

# Dell EMC Networking OS10 Enterprise Edition Storage Overview

High-level networking guide for storage technologies and storage networking topologies

## [Abstract](#)

This document provides high-level storage networking information for various storage technologies. Detailed deployment instructions are provided through linked documents where available.

May 2019

# Revisions

Date	Description
May 2019	Initial release

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.

© 2019 Dell Inc. or its subsidiaries. All Rights Reserved. Dell, EMC, Dell EMC, and other trademarks are trademarks of Dell Inc. or its subsidiaries. Other trademarks may be trademarks of their respective owners.

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

# Table of contents

Revisions.....	2
1 Introduction.....	5
2 Storage formats.....	6
2.1 Block.....	6
2.2 File.....	6
2.3 Object.....	6
3 Storage foundation.....	7
3.1 Direct attached storage.....	7
3.2 Network attached storage.....	7
3.3 Storage area network.....	7
3.4 Object storage.....	8
3.5 Software-defined storage.....	8
4 Storage benefits and limitations.....	9
4.1 Direct attached storage.....	9
4.2 Network attached storage.....	9
4.3 Storage area network.....	9
4.4 Object storage.....	10
4.5 Software-defined storage.....	10
5 Dell EMC product sets.....	11
5.1 PowerMax.....	11
5.2 VxRail appliance.....	11
5.3 ECS.....	11
5.4 VPLEX.....	11
5.5 Isilon.....	11
5.6 VMAX.....	12
5.7 XtremIO.....	12
5.8 Unity.....	12
5.9 PowerVault MD series.....	12
5.10 PowerVault ME series.....	12
5.11 SC series.....	12
5.12 NX series.....	13
6 Storage protocols.....	14
6.1 iSCSI.....	14
6.2 FC.....	14
6.3 FCoE.....	14

6.4	NFS.....	14
6.5	SMB.....	14
6.6	CIFS.....	15
7	Data center leaf-spine networks.....	16
7.1	Layer 2 leaf-spine network.....	16
7.2	Layer 3 leaf-spine network.....	17
8	Fibre Channel storage deployment options.....	18
8.1	Fibre Channel storage with S4148U-ON in F_port mode.....	18
8.2	FCoE with S4148U-ON in NPG mode.....	19
8.3	FCoE to Fibre Channel with S4148U-ON in F_Port mode.....	19
8.4	FCoE with FIP snooping bridge.....	20
9	Ethernet-based storage.....	22
9.1	iSCSI SAN with leaf-spine data center.....	22
9.1.1	iSCSI configuration.....	22
9.1.2	Leaf-Spine configuration.....	23
9.2	Shared leaf-spine data center with iSCSI or SDS.....	23
9.2.1	Leaf-spine deployment.....	23
9.2.2	vSAN deployment with SDS example.....	23
9.2.3	iSCSI configuration.....	24
A	Support and feedback.....	25

# 1 Introduction

Our vision at Dell EMC is to be the essential infrastructure company in the data center, campus, branch, or office. Dell EMC Networking ensures modernization for today's applications and the emerging cloud-native world. Dell EMC is committed to disrupting the fundamental economics of the market with an open strategy that gives you the freedom of choice for networking operating systems and top-tier merchant silicon. The Dell EMC strategy enables business transformations that maximize the benefits of collaborative software and standards-based hardware, including lowered costs, flexibility, freedom, and security. Dell EMC provides further customer enablement through validated deployment guides which demonstrate these benefits while maintaining a high standard of quality, consistency, and support.

Dell EMC Networking offers several switch models to support numerous storage technologies.

This document provides high-level storage networking information for various storage technologies. Detailed deployment instructions are provided through linked documents for a variety of scenarios.

## 2 Storage formats

Block, file, and object storage each has a limitation on how data is organized and presented. Block storage chunks data into arbitrarily organized, evenly sized volumes. File storage organizes and represents data as a hierarchy of files in folders. Object storage manages data and links it to associated metadata.

### 2.1 Block

Block storage is the most commonly used storage type for most applications. Data is stored in fixed-size blocks within a rigid arrangement that is ideal for enterprise databases. Block storage resources can be:

- LUNs
- Consistency groups
- VMware Virtual Machine File System Datastores.

### 2.2 File

File storage, also called file-level or file-based storage, stores data as files in hierarchically nested folders ideal for active documents. File storage resources can be file systems, such as Network File System (NFS) or Server Message Block (SMB), and VMware NFS datastores.

### 2.3 Object

Object storage sees devices and software that house data in objects and serves client data through RESTful HTTP APIs. Object storage is highly scalable for handling large amounts of data in the petabyte and exabyte range. The stateless nature of object storage is not recommended in high-performance or low latency applications but is ideal for significant unstructured data, analytics, and archiving.

## 3 Storage foundation

Elements including the storage hardware, servers, network, and the efficiency of the software, influence the storage system performance. The three main ways of connecting storage to servers are:

- Direct attached storage (DAS)
- Network attached storage (NAS)
- Storage area network (SAN)

With DAS, which is the most basic form of the three, the drive connects directly to the server and is often even in the same enclosure. SANs and NAS both separate data storage from servers which enables servers to share those resources. NAS devices have their own file system, so they work best as file servers. SANs give block-level access and appear to computers as normal drives, so they work better for applications such as databases.

### 3.1 Direct attached storage

DAS is storage that is directly attached to the computer accessing it, as opposed to storage that a NAS accesses. Hard drives, solid-state drives, optical disc drives, and storage on external drives are examples of DAS devices. A typical DAS consists of a data storage enclosure that holds several hard disk drives connected directly to a computer through a network adapter. There are no network devices like a hub, switch, or router between the two points. The main protocols that are used for DAS connections are:

- ATA
- SATA
- NVMe
- SCSI
- SAS
- USB
- IEEE 1394
- Fibre Channel

### 3.2 Network attached storage

NAS is a dedicated networked appliance. NAS uses its own operating system, integrated hardware, and software components to provide distributed file storage to other devices on an IP network. NAS provides access to files using network file sharing protocols such as SMB or CIFS for Microsoft Windows, NFS for Linux, and AFP for Apple. The benefits of dedicated a NAS, compared to general purpose servers also serving files, include faster data access, easier administration, and simpler configuration.

