

Dell EMC PowerScale: Ethernet Back-end Network Overview

Abstract

This white paper provides an introduction to the Ethernet back-end network for Dell EMC™ PowerScale™ scale-out NAS.

June 2020

Revisions

Date	Description
June 2020	Content and template update

Acknowledgments

Author: Abiy Mesfin

The information in this publication is provided “as is.” Dell Inc. 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 software described in this publication requires an applicable software license.

Copyright © 2020 Dell Inc. or its subsidiaries. All Rights Reserved. Dell Technologies, Dell, EMC, Dell EMC and other trademarks are trademarks of Dell Inc. or its subsidiaries. Other trademarks may be trademarks of their respective owners. [6/16/2020] [Technical White Paper] [H16346]

Table of contents

Revisions.....	2
Acknowledgments.....	2
Table of contents	3
Executive summary.....	4
1 Legacy Isilon back-end network.....	5
2 Isilon platform back-end network option.....	6
3 PowerScale platform back-end network option.....	7
3.1 Ethernet back end	7
3.2 Dell EMC switch support for Ethernet back end.....	8
3.3 Configuration and monitoring	8
3.4 Troubleshooting.....	10
A Technical support and resources	26
A.1 Related resources	26

Executive summary

This document provides design considerations for Dell EMC™ PowerScale™ back-end (internal) networking. This back-end network, which is configured with redundant switches for high availability, acts as the backplane for the Dell EMC Isilon™ cluster. This backplane enables each Isilon node to act as a contributor in the cluster and provides node-to-node communication with a private, high-speed, low-latency network.

1 Legacy Isilon back-end network

Prior to the introduction of the latest generation of Isilon scale-out NAS storage platforms, inter-node communication in an Isilon cluster has been performed using a proprietary, unicast (node-to-node) protocol known as Remote Block Manager (RBM). This inter-node communication uses a fast low-latency, InfiniBand (IB) network. This back-end network, which is configured with redundant switches for high availability, acts as the backplane for the Isilon cluster. This backplane enables each Isilon node to act as a contributor in the cluster and provides node-to-node communication with a private, high-speed, low-latency network. This back-end network utilizes Internet Protocol (IP) over IB (IPoIB) to manage the cluster. Sockets Direct Protocol (SDP) is used for all data traffic between nodes in the cluster.

2 Isilon platform back-end network option

Isilon scale-out NAS storage platforms offer increased back-end networking flexibility. With Isilon platforms, customers may choose to use either an InfiniBand or Ethernet switch on the back end. For customers electing to use an InfiniBand back-end network, the configuration and implementation will remain the same as previous generations of Isilon systems. Customers looking to add Isilon platforms (Isilon F800, H600, H5600, H500, H400, A200, and A2000) to an existing Isilon IB cluster comprised of earlier Isilon systems, will need to configure the nodes with an InfiniBand back-end interface. The Ethernet back-end network option is only supported in clusters that are comprised entirely of Ethernet back-end nodes. In these configurations, only Ethernet back-end switches that are provided and managed by Dell EMC will be supported.

The Isilon back-end Ethernet connection options are detailed in Table 1.

Table 1 Latest-generation Isilon back-end Ethernet options

Back-end options	Compute compatibility
10 GbE SFP+	Isilon H400, Isilon A200, or Isilon A2000
40 GbE QSFP+	Isilon F800/F810, Isilon H600, Isilon H5600, or Isilon H500

In general, high-performance platforms such as the Isilon F800 all-flash or Isilon H600 hybrid scale-out NAS platforms will typically utilize the bandwidth capabilities of 40 GbE ports. Lower-performance platforms such as the Isilon A200 or A2000 archive scale-out NAS platforms will typically be well supported with the bandwidth provided by 10 GbE ports.

