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

# DATA CENTERS

For over 40 years, CommScope's highly trained specialists have partnered with our customers to identify, design and build specialized solutions for data centers. Drawing upon this wealth of expertise and experience, CommScope developed this eBook to provide a holistic overview of the data center and share guidance about how to navigate through the ever-increasing challenges faced by data center managers both now and in the future.

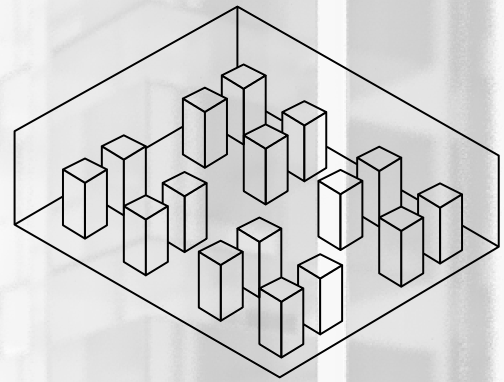

Explore the chapters to find out tips, answers and insights to demystify the technology, untangle the complexity and accelerate time to market so you can identify the challenges—and opportunities—in your own data center.

# DATA CENTERS

Chapter 1

## Data center standards

# Data center standards

## Setting the standards for higher speeds

Data centers and their contents must adhere to a wide range of standards, ranging from local building codes to guidelines from the American Society of Heating, Refrigerating and Air-Conditioning Experts (ASHRAE) on cooling to a number of requirements placed on the IT equipment. There are also a number of standards related to the structured cabling infrastructure that serves as the platform for IT equipment in the data center.

CommScope monitors data center trends and participates in standards organizations to help data center operators stay ahead of the industry. Given the relentless growth in data traffic and the need to provide high-bandwidth, low-latency connections, there has been a tremendous amount of activity within the standards bodies to define higher speeds. It is important to keep up with the latest developments to ensure the cabling infrastructure can support these higher speeds with minimal disruption.

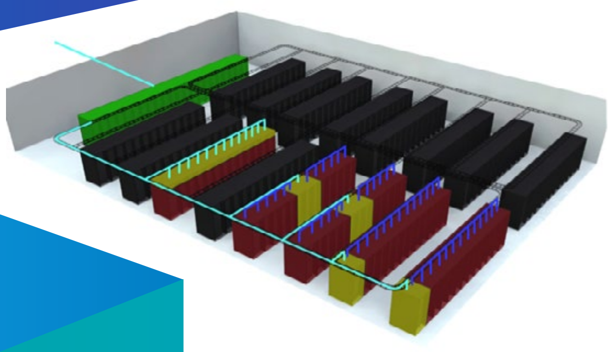

*Server to Switch Connectivity*

# How the standards define data center cabling

There are two main types of standards relevant to data center cabling infrastructure:

## Application Standards

Applications standards define the application that will run on the cabling infrastructure. There are three applications standards that are the most commonly deployed in data centers.

**IEE 802.3** (Ethernet standards) have been particularly active, and currently have draft standards underway for applications up to 400 Gb/s.

**INCITS T11** (Fibre Channel) covers storage area networks (SANs), with published standards for up to 128 Gb/s with a roadmap out to 1 Tb/s.

**Infiniband™** Infiniband Trade Association is used primarily for high-performance computing applications, with a roadmap with options for up to 600 Gb/s.

Applications standards also define the distance that an application can operate over a given media type. For example, under IEEE 802.3an, 10GBASE-T can operate at up to 100 meters over Category 6A cabling.

## Cabling Standards

Cabling standards provide more detail around the physical media and define the channel that supports the applications. There are three main cabling standard bodies.

| TIA® | CENELEC | ISO/IEC |
|---|---|---|
| North America | Europe | Global |

Each of these groups has a general standard which defines structured cabling, as well as a standard specifically for data center applications to reflect the need for higher speeds, increased density and an array of architectures. While there are differences between these standards, there is agreement around the minimum recommended cabling categories and connector types.

| | TIA® | CENELEC | ISO/IEC |
|---|---|---|---|
| Data center standard | TIA-942-B | EN 50173-5 | ISO-IEC 11801-5 |
| Fiber | OM4 | OM3 | OM3 |
| | OS1a | OS2 | OS2 |
| Connectors | LC (≤ 2 fibers) | | |
| | MPO (≥ 2 fibers) | | |

In addition to EN50173-5, CENELEC has also developed the EN 50600-2-4 standard "Telecommunication Cabling Infrastructure". It focuses primarily on design requirements for the different DC availability classes with strong emphasis on migration and growth.

# Keep up with the standards to future-proof the data center

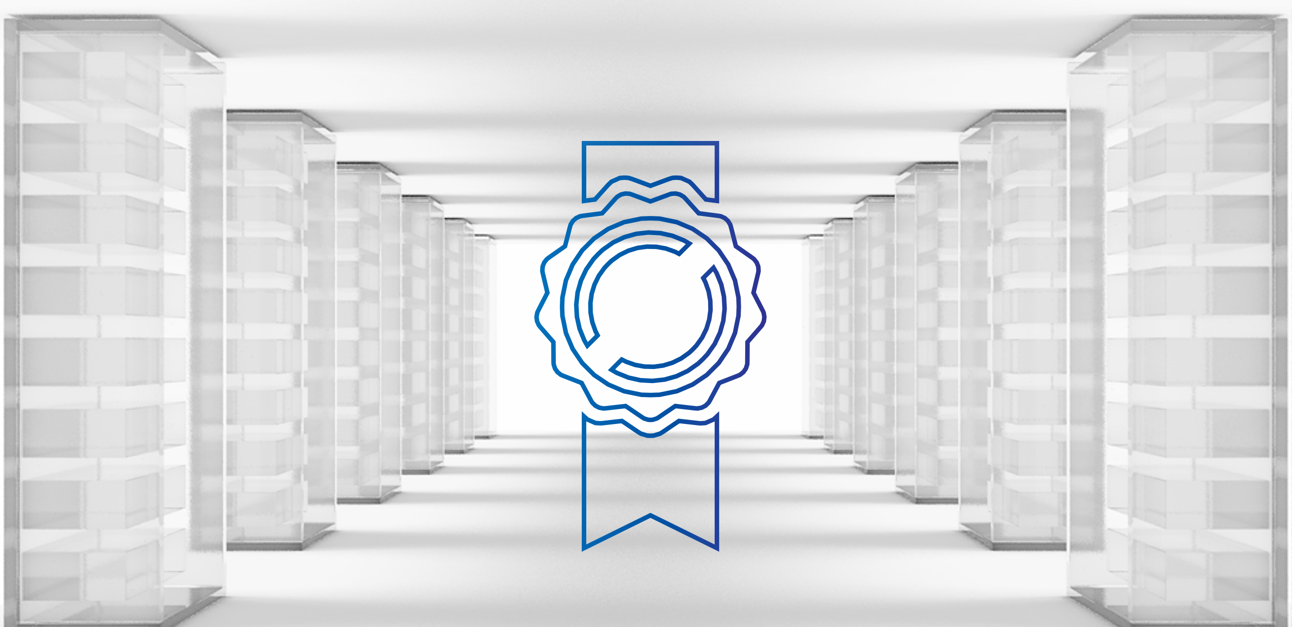

Since data center cabling infrastructure will likely need to support multiple generations of equipment and speeds in the future, keeping up with the latest standards developments is critical. For new builds, it is important to deploy the highest bandwidth cabling as per the data center cabling standards.

**White Paper**

Data Center Application Standards Reference Guide - Networking and Storage

**Standards updates**

CommScope Quarterly Standards Advisor

**White Paper**

Data Center Cabling Design Fundamentals: Telecommunication Cabling Infrastructure Requirements

# DATA CENTERS

Chapter 2

## Multi tenant data centers

# Multi tenant data centers

## Freeing enterprises to focus on their core business

With more than 130 million square feet of white space, multi tenant data centers (MTDCs) are one of the fastest-growing segments in the data center industry. Expanding at a rate of 16 percent annually, MTDCs enable enterprises and service providers to outsource their data center facilities. By leasing third-party data center whitespace, enterprises can remain focused on their core business while enjoying optimal data center availability, reliability and cost control.

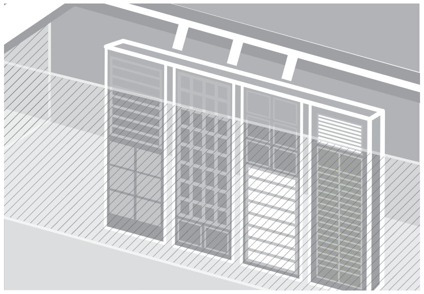

As a leading provider of multi tenant data center infrastructure solutions, CommScope has developed considerable expertise in this space, helping both operators and tenants maximize the value of these flexible, cost-effective facilities.

# KEY ADVANTAGES

## Flexibility

Multi tenant data center infrastructure makes advanced technology such as cloud computing and virtualized data centers available to small- and mid-sized businesses while also allowing easy expandability as the business grows.

## Reliability

Multi tenant data centers provide their own technicians to maintain the infrastucture and ensure that hosted functions operate at peak efficiency at all times. Multi tenant data center operators provide service level agreements (SLAs) to tenant clients to ensure commitment to uptime and operational parameters. Multi tenant data center operators typically offer 2N, N+1, N and hybrid mesh solutions for power redundancy with multiple POPs (point of presence)/ POEs (point of entrance), as well as multiple metro/WAN connectivity providers to provide redundancies that increase reliability. This enables clients to balance their redundancy/reliability needs against their cost options. Some clients may require lower levels of reliablility for certain applications, such as deploying a lab environment; Multi tenant data centers can match the reliability requirements to specific user requirements.

## Reduced Latency

By providing direct connectivity to service providers, content providers, cloud providers, high-frequency traders, financial transaction and peering partners also co-located at the multi tenant data center, latency can be significantly reduced.

# KEY ADVANTAGES

## Savings

By outsourcing data center services instead of building, hosting, maintaining and upgrading them themselves, multi tenant data center tenants can realize significant OpEx and CapEx savings. Most companies are not in the business of building and operating data centers. The expertise and efficiencies gained by multiple builds and design iterations have enabled Multi tenant data center operators to optimize their designs and operational efficiencies. Multi tenant data center operators can not only build a data center more cost-efficiently but are also able to operate it more cost-effectively, as well. Building a traditional data center is a significant capital expense for enterprises; Multi tenant data center operators offer conversion from CapEx to OpEx by leasing the data center to the client, and also offer savings from tenant improvement or asset amortization. Enabling direct connection between enterprises, vendors, content providers and cloud providers in the same facility eliminates the need for metro/WAN connections that have backhaul and bandwidth charges. Multi tenant data centers offer clients the ability to scale as they grow, and to deploy assets on a just in time basis.Most leases run from three to 15 years, which gives the customer the ability to dynamically manage their business versus trying to over-plan and build a traditional data center that is an up to 30-year depreciating asset.

## Security

A multi tenant data center offers multiple levels of security against external threats plus faster, more thorough recovery from disaster situations. The initial layer of security is at the entry points of the facility or campus, which are usually surrounded by high steel fences, gates and bollards, and equipped with a badge or biometric readers and security personnel. The facilities themselves are designed to restrict accessibility while maintaining a discrete appearance. Inside, there are security guards, restricted access and man traps that are designed to slow and restrict entry. Only authorized personnel are allowed entry to designated areas via badge or biometric access. In addition, the entire campus is under continuous monitoring via security cameras, and may often be subject to random security patrols.

# Factors to consider when choosing a multi tenant data center

The best multi tenant data center providers help clients through all stages of the project, from need assessment through migration through operation. They have the core competencies to understand, meet and exceed clients requirements.

High-quality, robust mechanical, electrical, and data transport infrastructure that performs optimally and deploys quickly, allowing for fast, simple changes.

Physical layer management that supports the tenant's differentiated services and allows visibility into the enterprise.

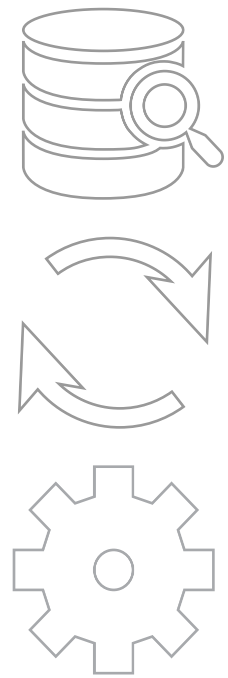

Scalability to readily expand capacity and functionality under the same roof to meet increasing data center demand as the tenant's business grows. This includes space, power and bandwidth scalability, and also the ability to scale down should there be a shift in public cloud utilization.

Direct access to cloud and content providers. Today's and tomorrow's data centers are and will continue to be connected with content and cloud providers in an effort to support internal and external customers. The ability to have direct access to these providers improves latency and cost objectives.

Effective in-building wireless (IBW) service for tenants working onsite, since data centers are often constructed with reinforced materials that make IBW coverage difficult. Providing mobile service to all mobile operators ensures that tenants have access, ensuring their productivity and efficiency while they are onsite.