Today, most storage capacity sold worldwide is shipped in support of file level or unstructured data. File-based information that does not fit efficiently in a column-and-row database, is an example of unstructured data. If you cannot fit the information neatly in a database, a file system is needed to make order sitting above the low-level disk operations.

### 3.3 Storage area network

SAN is a computer network that provides access to consolidated, block-level data storage. A SAN enhances the accessibility of storage devices, such as disk arrays and tape libraries, to servers so that the operating system recognizes the devices as locally attached devices. With the rise in unstructured data growth and file

content, SAN arrays account for most of the money that is spent on external disk storage worldwide. SAN connectivity options were traditionally based around Fibre Channel, but in recent years the broader adoption of 10 GbE has firmly planted iSCSI (Internet Small Computer System Interface) as a mainstream choice. Fibre Channel over Ethernet (FCoE) is another option that has had slow adoption in the market.

Virtualization has driven SAN growth in the last 15 years. With the ability to use compute resources effectively, virtualization provides the abstraction layer between hardware and workloads providing workload mobility. Moving applications across physical hosts for maintenance, workload balancing, or disaster recovery requires shared storage.

## 3.4 Object storage

Object storage is a technology that creates extra abstraction, often on top of and across local file systems. This means that the data in the system is managed as objects, instead of blocks or files. The objects include the data, a variable amount of metadata, and a globally unique identifier. Object storage can be implemented at multiple levels, including the device level, or object-storage device, the system level, and the interface level. This namespace can span hundreds of servers, enabling easier scaling of capacity than either SAN or NAS models.

## 3.5 Software-defined storage

Software-defined storage (SDS) is part virtualization software and part storage management software. SDS abstracts the bits of data that is included with hardware, formats the data into an object, block, or file format, and organizes the data for network use. SDS works well with unstructured data workloads such as the object and block storage systems that containers and microservices rely upon since it can scale in ways that non-virtualized storage solutions cannot.



## 4 Storage benefits and limitations

The general benefits and limitations of each storage type are outlined in this section.

### 4.1 Direct attached storage

DAS limits the use of shared storage. The benefits of DAS is mainly cost. Cloud providers use a DAS solution to deploy massive farms of storage servers and low-cost JBODs (just a bunch of disks) as an alternative to large dedicated disk arrays. DAS includes disks internal to a physical server and external JBOD cabinets attached to the server using a SAS cable. The DAS solution allows you to add low-cost storage easily.

### 4.2 Network attached storage

With NAS, data is stored as files and designed to interface with locally networked storage but provides more structure in the form of file systems and files. File systems have natural limits that are based on the local internal structures that are used to manage the file hierarchy and file access. The information within the managed file hierarchy provides awareness of the content in the system, however, it is entirely localized to a physical storage controller. By design, NAS systems are limited in scale and scope. Clustered NAS systems extend the expandability of the technology, however, NAS has natural limits that are tied to physical controllers, and to the central database that is used to track the file hierarchy and files.

The main benefits of NAS include:

- **Scale-out capacity** - The addition of more storage capacity to NAS is done by adding more hard disks. You do not have to upgrade or replace existing servers, and new storage can be made available without shutting down the network.
- **Performance** - Because NAS is dedicated to serving files, it removes the responsibility of file serving from other networked devices. Also, since NAS is tuned to specific use cases (like big data or multimedia storage), clients can expect better performance.
- **Easy setup** - NAS architectures are often delivered with simplified scripts, or even as appliances preinstalled with a streamlined operating system. This configuration reduces the time that it takes to set it up and manage the system.
- **Accessibility** - Every networked device has access to NAS.
- **Data protection**: NAS can be formatted to support replicated disks, a RAID, or erasure coding to ensure data integrity.

### 4.3 Storage area network

Data is stored as blocks and designed as the most basic way to interface with storage over a dedicated local network. SAN controls blocks of data in small logical volumes but has no context of what the data is. SAN depends entirely on the application to organize, catalog, and structure the data. By design, SANs are limited in scale, interfaces, and scope. SANs are typically higher in cost due to the dedicated network infrastructure.

The use of SSDs in a SAN array is functionally similar to traditional disk arrays. SSDs in a SAN array present shared capacity to multiple hosts but with the inclusion of solid-state drives, optimized platforms reach high IO points and low latency with shared resources.

## 4.4 Object storage

Data is stored as objects and a technology that creates more abstraction, often on top of and across local file systems. Instead of blocks or files, data in the system is managed as objects across a global namespace, with unique metadata for each object. This namespace can span hundreds of servers, enabling easier scaling of capacity than either SAN or NAS models.

## 4.5 Software-defined storage

SDS is a flexible storage model that separates storage functionality from the hardware, where storage becomes one of the services. Hardware independence architecture enables services and data to flow more freely when applications change, and systems scale up and down. Data protection and data recovery can then be implemented as another service limiting silos of functionality in a large business.

Self-service clouds that operate at application and infrastructure layers are multibillion dollar businesses.

## 5 Dell EMC product sets

Dell EMC provides a portfolio of storage and data protection offerings for both modern and legacy applications. The storage solutions are architected for varied datasets, such as structured, semi structured, or unstructured. The different datasets have differing I/O characteristics and service-level agreements that provide flexibility and enable data-driven decision making for the deployment.

### 5.1 PowerMax

Dell EMC PowerMax is the fastest storage array available that delivers up to 10 million IOPS with true end-to-end NVMe and integrated machine learning. Designed for performance, scale, and availability, PowerMax can handle even the toughest workloads including block, file, mainframe. With rich data services such as deduplication compression, nondestructive upgrades, and migrations, data is always protected and available.

