v and @240v	300W	
Typical thermal rating	400W (VA), 1350 BTU/hr	

Environmental specifications

Attribute	Value
Operating environment	50°F to 95°F (10°C to 35°C), 5% to 95% relative humidity, non-condensing
Non-operating environment	-40°F to 149°F (-40°C to 65°C), 5% to 95% relative humidity, non-condensing
Altitude	0 ft to 7,000 ft (0 m to 2,133.6 m)
Minimum service clearances	35" (88.9 cm), rear: 14" (35.6 cm)

A100 guidelines

Follow these guidelines to get optimal performance from the A100 accelerator node.

The A100 accelerator provides the most benefit to work flows where it can:

- Magnify cached read performance.

- Expand physical GigE network ports.
- Increase high single-stream throughput.
- Enhance a 10 GigE network interface.
- Reduce latency on cached, read-intensive work flows.

To ensure that the A100 accelerator is running in an optimal environment, note the version of OneFS for the Isilon account team.

The A100 accelerator does not provide significant performance improvements in work flows that include:

- Uncached concurrent reads.
- Concurrent writes.
- High IOPs.

With these types of work flow it is best to consider additional storage nodes to improve performance.

Racks and rails

You can secure Isilon nodes to standard storage racks with a sliding rail system.

Rail kits are included 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 pre-threaded holes

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

Titan HD racks

Titan HD is designed to support fully populated racks of A2000 chassis/nodes. However, all Generation 6 platforms can be installed in the Titan HD racks.

[Titan HD PDU jumper locations](#)

The following images illustrate the PDU jumper locations for the Titan-HD PDU. The jumpers should be in J1, J3, and J5 as shown in the images.

Figure 1 Single phase PDU




Figure 2 3 phase Delta PDU










Figure 3 3 Phase Wye PDU




Environmental requirements

Table 4 Titan HD environmental requirements

	+15°C to +32°C (59°F to 89.6°F) site temperature. A fully configured cabinet can produce up to 40,000 BTUs per hour.
	40% to 55% relative humidity is the recommended operating parameters. <i>Note:</i> Contents of the cabinet might be qualified outside these limits. The product-specific documentation for the system specifications provides complete details.
	The Titan HD cabinet weighs 249.48 KG (550 pounds); a cabinet fully configured with EMC products can weigh approximately 1,497 kg (3300 pounds). <i>Note:</i> Ensure that the flooring can safely support the configuration. Calculate the minimum load-bearing requirements for the site by using the product-specific weights for the system components at http://powercalculator.emc.com .
	0 meters to 2439 meters (0 ft to 8,000 ft) above sea level operating altitude.
	LAN and telephone connections for remote service and system operation.

Cabinet clearance

This Dell EMC cabinet ventilates from front to back. Provide adequate clearance to service and cool the system. Depending on component-specific connections within the cabinet, the available power cord length may be somewhat shorter than the 15-foot standard.



Stabilizing the casters

The cabinet bottom includes four caster wheels. The front wheels are fixed, and the two rear casters swivel in a 3.25-inch diameter. Swivel position of the caster wheels determines the load-bearing points on the site floor, but does not affect the cabinet footprint. Once you have positioned, leveled, and stabilized the cabinet, the four leveling feet determine the final load-bearing points on the site floor.

Cabinet specifications provide details about cabinet dimensions and weight information to effectively plan for system installation at the customer site.


Note: For installations that require a top of the cabinet power feed, a 3 m extension cord is provided. Do not move or invert the PDUs.

The following table provides cabinet load-rating information.

Load Type	Rating
Dynamic (with casters)	3,300 lb (1,497 kg) of equipment
Static (without casters)	3,300 lb (1,497 kg) of equipment

In addition, the pallets that are used to ship cabinets are specifically engineered to withstand the added weight of components that are shipped in the cabinet.

Figure 4 Cabinet component dimensions



NOTICE Customers are responsible for ensuring that the data center floor on which the Dell EMC system is configured can support the system weight. Systems can be configured directly on the data center floor, or on a raised floor supported by the data center floor. Failure to comply with these floor-loading requirements could result in severe damage to the Dell EMC system, the raised floor, subfloor, site floor, and the surrounding infrastructure. In the agreement between Dell EMC and the customer, Dell EMC fully disclaims all liability for any damage or injury resulting from a customer failure to ensure that the raised floor, subfloor and/or site floor can support the system weight as specified in this guide. The customer assumes all risk and liability that is associated with such failure.