3 PowerScale platform back-end network option

The new Dell EMC PowerScale all-flash storage platforms, powered by the Dell EMC PowerScale OneFS operating system, provide a powerful and simple scale-out storage architecture to speed up access to massive amounts of unstructured data. Powered by the new OneFS 9.0 operating system, the all-flash PowerScale platforms are available in 2 product lines:

PowerScale F200: Provides the performance of flash storage in a cost-effective form factor to address the needs of a wide variety of workloads.

PowerScale F600: With new NVMe drives, the F600 provides larger capacity with performance in a cost-effective compact form factor to power the most demanding workloads.

With the PowerScale platforms, currently the only back-end network option is Ethernet.

The PowerScale back-end Ethernet connection options are detailed in Table 2.

Back-end options	PowerScale nodes
10 GbE SFP+	F200
25 GbE SFP28	
40 GbE QSFP+	F600
100 GbE QSFP28+	

Note: The same NIC supports both 10 GbE and 25 GbE for the F200, and the same NIC supports both 40 GbE and 100 GbE for the F600.

New-generation PowerScale platforms with different back-end speeds can connect to the same switch with Isilon nodes (Isilon F800, H600, H5600, H500, H400, A200, and A2000) and not see performance issues. For example, in a mixed cluster where an archive node (such as A200 or A2000) with 10 GbE on the back end and PowerScale nodes with 40 GbE or 100 GbE on the back end, both node types can connect to a 100 GbE back-end switch without affecting the performance of other nodes on the switch. The 100 GbE back-end switch will provide 100 GbE to the ports servicing the high-performance PowerScale nodes and 10 GbE to the archive or lower performing nodes using breakout cables.

3.1 Ethernet back end

In legacy Isilon systems, back-end data traffic uses SDP and IPoIB for management. SDP has fast failover and incorporates a variety of InfiniBand-only features that ensures optimum performance. However, because SDP only works over InfiniBand, a new method was required to get optimal performance over the Ethernet back end. For this reason, the new generation of Isilon platforms now uses RBM over TCP on the back-end switches.

RBM now uses TCP, and the TCP stack has been enhanced to provide the performance required to support the cluster communication. All the modifications of the TCP stack have been made while conforming to the industry standard specification of the stack. The back-end and front-end networks will use the same TCP

stack and modifications to the performance of the back-end TCP stack should not affect TCP traffic on the front end. RBM over Ethernet will still provide fast failover.

3.2 Dell EMC switch support for Ethernet back end

Dell Ethernet switches to be used for the Isilon back end as a top-of-rack solution (TOR).

- Z9264-ON
- Z9100-ON
- S4148F-ON
- S4112-ON

These Ethernet switches will be zero-touch back-end switches that are used for inter-node communication in an Isilon cluster, and those are typically what are called **plug and play**. They are shipped with a fixed configuration and additional customer configuration is not necessary or allowed.

The Z9264-ON is a fixed 2U Ethernet switch which provides industry-leading density of either 64 ports of 100 GbE or 40 GbE in QSFP28 or 128 ports of 25 GbE or 10 GbE by breakout. Breakout cables are only used in the odd-numbered ports and using one in odd-numbered port disables the corresponding even-numbered port.

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

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

The S4112F-ON supports 10/100GbE with 12 fixed SFP+ ports to implement 10 GbE and three fixed QSFP28 ports to implement 4x10 or 4x25 using breakout. A total of 24 10GbE connection including the three fixed QSFP28 ports using 4x10 breakout cables.

Note: These switches are qualified to be used with currently available network cables (MPO, LC, QSFP+, SFP+ and breakout cables). These switches are shipped with a custom operating system that is built specifically to be compatible with Dell EMC PowerScale family.

3.3 Configuration and monitoring

When installing a new Isilon cluster, the Configuration Wizard has not changed. It still prompts you for int-a, int-b, and failover range. All configuration and setup steps will be the same regardless of InfiniBand or Ethernet option selected.