### 5.2 VxRail appliance

VxRail is an exclusive hyper-converged infrastructure appliance that extends and simplifies a VMware environment. VxRail provides consistent integration with existing VMware tools and also lets customers use and extend their current IT tools and processes.

### 5.3 ECS

Dell EMC Elastic Cloud Storage, or ECS, is a third-generation object platform. ECS unlocks data insights from both traditional and next-generation applications with unmatched expansion, flexibility, and resiliency. Seen as truly software-defined, ECS enables you to deploy this platform on your own terms:

- As a turnkey storage appliance
- As a software-only solution designed to run on industry-standard hardware
- Through public cloud solutions using Virtustream
- As a Dell EMC hosted ECS-as-a-Service

### 5.4 VPLEX

Dell EMC VPLEX storage delivers continuous availability and data mobility across arrays and data centers. VPLEX turns data storage environments into automated, business centers driven by service-level agreement, with analytics, performance, and configuration solutions to eliminate tedious tasks and allows for self-service delivery.

### 5.5 Isilon

Dell EMC Isilon is the industry-leading scale-out NAS system. Isilon supports a wide range of demanding file workloads:

- All-flash
- Hybrid
- Archive

Isilon is ideal for data consolidation and include the following benefits:

- Simple to manage no matter how large the data becomes
- Highly efficient with 80% utilization and automated tiering
- Scales from terabytes to petabytes
- Offers multiprotocol support for operational flexibility
- High-density design reduces the size of the system footprint

## 5.6 VMAX

Extreme performance meets mission-critical applications and scale with Dell EMC VMAX architected for flash storage. Dell EMC VMAX All Flash arrays are architected to solve the challenge of embracing a modernized flash-centric datacenter for mission-critical applications while simultaneously simplifying, automating, and consolidating IT operations. VMAX All Flash is engineered for the latest, high density flash technology and to specifically exploit the rich set of data services of VMAX All Flash.

## 5.7 XtremIO

Dell EMC XtremIO delivers predictable service-level agreement (SLAs) and operational efficiency with the enterprise-class all-flash array.

## 5.8 Unity

Dell EMC Unity gives you the ability to consolidate and manage data over your network. With simple, affordable, and reliable enterprise-level 10 GbE iSCSI technology, Unity streamlines storage consolidation for application deployments with simplicity and flexibility.

## 5.9 PowerVault MD series

The Dell EMC PowerVault MD series of arrays enable you to consolidate and manage data over your network. With simple, affordable, and reliable enterprise-level iSCSI technology, the MD series is scalable and consists of enterprise-class storage arrays.

## 5.10 PowerVault ME series

The affordable Dell EMC PowerVault ME series storage is designed for entry-level SAN and DAS workloads. The ME series is ideal for accelerated performance for HPC, Exchange, backup, VDI, and many more applications.

## 5.11 SC series

The Dell EMC SC series of systems consists of automated hybrid-to-all-flash storage solutions. The SC series offers intelligent, self-optimizing arrays with auto-tiering, deduplication, compression, and affordable performance for organizations of all sizes.

## 5.12 NX series

The Dell EMC NX series family of network attached storage uses Dell EMC PowerEdge servers that are paired with the advantages of Microsoft Windows Storage Server 2016 with NAS gateway support.

<b>Product Categories</b>	<b>Product Sets</b>
NAS and object storage	<ul style="list-style-type: none"><li>• Dell EMC ECS</li><li>• Dell EMC Isilon</li><li>• Dell EMC PowerMax</li><li>• Dell EMC Unity</li><li>• Dell EMC VMAX</li></ul>
Software-defined storage	<ul style="list-style-type: none"><li>• Dell EMC ECS</li><li>• Dell EMC VxRail</li><li>• Dell EMC Unity VSA</li></ul>
Direct attached storage	<ul style="list-style-type: none"><li>• MD Series</li><li>• ME Series</li></ul>
All-flash and hybrid storage	<ul style="list-style-type: none"><li>• Dell EMC Isilon</li><li>• Dell EMC PowerMax</li><li>• Dell EMC Unity</li><li>• Dell EMC XtremIO</li><li>• Dell EMC VMAX</li><li>• Dell EMC SC</li></ul>
Entry and midrange storage	<ul style="list-style-type: none"><li>• Dell EMC Unity</li><li>• Dell EMC SC Series</li><li>• Dell EMC NX Series</li><li>• Dell EMC FS Series</li></ul>

## 6 Storage protocols

### 6.1 iSCSI

Internet Small Computer System Interface, or iSCSI, is a transport layer protocol that uses the Transmission Control Protocol/Internet Protocol (TCP/IP) to transport Small Computer System Interface (SCSI) packets, enabling the use of Ethernet-based networking infrastructure as a SAN. The simple, yet powerful technology can help provide a high-speed, low-cost, long-distance storage solution for websites, service providers, enterprises, and other organizations.

### 6.2 FC

Fibre Channel (FC) was designed to extend the functionality of SCSI into point-to-point, loop, and switched topologies. Fibre Channel is an established low-latency, high-bandwidth, high-throughput protocol due to its deterministic nonblocking design. This high-link efficiency makes FC well-suited for storage traffic. FC is transparent to the applications, as the interface is provided by the server operating system or hypervisor. FC has low server requirements because its processing is offloaded from the server to the network adapter leaving more CPU cycles and memory for application services.

### 6.3 FCoE

Fiber Channel over Ethernet (FCoE) is a protocol that encapsulates FC frames over Ethernet networks. FCoE relies on SAN mapping without Ethernet forwarding dependence for robust and reliable SAN performance. FCoE uses 10 Gb Ethernet networks and fully adheres with FC protocol requirements.

### 6.4 NFS

Network File System (NFS) is a protocol regularly used on Linux and UNIX systems to access or share files between devices on a local area network (LAN).