Dell EMC 40U-P and 40U-D Titan racks

The Titan 40U-P rack supports all Generation 5 and earlier nodes, with the exception of the HD400 nodes. The Titan 40U-D supports all Generation 6 and earlier nodes, including the HD400 nodes. The Titan racks can be ordered from Dell EMC Isilon

Ensure that the facility can accommodate the rack, including:

- Delivery package dimensions
- Installed clearance requirements

- AC power input parameters


Also, the rack can be stabilized with optional stabilizer brackets.

Delivery dimensions

Ensure that all doorways and elevators can accommodate the shipping pallet and cabinet.


Use a mechanical lift or pallet jack to position the packaged cabinet in its final location.

Figure 5 Delivery dimensions



Leave approximately 2.43 meters (8 ft) of clearance at the back of the cabinet to unload the unit and roll it off the pallet.

Figure 6 Cabinet clearance



Installed clearance

The 40U-P and 40U-D Titan rack cabinets ventilate from front to back. Provide adequate clearance to service and cool the system.

Depending on component-specific connections within the cabinet, the available power cable length may be somewhat shorter than the 15-foot standard.

Figure 7 40U-P






Figure 8 40U-D

Caster wheels

The bottom of the 40U-P and 40U-D Titan rack cabinets includes four caster wheels. Of the four wheels on the bottom of the rack, the two front wheels are fixed, and the two rear casters swivel in a 1.75-inch diameter. The swivel position of the caster wheels determines the load-bearing points on the site floor, but does not affect the cabinet footprint. After you position, level, and stabilize the cabinet, the four leveling feet determine the final load-bearing points on the site floor.

Figure 9 Caster wheels



CL3627

WARNING The data center floor on which you configure the system must support that system. You are responsible to ensure that the data center floor can support the weight of the system. Whether, or not the system is configured directly on the data center floor, or on a raised floor supported by the data center floor. Failure to comply with these floor-loading requirements could result in severe damage to the system, the raised floor, sub-floor, site floor, and the surrounding infrastructure. Dell EMC fully disclaims any liability for damage or injury resulting from a failure to ensure that the raised floor, sub-floor, and/or site floor can support the system weight as specified. The customer assumes all risk and liability that is associated with such failure.

Stabilizer brackets


Optional brackets help to prevent the rack from tipping during maintenance or minor seismic events. If you intend to secure the optional stabilizer brackets to the site floor, prepare the location for the mounting bolts. Install an anti-tip bracket to provide an extra measure of security.

Note: There are two kits that can be installed. For cabinets with components that slide, it is recommended that you install both kits.


The following lists the stabilizer brackets:

- Anti-tip bracket - Use this bracket to provide an extra measure of anti-tip security. One or two kits may be used. For cabinets with components that slide, it is recommended that you use two kits.


Figure 10 Anti-tip bracket



- Anti-move bracket - Install an anti-move bracket to permanently fasten the unit to the floor.



- Seismic restraint bracket - Install a seismic restraint bracket to provide the highest protection from moving or tipping.

Figure 11 Seismic restraint bracket

Shock and vibration

The hardware enclosures have been tested to withstand shock and random vibration.

The response measurement levels apply to all three axes. Measure shock and vibration with an accelerometer on the equipment enclosure within the cabinet. The measurements should not exceed the levels in the following table.

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, 5-500Hz, 30 minutes
Operational random vibration	0.21 Grms, 5-500Hz, 10 minutes


AC power input

After the 40U-P and 40U-D Titan racks are positioned and loaded, connect power cords to the P1 and P2 connectors on the four power distribution units (PDU) within the cabinet.

Depending on the cabinet components and configuration, the 40U-P rack requires two, four, or six independent 200–240 V power sources. Power cords included with the shipment support the maximum configurations. There might be extra cords as part of the shipment.


CAUTION 40U-P cabinet PDUs do not include a power ON/OFF switch. Ensure that the four circuit breaker switches on each PDU are up, in the off position, until AC power is supplied to the unit. The power must be off before disconnecting jumper or power cords from a PDU.

Attach the power cords to the power distribution units on each side of the rack. The following image displays where to attach two AC source connections.


Figure 12 AC power input

The following image displays where to attach four AC source connections.

Figure 13 Four AC connections



The following image displays where to attach six AC source connections.

Figure 14 Six AC connections

Third party rack specifications for the Dell EMC Generation 6 deep chassis solution

The current Dell EMC rack solutions support up to 8 PDUs (4 on each side). The illustrations in this section provide the dimensions and guidelines for 3rd party rack solutions. The table following the illustrations lists the components and dimensions for the labels in the illustrations.