Figure 1 below shows the relative positioning of back-end ports provided in the Compute Assembly for each Dell EMC Isilon node in the following platforms: F800, F810, H600, H5600, H500, H400, A200, and A2000




Figure 1 F200



Figure 2 F600



Figure 3 Isilon back-end ports

Table 2 provides configuration information for the back-end ports in PowerScale platforms:

Table 2 Configuration for int-a, int-b, and failover

Setting	Description
Int-a network setting	The network settings used by the int-a network. The int-a network is used for communication between nodes.
Netmask	The int-a network must be configured with IPv4.
IP range	The int-a network must be on a separate/distinct subnet from an int-b/failover network.
Int-b and failover network setting	The network settings used by the optional int-b/failover network.
Netmask	The int-b network is used for communication between nodes and provides redundancy with the int-a network.
IP range	The int-b network must be configured with IPv4.
Failover IP range	The int-a, int-b and failover networks must be on separate or distinct subnets.

The monitoring capabilities on Isilon switches correspond to the FRU (field replaceable unit) components such as power supply, the fan, or others. Protocol and performance monitoring capability is not provided. Customers should not attempt to alter the back-end network configurations provided by Dell EMC. Any attempt to do so can result in a cluster-wide outage.

For SNMP capabilities, customer may send an SNMP alert through the CELOG system. In today's back-end Ethernet world, we no longer have **opensm** topology files to view all connected devices on the back-end network. If you want to know what is connected to the fabric of back-end Ethernet (int-a or int-b) you may use the **isi_dump_fabric int-a** (or int-b) command.

3.4 Troubleshooting

In the past, to get the back-end networking information, we issued the **isi_eth_mixer_d** command. As the result, you could determine the back-end interfaces, or which of Int-a or int-b interfaces was currently being used. This information is now available using the **sysctl isi.lbfo.config** command. This will tell you, for each failover address that exists on other nodes, which interface is primary. There is no preference for one or the other, the current connection shows the path that was last used. Failover occurs in under a half second to the other route.

```
# sysctl isi.lbfo.config
```

```
isi.lbfo.config:
```

```
Node: 169.254.3.75, Int-A: mlxen0 (P), Int-B: mlxen1 (A)(C)
```

```
Node: 169.254.3.76, Int-A: mlxen0 (P)(C) Int-B: mlxen1 (A)
```

(P)=Primary

(A)=Alternate


(C)=Current path to the node

mlxen0= Mellanox EN card for int-a

mlxen1= Mellanox EN card for int-b

Example 1:


All Performance Dell EMC Isilon 40 GbE back end: When using performance nodes, the back end must be 40 GbE (10 GbE is not supported).



In this example, your configuration will include:

- Two 100/40 GbE back-end switches
- 16 QSFP+/MPO back-end cables
- 16 Optics (If MPO cables used)

Example 2: Mixed Environment of Isilon 10 and 40 GbE back end: When mixing performance and archive nodes, use a 40 GbE infrastructure with 40 GbE connections to the performance nodes and 4 x 10 GbE breakout cables to the archive nodes.



In this example, your configuration will include:

- Two 100/40 GbE back-end switches
- 8 QSFP+/MPO back-end cables
- 8 Optics (If MPO cables used)
- 4 QSFP to SFP+ breakout cables

Table 3 40 GbE switch options

Vendor	Model	Legacy Isilon Model	Dell SKU	Back-end ports	Port type	Rack units	100 GbE nodes	Mixed environment (10, 25, 40, and 100 GbE)
Dell EMC	Z9264-ON	851-0318	210-AWOW	64	All 100 GbE	2	64 or less	Support breakout cables of 4 x 10 or 4 x 25. total 128 10 GbE or 25 GbE nodes
Dell EMC	Z9100-ON	851-0316	210-AWOV /210-AWOU	32	All 100 GbE	1	32 or less	Support breakout cables of 4 x 10 or 4 x 25. total 128 10 GbE or 25 GbE nodes
*Celestica	D4040	851-0259	N/A	32	All 40 GbE	1	100 GbE not supported	Support breakout cables, total 96 10 GbE nodes
*Arista	DCS-7308	851-0261	N/A	64	All 40 GbE	13	100 GbE not supported	No breakout support with FT, but you can add 10 GbE line card
*Arista		851-0282	N/A	leaf upgrade (32 ports)	All 40 GbE		252-node cluster will require 6 additional leaves	OneFS 8.2 is required to scale beyond 144 nodes using Arista switches.