### 6.5 SMB

In computer networking, Server Message Block (SMB), also known as Common Internet File System (CIFS), operates as an application-layer network protocol used to provide shared access to files, printers, and devices on a network. SMB protocol is a network file sharing protocol that consists of data packets in a set. Each set contains a request that a client sends or a response that the server sends. These packets can be broadly classified as follows:

- Session control packets establish and discontinue a connection to shared server resources
- File access packets Accesses and manipulates files and directories on the remote server
- General message packets to send data to print queues, mail slots, and named pipes, to provide data about the status of print queues

## 6.6 CIFS

Common Internet File System (CIFS) is a file-based storage system based on SMB. SMB is a shared storage protocol typically used in Microsoft environments for file sharing. Windows-based file shares rely on CIFS as the transfer protocol of the file level data. File-based storage relies on an underlying files system such as FAT32, XFS, NTFS or otherwise which differs from block-based storage which does not. File level storage is an excellent medium for some applications but is not traditionally effective in others. When an application needs direct block access to disk file-based storage is not appropriate. Deployments that fall into this category include some databases and most Operating Systems.

## 7 Data center leaf-spine networks

Modern data centers are designed to optimize east-west traffic due to the increase of server-to-server communication and the rising popularity of Software Defined Storage (SDS). Leaf-spine architecture is highly scalable. As administrators add racks to the data center, a pair of leaf switches are added to each new rack. Spine switches may be added as bandwidth requirements increase.

In a leaf-spine architecture, the access layer is seen as the leaf layer. Servers and storage devices connect to leaf switches at this layer. At the next level, the aggregation and core 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. This configuration minimizes latency and the likelihood of bottlenecks in the network.

### 7.1 Layer 2 leaf-spine network

In a Layer 2 leaf-spine network, traffic between leaves and spines is switched (except for a pair of edge leaves) as shown in Figure 1. VLT is used for multipathing and load balancing traffic across the Layer 2 leaf-spine fabric. Connections from hosts to leaf switches are also Layer 2.

For connections to external networks, the links between the spines and edge leaves are Layer 3.

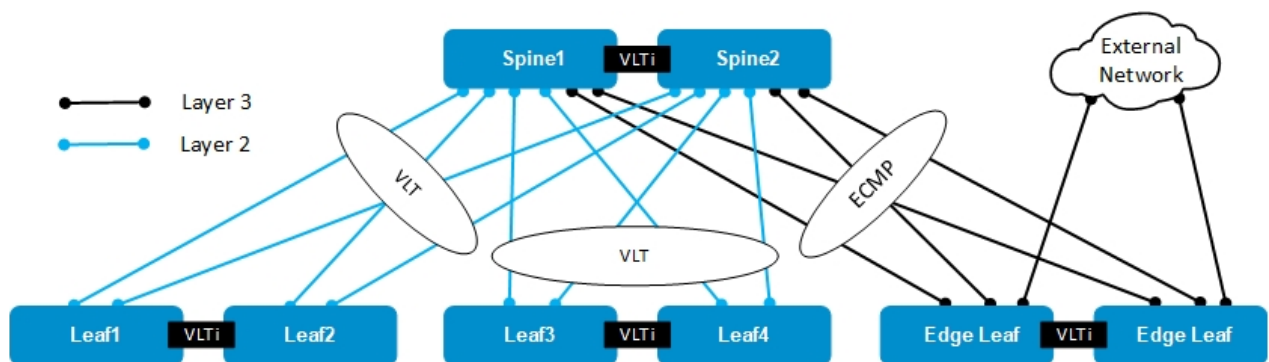


Figure 1 Layer 2 leaf-spine network

Layer 2 topology limitations:

- For each VLAN, the Layer 2 topology creates one large broadcast domain across the fabric. The Layer 3 topology has the benefit of containing broadcast domains to each rack.
- The Layer 2 topology is limited to 4094 VLANs across the fabric. The Layer 3 topology allows up to 4094 VLANs per rack.
- The Layer 2 topology is limited to two physical switches at the spine layer (configured as VLT peers). In a Layer 3 topology, extra spines may be added as needed to provide more paths and bandwidth. A Layer 3 topology is more scalable and is better suited for larger networks.
- Overlay networks using VXLAN, such as VMware NSX, require a Layer 3 underlay network.

The Layer 2 design is relatively simple. The configurations for the leafs and spines are generated by the Dell EMC Fabric Design Center. The Build-Your-Own Network design feature provides a base design for any Layer 2 network. To access the Dell EMC Fabric Design Center, go to <https://fdc.emc.com>.



## 7.2 Layer 3 leaf-spine network

In a Layer 3 leaf-spine network, traffic between leaf and spine switches is routed. The Layer 3, Layer 2 boundary is at the leaf switches. This means at the leaf layer and below (leaf switches and hosts), communication is achieved at Layer 2. However, communication at and above the leaf switches is achieved at Layer 3. Spine switches are never connected to each other in a Layer 3 topology. Equal cost multipath routing (ECMP) is used to load balance traffic across the Layer 3 network. Connections within racks from hosts to leaf switches are Layer 2. Connections to external networks are made from a pair of edge or border leaf switches as shown in Figure 2. Also, connections from the data center core can be made directly to the spines.