Figure 15 Rear facing rack




Figure 16 Center facing rack




Table 5 3rd party rack components and dimensions

Rack component label	Description
a	Distance between front surface of the rack and the front NEMA rail
b	Distance between NEMA rails, minimum=24in (609.6mm), max=34in (863.6mm)
c	Distance between the rear of the chassis to the rear of the rack, min=2.3in (58.42mm)
d	Distance between inner front of the front door and the NEMA rail, min=2.5in (63.5mm)

Table 5 3rd party rack components and dimensions (continued)

Rack component label	Description
e	Distance between the inside of the rear post and the rear vertical edge of the chassis and rails, min=2.5in (63.5mm)
f	Width of the rear rack post
g	19in (486.2mm)+(2e), min=24in (609.6mm)
h	<p>19in (486.2mm) NEMA+(2e)+(2f)</p> <p>Note:</p> <ul style="list-style-type: none"> Width of the PDU+0.5in (13mm) <=e +f If j=i+c+PDU depth+3in (76.2mm), then h=min 23.6in (600mm) <p>Assuming the PDU is mounted beyond i+c.</p>
i	<p>Chassis depth:</p> <ul style="list-style-type: none"> Normal chassis=35.80in (909mm) Deep chassis=40.40in (1026mm) <p>Switch depth (measured from the front NEMA rail):</p> <ul style="list-style-type: none"> Arista 7304x and 7308x = 23.74in (909mm) Celestica D2024, D2060 and D4040 = 35in (889mm) to the switch port plate <p>Note:</p> <ul style="list-style-type: none"> The inner rail is fixed at 36.25in (921mm) Allow up to 6in (155mm) for cable bend radius when routing up to 32 cables to one side of the rack Select the greater of the installed equipment
j	Minimum rack depth=i+c
k	Front
l	Rear
m	Front door
n	Rear door
p	Rack post
q	PDU
r	NEMA

Table 5 3rd party rack components and dimensions (continued)

Rack component label	Description
s	NEMA 19 inch
t	Rack top view
u	Distance from front NEMA to chassis face: <ul style="list-style-type: none"> • Dell EMC Generation 6 deep and normal = 0in • Arista switches = 0in • Celestica D2024 = 19.2in (488mm) • Celestica D2060 and D4040 = 14.5in (368mm)

Switches and cables


Select network switches and cables that are compatible with the Isilon nodes, and that support the network topology.


Isilon Generation 4 and Generation 5 nodes use standard copper Gigabit Ethernet (GigE) switches for the front-end (external) traffic and InfiniBand for the back-end (internal) traffic. Generation 6 uses 10GbE and 40GbE for the back-end, and InfiniBand is optional for existing clusters only.

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

If you do not use a front end switch that is recommended by Isilon, the switch must meet the following minimum specifications:

- 10/40 GbE support
- Nonblocking fabric switch
- Minimum of 1 MB per port of packet buffer memory
- Support for jumbo frames (if you intend to use this feature)

 **CAUTION** Isilon requires that separate switches are used for the external and internal interfaces. Isilon Nodes support Ethernet traffic on the front end and InfiniBand for Generation 5 and older nodes and Ethernet and InfiniBand for Generation 6 nodes.

 **CAUTION** Only Isilon nodes should be connected to the back-end (internal) switch. Information that is exchanged on the back-end network is not encrypted. Connecting anything other than Isilon nodes to the back-end switch create 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 Isilon nodes.

Protect cables


Damage to the InfiniBand or Ethernet cables (copper or optical Fibre) can affect the Isilon 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 17 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 μ m cables for all optical connections.

Table 6 Optical cables

Cable type	Operating speed	Length
50 μ m	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

Table 6 Optical cables (continued)

Cable type	Operating speed	Length
	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 μ m	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 μ m or 62.5 μ m optical cables includes two connections or splices between the source and the destination.

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

Rail installation for switches

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

The following internal network switches ship with Dell EMC 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.

NOTE: The Celestica Ethernet rails are designed to overhang the rear NEMA rails to align the switch with the Generation 6 chassis at the rear of the rack. The rails require a minimum clearance of 36 in from the front NEMA rail to the rear of the rack to ensure that the rack door can be closed.

Table 7 Infiniband switches

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

Table 8 Dell Z9100-ON Ethernet switch

Switch	Ports	Network
Dell Z9100-ON	32-port	32x100GbE, 32x40GbE, 128x10GbE (with breakout cables)

Note: In OneFS 8.2.1, the Dell Z9100-ON switch is required if you plan to implement Leaf-Spine networking for large clusters.