Note: For the Celestica 851-0259, you can use 24 breakout cables to connect 96 nodes at 10G, though only ports 1 - 12 and 17 - 28 can break out (this is a Celestica design limitation). Breakout cables do not require manual configuration on the switch – they are plug and play.

Table 4 Cable options for F800 (also H600 and H500)

Cable type (passive)	Legacy Isilon Model	Connector	Length	EMC P/N	Dell SKU number	Reason
Copper	851-0253	QSFP+	1m	038-002-064-01	470-AEGB	Ethernet cluster
Copper	851-0254	QSFP+	3m	038-002-066-01	470-AEGG	Ethernet cluster
Copper	851-0255	QSFP+	5m	038-002-139-01	470-AEGM	Ethernet cluster

Cable type (passive)	Legacy Isilon Model	Connector	Length	EMC P/N	Dell SKU number	Reason
Optical	851-0274	MPO	1m	038-004-214	407-BCIV	Ethernet/IB cluster
Optical	851-0275	MPO	3m	038-004-216	407-BCIW	Ethernet/IB cluster
Optical	851-0276	MPO	5m	038-004-227	407-BCJD	Ethernet/IB cluster
Optical	851-0224	MPO	10m	038-004-218	407-BCIY	Ethernet/IB cluster
Optical	851-0225	MPO	30m	038-004-219	407-BCJB	Ethernet/IB cluster
Optical	851-0226	MPO	50m	038-004-220	407-BCJC	Ethernet/IB cluster
Optical	851-0227	MPO	100m	038-004-221	407-BCIZ	Ethernet/IB cluster
Optical	851-0277	MPO	150m	038-000-139	407-BCIX	Ethernet/IB cluster


Note: QSFP+ cables for Ethernet use do not require optics. MPO cables for Ethernet use require passive optics. The model is 851-0285 (019-078-046)(Dell SKU 407-BCJE). MPO optics are added automatically when MPO cables are quoted.

Table 5 Cable options for Isilon A200 (also Isilon A2000 and H400)

Cable type	Legacy Isilon Model	Length	Connector	Optic part #	Dell SKU number	EMC P/N	Reason
Copper	851-0278	1m	(1) QSFP to (4) SFP+	N/A	470-AEGC	038-004-506-03	Breakout: 40Ge/10Ge (4)
Copper	851-0279	3m	(1) QSFP to (4) SFP+	N/A	470-AEGH	038-004-507-03	Breakout: 40Ge/10Ge (4)
Copper	851-0280	5m	(1) QSFP to (4) SFP+	N/A	470-AEGN	038-004-508-03	Breakout: 40Ge/10Ge (4)

Note: Breakout cables do not require optics. For the Celestica 851-0259, you can use 24 breakout cables to connect 96 nodes at 10G, though only ports 1 - 12 and 17 - 28 can break out (this is due to a Celestica design factor).

Example 3: All Archive Isilon 10 GbE nodes



In this example, your configuration will include:

- Two 10 GbE SFP+ switches
- 16 SFP+/LC cables
- 16 optics (If you go with LC cables)

