

# Dell EMC PowerScale: Leaf-Spine Network Best Practices

## Abstract

This white paper provides the best practices for deploying a scalable and resilient back-end network infrastructure for Dell EMC™ PowerScale™ clusters. Dell EMC Isilon™ OneFS™ 8.2 enables the deployment of a leaf-spine back-end network switch architecture that increases the size, scale, and performance of PowerScale clusters.

June 2020

## Revisions

Date	Description
June 2020	Content and format updates

## 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. [6/12/2020] [Best Practices] [H17685]

# Table of contents

Revisions.....	2
Acknowledgments.....	2
Table of contents .....	3
Executive summary.....	4
1 Introduction.....	5
2 Leaf-spine network architecture .....	6
3 General leaf-spine switch design considerations .....	8
4 Network design examples .....	10
4.1 Example 1: Network design with 12 Isilon nodes.....	10
4.2 Example 2: Network design with 40 Isilon nodes.....	11
4.3 Example 3: Network design with 88 PowerScale nodes .....	14
4.4 Example 4: Network design with 128 Isilon nodes.....	16
4.5 Example 5: Network design with 252 Isilon nodes.....	18
4.6 Example 6: Network design with 150 PowerScale nodes .....	20
A Technical support and resources .....	26
A.1 Related resources .....	26

## Executive summary

Dell EMC™ Isilon™ OneFS™ 8.2 introduced support for a back-end network infrastructure that can be configured as a set of leaf switches connected to spine switches. With the introduction of Dell EMC PowerScale™ OneFS 9.0, we support 252 nodes in a leaf-spine architecture. This document includes best practices for configuring your architecture to account for data-center space usage, rack positioning, cabling, and scale.

# 1 Introduction

Dell EMC PowerScale scale-out NAS nodes use InfiniBand switches as the private network for the back-end, intra-cluster, node-to-node communication. OneFS 8.1 operating system introduced the use of Ethernet switches for the back-end node-to-node communication. OneFS 8.1.1 introduced a choice of Dell EMC Ethernet switches for the back end to simplify configurations and provide a full Dell EMC solution.

Next-generation, multi-rack, data center solutions require performance, scale, and capacity to drive new and demanding workloads. A leaf-spine back-end network architecture facilitates much larger deployments. As a first step in designing a leaf-spine back-end network architecture using a Dell EMC switch solution, the solution includes a PowerScale cluster which can scale up to 144 nodes.

## 2 Leaf-spine network architecture

In a leaf-spine network switch architecture, the access layer of the network is referred to as the leaf layer. The Isilon nodes connect to leaf switches at this layer. At the next level, the aggregation and core network layers are condensed into a single spine layer. Every leaf switch connects to every spine switch to ensure that all leaf switches are no more than one hop away from one another. In addition, leaf switch to spine switch connections need to be evenly distributed meaning there should be the same number of connections to each spine switch from each leaf switch. This minimizes latency and the likelihood of bottlenecks in the back-end network. A leaf spine network architecture is highly scalable and built with redundancy.

Leaf spine network deployments can have a minimum of two leaf switches and one spine switch. For small to medium clusters, this back-end network is comprised of a pair redundant top-of-rack switches. Only the Dell EMC Z9100 Ethernet switch is supported in the leaf-spine architecture.

Legacy Isilon model	Dell SKU	Back-end ports	Port type	Rack units	100 GbE and 40 GbE nodes	Mixed environment (10, 25, 40, and 100 GbE)
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 as ToR back-end switch.

The Z9100-ON switch is compatible with all Isilon platforms and PowerScale platforms. While the Z9100-ON supports many features, not all capabilities of the switch are exposed or used when the switch is being used as an Isilon back-end switch.

Component	Description	Connection considerations
Network Spine Switch	Dell Z9100-ON 32-port switch	Back-end network with 100 GbE (uplink) connects to the leaf switch.
Network Leaf Switch	Dell Z9100-ON 32-port switch	Downlink from the leaf switch to the nodes. Supported connection type 100 GbE, 40 GbE, 25 GbE and 10 GbE back-end nodes.
Dell EMC PowerScale nodes	F200 and F600	F200: nodes support a 10 GbE or 25 GbE connection to the leaf using the same NIC.  F600: nodes support a 40 GbE or 100 GbE connection to the leaf using the same NIC.
Dell EMC Isilon Performance nodes (Flash and Hybrid)	F810, F800, H600, H500, and H5600	Performance nodes support a 40GbE connection to the leaf switch
Isilon Archive and Hybrid Nodes	A200, A2000, and H400	Archive nodes support a 10GbE connection to the leaf switch using breakout cable