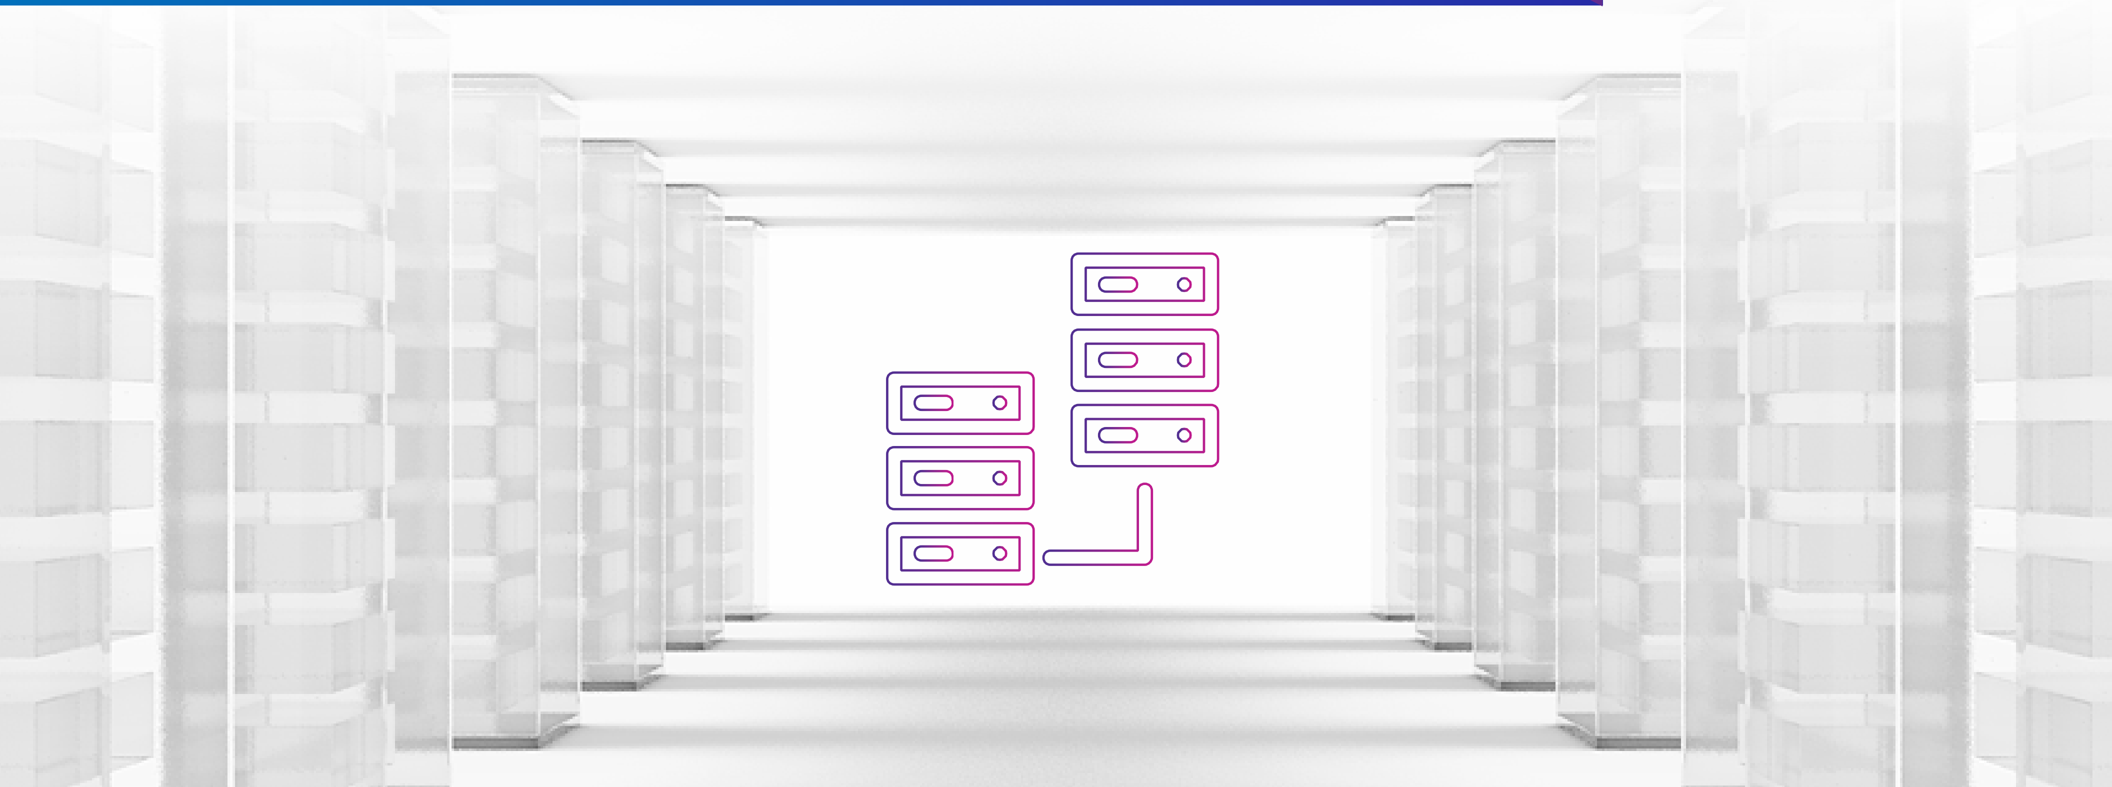

When designing the cabling infrastructure for deployment in an MTDC, it is critical that the infrastructure be able to scale up and down, and adapt to technology changes. It must be able to integrate peer, cloud and hybrid strategies, and must also ensure that adequate space is allocated for horizontal and vertical cable management so that moves, adds and changes can be carried out easily. There is often the temptation to fully load server and switch cabinets without consideration of cable management. While this may reduce the amount of leased rack space, it also greatly increases the risk of downtime due to manual error. Proper cabling design, along with the use of an automated infrastructure management (AIM) system, can reduce this risk.

# DATA CENTERS

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:

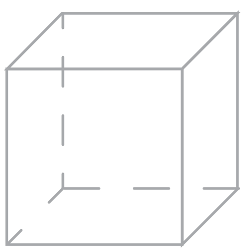

The size of the current data center

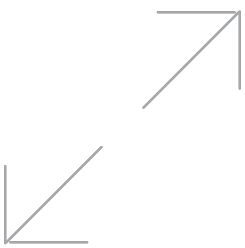

Anticipated growth of the data center

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.

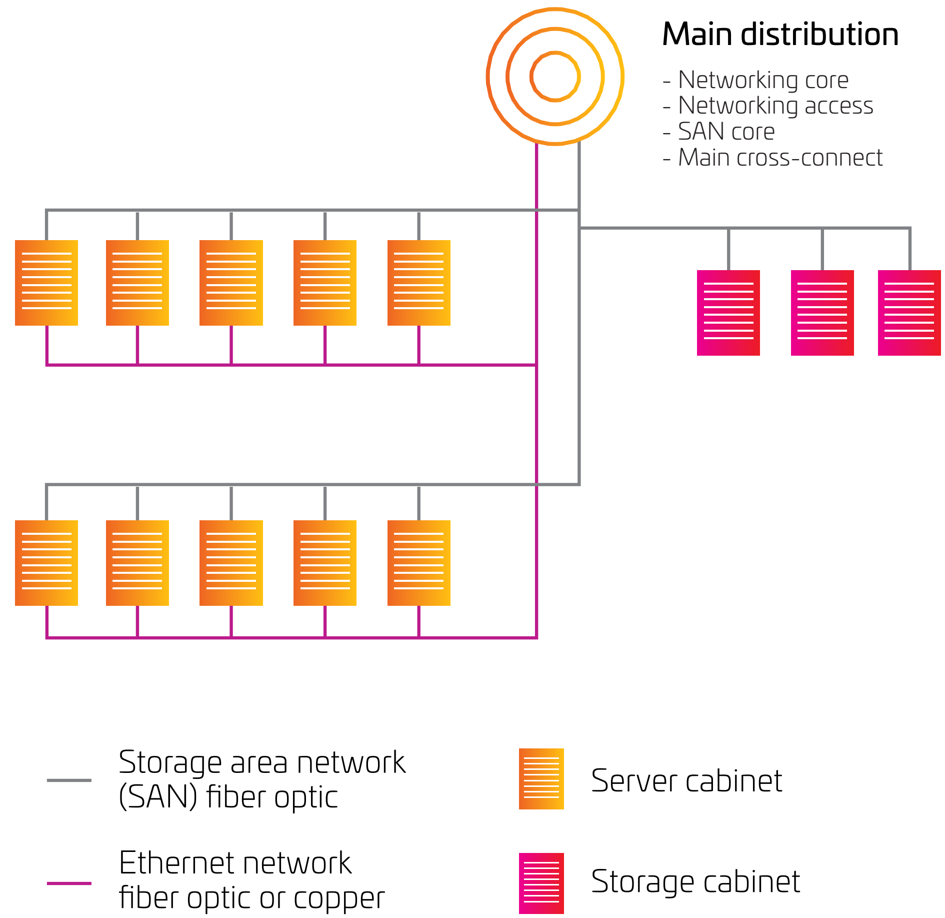

## 2. ZONED

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.

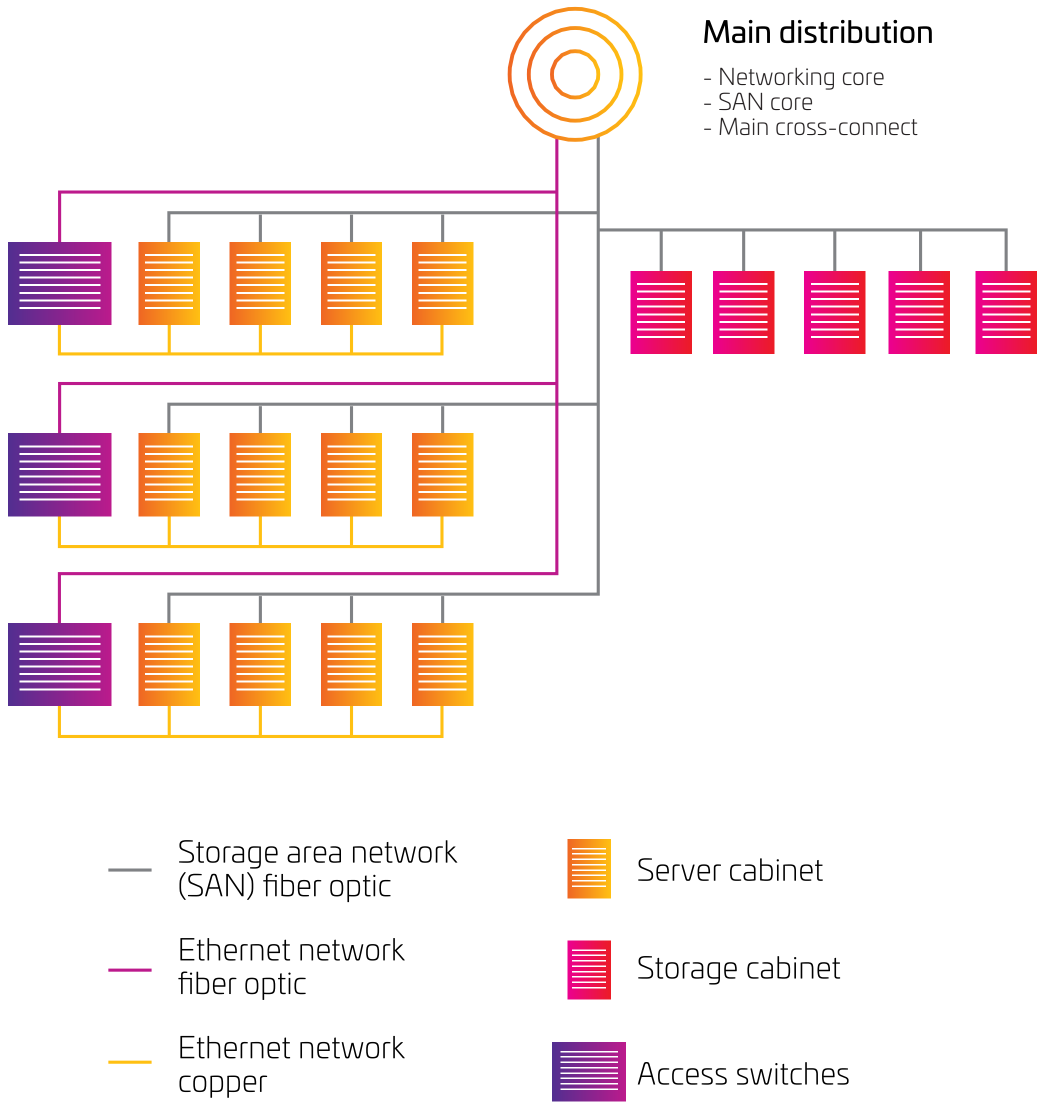

## 3. TOP OF RACK

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.