Table 6 10 GbE switch option

Vendor	Model	Legacy Isilon Model	Dell SKU	Back-end ports	Port type	Rack units	All 10 GbE nodes	Mixed environment (10 GbE and 40 GbE)
Dell EMC	S4148F-ON	851-0317	210-AWOT	48	48 port 10 GbE, 2 port 40 GbE	1	Less than 48	Not supported
Dell EMC	S4112F-ON	851-0334	210-AWOS	15	12 port 10 GbE, 12 additional 10 GbE ports using 4x10 breakout	1/2	Less than 24	Not supported
Celestica	D2024	851-0258	N/A	24	24 port 10 GbE, 2 port 40 GbE	1	Less than 24	Not supported
Celestica	D2060	851-0257	N/A	48	48 port 10 GbE, 6 port 40 GbE	1	Greater than 24 and less than 48	Not supported
Arista	DCS-7304	851-0260	N/A	96	48 port 10 GbE, 4 port 40 GbE	8	Greater than 48 and less than 96 (included two 48 ports line cards)	40 GbE line card can be added
Arista		851-0283	N/A	Leaf upgrade (48 ports)			Greater than 96 and less than 144 (max 1 leaf upgrade)	

Note: For Celestica D2024, the two 40 GbE ports are not supported. For Celestica D2060, the six 40 GbE ports have been tested and can breakout to 4x10 GbE mode. For Arista DCS 7304, the four 40 GbE ports are not supported.

Table 7 Cable options for H400 and A2000 (A200)

Cable type	Legacy Isilon Model	Connector	Length	EMC P/N	Dell SKU number
Copper	851-0262	SFP+	1 m	038-003-728-01	470-AEFZ
Copper	851-0263	SFP+	3 m	038-003-729-01	470-AEGF
Copper	851-0264	SFP+	5 m	038-004-730-01	470-AEGL
Optical	851-0266	LC	10 m	038-004-153	470-AEFX

PowerScale platform back-end network option

Optical	851-0267	LC	30 m	038-004-154	470-AEGD
Optical	851-0268	LC	50 m	038-004-155	470-AEGK
Optical	851-0269	LC	100 m	038-004-156	470-AEHO
Optical	851-0270	LC	150 m	038-004-591	470-AEHN


Note: The optics for the LC-LC cables are bundled with the cable BOM and not listed separately on the quoting tool.

Example 4: All PowerScale F200 and F600: When mixing F600 and F200 nodes, use a 100 GbE infrastructure with 100 GbE connections to the F600 nodes and 4 x 25 GbE breakout cables to the F200 nodes.

3 x F600

3 x F200

2 x 100GbE Dell




In this example, your configuration will include:

- Two 100 GbE back-end switches
- 6 100G back-end cables
- 6 Optics (If MPO cables used)
- 2 QSFP28 to 4x SFP28 breakout cables (you can also use SFP28 Direct attached or LC-LC cables)

Example 5: Mixed Environment of PowerScale and Isilon:

Assuming you are using Dell EMC switches on the back end, our recommendation is to use the native NIC speed of the PowerScale nodes (100 GbE for F600 and 25 GbE for F200). If you are using non Dell EMC switches on the back end, you will have to use 40 GbE connection for F600 and 10 GbE connection for F200.

- 3 x F600**
- 3 x F200**
- 4 x F800**
- 2 x 100GbE Dell Ethernet**




In this example, your configuration will include:

- Two 100 GbE back-end switches
- 8 QSFP+/MPO back-end cables
- 8 Optics (If MPO cables used)
- 6 100G back-end cables
- 6 Optics (If MPO cables used)
- 2 QSFP28 to 4x SFP28 breakout cables (you can also use SFP28 Direct attached or LC-LC cables)

Example 6: Mixed Environment of PowerScale and Isilon: Assuming you are using non Dell EMC switches on the back end, our recommendation is to use 40 GbE for F600 and 10 GbE for F200.