**Note:** The 32 nodes limit per leaf switch (using breakout cable) has been removed due to a fix in Dell EMC Networking OS version to 10.5.0.6. Now with a 22-downlink connection, you can have up to 88 10 GbE nodes using 4 x 10 GbE breakout cables. OneFS version 9.0 requires the switch operating system version to be

10.5.0.6. The Dell EMC Networking OS version 10.5.0.6 is also supported in OneFS version 8.2.2. 144-node L/S clusters can still work with older versions of Dell EMC Networking OS version 10.4.x. Dell EMC Networking OS version 10.5.0.6 requires manual designation of leaf and spine switches through the command line from the switches. For detailed instruction on upgrading to Dell EMC Networking operating system version 10.5.0.6, see the [Leaf-Spine Installation Guide](#) and [Switch OS Upgrade Guide](#).

---

### 3 General leaf-spine switch design considerations

Here are some general design considerations to simplify the setup and management of your Isilon environment.

You must avoid network oversubscription between your uplink and downlink connections between the leaf and spine switches. For example, if majority of nodes in the cluster consist of 100 GbE back-end nodes; you are required to have up to 16 x 100 GbE uplink connection per leaf.

If you have greater than 1 spine switch in the architecture, make sure connections between leaf and spine switches are equally distributed among all leaf switches.

Both Int-a and Int-b should be identical in terms of configuration and leaf spine network architectural design.

You should strategically locate the spine switches within a data center. This is to ensure cabling is planned, organized and manageable when scaling out the nodes and switches within the cluster.

Deploy a leaf-spine network topology for the expected growth of that cluster, rather than the initial configuration.

Live migration from a ToR back end to an L/S back end is supported. For detailed steps, see the Best Practices Guide for Live Migration document.

Maximum nodes	Spines	Leaves	Cables between each pair of leaves and spines
<b>All 40G ports:</b>			
44	1	2	9
66	1	3	9
88	2	4	5
110	2	5	5
132	2	6	5
154	3	7	3
176	3	8	3
198	3	9	3
220	5	10	2
242	5	11	2
252	5	12	2
<b>All 100G ports:</b>			
32	1	2	16
64	2	4	8



Maximum nodes	Spines	Leaves	Cables between each pair of leaves and spines
112	4	7	4
128	4	8	4
135	5	9	3
150	5	10	3

---

**Note:** Maximum number of leaves and spines in the cluster should not exceed 17 switches per side (int-a and int-b combined 34 switches per cluster).

---

## 4 Network design examples

### 4.1 Example 1: Network design with 12 Isilon nodes

Proposed solution of 12 Isilon nodes in the cluster:

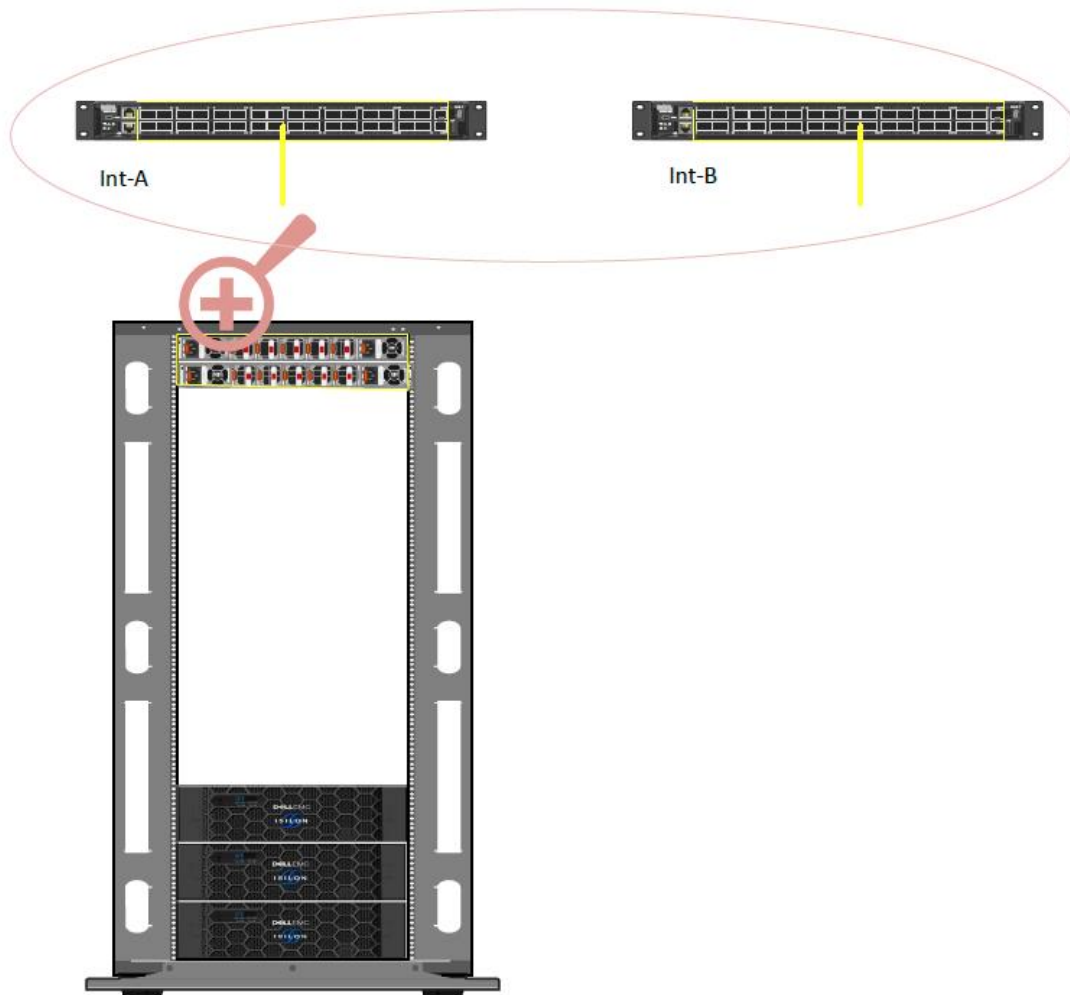
- 8 Performance nodes (40 GbE back end)
- 4 Archive nodes (10 GbE back end)

**Assumptions/requirements:**

- The customer has confirmed that this cluster will **NEVER** grow beyond 32 nodes of performance nodes with 40 GbE or 128 archive nodes with 10 GbE using breakout cables.
- A leaf-spine back-end network configuration is **NOT** needed.

However, if the customer changes their mind and decides to grow this cluster beyond what is supported in the ToR solution, we need to add 2 more Z9100 switches per each side (Int-a and Int-b). See example 2, to see what it takes to expand this ToR solution to a leaf-spine architecture.

Top of rack Z9100



In this example, your configuration will include:

- 2 Dell EMC Networking Z9100 switches
- 16 QSFP+ or MPO back-end cables
- 16 Optics (if MPO cables used)
- 2 QSFP to SFP+ breakout cables

## 4.2 Example 2: Network design with 40 Isilon nodes

Proposed solution of 40 Dell EMC Isilon nodes in the cluster:


- 20 Performance nodes (40 GbE back end)
- 20 Archive nodes (10 GbE back end)

**Assumptions/requirements:**

- The customer has confirmed that this cluster will **never** grow beyond 44 performance nodes with 40 GbE.

Although in this example, the starting point does not require a leaf-spine switch architecture, the target growth of the cluster exceeds what a single Z9100 network switch can support.

Leaf switch	Spine switch connection
L1	Port 1, 2, 3, 4, 5, 6, 7, 8, 9 and 10 on Spine switch 1
L2	Port 11, 12, 13, 14, 15, 16, 17, 18, 19 and 20 on Spine switch 1



In this example, your configuration will include:

- 6 Dell EMC Z9100 switches (3 per side)
  - 1 spine switches
  - 2 leaf switches
- 36 QSFP28 100 Gb uplink cables (9 uplink cable per leaf switch)
- 40 QSFP+ or MPO back-end cables
- 40 Optics (if MPO cables used)
- 10 QSFP to SFP+ breakout cables

**Design considerations:**

Organize the nodes per nodepools/diskpools: In this example above all performance nodes are on rack #1 and all archive nodes are on rack #2.

To simplify and organize cabling, place both an Int-a and Int-b leaf switch into the same rack allows all nodes to connect to switches only within the same rack, and rack to rack cabling to only be uplinks (leaf switch to spine switches).