Table 9 Celestica Ethernet switches

Switch	Ports	Network
Celestica D2024	24-port	10GbE
Celestica D4040	32-port	40GbE
Celestica D2060	48-port	10GbE

Note: There is no breakout cable support for Arista switches. However, you can add a 10GbE or 40GbE line card depending on the Arista switch model. Details are included in the following table.

Table 10 Arista Ethernet switches

Switch	Nodes
Arista DCS-7304	Shipped with 2 line cards, each with 48 10GbE ports to a maximum of 144 nodes, you can add either of the following: <ul style="list-style-type: none"> 1 additional line card of 48 ports 10GbE 2 additional 32 port line card of 40GbE
Arista DCS-7308	Shipped with 2 line cards. Each with 32 port of 40GbE to a maximum of 144 nodes, you can add either of the following: <ul style="list-style-type: none"> 3 additional 32GbE line cards 2 additional 48 ports of 10GbE line cards (or combination of 10 and 40GbE line cards)

The following table lists rack and power requirements for Arista switches.

Table 11 Arista switch requirements

Switch	Rack enclosure	PDU
Arista DCS-7304	8U	C19
Arista DCS-7308	13U	

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 OU cable management option available at this time.

Procedure

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 18 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 19 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 20 Secure the switch to the rail



8. Snap the bezel in place.

Twinaxial cabling for 10 Gigabit and 40 Gigabit ethernet environments

Active twinaxial copper cables include small form factor pluggable (SFP) modules at either end. They connect Fibre Channel over Ethernet (FCoE) modules to FCoE switches, 40GbE, or 10GbE/iSCSI modules to an external iSCSI environment.

The twinaxial cables are shielded, quad-construction style, 100Ω differential. The SFP connectors comply with SFF-8431 and SFF-8472 standards. To facilitate 40GbE or 10GbE transmission, transmit and receive ends include active components. Various lengths of twinaxial cable are available, from 1 meter to 15 meters and wire gauges range from 24 to 30 with cable length.

HD400 cable management accessories

A cable management accessory must accompany every HD400 node that you install.

To ensure proper airflow around the HD400 node and to aid with maintenance, you must install one of two cable management accessories with every node. You can either install cable management arms on the back of the node, or a 1U cable management tray in the one unit space directly beneath the node.

Note: The 35 in depth of the HD400 does not include additional space required for cable management arms.

Installation instructions for these accessories are available in the *HD400 Installation Guide*.

Leaf-Spine topology

OneFS 8.2 supports Leaf-Spine network topology for internal networks that communicate with the nodes that form clusters. For large Generation 6 clusters or clusters that will grow significantly over time, the Leaf-Spine topology is recommended.

Architecture

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

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

Table 12 Leaf-Spine network components

Component	Description	Connection considerations
Spine	Dell Z9100-ON 32-port switch	Back-end network with 100GbE (uplink) connects to the Leaf
Leaf	Dell Z9100-ON 32-port switch	<ul style="list-style-type: none"> 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 node	F800, H500, and H600 Isilon nodes	Performance nodes support a 40GbE connection to the Leaf
Archive node	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
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 Isilon node connection	


Network topology

External networks connect the cluster to the outside world.

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, if a network topology has nodes that connect to one external IP subnet, nodes that connect to a second IP subnet, and nodes that do not connect at all externally.

 **Note:** 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 10GbE or 40GbE, and InfiniBand is optional for older clusters.


To configure the cluster, set up an initial network. Optionally, set up an alternate interface as a fail over network. The Internal-A port (int-a) is the initial network. Configuration of Internal-A is required for proper operation of the cluster. The Internal-B port (int-b) is the alternate interface for internal communications and can also be used for failover.

The Internal-A Port

When setting up the cluster, connect the Internal-A port of each node to the switch that supports the Internal-A segment of the internal network.

The Internal-B Failover Port

Optionally, an Internal-B/failover network can be configured to provide the cluster with continued internal communications in the event of a faulty switch or other networking infrastructure failure.

 **CAUTION** Only Isilon nodes should be connected to the back-end internal switch. Information that is exchanged on the back-end network is not encrypted. Connecting anything other than Isilon nodes to the back-end switch creates a security risk.

Assisting with installation

Dell EMC Isilon Professional Services can help plan a cluster that best supports the workflow. Contact an Isilon Account Manager for more information.

Installation and implementation details

Three to four weeks prior to your installation date, the Dell EMC Isilon Professional Services team will help to gather the information necessary to configure the Isilon cluster.