- 3 x F600
- 3 x F200
- 4 x F800
- 2 x 40GbE Celestica



In this example, your configuration will include:

- Two 40 GbE Celestica back-end switches
- 14 40G back-end cables (QSFP+/MPO)
- 14 Optics (If MPO cables used)
- 2 QSFP to 4 x SFP+ breakout cables (you can also use SFP Direct attached or LC-LC cables)

Table 8 PowerScale Ethernet back-end switch options

Vendor	Model	Legacy Isilon Model	Dell SKU	Back-end ports	Port type	Rack units	100 GbE nodes	Mixed environment (10, 25, 40, and 100 GbE)
Dell EMC	Z9264-ON	851-0318	210-AWOW	64	All 100 GbE	2	Less than 64	Support breakout cables of 4 x 10 or 4 x 25. total 128 10 or 25 GbE nodes
Dell EMC	Z9100-ON	851-0316	210-AWOV /210-AWOU	32	All 100 GbE	1	Less than 32	Support breakout cables of 4 x 10 or 4 x 25. total 128 10 GbE or 25 GbE nodes
*Celestica	D4040	851-0259	N/A	32	All 40 GbE	1	100 GbE not supported	Support breakout cables, total 96 10 GbE nodes
*Arista	DCS-7308	851-0261	N/A	64	All 40 GbE	13	100 GbE not supported	No breakout support with FT, but you can add 10 GbE line card
*Arista		851-0282	N/A	leaf upgrade (32 ports)	All 40 GbE		252 nodes cluster will require 6 additional leaf.	OneFS 8.2 is required to scale beyond 144 nodes using Arista switches.

*Since 100 GbE is not supported in Celestica or Arista switches, you will have to connect to those switches using the appropriate 40 GbE and 10 GbE optics and cable combination.

10 GbE	SKU
Copper Cables	
Dell Networking, Cable, SFP+ to SFP+, 10GbE, Copper Twinax Direct Attach Cable, 1 Meter,CusKit	470-AAVH
Dell Networking Cable, SFP+ to SFP+, 10GbE, Passive Copper Twinax Direct Attach, 2 Meter,Cust Kit	470-ABPS
Dell Networking, Cable, SFP+ to SFP+, 10GbE, Copper Twinax Direct Attach Cable, 3 Meter,CusKit	470-AAVJ
Dell Networking, Cable, SFP+ to SFP+, 10GbE, Copper Twinax Direct Attach Cable, 5 Meter,CusKit	470-AAVG
Dell Networking, Cable, SFP+ to SFP+, 10GbE, Copper Twinax Direct Attach Cable, 7 Meter,CusKit	470-AAVI
Breakout Cables	
Dell Networking Cable, 40GbE, QSFP+ to 4x10GbE SFP+, Passive Copper Breakout Cable, 1M, Cust Kit	470-AAVO

10 GbE	SKU
Dell Networking,Cable,40GbE (QSFP+) to 4 x 10GbE SFP+ Passive Copper Breakout Cable, 2 Meter Customer Kit	470-ABXO
Dell Networking Cable 40GbE (QSFP+) to 4 x 10GbE SFP+ Passive Copper Breakout Cable 3 Meters, Customer Install	470-AAXG
Dell Networking 40GbE (QSFP+) to 4x10GbE SFP+ Passive Copper Breakout Cable, 5 Meters, Cust Kit	470-AAXH
Dell Networking Cable, 40GbE, QSFP+ to 4x10GbE SFP+, Passive Copper Breakout, 7 Meter, Cust Kit	470-AAWU
Optics	
Node Side: Dell EMC PowerEdge 10/25GbE Dual Rate SFP28 SR 85C Optic for all SFP28 ports, Customer Install	407-BCHI
Node Side: SFP+ SR 10GbE Optical Transceiver, High Temperature, Dell, Customer Kit	407-BBZM
Switch Side: Dell Networking, Transceiver, SFP+, 10GbE, SR, 850nm Wavelength, 300m Reach - Kit	407-BBOU
LC OM4 passive optical cables	
Dell Networking Cable, OM4 LC/LC Fiber Cable, (Optics required), 1 Meter, Customer kit	470-ACLV
Dell Networking Cable, OM4 LC/LC Fiber Cable, (Optics required), 2 Meter, Customer kit	470-ACLT
Dell Networking Cable, OM4 LC/LC Fiber Cable, (Optics required), 3 Meter, Customer kit	470-ACMO
Dell Networking Cable, OM4 LC/LC Fiber Cable, (Optics required), 5 Meter, Customer kit	470-ACLY
Dell Networking Cable, OM4 LC/LC Fiber Cable, (Optics required), 10 Meter, Customer kit	470-ACMN
Dell Networking Cable, OM4 LC/LC Fiber Cable, (Optics required), 30 Meter, Customer kit	470-ACLM
25GbE	
Copper Cables	
Dell Networking, Cable, SFP28 to SFP28, 25GbE, Passive Copper Twinax Direct Attach, 1M, Cust Kit	470-ACEX

