

DATA CENTERS

A holistic view of the data center and the opportunities to enhance its infrastructure to meet current and future demands



Chapter 3


Data center topologies
and architectures

Data center topologies and architectures


Creating a blueprint for better data center performance

A data center networking architecture—the layout of the cabling infrastructure and the way servers are connected to switches—must strike a balance between reliability, performance, agility, scalability and cost. It must also support both current and future applications and speeds.

Key factors in selecting a data center architecture include:



The size of the current data center



Anticipated growth of the data center



Whether it's a new installation or an upgrade of a legacy system

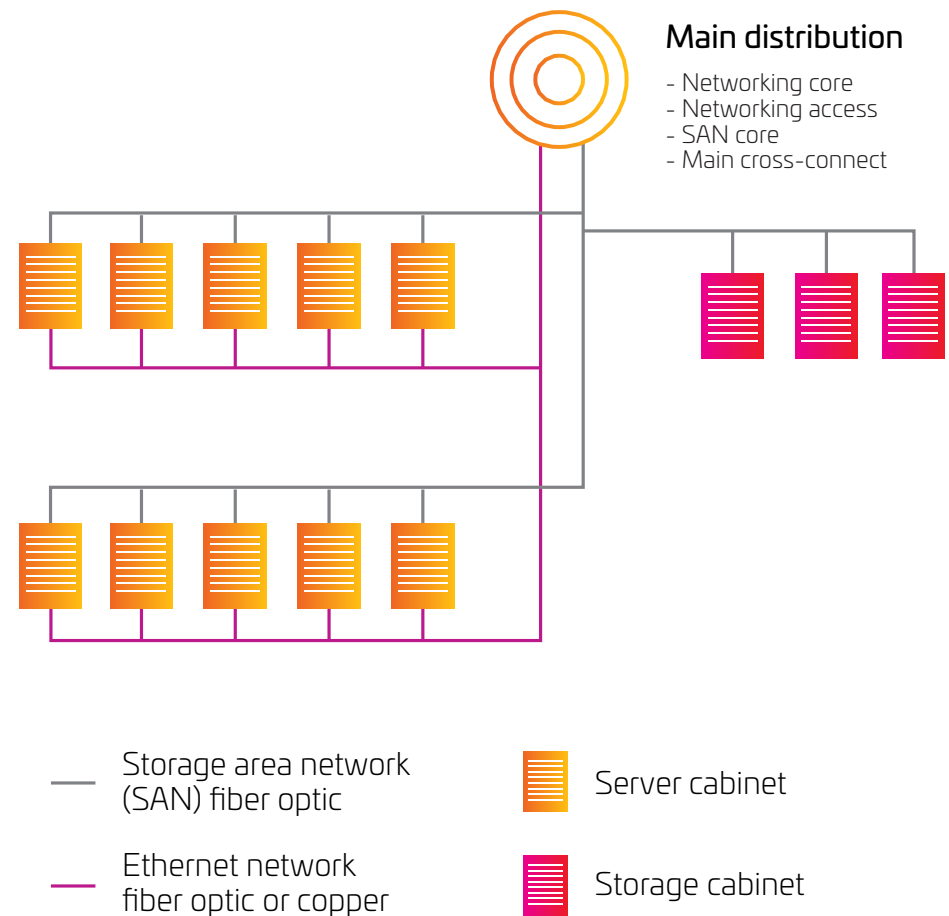
TOPOLOGIES

There are three main data center topologies in use today—and each has its advantages and trade-offs. In fact, some larger data centers will often deploy two or even all three of these topologies in the same facility.

1. CENTRALIZED

The centralized model is an appropriate topologies for smaller data centers (under 5,000 square feet). As shown, there are separate local area network (LAN)/ storage area network (SAN) environments and each one has home run cabling that goes to each of the server cabinets and zones. Each server is effectively cabled back to the core switches, which are centralized in the main distribution area.