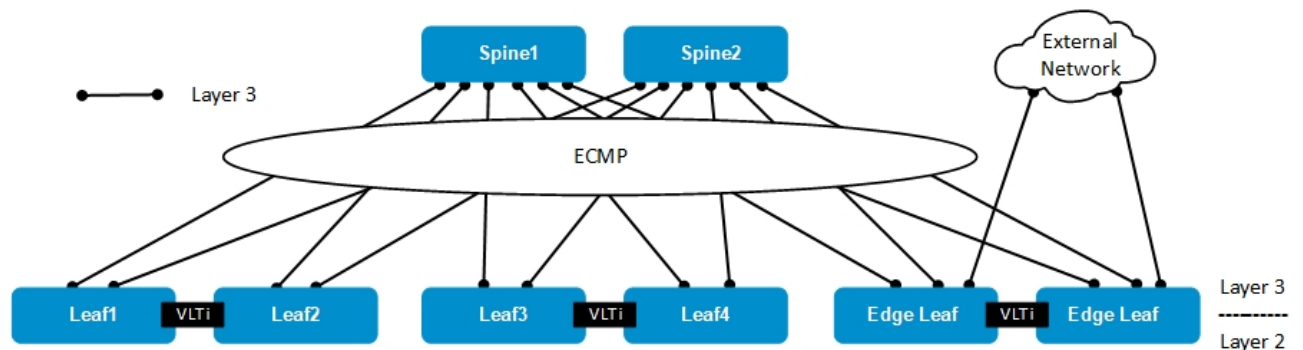


Figure 2 Layer 3 leaf-spine network

The following concepts apply to Layer 3 leaf-spine topologies:

- Each leaf switch connects to every spine switch in the topology
- Servers, storage arrays, edge routers, and similar devices, are always connected to leaf switches and never to spine switches

Layer 3 topologies use two leaf switches at the top of each rack that is configured as a Virtual Link Trunking (VLT) pair. VLT allows all connections to be active while also providing fault tolerance. As administrators add racks to the data center, two leaf switches that are configured for VLT, are added to each new rack.

The total number of leaf-spine connections is equal to the number of leaf switches multiplied by the number of spine switches. To increase fabric bandwidth, implement more connections between leaf and spine switches if the spine layer has the capacity for the additional connections.

To download detailed deployment and best practices guides, go to [Dell EMC Networking Guides](#). To download the document, see [Dell EMC Networking Layer 3 Leaf-Spine Deployment and Best Practices with OS10](#).

## 8 Fibre Channel storage deployment options

Fibre Channel is an established technology and solution for many data centers. The advantages of Fibre Channel provide customers with a high performing, stable solution. Dell EMC Networking provides several options to service and transports Fiber Channel storage traffic through the topologies that are detailed in this section.

### 8.1 Fibre Channel storage with S4148U-ON in F\_port mode

The solution that is provided in this section uses a Dell EMC PowerSwitch S4148U-ON switch for both Ethernet and FC traffic. The S4148U-ON acts as a leaf in a leaf-spine data center. Fibre Channel services are managed on the S4148U-ON switch and link directly to FC HBAs in the servers. An FC storage array is deployed within the rack and FC traffic is contained within the rack. A separate Ethernet NIC is used in the server for data traffic.

This solution is an efficient design that makes full use of the performance advantages of the Fibre Channel protocol at a reduced cost. The S4148U-ON replaces the need for a dedicated FC switch, saving on space and power within the rack.

The combined FC SAN and leaf-spine topology is shown in Figure 3.

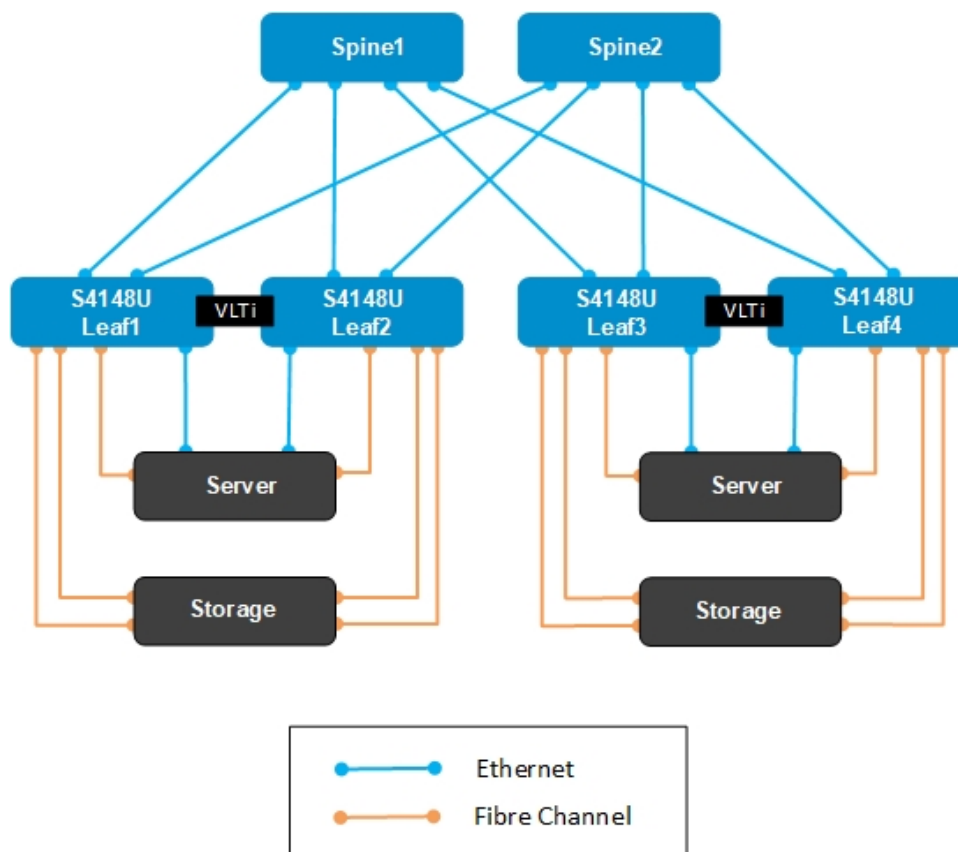


Figure 3 FC with S4148U-ON in F\_Port mode

To obtain the detailed deployment guide with step-by-step configuration of the S4148U-ON switches, see [Dell EMC Networking Fibre Channel Deployment with S4148U-ON in F\\_port Mode](#)

## 8.2 FCoE with S4148U-ON in NPG mode

The solution in this section uses FCoE and the S4148U-ON in NPG mode to manage storage traffic to a dedicated storage network. Fibre Channel services are managed on a traditional FC switch to which the FC storage arrays are connected. The servers contain FCoE CNAs to pass both storage and data traffic to the S4148U-ON switches deployed as a leaf pair in a traditional leaf-spine network. In NPG mode, the S4148U-ON provides FCoE to FC bridging for storage traffic.