Group your uplink cables connected to spine in a logical way. As you can see from the magnified drawing above, the first 9 ports on the spine switch were used by leaf switch #1 then the next 9 ports by leaf switch #2.

Reserve the 10 ports on leaf switches to be used for uplink. Or reserve 16 ports instead if you were to add majority of 100 GbE nodes in the future.

### 4.3 Example 3: Network design with 88 PowerScale nodes


Proposed solution of 80 Isilon nodes in the cluster:

- 66 Dell EMC PowerScale F600 nodes (40 10 GbE)
- 22 Dell EMC PowerScale F200 nodes (10GbE)

**Assumptions/requirements:**

- All nodes are connected to back-end leaf switch on the same rack and **only** uplink cables to spine switches are connected to a different rack.

Leaf switch	Spine connection
L1	Port 1, 2, 3, 4 and 5 on Spine switch 1 and 2
L2	Port 6, 7, 8, 9 and 10 on Spine switch 1 and 2
L3	Port 11, 12, 13, 14 and 15 on Spine switch 1 and 2
L4	Port 16, 17, 18, 19 and 20 on Spine switch 1 and 2



In this example, your configuration will include:

- 12 Dell EMC Z9100 switches (6 per side)
  - 2 spine switches
  - 4 leaf switches
- 80 QSFP28 100 Gb uplink cables (10 uplink cables per leaf)
- 132 QSFP+ or MPO back-end cables
- 132 Optics (if MPO cables used)
- 12 25 GbE breakout cables

**Design considerations:**

As you can see from the drawing, all spine switches for both networks (Int-a and Int-b) are spread across rack #2 and #3. Rack #2 has one of each spine switch from Int-a and Int-b. Rack #3 has one of each spine switch from Int-a and Int-b.

To simplify and organize cabling, place your leaf switches accordingly (see the drawing above). As you can see, one leaf switches from Int-a and one leaf switches from Int-b are spread across all racks. This way all the downlink cables are contained within the rack and uplink cables connect to the spine switch on rack #2 and rack #3.

Reserve the first 10 ports on leaf switches to be used for uplink. Keep in mind that if you were to change your mind and deploy these nodes by 100 GbE and 25 GbE connection on the back end, you will be required to recable your uplinks. To avoid re-cabling your uplink, reserve the first 16 ports on the leaf switches to be used for uplink.

On rack 1, 2 and 3 we have all our F600 nodes (22 nodes per rack). Since 10 ports are used for uplink and 22 for downlink, no more downlink ports are available on these 3 racks.

All 22 F200 nodes with 25 GbE back end are on rack #4. Out of the 22 downlink ports available on rack #4, we're only using 6 ports using 4 x 25 GbE breakout to connect the 22 F200 nodes. So, we have remaining of 16 ports. Without oversubscription, we can add 18 25GbE F200 nodes.

## 4.4 Example 4: Network design with 128 Isilon nodes

Proposed solution of 128 Isilon nodes in the cluster:


128 performance only nodes (40 GbE back end)

### Assumptions/requirements:

- The customer has confirmed that this cluster will **never** grow beyond 128 performance nodes.
- Not just uplink cables but also downlink cables can connect to nodes on a different rack.

Leaf switch	Spine switch connection
L1	Port 1, 2, 3, 4 and 5 on Spine switch 1 and 2
L2	Port 6, 7, 8, 9 and 10 on Spine switch 1 and 2
L3	Port 11, 12, 13, 14 and 15 on Spine switch 1 and 2
L4	Port 16, 17, 18, 19 and 20 on Spine switch 1 and 2
L5	Port 21, 22, 23, 24 and 25 on Spine switch 1 and 2
L6	Port 26, 27, 28, 29 and 30 on Spine switch 1 and 2





In this example, your configuration will include:

- 16 Dell EMC Z9100 switches (8 per side)
  - 2 spine switches
  - 6 leaf switches
- 120 QSFP28 100 Gb uplink cables (10 uplink cables per leaf)
- 256 QSFP+ or MPO back-end cables
- 256 Optics (if MPO cables used)

**Design considerations:**

Place the spine switches for both networks (Int-a and Int-b) in the center (in this example they're on rack #2 and #3).

To simplify and organize cabling, place your leaf switches accordingly (see the drawing above). As you can see, the leaf switches from Int-a and Int-b are spread across all racks. And rack #1 and #4 downlink cables are contained within the rack since there are enough downlinks from both networks (Int-a and Int-b). Only downlink cables leaving racks would be from rack #2 and #3.