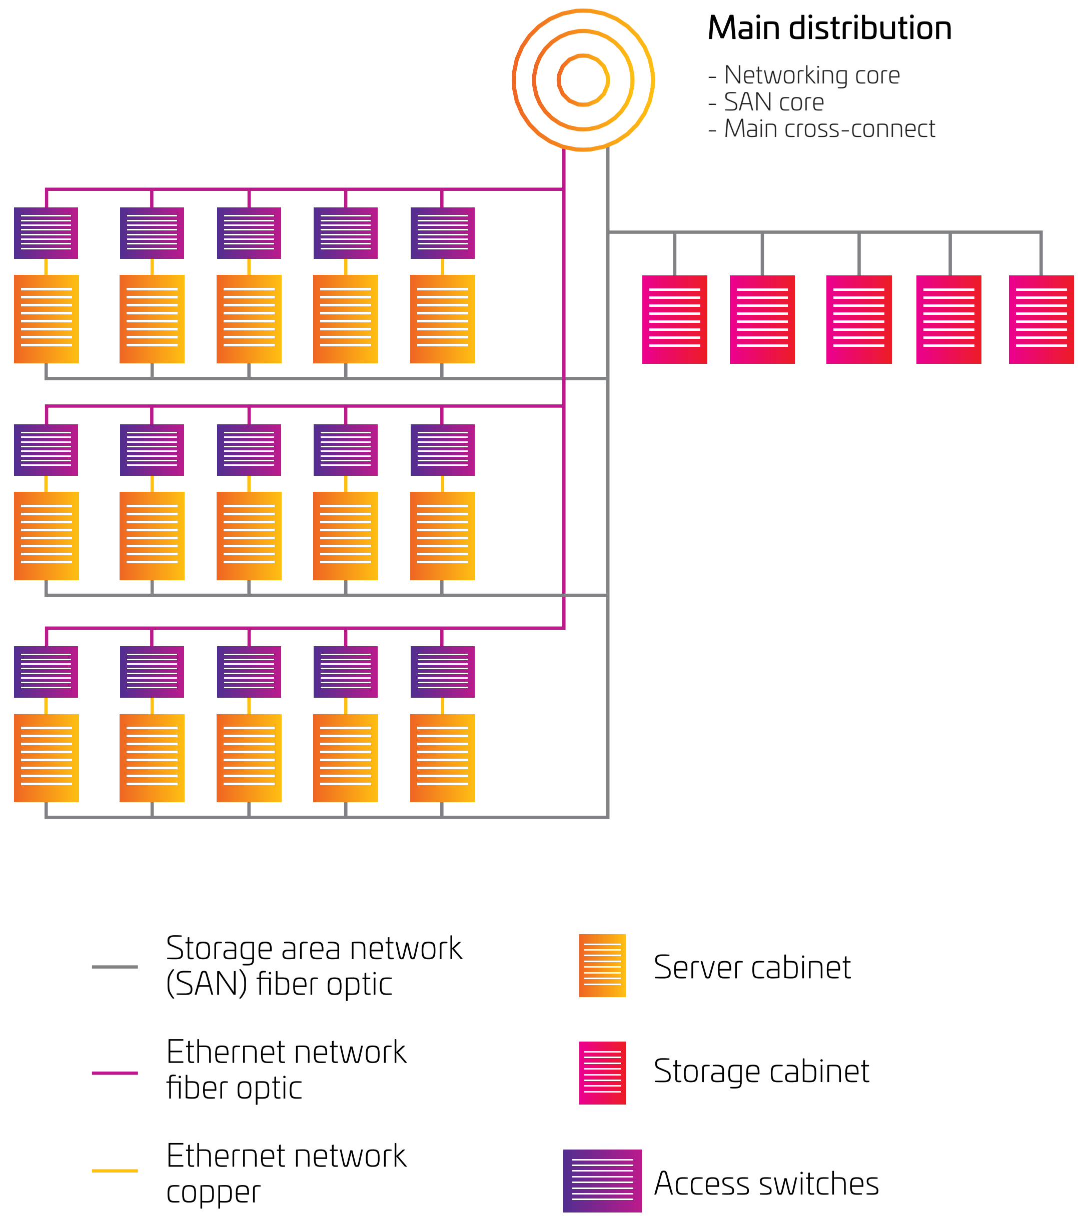

# ARCHITECTURES

## 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.

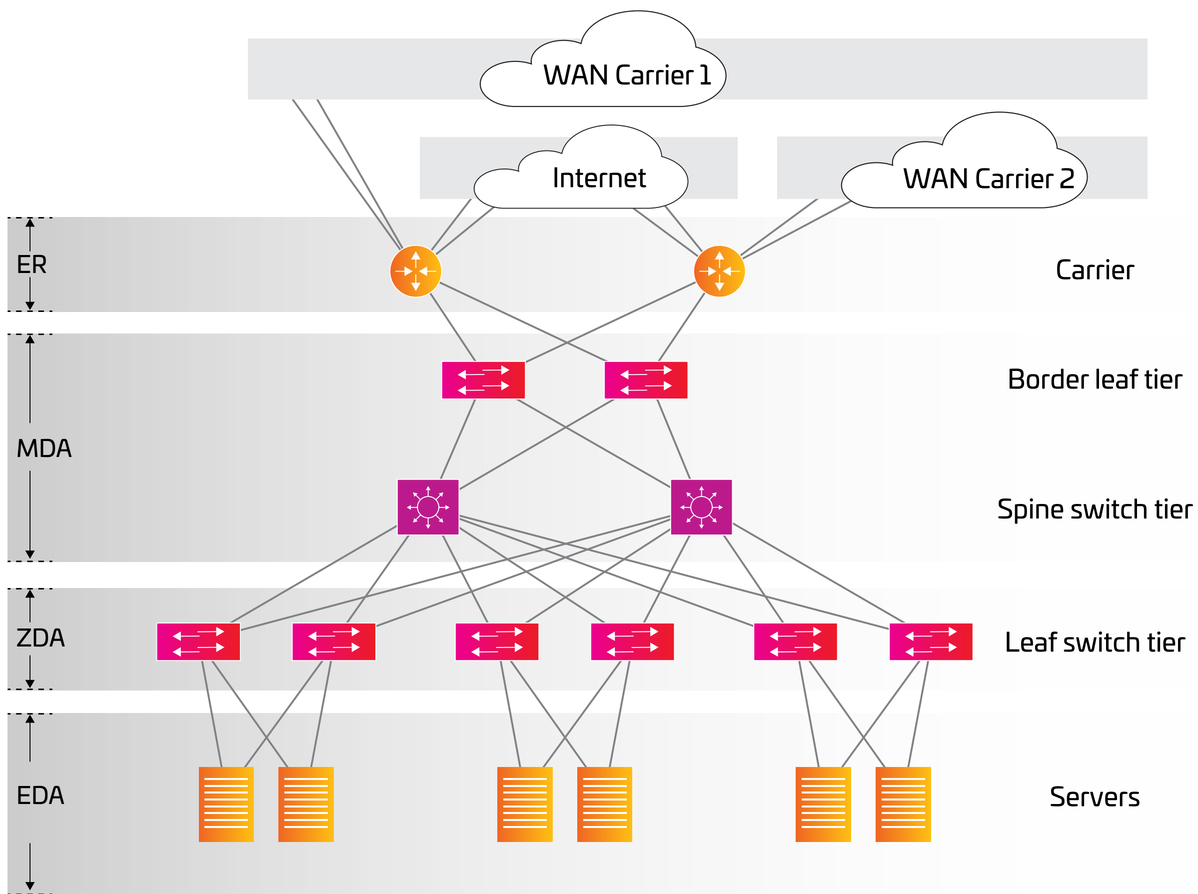

## 2. THREE-TIER OR MULTI-TIER MODEL

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.

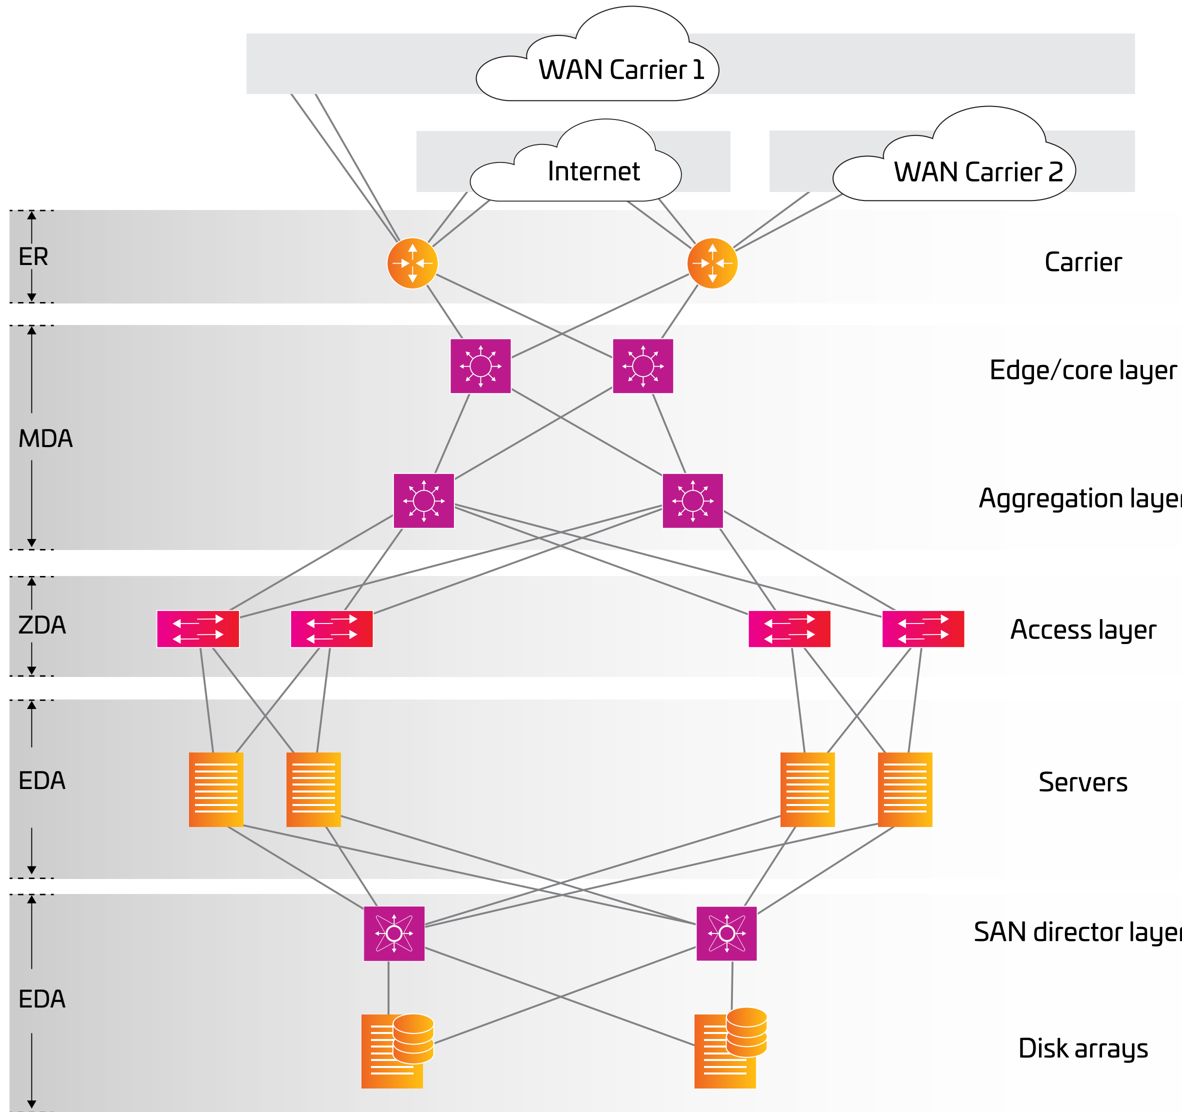

## 3. MESH POINT OF DELIVERY (PoD)

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

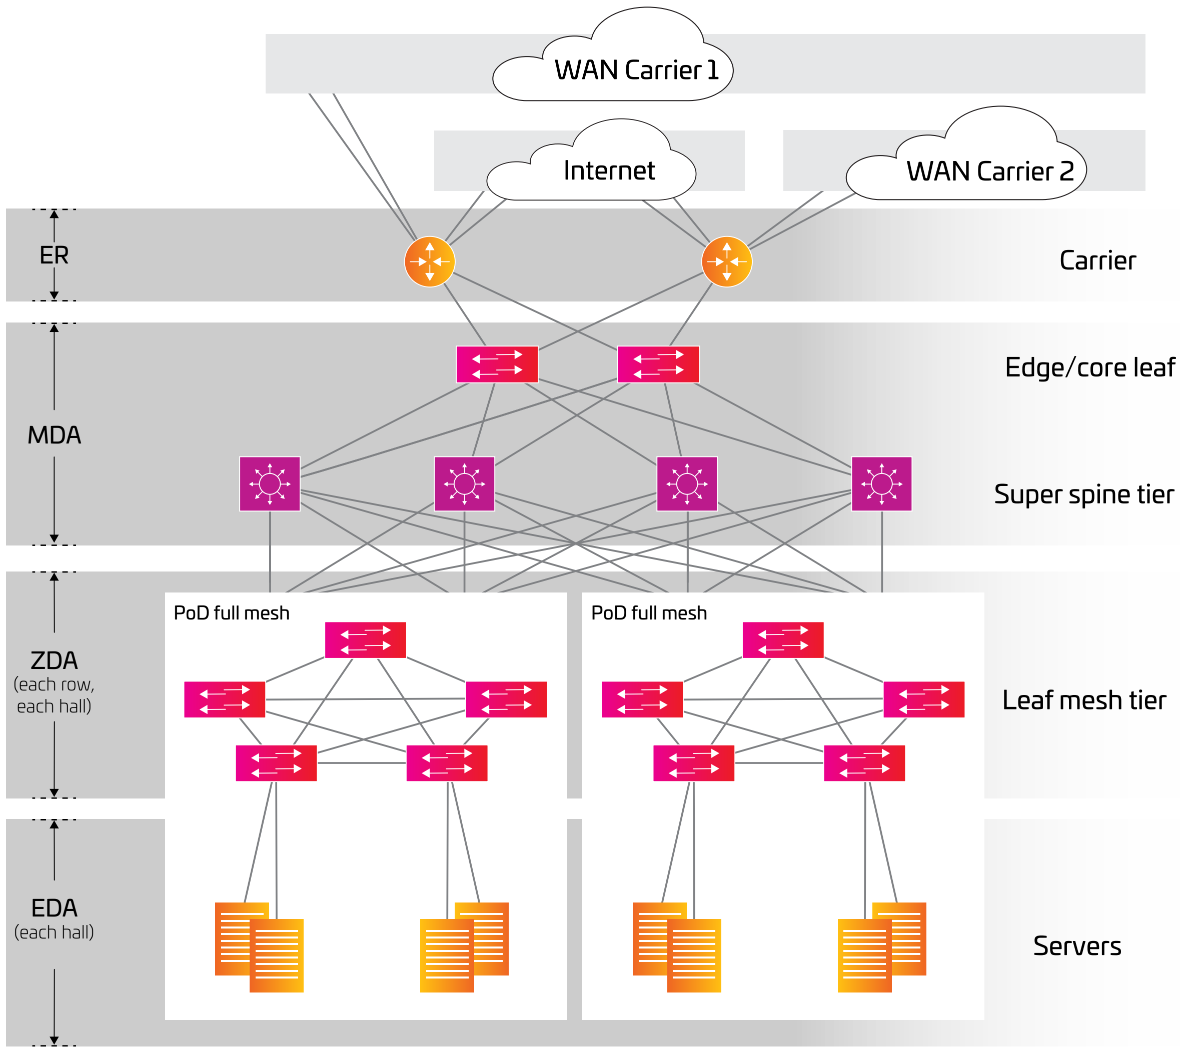

Super spine architecture is commonly deployed by hyperscale organizations deploying large-scale data center infrastructures or campus-style data centers.

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

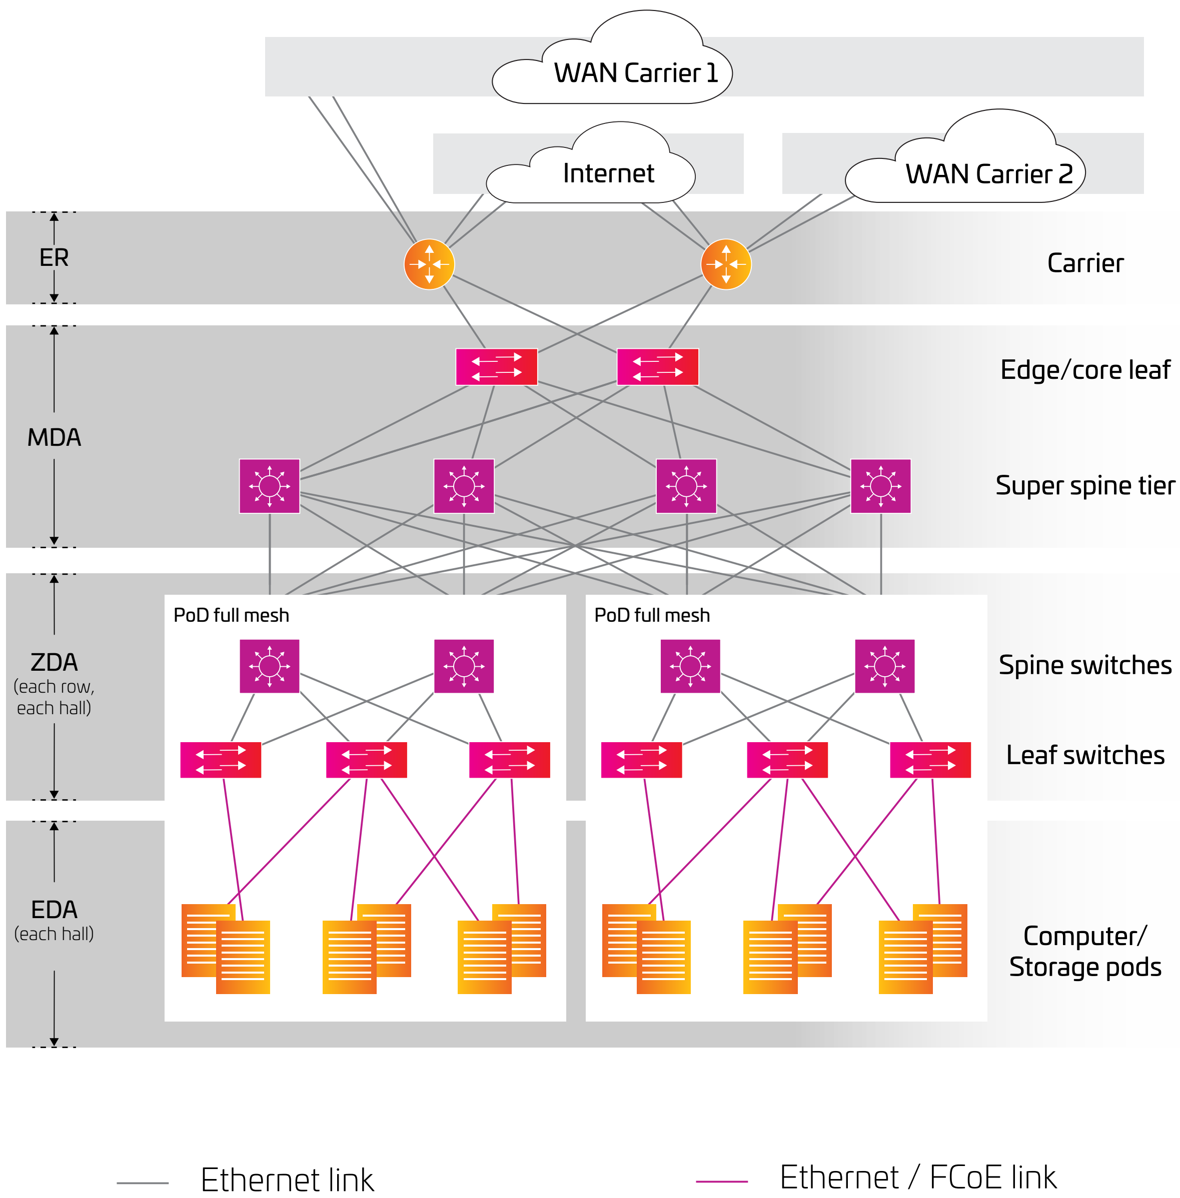

# Data center equipment connection methods

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

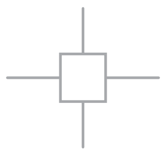

## Cross-connect

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.