This solution provides a scalable design that uses established FC SAN deployments. The S4148U-ON provides flexibility in operating as both a leaf and FCoE to FC bridge to a traditional FC switch.

The leaf-spine topology with a dedicated FC SAN is shown in Figure 4.




Figure 4 FCoE with S4148U-ON in NPG mode

To obtain the detailed deployment guide with step-by-step configuration of the S4148U-ON switches, see [Dell EMC Networking FCoE Deployment with S4148U-ON in NPG Mode](#)

## 8.3 FCoE to Fibre Channel with S4148U-ON in F\_Port mode

The solution in this section uses a Dell EMC PowerSwitch S4148U-ON switch for both Ethernet and FC traffic. The S4148U-ON acts as a leaf in a leaf-spine data center. Fibre Channel services are managed on the S4148U-ON switch. A FC storage array is deployed within the rack and FC traffic is contained within the rack. FCoE CNAs are used in the server for both server and data traffic. In F\_Port mode, the S4148U-ON provides FCoE to FC bridging for storage traffic.

This solution is an efficient design that takes advantage of the performance advantages of the Fibre Channel protocol at a reduced cost. The S4148U-ON replaces the need for a dedicated FC switch, saving on space and power within the rack.

The combined FC SAN and leaf-spine topology is shown in Figure 5.

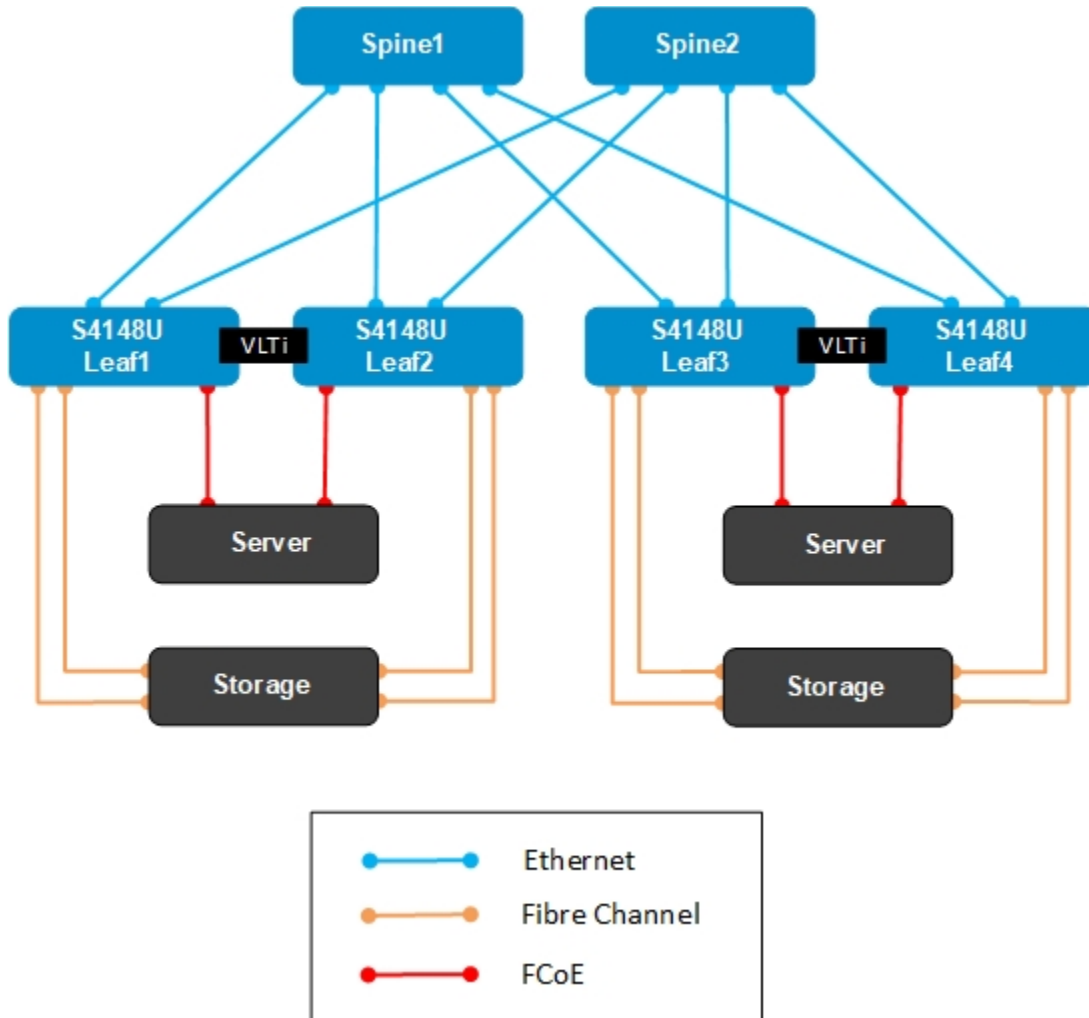


Figure 5 FCoE to Fibre Channel with S4148U-ON in F\_Port mode

To obtain the detailed deployment guide with step-by-step configuration of the S4148U-ON switches, see [Dell EMC Networking FCoE-to-Fibre Channel Deployment with S4148U-ON in F\\_port Mode](#).

## 8.4 FCoE with FIP snooping bridge

The solution that is provided in this section uses two pairs of Dell EMC Networking switches. A pair of leaf switches act as a FIP snooping bridge to a pair of S4148U-ON switches that manage Fibre Channel services for the dedicated storage network. The servers contain FCoE CNAs to pass both storage and data traffic to the switches deployed as a leaf pair in a traditional leaf-spine network.