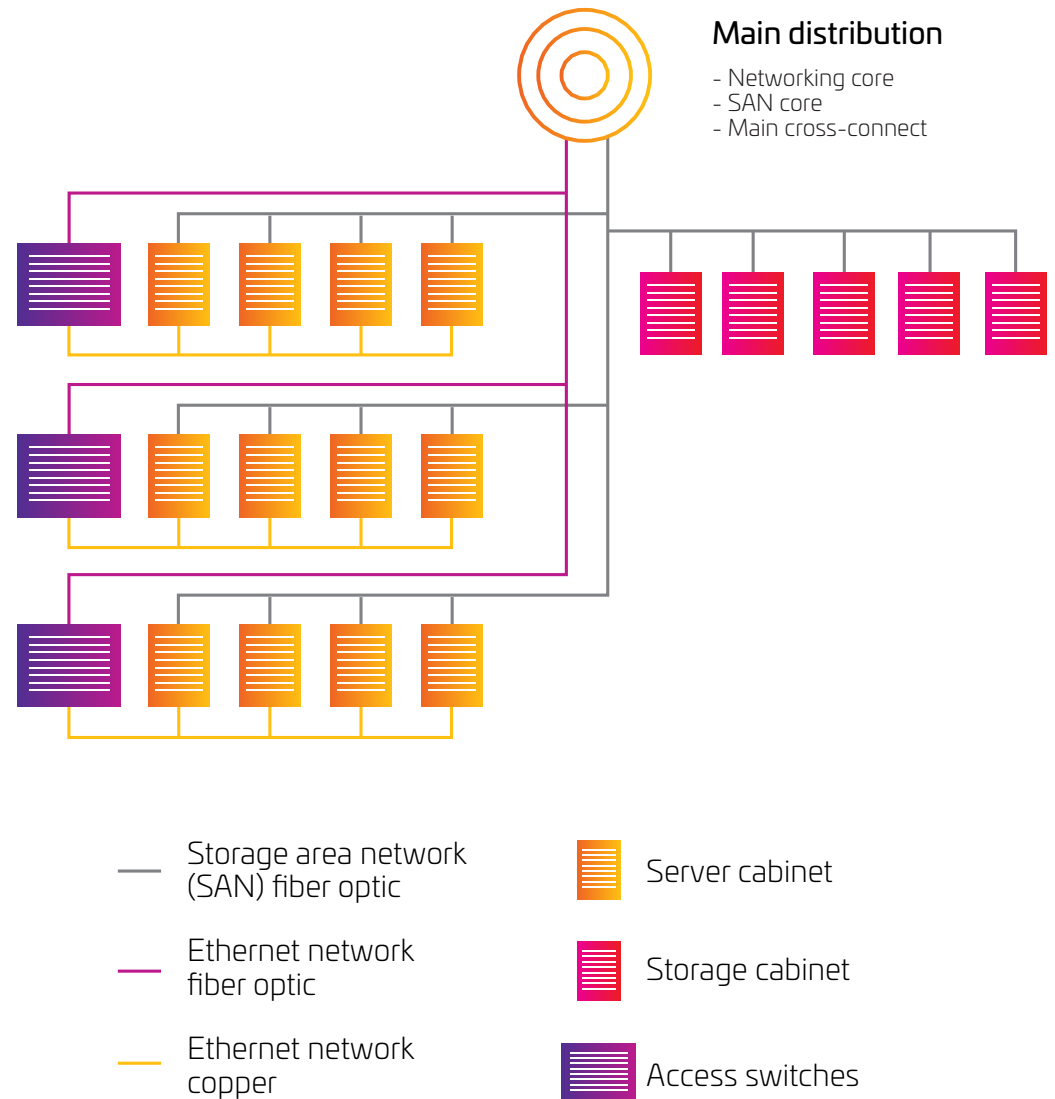
This provides very efficient utilization of port switches and makes it easier to manage and add components. The centralized topology works well for smaller data centers but does not scale up well, which makes it difficult to support expansions. In larger data centers, the high number of extended-length cable runs required causes congestion in the cable pathways and cabinets, and increases cost. While some larger data centers use zoned or top-of-rack topologies for LAN traffic, they may also utilize a centralized architecture for the SAN environments. This is especially true where the cost of SAN switch ports is high and port utilization is important.



2. ZONED

Zoned topology consists of distributed switching resources. As shown below, the switches can be distributed among end-of-row (EoR) or middle-of-row (MoR) locations, with chassis-based switches typically used to support multiple server cabinets. This solution is recommended by the ANS/TIA-942 Data Center Standards and is very scalable, repeatable, and predictable. Zoned architecture is usually the most cost-effective design, providing the highest level of switch and port utilization while minimizing cabling costs.