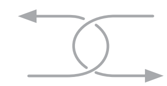

## Interconnect

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.

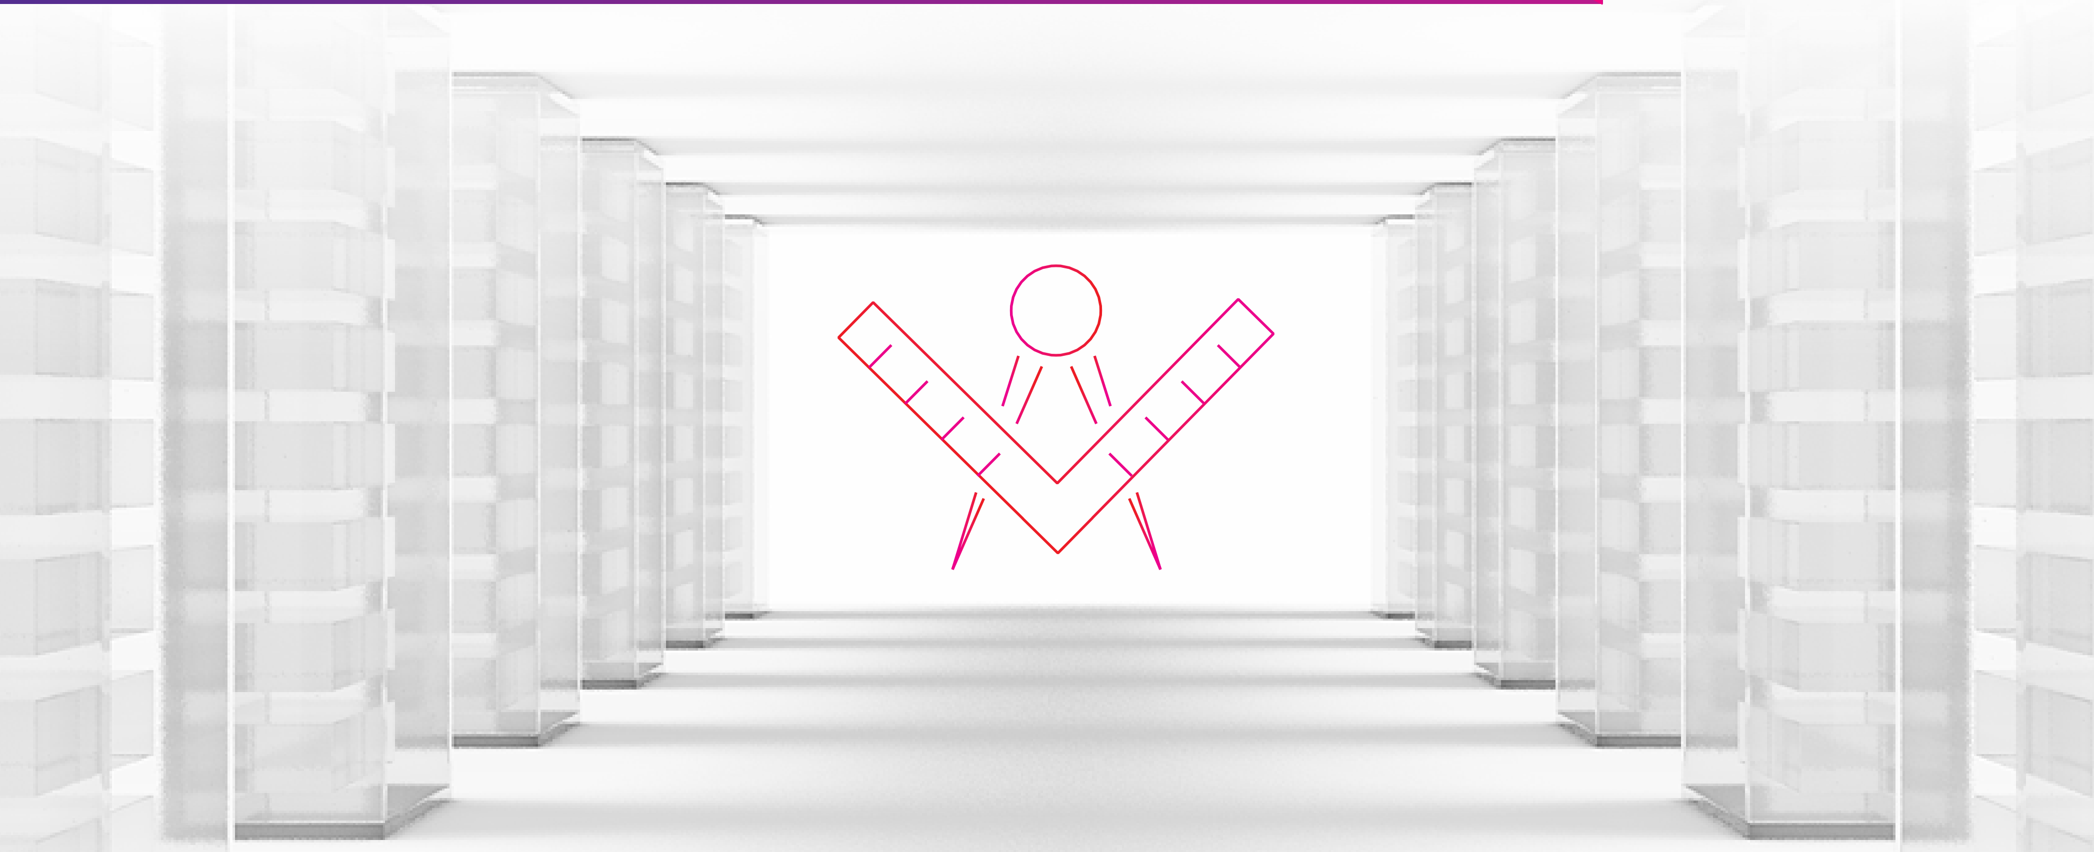

## Architecture for a higher-speed future

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.

# DATA CENTERS

Chapter 4

## High Speed Migration

# High Speed Migration

## Redesigning data center connectivity for a higher-speed future

To support the rapid growth of cloud-based storage and compute services, data centers are adapting their traditional three-layer switching topologies to accommodate highly agile, low-latency performance. These new architectures resemble "warehouse scale" facilities that are designed to support many different enterprise applications.

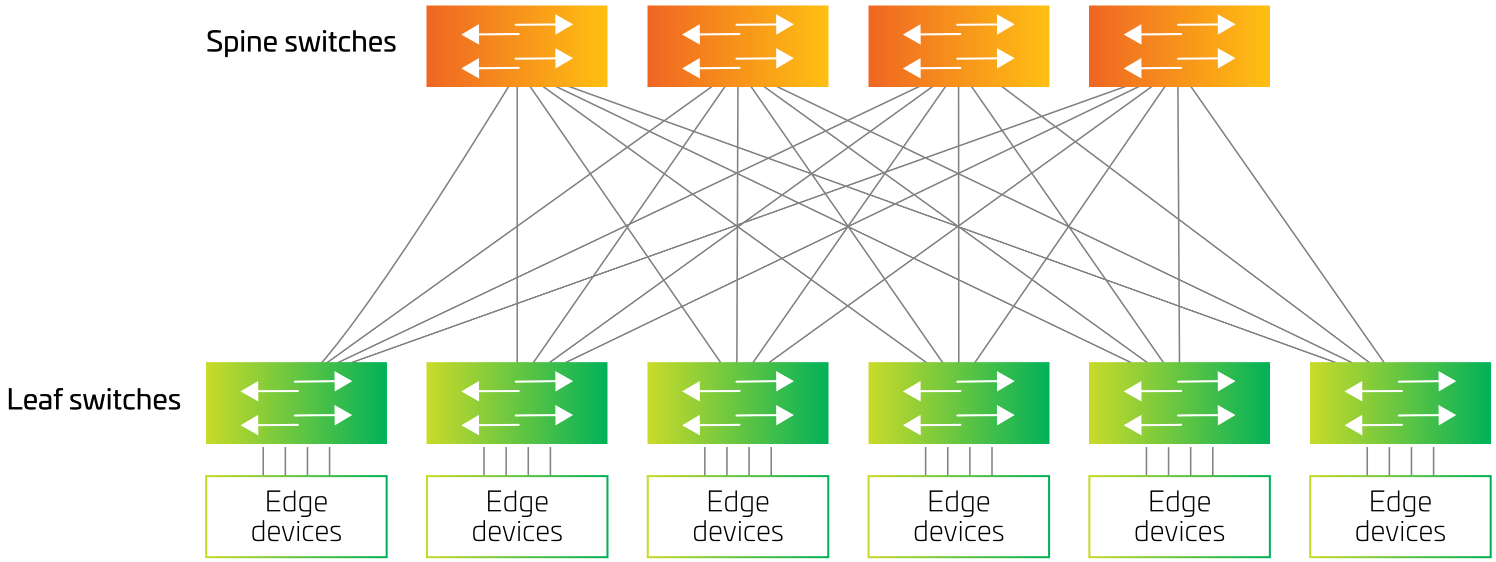


Leaf-spine architectures, for example, create an optimized path for server-to server communication that can accommodate additional nodes, as well as higher line rates, as the network grows. The meshed connections between leaf-and-spine switches allow applications on any compute and storage device to work together in a predictable, scalable way regardless of their physical location within the data center.

# Future-ready fabric network technology

Demand for lower costs and higher capacities in the data center is growing. New fabric network systems that can better support cloud-based compute and storage systems are becoming the architecture of choice. Their ability to deliver any-to-any connectivity with predictable capacity and lower latency makes today's fabric networks a key to enabling universal cloud services.

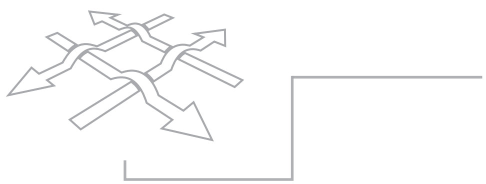

These fabric networks can take many forms: fabric extensions in a top-of-rack deployment, fabric at the horizontal or intermediate distribution area, and fabric in a centralized architecture. In all cases, consideration must be given to how the physical layer infrastructure is designed and implemented to ensure the switch fabric can scale easily and efficiently.

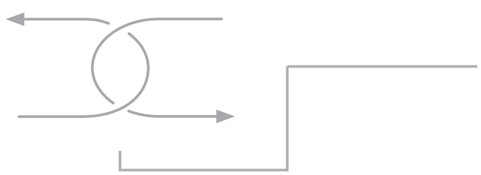

The fabric has inherent redundancy, with multiple switching resources interconnected across the data center to help ensure better application availability. These meshed network designs can be much more cost-effective to deploy and scale when compared to very large, traditional switching platforms.

The design of high-capacity links is more complex since the number of links and link speeds is increasing. Providing more data center capacity means pushing the limits of existing media and communication channel technologies. As shown below, the Ethernet Alliance Roadmap illustrates existing application standards and future application rates beyond one terabit per second. This will further challenge complexity as application speeds move from duplex transmission to parallel transmission. The advent of new technologies—shortwave wavelength division multiplexing (SWDM), OM5 wideband multimode fiber (WBMMF), bi-directional (BiDi) transmission, coarse wavelength division multiplexing (CWDM) and more efficient line coding—is expected to delay the transition to parallel optics.

## ETHERNET ROADMAP

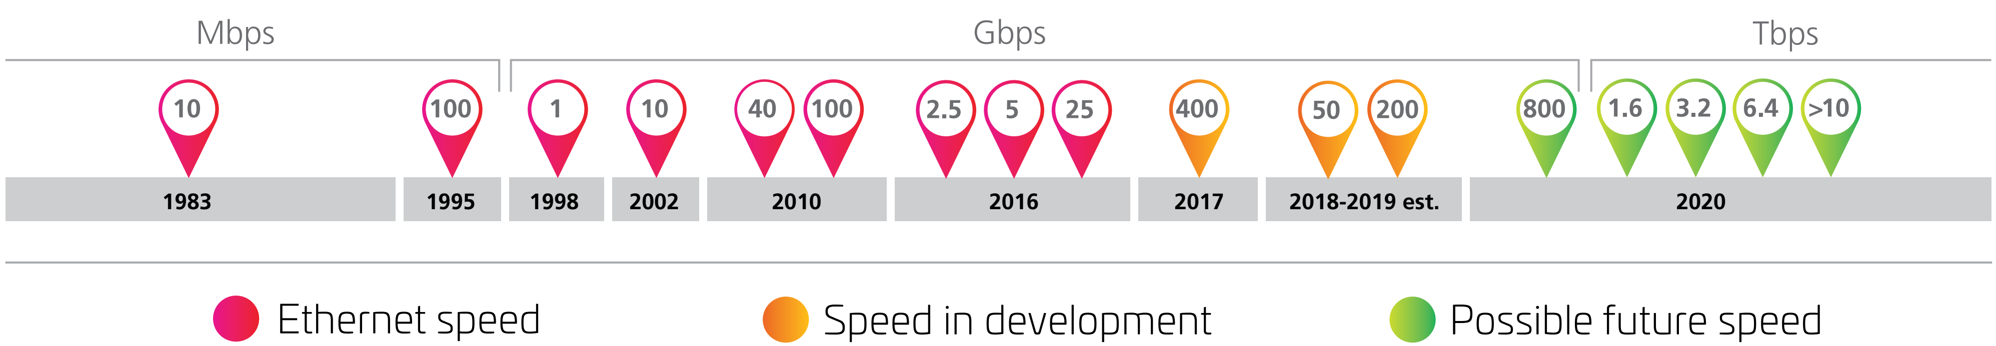

## Tomorrow's possible interfaces

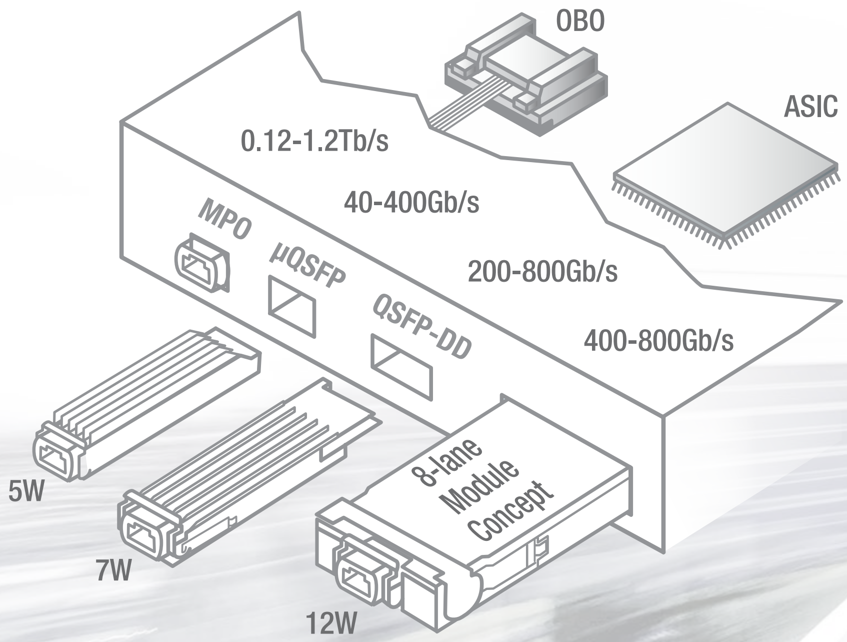


*Image Courtesy of the Ethernet Alliance*

The combination of shortwave wavelength division multiplexing (SWDM) and OM5 wideband multimode fiber (WBMMF) provides the opportunity to extend the use of multimode technology, which continues to be the most prevalent fiber technology deployed in data centers.