This solution provides a flexible design that uses established FC storage designs. The switch model for the leaf pair can be any switch with the Dell EMC OS10 Enterprise Edition operating system that allows for the optimal switch port and speeds to be deployed. All Dell EMC Networking switches running OS10EE can be used as an FSB.

Figure 6 shows the leaf-spine topology with a dedicated FC SAN.

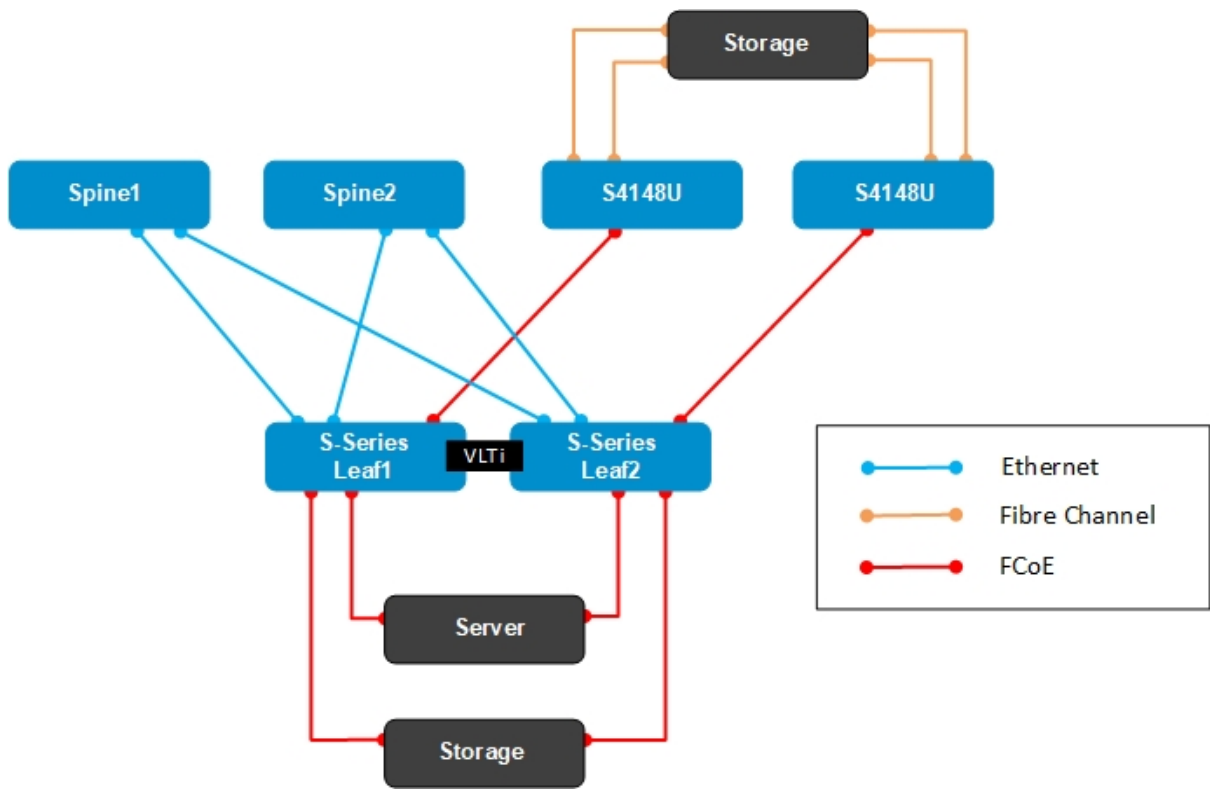


Figure 6 FCoE with FIP Snooping Bridge

To obtain the detailed deployment guide with step-by-step configuration of S5248-ON and S4148U-ON switches, see [Dell EMC Networking OS10EE FCoE Deployment with FSB](#)

## 9 Ethernet-based storage

Ethernet-based storage has the advantage of using the same technology that is established throughout the data center. iSCSI storage technology can be implemented either with a dedicated SAN or through a shared or distributed topology. Advancements in SDS technology allow for great flexibility and scalability through distributed storage with hyper-converged appliances. The topologies in this section show how Dell EMC Networking products can be used to implement both iSCSI and SDS storage solutions.

### 9.1 iSCSI SAN with leaf-spine data center

This section uses a dedicated switch pair to provide an iSCSI session from server CNAs to an iSCSI storage array. Figure 7 shows how storage traffic is serviced with a dedicated storage network while data traffic is serviced on the standard leaf-spine topology.

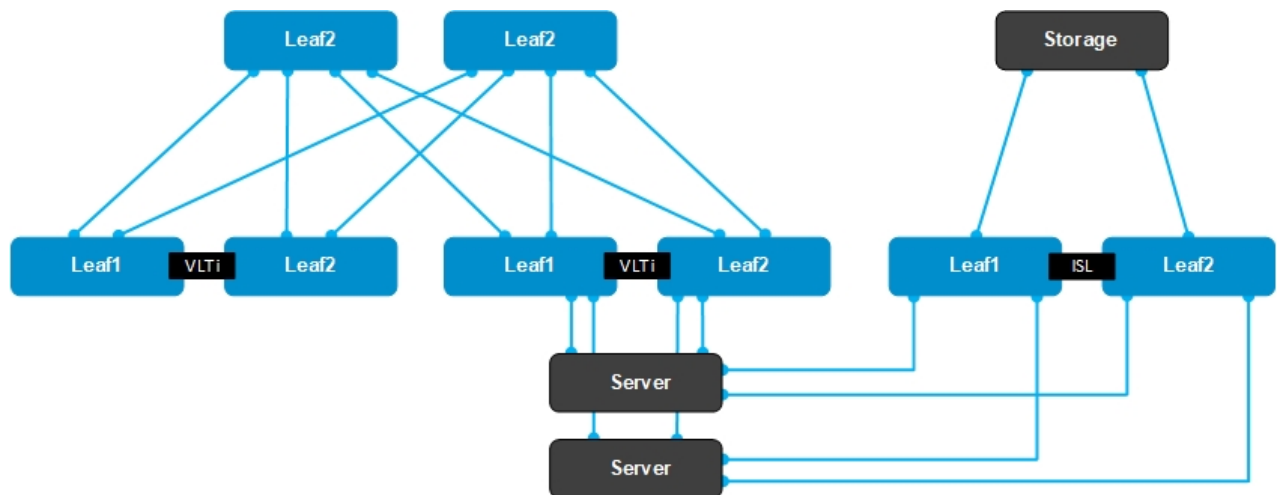


Figure 7 Dedicated iSCSI SAN with leaf-spine data center

A dedicated storage network has been a standard for iSCSI-based storage. The two main reasons to deploy a dedicated storage network are:

- Performance:
  - Low latency
  - Predictable performance that does not compete with application traffic
- Administration:
  - Simplifies management by allowing storage administrators to manage the storage network
  - Streamlines storage troubleshooting using simple fault isolation

#### 9.1.1 iSCSI configuration

Although iSCSI optimization is globally enabled by default on the OS10EE operating system (release 10.4.1.1 and later), manual configuration is often required. To support the iSCSI optimization profile, enable the interface directly connected to the iSCSI storage device. For manual configuration, see the *Configure iSCSI optimization* section in the latest [OS10 Enterprise Edition User Guide](#).

## 9.1.2 Leaf-Spine configuration

The data center topology recommended for modern applications to support large amounts of east-west traffic is the leaf-spine topology. Leaf-spine networks are highly scalable and minimize latency by providing all devices no more than one hop to each destination within the data center.

For detailed deployment instructions see [Dell EMC Networking Layer 3 Leaf-Spine Deployment and Best Practices with OS10](#).

## 9.2 Shared leaf-spine data center with iSCSI or SDS

The solution in this section provides flexibility in the deployment of traditional iSCSI storage arrays and hyper-converged SDS-based storage solutions. This topology allows for storage to be distributed throughout the data center. It is highly scalable and supports large amounts of east-west traffic that is common in modern data centers.

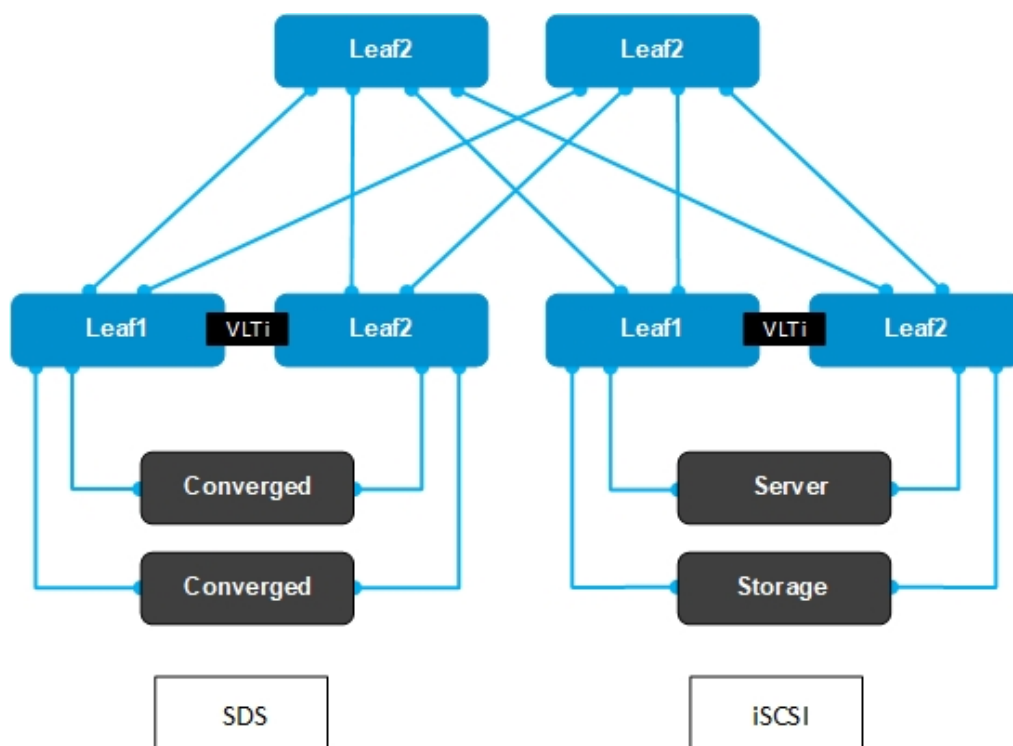


Figure 8 Shared leaf-spine data center

### 9.2.1 Leaf-spine deployment

This solution provided is the deployment of a leaf-spine architecture. For detailed deployment instructions, see [Dell EMC Networking Layer 3 Leaf-Spine Deployment and Best Practices with OS10](#).

### 9.2.2 vSAN deployment with SDS example

Software-defined storage can be implemented through several solutions. VMware vSAN is an example of a SDS solution that can be deployed in various data center architectures. Other examples of SDS include Dell EMC VxFlex OS and Microsoft Storage Spaces Direct.

For detailed vSAN deployment instructions, see [Dell EMC Networking Deploying VMware vSAN with OS10 Enterprise Edition](#).

### 9.2.3 iSCSI configuration

Although iSCSI optimization is globally enabled by default on the OS10EE operating system (release 10.4.1.1 and later), manual configuration is often required. The interface directly connected to the iSCSI storage device must be enabled to support the iSCSI optimization profile. For manual configuration, see the *Configure iSCSI optimization* section in the latest [OS10 Enterprise Edition User Guide](#).



# A Support and feedback

## Contacting Technical Support

Support Contact Information

Web: <http://www.dell.com/support>

Telephone: USA: 1-800-945-3355

## Feedback for this document

To provide feedback on the quality and usefulness of this publication, you are encouraged to send an email to [Dell\\_Networking\\_Solutions@Dell.com](mailto:Dell_Networking_Solutions@Dell.com).