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

Be prepared to discuss the following information with the Dell EMC Isilon project team:

- Data workflow, including the role of the cluster in that workflow: production, test, or disaster recovery.
- The OneFS version to install on the cluster.
- Network connectivity details, including IP ranges, for the client and InfiniBand networks.
- DNS configuration details, including name servers and search lists.
- 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.
- Data protection levels, antivirus solutions, and NDMP backup.
- Cluster alert solutions such as SupportIQ and SNMP monitoring.

CHAPTER 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 an EMC Isilon sales representative.

- [Data management modules](#)..... 58
- [Data protection modules](#)..... 60

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.

SmartPools

The SmartPools module enables you to administer node pools and storage tiers, and to create policies to manage files on a granular level.

OneFS provides functions such as autoprovisioning, compatibilities, virtual hot spare (VHS), SSD strategies, global namespace acceleration (GNA), L3 cache, and storage tiers.

The following table compares storage pool features based on whether a SmartPools license is active.

Function	Inactive SmartPools license	Active SmartPools license
Automatic storage pool provisioning	Yes	Yes
Node class compatibilities (node equivalency)	Yes	Yes
SSD capacity compatibilities	Yes	Yes
SSD count compatibilities	Yes	Yes
Virtual hot spare	Yes	Yes
SSD strategies	Yes	Yes
L3 cache	Yes	Yes
Tiers	Yes	Yes
GNA	Yes	Yes
File pool policies	No	Yes
Spillover management	No	Yes

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.

InsightIQ

The InsightIQ module provides advanced monitoring and reporting tools to help streamline and forecast cluster operations.

InsightIQ helps to create customized reports containing key cluster performance indicators such as:

- Network traffic on a per-interface, per-node, per-client, and per-protocol basis.
- Protocol operation rates and latencies that are recorded on a per-protocol, per-client, and per-operation class basis.
- Per-node CPU utilization and disk throughput statistics.

To run the Isilon OneFS virtual appliance, the environment must meet the following minimum system requirements.

Isilon cluster

The monitored cluster must be running version 5.5.3 or later of the operating system. The InsightIQ File System Analytics functionality requires OneFS 6.0 or later. The available InsightIQ features depend on the OneFS version that the monitored system is running.

For monitored clusters running OneFS 7.0 and later, enable HTTPS port 8080. For monitored clusters running an earlier version of OneFS, enable HTTPS port 9443. For the File System Analytics feature, enable the NFS service, HTTPS port 111, and HTTPS port 2049 on all monitored clusters.

InsightIQ product license

Before installing the InsightIQ virtual appliance, obtain an InsightIQ license key for each cluster that to monitor. For more information, contact Dell EMC Isilon Customer Support. After you obtain the license keys, activate the licenses on each cluster that you want to monitor.

Web browser

You can access the Isilon InsightIQ application through any web browser that supports sophisticated graphics. Examples of supported browsers include

Microsoft Internet Explorer 9, Mozilla Firefox, Apple Safari 5, and Google Chrome.

To access the InsightIQ application through Internet Explorer, you must first install the Google Chrome Frame plug-in. If InsightIQ cannot detect the plug-in, an installation prompt appears the first time that you try to access InsightIQ through Internet Explorer.

The InsightIQ web application interface is optimized for a screen that is 1,280 pixels wide.

Storage space for the InsightIQ data store

InsightIQ can store monitored-cluster data either locally or on an NFS-mounted server. If InsightIQ is installed as a virtual appliance, InsightIQ by default stores the monitored-cluster data locally on a virtual hard drive that is included with the InsightIQ virtual appliance. If you want to use this local virtual hard drive as the InsightIQ data store, ensure that the virtualization host contains at least 70 GB of free disk space. If InsightIQ is installed as an RPM, InsightIQ stores data on the local computer by default. If you store InsightIQ data on the local computer, ensure that the computer contains at least 64 GB of free disk space.

As an alternative to storing InsightIQ data locally, you can configure InsightIQ to store monitored-cluster data on an Isilon cluster—either the monitored cluster itself or a different cluster—or on any NFS-mounted server. If you want to store InsightIQ monitored-cluster data on an Isilon cluster or other NFS-based server, you must provide adequate space in which to store the data. Verify that an NFS export rule is correctly configured on the cluster or NFS-based server. For information about NFS data store requirements, see the *InsightIQ Installation Guide*.

Data protection modules

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

You can install advanced data protection modules to enable enhanced replication, restoration, dynamic cluster load balancing, and technical support.

SnapshotIQ

The SnapshotIQ module provides the ability to create and manage snapshots on the Isilon cluster.