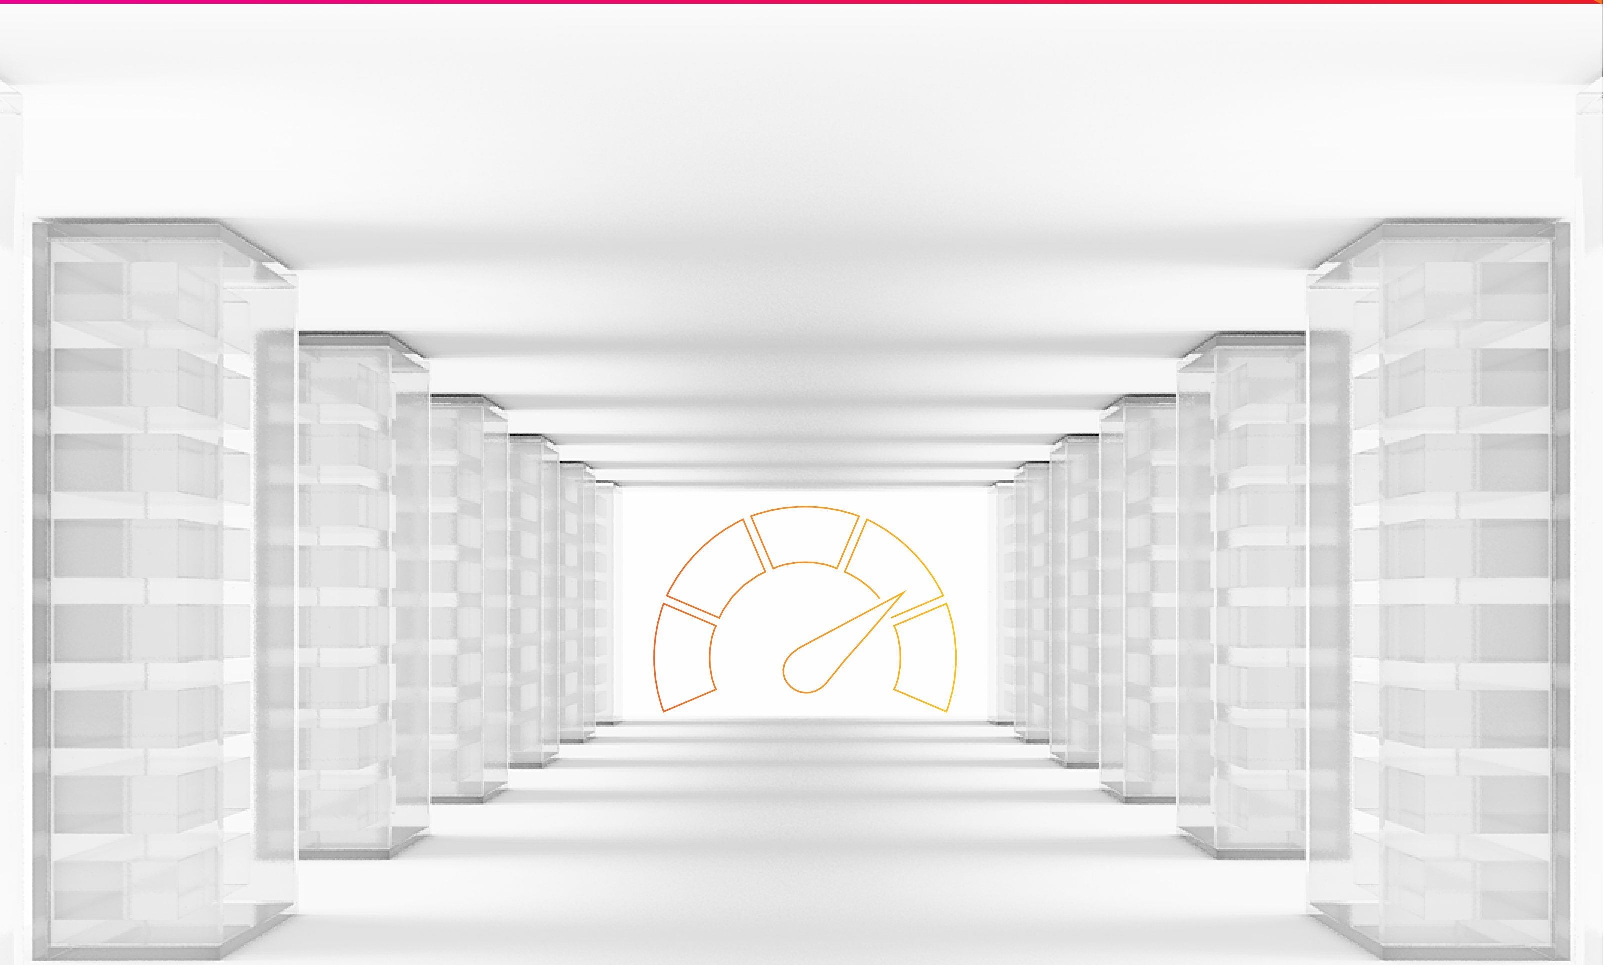

Engineered solutions make complex fabric networks easy to design, implement and manage. Preterminated high-performance systems support next-generation network media, duplex and multifiber modular applications while reducing deployment management time and expense.

**Design guide:**

Fabric networks: Designing your network for the future—from 10G through 400G and beyond

**Brochure:**

High Speed Migration

**Video:**

SYSTIMAX Application Assurance: High-speed application support—guaranteed

**Application note:**

High Speed Migration: Choosing the right multimode multi-fiber push on (MPO) system for your data center

# DATA CENTERS

Chapter 5

## Multisource agreements

# Multisource agreements

## Bringing options to a changing environment

The data center is a complex environment, comprising a wide range of equipment and technology manufactured by many different companies. Ever-increasing bandwidth and line rates have led to optical fiber being the preferred technology to enable higher speeds. To ensure proper operation and maximum efficiency of the data center networks, optical transceivers of the same type must be interchangeable and interoperable so replacements and upgrades can be performed quickly and easily, without the need to replace or modify other network equipment.

The solution is a multisource agreement (MSA)—an agreement among multiple manufacturers to make equipment consistent and interchangeable by defining a common physical form for various devices and components. In the case of data center connectivity, there are MSAs that cover both the specification and implementation of the optical transceivers made by various manufacturers.

The phenomenal growth in data, voice and video has triggered the need for higher and higher speeds in the data center and across data centers. This has driven the standards bodies to develop higher application speeds, which in turn have driven the need for new MSAs. Per the latest version of the Ethernet Roadmap, there are currently seven new applications in progress—most of which involve fiber optics.

There are now many different MSAs reflecting the variety of applications we see in the data center:

10G, 25G, 40G, 50G, 100G, 200G, singlemode, multimode, and others.

## Examples of optical MSAs

The optical transceiver multisource agreement (MSA) environment is very dynamic, with numerous MSAs—too many to list in this publication. These MSAs cover everything from form factor, application (standard, prestandard or proprietary), maximum power consumption, fiber connector type, strand count, wavelength and cable reach. Examples of current and future MSAs are shown below:

| TODAY'S INTERFACES | TOMORROW'S POSSIBLE INTERFACES |
|---|---|

*Images Courtesy of the Ethernet Alliance*

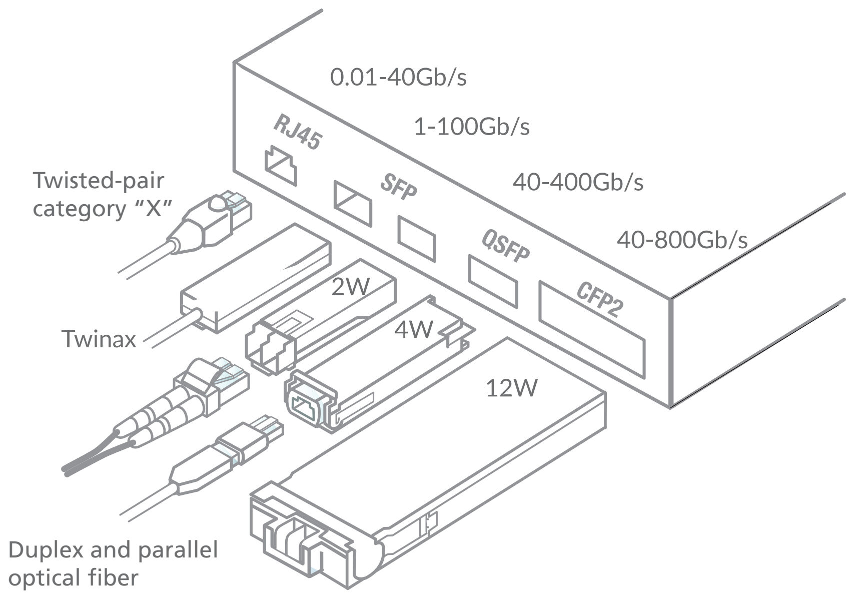


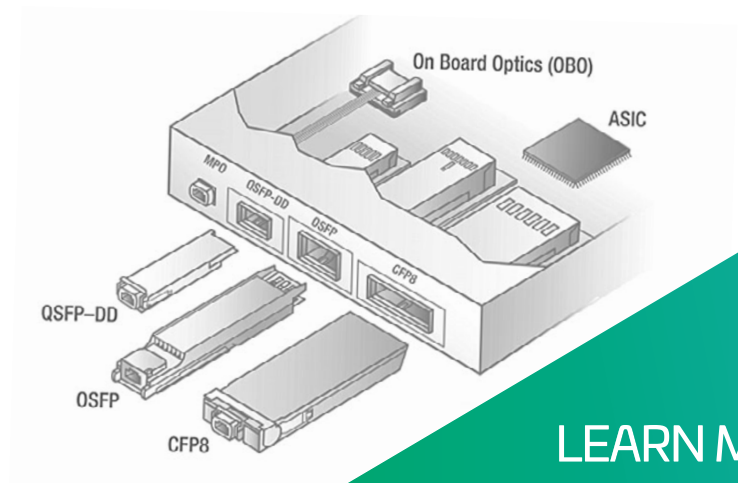


**LEARN MORE** about tomorrow's possible interfaces on the next page >>>

## QSPF-DD

QSFP-DD—"Quad Small Factor Pluggable—Double Density."

The smallest 400 Gbps module will provide backwards compatibility to 40GbE and 100GbE QSFP modules. Will support Ethernet, FibreChannel or InfiniBand protocols.

**Designed to support:**

- 200 Gbps or an aggregate of 400 Gbps, using 25 Gbps NRZ modulation per lane or 50 Gbps PAM4 per lane

## OSPF

OSFP—"Octal Small Factor Pluggable."

The smallest 400 Gbps module will provide backwards compatibility to 40GbE and 100GbE QSFP modules. Will support Ethernet, FibreChannel or InfiniBand protocols.

**Designed to support:**

- 400GBASE-DR4 parallel SMF (4x100G PAM4)
- 400GBASE-SR8 parallel MMF (8x50G PAM4)
- 400GBASE-FR4 duplex SMF (4x100G PAM4 WDM)
- 400GBASE-FR8/LR8 duplex SMF (8x50G PAM4 WDM)
- 2x200GBASE-FR4 parallel SMF
- 2x100GBASE-CWDM4 parallel SMF

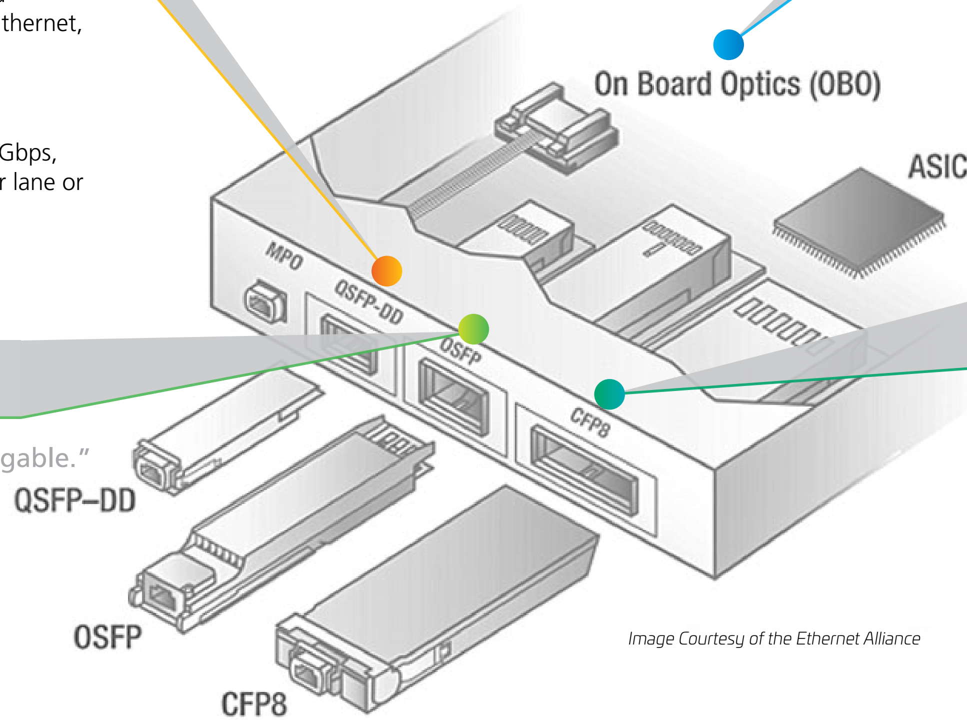


*Image Courtesy of the Ethernet Alliance*

## OBO

OBO, or Consortium for On-Board Optics (COBO)

Eliminates the E/O function traditionally performed by transceivers, meaning the bandwidth density at the faceplate can be dramatically increased. Data applications to be supported by OBO are currently to be defined, but this technology is primarily targeted at data rates from 400 Gbps to more than 800 Gbps.