In certain scenarios, end-of-row switching provides performance advantages. For example, the local area network (LAN) ports of two servers (that exchange large volumes of information) can be placed on the same end-of-row switch, for low-latency port-to-port switching. A potential disadvantage of end-of-row switching is the need to run cable back to the end-of-row switch. Assuming every server is connected to redundant switches, this cabling can exceed what is required in top-of-rack architecture.



3. TOP OF RACK

Top-of-rack (ToR) switching typically consists of two or more switches placed at the top of the rack in each server cabinet, as shown below. This topology can be a good choice for dense one rack-unit (1RU) server environments. All servers in the rack are cabled to both switches for redundancy. The top-of-rack switches have uplinks to the next layer of switching. Top of rack significantly simplifies cable management and minimizes cable containment requirements. This approach also provides fast port-to-port switching for servers within the rack and predictable oversubscription of the uplink.


A top-of-rack design utilizes cabling more efficiently. The tradeoffs are often an increase in the cost of switches and the high cost for under-utilization of ports. Top-of-rack switching may be difficult to manage in large deployments, and there is also the potential for overheating of local area network (LAN) switch gear in server racks. As a result, some data centers deploy top-of-rack switches in a middle-of-row or end-of-row architecture to better utilize switch ports and reduce the overall number of switches used.



ARCHITECTURES


1. MESH NETWORK

The mesh network architecture, often referred to as a “network fabric,” or leaf-spine, consists of meshed connections between leaf-and-spine switches. The mesh of network links enables any-to-any connectivity, with predictable capacity and lower latency—making this architecture well suited for supporting universal “cloud services.” With multiple switching resources spread across the data center, the mesh network is inherently redundant for better application availability. These distributed network designs can be much more cost-effective to deploy and scale when compared to very large, traditional centralized switching platforms.