A snapshot contains a directory on a cluster, and includes all data that is stored in the specified directory and its subdirectories. If data contained in a snapshot is modified, the snapshot stores a physical copy of the original data and references the copied data. Snapshots are created according to user specifications, or generated automatically by OneFS to facilitate system operations.

To create and manage snapshots, you must activate a SnapshotIQ license on the cluster. Some applications must generate snapshots to function, but do not require you to activate a SnapshotIQ license by default. The snapshots are automatically deleted when the system no longer needs them. However, if a SnapshotIQ license is active on the cluster, some applications can retain snapshots. You can view auto-generated snapshots regardless of whether a SnapshotIQ license is active.

The following table lists the available snapshot functionality depending on whether a SnapshotIQ license is active.

Functionality	Inactive	Active
Create snapshots and snapshot schedules	No	Yes
Configure SnapshotIQ settings	No	Yes
View snapshot schedules	Yes	Yes
Delete snapshots	Yes	Yes
Access snapshot data	Yes	Yes
View snapshots	Yes	Yes

SyncIQ

The SyncIQ module enables you to replicate data from one Isilon cluster to another.

With SyncIQ, you can replicate data at the directory level while optionally excluding specific files and sub-directories from being replicated. SyncIQ creates and references snapshots to replicate a consistent point-in-time image of a root directory. Metadata such as access control lists (ACLs) and alternate data streams (ADS) are replicated along with data.

SyncIQ enables you to retain a consistent backup copy of your data on another Isilon cluster. SyncIQ offers automated failover and fallback capabilities that enable you to continue operations on another Isilon cluster if a primary cluster becomes unavailable.


SmartLock

The SmartLock module allows you to prevent users from modifying and deleting files on protected directories.

Use the SmartLock tool to create SmartLock directories and commit files within those directories to a write once, read many (WORM) state. You cannot erase or re-write a file that is committed to a WORM state. You can delete a file that has been removed from a WORM state, but you cannot modify a file that has ever been committed to a WORM state.

Note the following SmartLock considerations:

- Create files outside of SmartLock directories and transfer them into a SmartLock directory only after you finish working with the files.
Upload files to a cluster in two steps.
 1. Upload the files into a non-SmartLock directory.
 2. Transfer the files to a SmartLock directory.

 **Note:** Files committed to a WORM state while being uploaded will become trapped in an inconsistent state.

Files can be committed to a WORM state while they are still open. If you specify an autocommit time period for a directory, the autocommit time period is calculated according to the length of time since the file was last modified, not when the file was closed. If you delay writing to an open file for more than the autocommit time period, the file will be committed to a WORM state the next time you attempt to write to it.

- In a Microsoft Windows environment, if you commit a file to a WORM state, you can no longer modify the hidden or archive attributes of the file. Any attempt to modify the hidden or archive attributes of a WORM committed file will generate an


error. This can prevent third-party applications from modifying the hidden or archive attributes.

SmartLock compliance mode

SmartLock compliance mode enables you to protect your data in compliance with the regulations defined by U.S. Securities and Exchange Commission rule 17a-4.

You can upgrade a cluster to SmartLock compliance mode during the initial cluster configuration process, before you activate a SmartLock license. To upgrade a cluster to SmartLock compliance mode after the initial cluster configuration process, contact Isilon Technical Support.

If you set a cluster to compliance mode, you will not be able to log in to that cluster through the root user account. Instead, you can log in to the cluster through the compliance administrator account. You must configure the compliance administrator account during the initial cluster configuration process. If you are logged in through the compliance administrator account, you can perform administrative tasks through the `sudo` command.

 **Note:** Compliance mode is not compatible with Isilon for vCenter, VMware vSphere API for Storage Awareness (VASA), or the vSphere API for Array Integration (VAAI) NAS Plug-In for Isilon.

SmartConnect Advanced

The SmartConnect Advanced module adds enhanced balancing policies to evenly distribute CPU usage, client connections, or throughput.

If you activate a SmartConnect Advanced license, you are also able to:

- Enable dynamic IP allocation and IP failover in your cluster.
- Define IP address pools to support multiple DNS zones in a subnet.
- Establish multiple pools for a single subnet.

SupportIQ

The SupportIQ module allows Isilon Technical Support, with your permission, to securely upload and analyze your OneFS logs to troubleshoot cluster problems.

When SupportIQ is enabled, Isilon Technical Support personnel can request logs through scripts that gather cluster data and then upload the data to a secure location. You must enable and configure the SupportIQ module before SupportIQ can run scripts to gather data.