## CFP8

CFP8—"C Form Factor Pluggable"

Primarily aimed at supporting 400 Gbps with a claim to offer a path to support 800 Gbps in the future.

**Designed to support:**

- 400GBASE-SR16 parallel MMF (16x25G NRZ)
- 400GBASE-FR8/LR8 duplex SMF (8x50G PAM4 WDM)
- 400GBASE-DR4 parallel SMF (4x100G PAM4)

A networking technology may come to market with multiple choices or generations of optical transceivers. The market will eventually identify the winning solution based on cost, size, power consumption, vendor support and other factors.

# IMPLICATIONS FOR FIBER CABLING INFRASTRUCTURE DESIGN

The clear trend in the development of new multisource agreements (MSAs) has been toward both higher speeds and increased densities. Higher speeds are the result of new applications standards that specify higher line rates. Higher densities have been driven largely by technology advances that enable the transceiver to make use of lower power, which allows for smaller packaging. As shown, the physically larger MSAs are designed to accommodate higher power transceivers, while reduced power transceivers can make use of smaller MSAs for more ports or higher density communication hardware.

Fortunately, each of the data center cabling standards (TIA 942, ISO/IEC 11801-5 and CENELEC 50173-5) has standardized on two optical connectors for use in the data center: the LC for single or duplex applications and the MPO for applications requiring more than two fibers. This has simplified the fiber connectivity as the MSAs that are relevant in the data center environment also have made use of the LC and MPO connectors. And, while the standardization of connectors has helped simplify cabling, it has also become very important to provide very flexible, agile connectivity that can accommodate the ever-increasing speeds and the higher densities that are being driven by higher densities at the equipment faceplate.

Standards:

Multisource agreements

# DATA CENTERS

Chapter 6

## Fiber selection

# Fiber selection

## Vital connections for today's data centers

The data center is at the core of today's business—and fiber-optic connectivity is the fabric, carrying vital data to drive critical business processes and providing connectivity to link servers, switches and storage systems.

Data center designers have two high-level choices when it comes to fiber types: multimode fiber and singlemode fiber. In this chapter, we'll discuss the development, deployment and advantages of each fiber type, as well as the connectors that pull it all together.

## Multimode fiber: the low-cost platform

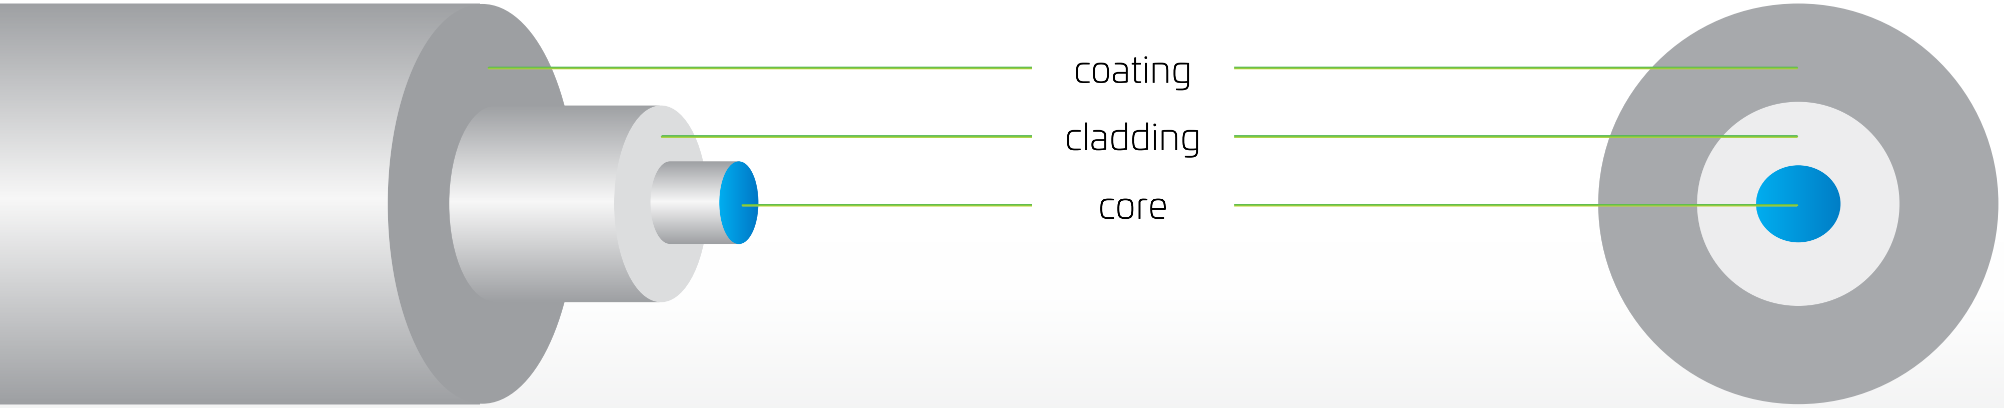

Today, multimode fiber (MMF) is the workhorse medium for data centers because it is the lowest-cost way to transport data at high rates over the relatively short distances in these environments. MMF has evolved from being optimized for multimegabit-per-second transmission using light-emitting diode (LED) light sources to being optimized to support multigigabit transmission using 850 nm vertical cavity surface emitting laser (VCSEL) sources, which tend to be less expensive than their singlemode counterparts.

This leap in performance is reflected in the classifications given by the standards bodies. OM1 and OM2 represented the earlier MMF types with low modal bandwidth and very limited support for higher-speed optics. OM3 and OM4 represent the newer, laser-optimized MMFs typically installed in data centers today. The following table provides examples of some of the current data center applications and the maximum channel lengths over different fiber types.

| Application | Standard, MSA or Mfg. Specification | IEEE reference | Media | Speed | Target distance |
|---|---|---|---|---|---|
| 40 Gigabit Ethernet | 40GBASE-SR4 | IEEE 802.3ba | 4-pair MMF | 40 Gb/s | 150 m (OM4/OM5) |
| | 40GBASE-eSR4 | - | 4-pair MMF | | 400 m (OM4/OM5) |
| | 40G-BiDi | - | 1-pair MMF | | 150 m (OM4) 200 m (OM5) |
| | 40G-SWDM | - | 1-pair MMF | | 350 m (OM4) 440 m (OM5) |
| | 40GBASE-FR | IEEE 802.3bg | 1-pair SMF | | 2 km |
| | 40GBASE-LR4 | IEEE 802.3ba | 1-pair SMF | | 10 km |
| | 40GBASE-ER4 | IEEE 802.3ba | 1-pair SMF | | 40 km |
| 100 Gigabit Ethernet | 100GBASE-SR4 | IEEE 802.3bm | 4-pair MMF | 100 Gb/s | 100 m (OM4/OM5) |
| | 100GBASE-eSR4 | - | 4-pair MMF | | 300 m (OM4/OM5) |
| | 100GBASE-SR10 | IEEE 802.3ba | 10-pair MMF | | 150 m (OM4/OM5) |
| | 100G-SWDM4 | - | 1-pair MMF | | 100 m (OM4) 150 m (OM5) |
| | 100G-eSWDM4 | - | 1-pair MMF | | 300 m (OM4) 400 m (OM5) |
| | 100G-PSM4 | - | 4-pair SMF | | 500 m |
| | 100G-CWDM4 Lite | - | 1-pair SMF | | 500 m |
| | 100G-CWDM4 | | 1-pair SMF | | 2 Km |
| | 100GBASE-LR4 | IEEE 802.3ba | 1-pair SMF | | 10 km |
| | 100GBASE-ER4 | IEEE 802.3ba | 1-pair SMF | | 40 km |

# Introducing wideband multimode fiber (WBMMF)

OM3 and OM4 provide very high modal bandwidth at 850 nm—the predominant wavelength that can be efficiently supported by VCSEL transmitters. To support an increase in performance over a single pair of multimode fibers, additional wavelengths need to be transmitted alongside 850 nm, achieved via a new technology—shortwave wavelength division multiplexing (SWDM). Because the modal bandwidth of OM3 and OM4 fibers was specified for laser operation at 850 nm only, a new specification for optical fiber was required. Many data center managers are now considering wideband multimode fiber (WBMMF), which optimizes the reach of SWDM transmission that delivers four times more information with the same number of fiber strands over practical distances. Being optimized to support the additional wavelengths required for SWDM operation (in the 850 nm to 950 nm range), WBMMF ensures not only more efficient support for future applications across the data center fabric, but also full compatibility with legacy applications because it remains fully compliant with OM4 specifications.

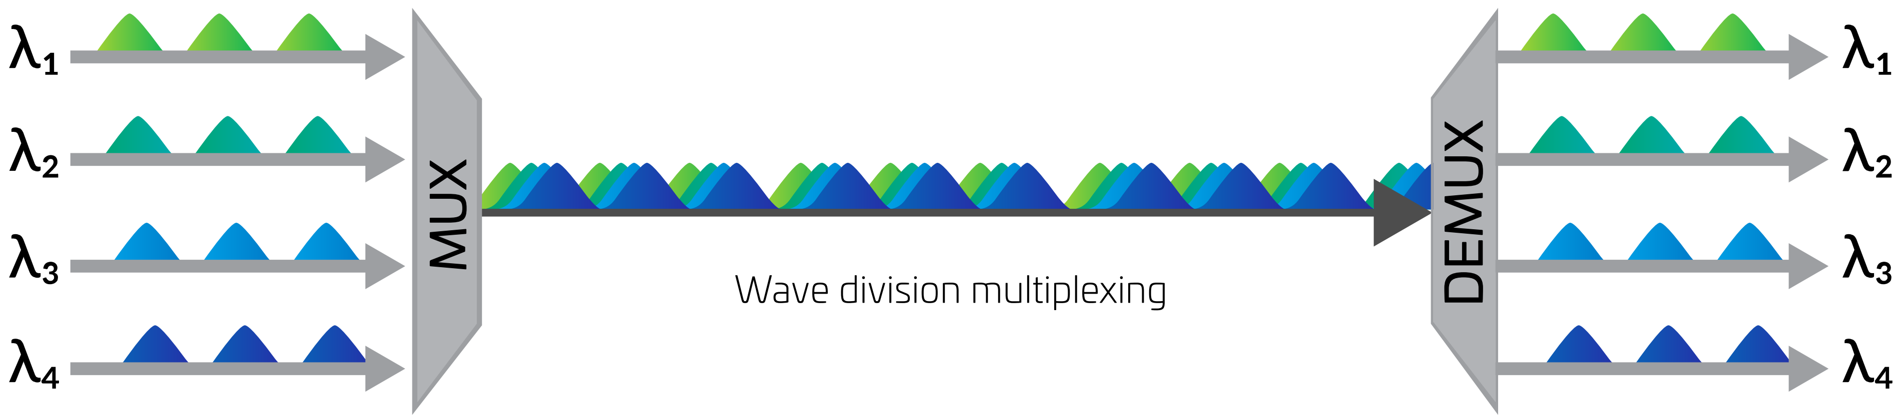

By the middle of 2017, the journey to standardization of WBMMF cabling was complete, having been recognized by ISO/IEC and TIA standard bodies. The OM5 designation was adopted for inclusion of this new cabled optical fiber category in the third edition of the ISO/IEC 11801 standard. Once again, CommScope led the market in next-generation standards development as well as product availability and was one of the first manufacturers to deliver a commercially available OM5 end-to-end solution with the distinctive lime green color that is also being recognized by standards bodies. Well ahead of standards ratification, CommScope introduced the LazrSPEED® OM5 Wideband solution in 2016, knowing that the support of higher data throughput using low-cost optics is exactly what data center managers require to enable next-generation networks today and in the future.

Indeed, the future of OM5 is very bright. At the end of 2017, the IEEE agreed to initiate a project to define next-generation multimode transmission using shortwave division multiplexing—the transmission technology OM5 was designed to support.

## GOALS AND BENEFITS

- Retain legacy application support of OM4
- Increase capacity to > 100 Gbps per fiber
- Reduce fiber count by four
- Boost array cabling capacity for parallel applications
- Enable Ethernet
  40G-SR, 100G-SR, 200G-SR, 400G-SR4
- Enable Fiber Channel:
  128G-SWDM, 256GFC-SWDM
- Extend MMF utility
  as universal data communication medium

## Singlemode fiber: supporting longer distances

Designed with a much narrower core, singlemode fiber (SMF) is the technology of choice for longer-reach applications in the data center, such as extended runs in the fabric between leaf-and-spine switches, spine and routers, and into the transport network to connect data centers in different locations. SMF provides higher bandwidth and does not have the modal dispersion limitations inherent in MMF. For this reason, SMF is used in applications where support for higher and next-generation bandwidths can be absolutely guaranteed to be supported. This makes it a perfect medium of choice for hyperscale and service provider data center owners to deploy.

Very large data centers as well as hyperscale data centers typically deploy SMF to connect multiple halls and extended equipment zones using a centralized cross-connects architecture at the MDA. They typically use a dedicated optical distribution frame (ODF). Deploying an ODF can help ensure that cables are kept to an optimum length for transmission, while equipment zones and other data halls can be quickly and efficiently patched to one another with minimal disruption to service and networking equipment.

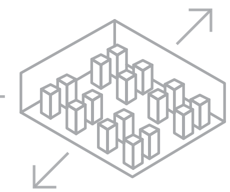

Singlemode fiber also enables duplex transmission at higher speeds because it is able to transport multiple wavelengths, thus reducing fiber counts. It is anticipated that one of the 200GE and 400 GbE applications will utilize four-pair parallel optics over SMF—taking advantage of the lower overall system cost parallel optics can offer. The PSM4 multisource agreement (MSA) also defines a four-pair transceiver for 100G applications.