Group your uplink cables connected to spine switches in a logical way. As you can see from the magnified drawing above, the first 5 ports on the spine switch were used by leaf switch #1, then the next 5 ports by leaf switch #2, and so on.

Reserve the first 10 ports on leaf switches to be used for uplink. Alternately, reserve 16 ports instead if you were to add mostly 100 GbE nodes in the future.

This solution is built to the maximum that the leaf spine architecture can support.

Recabling will be required to add additional spine switches.

## 4.5 Example 5: Network design with 252 Isilon nodes

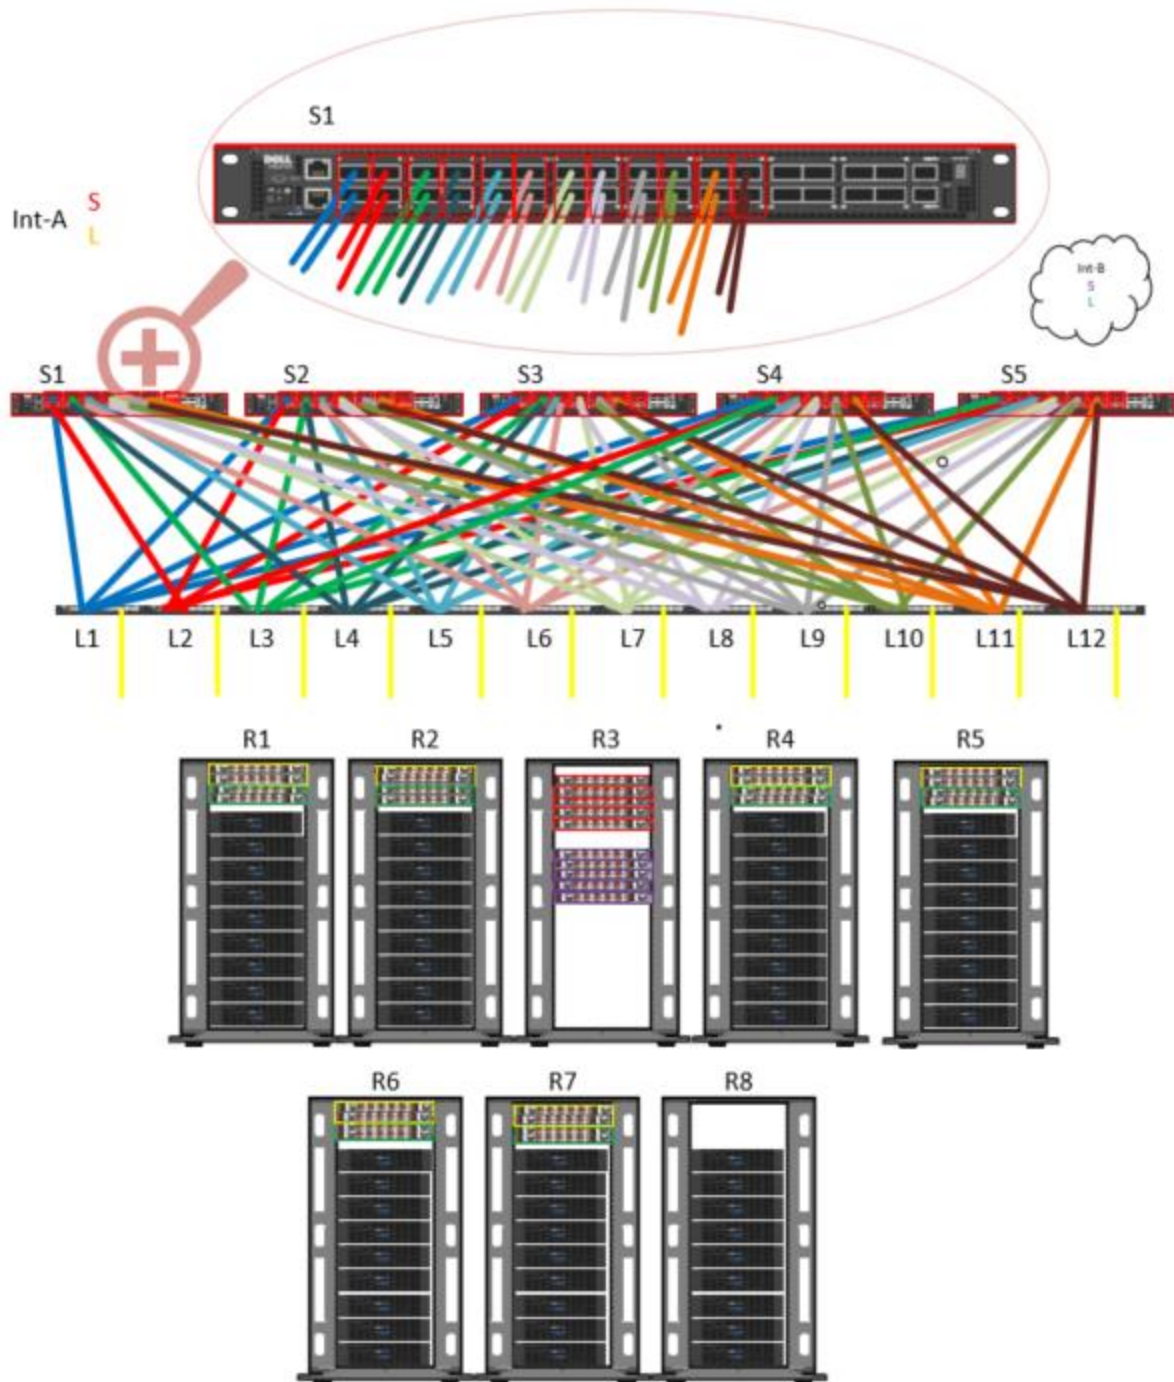
Proposed solution of 252 Isilon nodes in the cluster:

252 Performance only nodes (40 GbE back end)

### Assumptions/requirements:

- Not just uplink cables but also downlink cables can connect to nodes on a different rack.
- Due to limited rack space, customer requested to minimize space in the data center.

Leaf switch	Spine switch connection
L1	Port 1 and 2 on Spine switch 1, 2, 3, 4 and 5
L2	Port 3 and 4 on Spine switch 1, 2, 3, 4 and 5
L3	Port 5 and 6 on Spine switch 1, 2, 3, 4 and 5
L4	Port 7 and 8 on Spine switch 1, 2, 3, 4 and 5
L5	Port 9 and 10 on Spine switch 1, 2, 3, 4 and 5
L6	Port 11 and 12 on Spine switch 1, 2, 3, 4 and 5
L7	Port 13 and 14 on Spine switch 1, 2, 3, 4 and 5
L8	Port 15 and 16 Spine switch 1, 2, 3, 4 and 5
L9	Port 17 and 18 on Spine switch 1, 2, 3, 4 and 5
L10	Port 19 and 20 on Spine switch 1, 2, 3, 4 and 5
L11	Port 21 and 22 on Spine switch 1, 2, 3, 4 and 5
L12	Port 23 and 24 on Spine switch 1, 2, 3, 4 and 5



In this example, your configuration will include:

- 34 Dell EMC Networking Z9100 switches (17 per side)
  - 5 spine switches
  - 12 leaf switches
- 240 QSFP28 100 Gb uplink cables (10 uplink cables per leaf)
- 504 QSFP+ or MPO back-end cables

- 504 Optics (if MPO cables used)

---

**Note:** This example will be accurate if you were to size all F600 with 40 GbE back-end connectivity.

---

**Design considerations:**

To simplify and organize cabling, place your leaf switches accordingly (see the drawing above). As you can see, two leaf switches from Int-a and two leaf switches from Int-b are spread across all racks except R8. No leaf switches on R8, so the chassis on R8 needs to connect to a different rack

## 4.6 Example 6: Network design with 150 PowerScale nodes


Proposed solution of 150 Dell EMC PowerScale nodes in the cluster:

150 x 100 GbE only nodes

**Assumptions/requirements:**

- Not just uplink cables but also downlink cables can connect to nodes on a different rack.