You can also enable remote access, which allows Isilon Technical Support personnel to troubleshoot your cluster remotely and run additional data-gathering scripts. Remote access is disabled by default. To enable remote SSH access to your cluster, you must provide the cluster password to a Technical Support engineer.

Antivirus planning

You can scan the OneFS file system for computer viruses and other security threats by integrating with third-party scanning services through the Internet Content Adaptation Protocol (ICAP). This feature does not require you to activate a license.

If an ICAP server detects a threat it notifies OneFS. OneFS creates an event to inform system administrators, displays near real-time summary information, and documents the threat in an antivirus scan report. You can configure OneFS to request that ICAP


servers attempt to repair infected files. You can also configure OneFS to protect users against potentially dangerous files by truncating or quarantining infected files.

ICAP servers

The number of ICAP servers that are required to support an Isilon cluster depends on how virus scanning is configured, the amount of data a cluster processes, and the processing power of the ICAP servers.

If you intend to scan files exclusively through antivirus scan policies, it is recommended that you have a minimum of two ICAP servers per cluster. If you intend to scan files on access, it is recommended that you have at least one ICAP server for each node in the cluster.

If you configure more than one ICAP server for a cluster, it is important to ensure that the processing power of each ICAP server is relatively equal. OneFS distributes files to the ICAP servers on a rotating basis, regardless of the processing power of the ICAP servers. If one server is significantly more powerful than another, OneFS does not send more files to the more powerful server.

 **CAUTION** When files are sent from the cluster to an ICAP server, they are sent across the network in cleartext. Make sure that the path from the cluster to the ICAP server is on a trusted network.

Supported ICAP servers

OneFS supports anti-virus scan applications that follow the ICAP standard ([ICAP RFC 3507](#)).

For information about implementing anti-virus solutions with OneFS, see the [Anti-virus Solutions with EMC Isilon Scale-out NAS](#) white paper (requires login to [EMC Online Support](#)).

CHAPTER 4





Preparing your facility

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

- [Generation 6 environmental requirements](#)..... 66
- [Pre-Generation 6 environmental requirements](#)..... 71
- [Fire suppressant disclaimer](#)..... 75

Generation 6 environmental requirements

To support the recommended operating parameters for Isilon equipment, prepare the site

Environmental item	Description
	Operate at ≥ 90 percent of the time between 10 degrees celsius to 35 degrees celsius, and ≤ 10 percent of the time between 5 degrees celsius to 40 degrees celsius. Note: There are four independent cooling zones so that a cooling failure in one node, does not impact other nodes.
	40 to 55 percent relative humidity. Note: The cluster can be qualified to operate outside these limits. Product-specific documentation for system specifications provides more information.
	A fully configured cabinet must sit on at least two floor tiles, and can weigh approximately 1588 kilograms (3500 pounds).
	0 meters to 2439 meters (0 to 8,000 ft) above sea level operating altitude.

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:

- Dell EMC recommends 24 x 24 inches or (60 x 60 cm) heavy-duty, concrete filled steel floor tiles.
- 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 as 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 than 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 using the tools available at <http://powercalculator.emc.com>.

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. For Dell EMC products, use the Dell EMC Power Calculator available at, <http://powercalculator.emc.com>.

Table 13 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
Power zones	Two	Two
Power requirements at site (minimum to maximum)	<ul style="list-style-type: none"> 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 <p>Note: The options for the single-phase PDU interface connector are listed in the following table, Single-phase AC power input connector options.</p>	

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

Table 14 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

Table 14 Single-phase AC power input connector options (continued)








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

Table 15 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)	<p>North America (Delta):</p> <ul style="list-style-type: none"> 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. <p>International (Wye):</p> <ul style="list-style-type: none"> One 32 A, three-phase drop per zone Each Wye rack requires two 32 A drops. <p> Note: The interface connector options for the Delta and Wye three-phase PDUs are listed in the following tables, and .</p>	

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

Table 16 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 17 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

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 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.

Air quality requirements

Dell EMC products are designed to be consistent with the requirements of the American Society of Heating, Refrigeration, and Air Conditioning Engineers (ASHRAE) Environmental Standard Handbook and the most current revision of Thermal Guidelines for Data Processing Environments, Second Edition, ASHRAE 2009b.

Dell EMC cabinets are best suited for Class 1 datacom environments, which consist of tightly controlled environmental parameters, including temperature, dew point, relative humidity, and air quality. The 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. In addition, efforts should be maintained 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%. For data centers with gaseous contamination, such as high sulfur content, lower temperatures, and humidity are recommended to minimize the risk of hardware corrosion and degradation. In general, the humidity fluctuations within the data center should be minimized. It is also recommended that the data center be positively pressured and have air curtains on entry ways to prevent outside air contaminants and humidity from entering the facility.