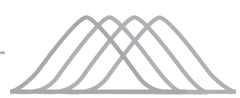

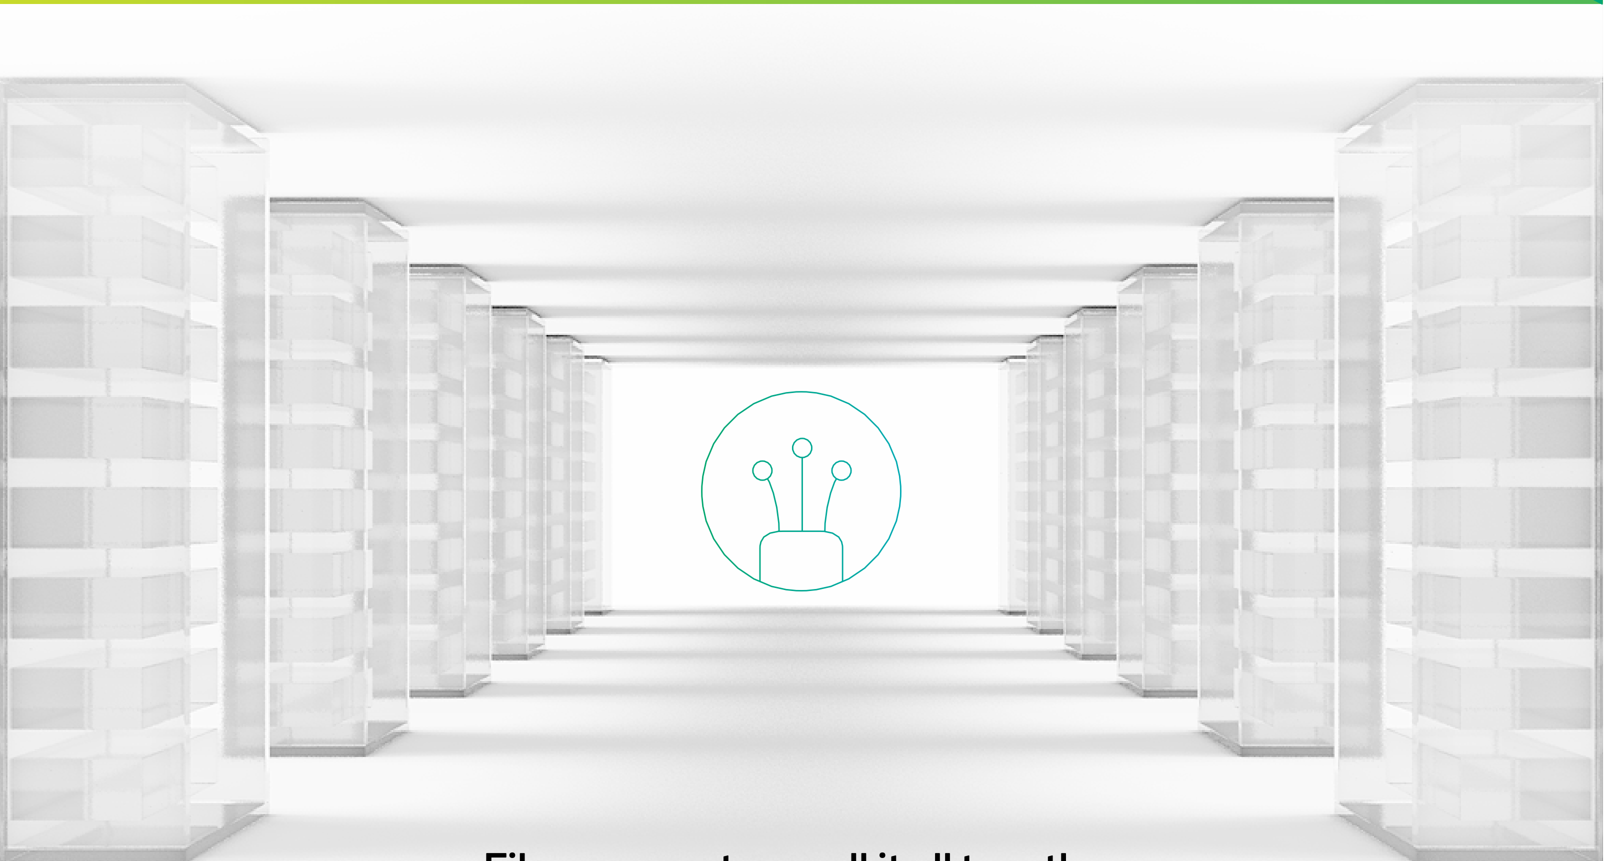

## Fiber connectors pull it all together

Fiber connectors have evolved along with fiber-optic cabling, driven by increasing fiber density. The duplex LC connector emerged during the early 2000s as the predominant two-fiber type—and remains so today. While the evolution of the duplex connector was underway, array connectors (parallel optics) were also emerging. First deployed in public networks, the multifiber push-on (MPO) connector has become a preferred choice for rapidly deploying cabling into data centers. The compact form of the MPO allows 12 or more fibers to be terminated in a compact plug, occupying the same space as a duplex LC. The MPO's high density enables installation of preterminated, high-strand-count cables, while eliminating the time-consuming process of field installing connectors on site.

**White paper:**

Wideband multimode fiber—
What is it and why does it make sense?

**Design guide:**

Fabric networks:
Designing your network for the future—
from 10G through 400G and beyond

# DATA CENTERS

Chapter 7

## Optical distribution frames

# Optical distribution frames

## Take control of data center cabling for optimal performance

The ever-growing demand for more bandwidth to accommodate a wide array of new applications in the data center is driving higher fiber counts, as more and more data centers are being designed and built to run high-speed applications for LAN/SAN. These high-speed applications are based on fiber-optic transmission, making fiber-optic cabling the predominant transmission medium in the data center now and into the future.

With more and more optical connections to contend with, the challenge becomes how to add optical density to the fiber frame while still maintaining proper accessibility, flexibility and manageability at the lowest possible cost. As data center operators add more fiber-optic cabling, they often face an out-of-control situation in terms of fiber count, density and space—resulting in potentially reduced availability and higher cost of operation.

The attempt to address these issues by using high-density patch panels can make the problem worse, if not done correctly. Trying to fit high-density cabling into cabinets that are designed to house active equipment can result in a tangled "spaghetti bowl" of cabling—especially in configurations where cable management is essentially non-existent.

The solution to the problem has two parts:

## Choosing a different cabling architecture

A centralized cross-connect configuration in the main distribution area (MDA) eliminates patching from the core equipment cabinets. All active core ports from LAN/SAN are mirrored in central cross-connect cabinets—resulting in safer operation and simplified design for future growth.

## Having the right cross-connect solution

However, they often provide limited cable management for the patch cords connecting to the active equipment. This scenario may be adequate for equipment connectivity, but does not provide cable management needed for cross-connects. A best-in-class cross-connect solution consists of frames or cabinets that have been designed around the fiber patch panels along with providing patch cord management to accommodate the quantity and types used today and those that will be used in the future. Only when all of this is considered can the data center designer design the cabling infrastructure properly.

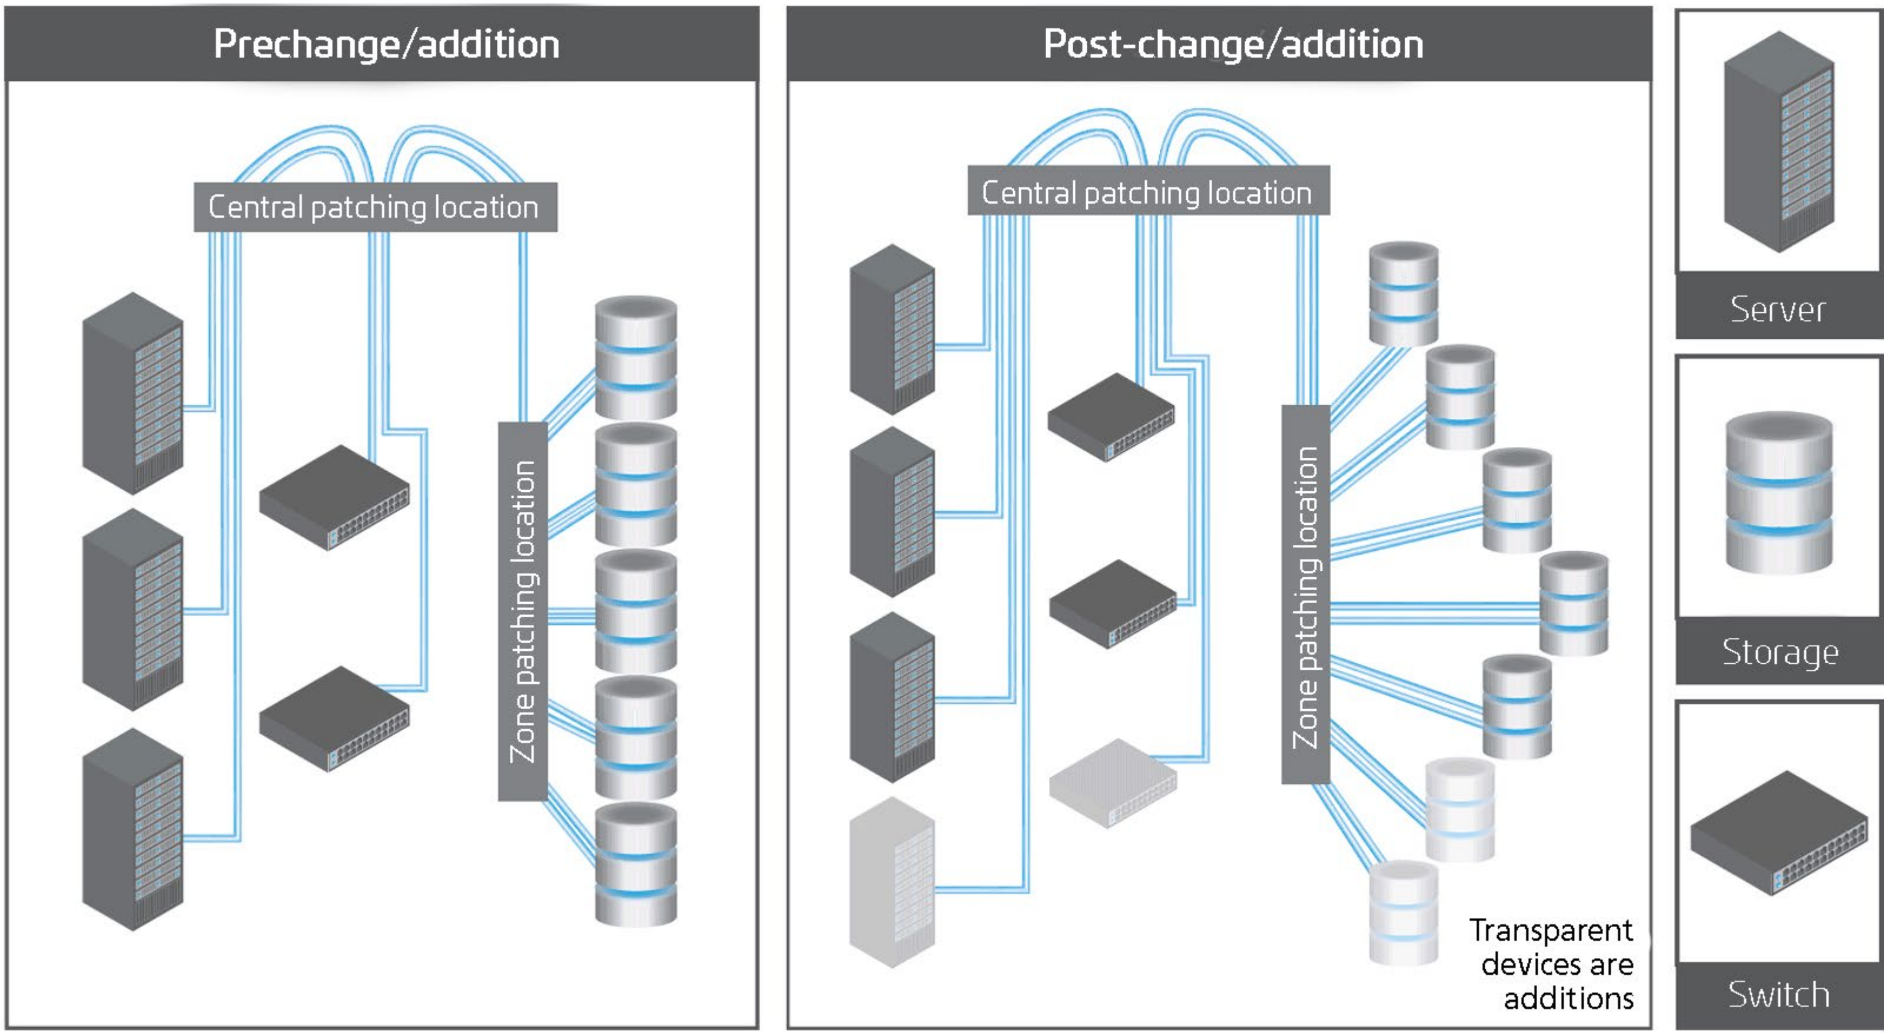


To meet these challenges today and equip their facilities for future growth, data centers must be designed with optical distribution frames (ODFs) functioning as cross-connects in the main distribution area (MDA).

# A proven telecom solution comes to the data center

Optical distribution frames (ODFs) have been available for years, used primarily in telecommunication providers' central offices, where tens of thousands of optical fibers converge at a single location. With similar challenges now facing data center operators, using ODFs to manage data center cabling has become an effective option.

Since cross-connect ODFs are optimized for cabling, not for equipment, they are able to solve the two largest data center cable management problems: those caused by application migration toward parallel fiber-optic applications, and those caused by the expected growth of the data center itself.

Both of these trends require deploying much more fiber in the data center, resulting in massive patch cord changes in both number and size. ODFs can easily deal with these challenges because they are optimized for cable management, offering bend radius protection for fiber patch cords and over-length storage for efficient use of the ODF—even with thousands of patch cords in it.

Correctly designed, cross-connect ODFs function very effectively as the single point of distribution for all LAN, SAN and telecommunication services in the data center, delivering best-in-class cable management and reduced operations costs, with these advantages:

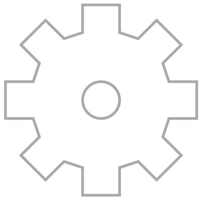

### Easy servicing

Precabled ODFs allow fast moves, adds and changes

### Increased availability

Cabling can be added or changed without disrupting running systems

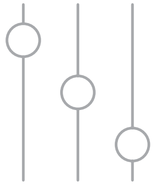

### Optimal flexibility

Equipment can be connected regardless of its location

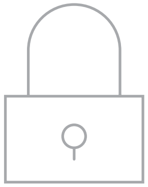

### Enhanced security

Requires no direct patching at the switch/SAN director

# DATA CENTERS

Chapter 8

## Automated infrastructure management

# Automated infrastructure management

## The data center connectivity challenge

In today's fast-evolving data centers, expanding the fiber-optic infrastructure is vital for providing the bandwidth and speed needed to transmit large amounts of data to and from multiple sources. As switches with 40G and 100G ports become commonplace, data center infrastructure becomes more complex—and it is becoming increasingly clear that traditional, manual methods for managing fiber connectivity may not be sufficient.