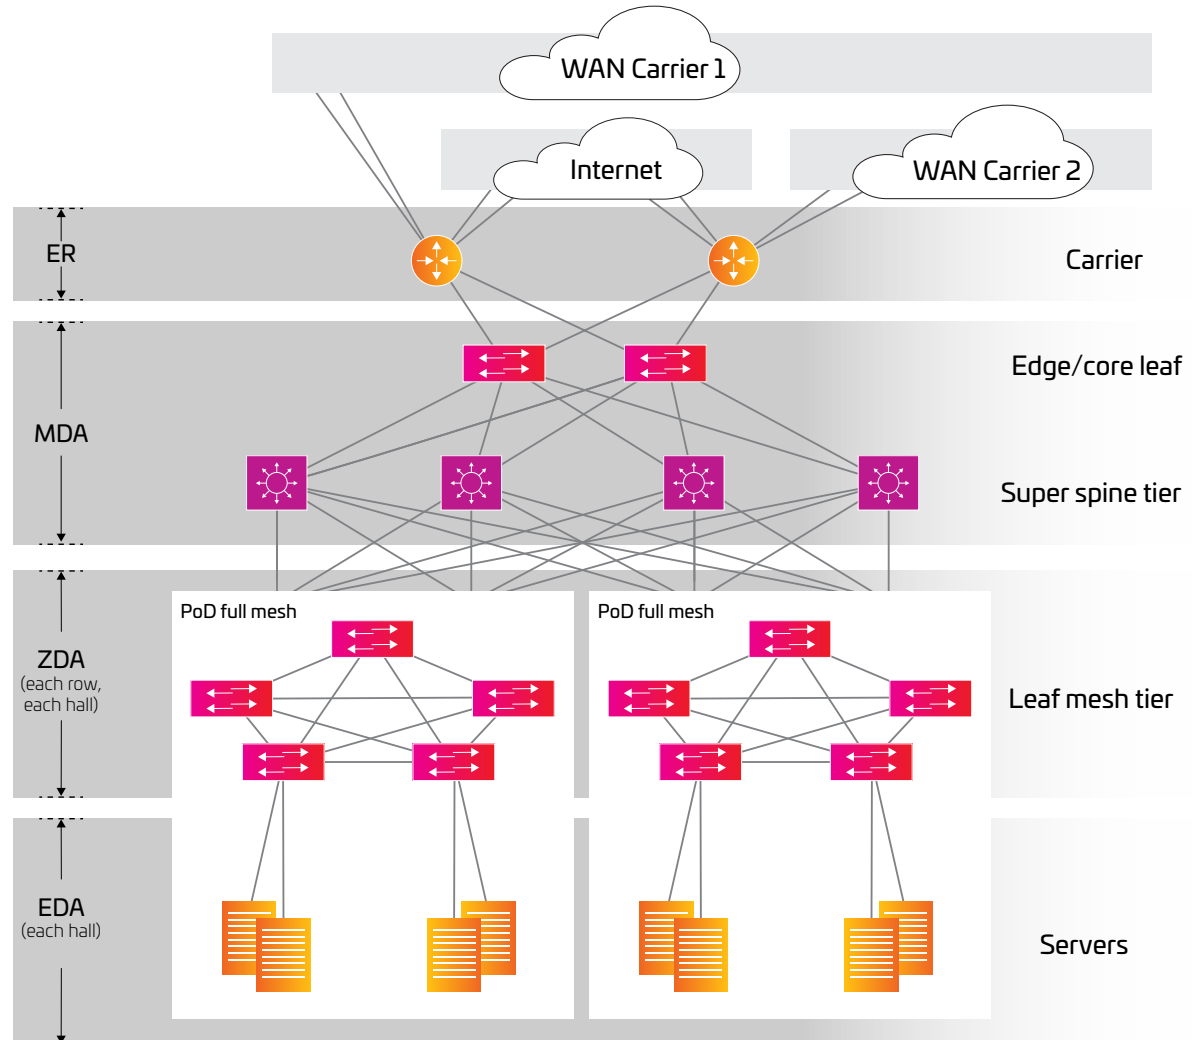
2. THREE-TIER OR MULTI-TIER MODEL

The multi-tier architecture has been the most commonly deployed model used in the enterprise data center. This design consists primarily of web, application and database server tiers running on various platforms, including blade servers, 1RU servers and mainframes.



3. MESH POINT OF DELIVERY (PoD)

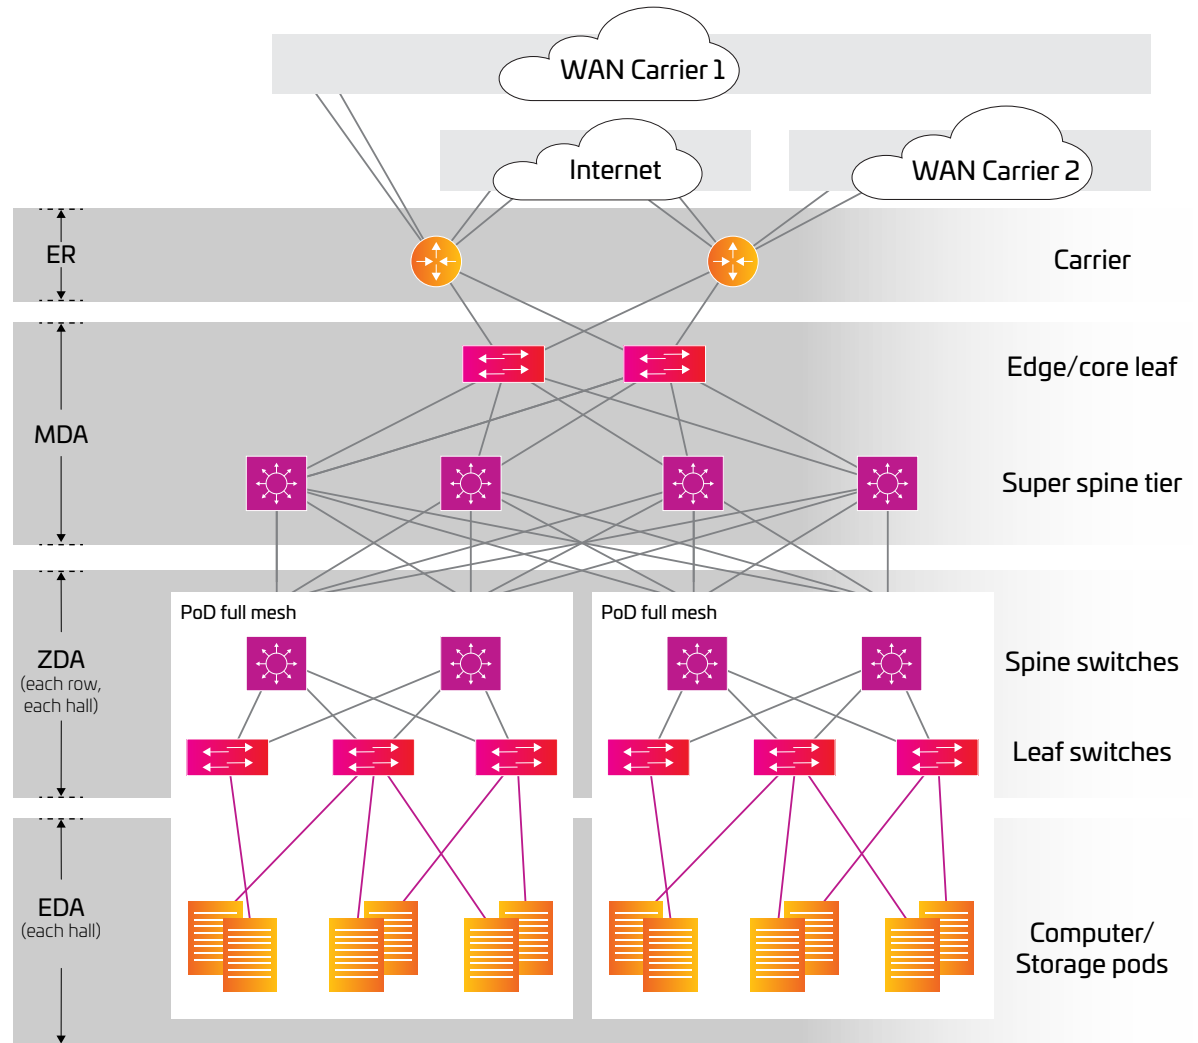
The mesh point of delivery (PoD) architecture features multiple leaf switches interconnected within in the PoDs, with spine switches typically aggregated in a central main distribution area (MDA). Among other advantages, this architecture enables multiple PoDs to connect efficiently to a super-spine tier. Data center managers can easily add new infrastructure to their existing three-tier topology to support the low-latency east-west data flow of new cloud applications. Mesh PoD networks can provide a pool of low-latency compute and storage for these applications that can be added without disrupting the existing environment