For facilities below 40% relative humidity, it is recommended to 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, it is recommended to place copper and silver coupons (per ISA 71.04-1985, Section 6.1 Reactivity), in airstreams representative of what is in the data center. The monthly reactivity rate of the coupons should be fewer than 300 Angstroms. When monitored reactivity rate is exceeded, the coupon should be analyzed for material species and a corrective mitigation process that is put in place.

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	<ul style="list-style-type: none"> 1.15 Grms 1 hour Frequency range 1–200 Hz

Pre-Generation 6 environmental requirements

To support the recommended operating parameters of Isilon equipment, prepare the site.

	+15°C to +32°C (59°F to 89.6°F) site temperature. A fully configured cabinet can produce up to 16,400 BTUs per hour
	40% to 55% relative humidity
	A fully configured cabinet sits on at least two floor tiles, and can weigh approximately 1,182 kilograms (2600 pounds)
	0 meters to 2439 meters (0 ft to 8,000 ft) above sea level operating altitude
	LAN and telephone connections for remote service and system operation

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

Shipping and storage requirements

NOTICE 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 18 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.

Floor load bearing requirements

Install node racks in raised or non-raised floor environments capable of supporting at least 2,600 lbs (1,180 kg) per rack.

Although the system configuration might weigh less, floor support that is rated at a minimum of 2,600 lbs (1,180 kg) per rack is required to accommodate any equipment upgrades or reconfiguration.

General floor requirements:

- Position the cabinet as 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 or the subfloor (non-raised floor).
- Ensure that the floor can support the system and surrounding weight by having a certified data center design consultant inspect the site. The overall weight of the equipment depends on the type and quantity of nodes, switches, and racks.

Raised floor requirements

- Dell EMC recommends that 24 in. x 24 in. (60 cm x 60 cm) or heavy-duty, concrete-filled steel floor tiles are used.
- Use only floor tiles and stringers rated to withstand:
 - concentrated loads of two casters or leveling feet, each weighing up to 1,000 lb (454 kg).
 - minimum static ultimate load of 3,000 lb (1,361 kg).
 - rolling loads of 1,000 lb (454 kg). On floor tiles that do not meet the 1,000 lb rolling load rating, use coverings such as plywood to protect floors during system roll.
- Position adjacent cabinets with no more than two casters or leveling feet on a single floor tile.
- Cutouts in 24 in. x 24 in. (60 cm x 60 cm) tiles must be no more than 8 in. (20.3 cm) wide by 6 in. (15.3 cm) deep, and centered on the tiles, 9 in. (22.9 cm) from the front and rear and 8 in. (20.3 cm) from the sides. Cutouts weaken the tile, but 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.

Air quality requirements

Dell EMC products 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.

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
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.

Radio Frequency Interference (RFI) requirements

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

Dell EMC products are 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 Dell EMC 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 Dell EMC equipment, use the following table as a guide.

Table 19 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

Isilon equipment power requirements

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

Plan to set up redundant power for each rack that contains Isilon nodes. Supply the power with a minimum of two separate circuits on the building's electrical system. If one of the circuits fails, one or more remaining circuits should be able to 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 Dell EMC Isilon 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 as well as 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 will sustain 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.

Power requirements

Power cords and connectors depend on the type ordered with your system, and must match the supply receptacles at your site.

Power cord connector	Operating voltage/ frequency	Service type	Site
 NEMA L6-30P	200-240 V ac, 50/60 Hz	30-amp service, single phase	North America, Japan
 IEC-309-332P6	200-240 V ac, 50/60 Hz	32-amp service, single phase	International
 56PA332 Right Angle	240 V ac, 50/60 Hz	32-amp service, single phase	Australia

Each AC circuit requires a source connection that can support a minimum of 4800 VA of single phase, 200-240 V AC input power. For high availability, the left and right sides of any rack or cabinet must receive power from separate branch feed circuits.

Note: Each pair of power distribution panels (PDP) in the 40U-C cabinet can support a maximum of 24 A AC current draw from devices connected to its power distribution units (PDU). Most cabinet configurations draw less than 24 A AC power, and require only two discrete 240 V AC power sources. If the total AC current draw of all the devices in a single cabinet exceeds 24 A, the cabinet requires two additional 240 V power sources to support a second pair of PDPs. Use the published technical specifications and device rating labels to determine the current draw of each device in your cabinet and calculate the total.

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 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.