Leaf switch	Spine switch connection
L1	Port 1, 2 and 3 on Spine switch 1, 2, 3, 4 and 5
L2	Port 4, 5 and 6 on Spine switch 1, 2, 3, 4 and 5
L3	Port 7, 8 and 9 on Spine switch 1, 2, 3, 4 and 5
L4	Port 10, 11 and 12 on Spine switch 1, 2, 3, 4 and 5
L5	Port 13, 14 and 15 on Spine switch 1, 2, 3, 4 and 5
L6	Port 16, 17 and 18 on Spine switch 1, 2, 3, 4 and 5
L7	Port 19, 20 and 21 on Spine switch 1, 2, 3, 4 and 5
L8	Port 22, 23 and 24 Spine switch 1, 2, 3, 4 and 5
L9	Port 25, 26 and 27 on Spine switch 1, 2, 3, 4 and 5
L10	Port 28, 29 and 30 on Spine switch 1, 2, 3, 4 and 5



In this example, your configuration will include:

- 30 Dell EMC Z9100 switches (15 per side)
  - 5 spine switches
  - 10 leaf switches



- 300 QSFP28 100 Gb uplink cables (15 uplink cables per leaf)
- 300 QSFP+ or MPO back-end cables
- 300 Optics (if MPO cables used)

**Note:** Since these nodes are 100 GbE back end, they require a lot more uplink than your typical 40 GbE back-end nodes. In this example, we have 15 x 100 GbE uplink per leaves. That means we can only connect 15 x 100 GbE nodes per leaf switch.

Table 1 Legacy Dell EMC Isilon 100 GbE Uplink Cable options for Z9100

Cable type	Legacy Isilon Model	Dell SKU number	Connector	Length
Pass Copper	851-0320	470-AEFW	QSFP28	1 m
Pass Copper	851-0321	470-AEGI	QSFP28	3 m
Pass Copper	851-0322	470-AEGO	QSFP28	5 m
Active Optical	851-0323	470-AEGJ	QSFP28	3 m
Active Optical	851-0324	470-AEGP	QSFP28	7 m
Active Optical	851-0325	470-AEFY	QSFP28	10 m
Active Optical	851-0326	470-AEGE	QSFP28	30 m

**Note:** Optics are added automatically when MPO cables are quoted.

Table 2 Legacy Isilon downlink 40 GbE Cable options for Performance nodes (F810, F800, H600 and H500)

Cable type (passive)	Legacy Isilon (model)	Dell SKU number	Connector	Length	EMC P/N
Copper	851-0253	470-AEGB	QSFP+	1m	038-002-064-01
Copper	851-0254	470-AEGG	QSFP+	3m	038-002-066-01
Copper	851-0255	470-AEGM	QSFP+	5m	038-002-139-01
Optical	851-0274	407-BCIV	MPO	1m	038-004-214
Optical	851-0275	407-BCIW	MPO	3m	038-004-216
Optical	851-0276	407-BCJD	MPO	5m	038-004-227
Optical	851-0224	407-BCIY	MPO	10m	038-004-218
Optical	851-0225	407-BCJB	MPO	30m	038-004-219
Optical	851-0226	407-BCJC	MPO	50m	038-004-220

Cable type (passive)	Legacy Isilon (model)	Dell SKU number	Connector	Length	EMC P/N
Optical	851-0227	407-BCIZ	MPO	100m	038-004-221
Optical	851-0277	407-BCIX	MPO	150m	038-000-139

**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). MPO optics are added automatically when MPO cables are quoted and appear as a separate line item. Legacy Isilon downlink 10 GbE breakout cable options for archive nodes (A2000, A200 and H400)

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

Table 3 Dell EMC PowerScale supported cables and optics

PowerScale 10 GbE Breakout cables	SKU
Dell Networking Cable, 40GbE, QSFP+ to 4x10GbE SFP+, Passive Copper Breakout Cable, 1M, Cust Kit	470-AAVO
Dell Networking, Cable, 40GbE (QSFP+) to 4 x 10 GbE 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
PowerScale 25GbE 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
<b>PowerScale 100GbE cables</b>	
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	470-BCEX
Switch Side: Dell Networking, Transceiver, 100GbE QSFP28 SR4, No FEC Capable, MPO, MMF, Customer Kit	470-BBWV
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
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
<b>40GbE * for Back-End compatibility with existing Isilon Gen 6 clusters</b>	
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	470-BBOI
Switch Side: Dell Networking, Transceiver, 40GbE QSFP+ SR4 Optics, 850nm Wavelength, 100-150m Reach on OM3/OM4, CK	470-BBOZ

## 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 ToR network best practices:

<https://www.dell.com/resources/en-us/asset/white-papers/products/storage/h16346-dell-emc-powerstore-back-end-network-overview.pdf>

Dell Switch OS Upgrade Guide:

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

PowerScale Leaf-Spine Installation Guide:

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