4. SUPER SPINE MESH

Super spine architecture is commonly deployed by hyperscale organizations deploying large-scale data center infrastructures or campus-style data centers.

This type of architecture services huge amounts of data passing east to west across data halls.




— Ethernet link

— Ethernet / FCoE link


Data center equipment connection methods

There are two methods typically used to connect data center electronics via structured cabling: cross-connect and interconnect.




Cross-connect

A cross-connect uses patch cords or jumpers to connect cabling runs, subsystems and equipment to connecting hardware at each end. It enables connections to be made without disturbing the electronic ports or backbone cabling. A cross-connect provides excellent cable management and design flexibility to support future growth. Designed for “any-to-any” connectivity, this model enables any piece of equipment in the data center to connect to any other regardless of location. A cross-connect also offers operational advantages, as all connections for moves, add and changes are managed from one location. The major disadvantage is higher implementation costs due to increased cabling requirements.



Interconnect

An interconnect uses patch cords to connect equipment ports directly to the backbone cabling. This solution requires fewer components and is, therefore, less expensive. However, it reduces flexibility and introduces additional risk, as users must directly access the electronics ports in order to make the connection. Therefore, CommScope generally recommends utilizing cross-connects for maximum flexibility and operational efficiency in the data center.



Architecture for a higher-speed future

With newer technologies—25/50/100GbE, 32G and 128G Fibre Channel—limitations on bandwidth, distance and connections are more stringent than with lower-speed legacy systems. In planning the data center’s local area network (LAN)/storage area network (SAN) environment, designers must understand the limitations of each application being deployed and select an architecture that will not only support current applications but also have the ability to migrate to higher-speed future applications.



Return back to full
Data Center eBook



For more information
on enhancing your data
center, reach out to one
of our experts now.



CommScope pushes the boundaries of communications technology with game-changing ideas and ground-breaking discoveries that spark profound human achievement. We collaborate with our customers and partners to design, create and build the world's most advanced networks. It is our passion and commitment to identify the next opportunity and realize a better tomorrow.

Discover more at commscope.com

www.commscope.com

Visit our website or contact your local CommScope representative for more information.

© 2018 CommScope, Inc. All rights reserved.

All trademarks identified by @ or ™ are registered trademarks or trademarks, respectively, of CommScope, Inc.

This document is for planning purposes only and is not intended to modify or supplement any specifications or warranties relating to CommScope products or services.

CO-110101.3-EN (08/18) CHAPTER 3