10 GbE	SKU
Dell Networking, Cable, SFP28 to SFP28, 25GbE, Passive Copper Twinax Direct Attach, 2M, Cust Kit	470-ACFB
Dell Networking, Cable, SFP28 to SFP28, 25GbE, Passive Copper Twinax Direct Attach, 3M, Cust Kit	470-ACEU
Dell Networking, Cable, SFP28 to SFP28, 25GbE, Passive Copper Twinax Direct Attach, 5M, Cust Kit	470-ACEY
Breakout Cables	
Dell Networking Cable, 100GbE QSFP28 to 4xSFP28 Passive DirectAttachBreakout Cable, 1 Meter, Customer Kit	470-ABPR
Dell Networking Cable, 100GbE QSFP28 to 4xSFP28 Passive DirectAttachBreakout Cable, 2 Meter, Customer Kit	470-ABQF
Dell Networking Cable, 100GbE QSFP28 to 4xSFP28 Passive DirectAttachBreakout Cable, 3 Meter, Customer Kit	470-ABQB
Dell Networking Cable QSFP28-4XSFP28, 25G, Passive Copper DAC, Breakout, 5 Meter Customer Kit	470-AECY
Optics	
Node Side: Dell EMC PowerEdge 10/25GbE Dual Rate SFP28 SR 85C Optic for all SFP28 ports, Customer Install	407-BCHI
Switch Side: Dell Networking, Transceiver, 25GbE SFP28 SR, No FEC, MMF, Duplex LC, Customer Kit	407-BBXU
LC OM4 passive optical cables	
Dell Networking Cable, OM4 LC/LC Fiber Cable, (Optics required), 1 Meter, Customer kit	470-ACLV
Dell Networking Cable, OM4 LC/LC Fiber Cable, (Optics required), 2 Meter, Customer kit	470-ACLT
Dell Networking Cable, OM4 LC/LC Fiber Cable, (Optics required), 3 Meter, Customer kit	470-ACMO
Dell Networking Cable, OM4 LC/LC Fiber Cable, (Optics required), 5 Meter, Customer kit	470-ACLY
Dell Networking Cable, OM4 LC/LC Fiber Cable, (Optics required), 10 Meter, Customer kit	470-ACMN
Dell Networking Cable, OM4 LC/LC Fiber Cable, (Optics required), 30 Meter, Customer kit	470-ACLM

10 GbE	SKU
100GbE	
Copper Cables	
Dell Networking Cable 100GbE, QSFP28 to QSFP28, Passive Copper Direct Attach Cable, 1 Meter, Customer Kit	470-ABPY
Dell Networking Cable, 100GbE QSFP28 to QSFP28, Passive Copper Direct Attach Cable, 2 Meter, Customer Kit	470-ADDP
Dell Networking Cable, 100GbE QSFP28 to QSFP28, Passive Copper Direct Attach Cable, 3 Meter, Customer Kit	470-ABQE
Dell Networking Cable, 100GbE QSFP28 to QSFP28, Passive Copper Direct Attach Cable, 5 Meter, Customer Kit	470-ABPU
Active Optical Cables	
Dell Networking Cable, QSFP28 to QSFP28, 100GbE, Active Optical (Optics included), 3 Meter, Cust Kit	470-ACLU
Dell Networking Cable, QSFP28 to QSFP28, 100GbE, Active Optical (Optics included) Cable, 7 Meter, Customer Kit	470-ABPI
Dell Networking Cable, QSFP28 to QSFP28, 100GbE, Active Optical (Optics included) Cable, 10 Meter, Customer Kit	470-ABPM
Dell Networking Cable, QSFP28 to QSFP28, 100GbE, Active Optical (Optics included), 30 Meter, Customer Kit	470-ABPJ
Optics	
Node Side: Dell EMC PowerEdge QSFP28 SR4 100GbE 85C optic Customer Install	407-BCEX
Switch Side: Dell Networking, Transceiver, 100GbE QSFP28 SR4, No FEC Capable, MPO, MMF, Customer Kit	407-BBWW
MPO/MPT passive optical cables	
Dell Networking MPO Type B Crossover Cable, Multi Mode Fiber OM4, 1 Meter, Customer kit	470-ABPO
Dell Networking MPO Type B Crossover Cable, Multi Mode Fiber OM4, 3 Meter, Customer kit	470-ABPN
Dell Networking MPO Type B Crossover Cable, Multi Mode Fiber OM4, 5 Meter, Customer kit	470-ABPQ
Dell Networking MPO Type B Crossover Cable, Multi Mode Fiber OM4, 7 Meter, Customer kit	470-ABPP

10 GbE	SKU
Dell Networking MPO Type B Crossover Cable, Multi Mode Fiber OM4, 10 Meter, Customer kit	470-ABPV
Dell Networking MPO Type B Crossover Cable, Multi Mode Fiber OM4, 25 Meter, Customer kit	470-ABPT
40GbE * for Back-End compatibility with existing Isilon Gen 6 clusters	
Dell Networking Cable QSFP+ to QSFP+ 40GbE Passive Copper Direct Attach Cable 1 Meter, Cust Kit	470-AAVR
Dell Networking Cable, QSFP+ to QSFP+, 40GbE Passive Copper Direct Attach Cable, 2 Meter, Customer Kit	470-ACIW
Dell Networking Cable QSFP+ to QSFP+ 40GbE Passive Copper Direct Attach Cable 3 Meters, CK	470-AAWN
Dell Networking Cable QSFP+ to QSFP+ 40GbE Passive Copper Direct Attach Cable 5 Meters, CK	470-AAWE
Dell Networking Cable, QSFP+, 40GbE Active Optical (no optics required), 3 Meters, Customer Kit	470-ACOR
Dell Networking, Cable, QSFP+, 40GbE, Active Fiber Optical, 10 Meters (No optics required), Cust Kit	470-AAZM
Optics	
Node Side: Mellanox, Transceiver, QSFP, 40Gb, Short-Range, for use in Mellanox NW Adpt Only, CusKit	407-BBOI
Switch Side: Dell Networking, Transceiver, 40GbE QSFP+ SR4 Optics, 850nm Wavelength, 100-150m Reach on OM3/OM4, CK	407-BBOZ

Note: When ordering an Isilon F200 and F600 system in the new Dell EMC sales tool (DSA/Gii/DOMs), optics are not automatically added. Make sure to add the appropriate optics for the node side as well as for the switch side.

A Technical support and resources

[Dell.com/support](https://www.dell.com/support) is focused on meeting customer needs with proven services and support.

[Storage technical documents and videos](#) provide expertise that helps to ensure customer success on Dell EMC storage platforms.

A.1 Related resources

Dell EMC PowerScale Leaf and Spine network best practices:

<https://www.dell.com/resources/en-us/asset/white-papers/products/storage/h17682-dell-emc-powerscale-leaf-spine-network-best-practices.pdf>

Dell Switch OS Upgrade Guide:

<https://support.emc.com/docu93726>

PowerScale Leaf-Spine Installation Guide:

<https://support.emc.com/docu93725>