Demand for fast data transmission and efficient network performance has been fueled by requirements to support virtualization, convergence and cloud computing, as well as high-bandwidth applications like streaming video. But, while supporting more bandwidth is important, there are additional trends impacting data center fiber infrastructure management:

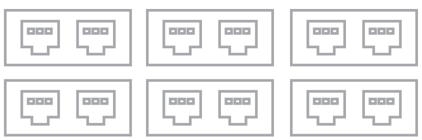

**Higher port density**
Space is at a premium in the data center, which has led to higher densities of fiber ports on equipment and fiber shelves. With increased density comes the increased risk of making or removing the wrong connection— potentially causing widespread disruption in network services.

**Increased complexity of cabling topology**
Point-to-multipoint connections have become commonplace with the advent of 40G and 100G technology—making manual record-keeping nearly impossible.

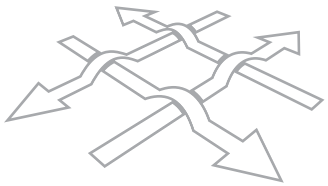

**Increased complexity of network architecture**
The move to heavily-meshed leaf-spine network architecture has greatly increased the number of connections along with a need for the consistent and accurate deployment of connectivity pattern/mesh for orderly network expansions in the data center.

## Taking AIM at data center downtime

In its simplest terms, automated infrastructure management (AIM) is an integrated hardware and software platform that manages the entire physical layer. It fully documents the cabling infrastructure, including connected equipment, to provide a complete view of where devices are located and how they are connected.

By capturing information about every physical connection in the network and relaying it to higher-level network management systems, the AIM system provides an accurate, real-time view of the physical network connectivity and can issue alarms when an unplanned or unauthorized change occurs. AIM streamlines the provisioning and monitoring of data center connectivity; produces up-to-date reports on the status and capacity of the network infrastructure; and ultimately can reduce data center downtime and mean time to repair through real-time, precision notification of connectivity outages.

AIM systems also improve other aspects of data center operations, including:

### Capacity management

AIM provides an accurate view of available panel, switch and server ports, which helps address network capacity challenges by eliminating dormant ports—resulting in more efficient planning and reduced CapEx.

### Troubleshooting

The precise location of a connectivity problem is documented so the technician doesn't have to spend time verifying manual documentation or hunting for the location of a problem.

### Security

AIM can raise alerts when a port is disconnected or connected in an unauthorized location—for example, if someone has moved a server without following approved change management processes.

### Automated workflow

AIM reduces time-intensive manual processes by generating electronic work orders and enabling guided administration of connectivity changes. This helps minimize human errors and unplanned network downtime.

# Choosing an AIM system that exceeds the standard


The IT industry has recognized the important role intelligent infrastructure solutions can play in data center management and has established standards for automated infrastructure management (AIM) capabilities and functions.

In early 2017, the ISO/IEC WG3 SC25 group is expected to publish the ISO/IEC 18598 Standard for Automated Infrastructure Management Systems—Requirements, Data Exchange and Application.

## To meet the key requirements of ISO/IEC 18598, an AIM solution must:

- Automatically detect and monitor connectivity
- Automatically detect, document and monitor the presence of network devices
- Automatically update records when any monitored connections are modified
- Document connectivity between non-AIM enabled ports and other equipment
- Identify and track the physical location of end devices connected to the network
- Generate a graphic representation of end-to-end connectivity (circuit trace)
- Generate electronic work orders and automatically monitor the accuracy of implementation of work order tasks

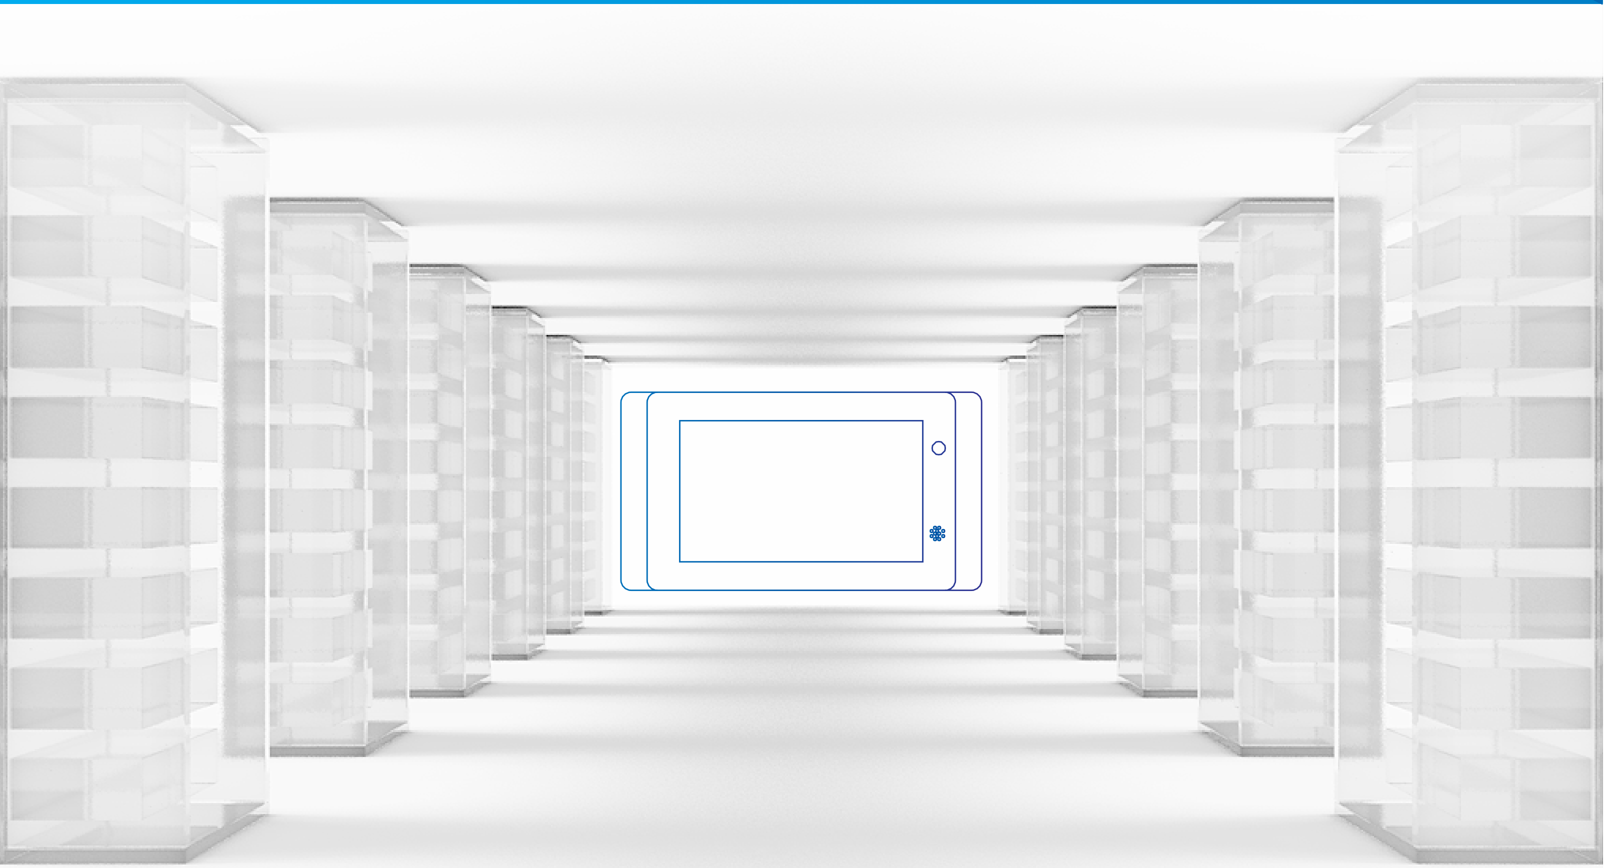

The emergence of AIM standards
is making intelligent connectivity
a mission-critical technology for data centers.
Since cabling infrastructure migration often requires
replacement of fiber-optic modules, now is the time to
upgrade to an AIM-driven intelligent connectivity system.

## RESOURCES

**Standards:**

ISO/IEC AIM Document
(18598/DIS draft)

**Brochure:**

imVision®
automated infrastructure
management

**Video:**

Managing critical
data center fiber
connectivity with imVision

**Video:**

imVision.
Infrastructure management.
Made easy.

# DATA CENTERS

Chapter 9

## Designing for fiber TAPs

## Real-time network monitoring with no service interruptions

The need for real-time network traffic monitoring in today's data center has become compelling. Data center network administrators need to gain better visibility of their networks to optimize the performance of mission-critical applications and keep their networks secure.

In fiber-optic data center networks, a traffic access point (TAP) is a critical tool for data center monitoring and management. A TAP module can be integrated into the fiber cabling infrastructure to enable network traffic monitoring from the physical layer (layer 1) and above in real time—without interrupting network service.

A TAP module is a compact package of fiber-optic couplers or splitters that passively diverts a fixed percentage of light energy away from main transportation channels to monitor the traffic status or content without disrupting the main channel traffic. The optical couplers or splitters inside a TAP module split the light energy from the input port into two output ports according to a designed split percentage—usually diverting from 10 to 50 percent to the TAP.

Because TAPs continuously pass all traffic running between the endpoint network devices with zero latency—while duplicating that exact same traffic to the monitor ports simultaneously—they are one of the most efficient ways to monitor traffic and network link quality in data center networks.

TAP modules help improve managers' understanding of how applications perform and how to measure their performance, and ensure that it meets the required standard. They are also being used to meet compliance or legal requirements that require a business to deploy reasonable tools to secure the data center network.

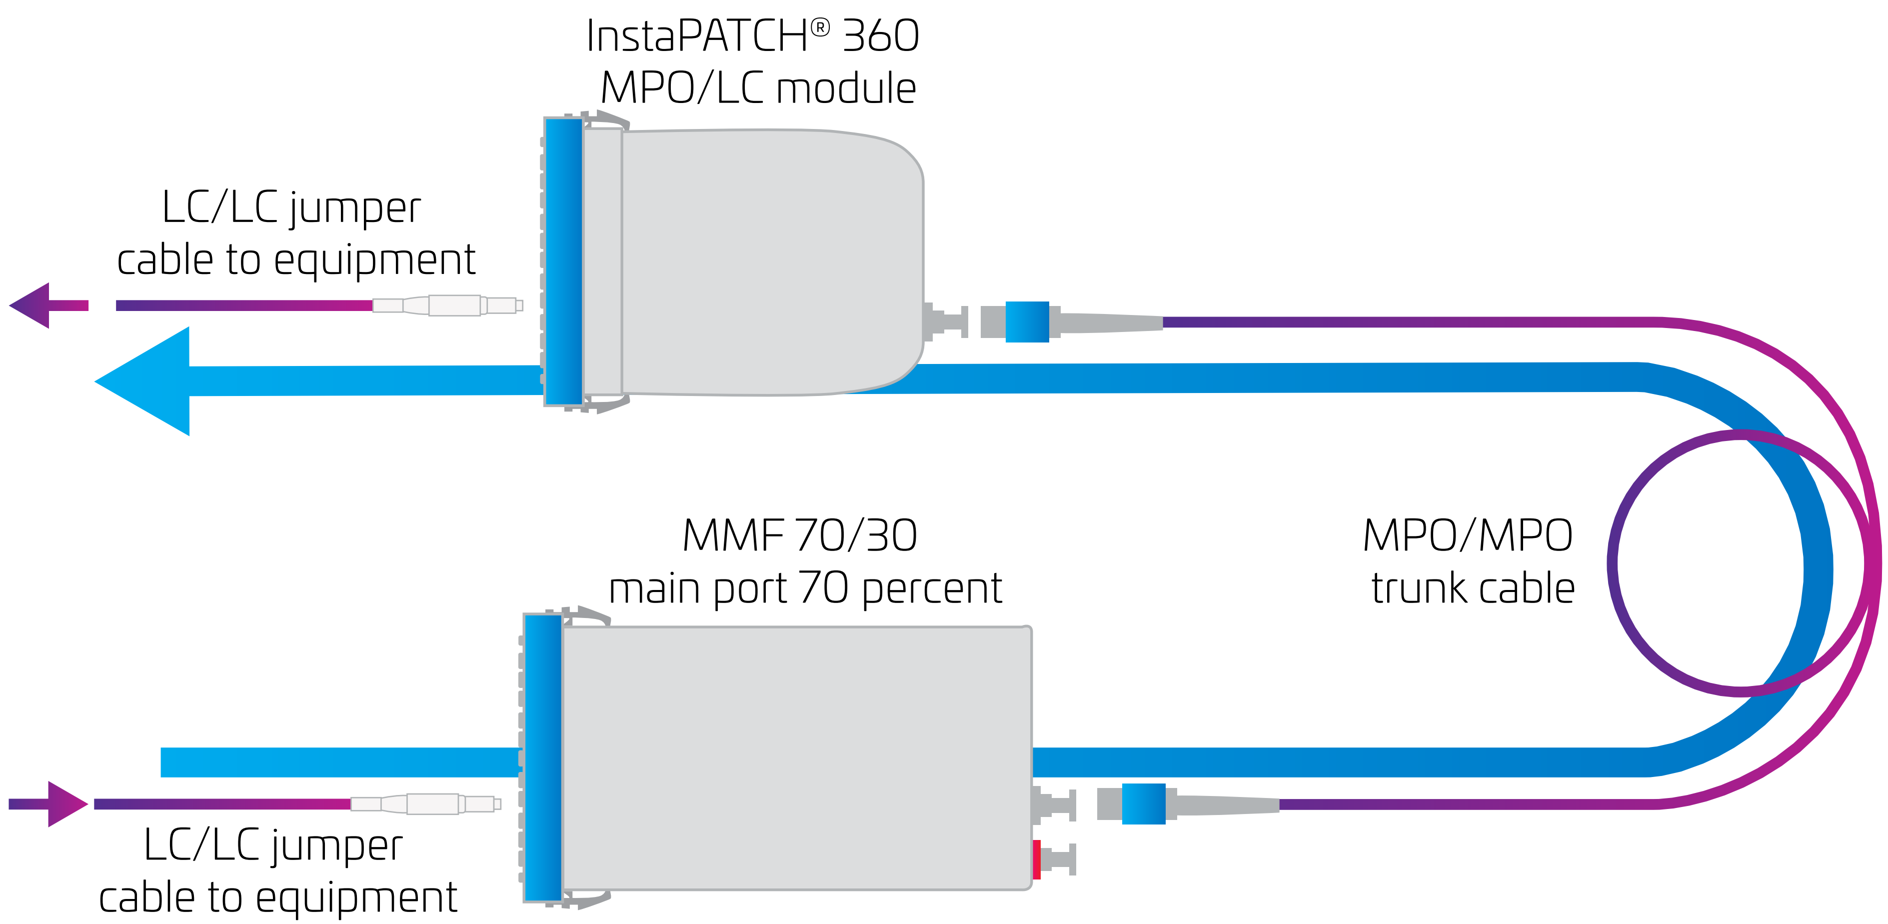


Fiber TAP with 70/30 split

## Designing a TAP solution to mitigate insertion loss

By diverting network traffic for monitoring, traffic access points (TAPs) can introduce additional insertion loss into the network. While industry standards for Ethernet and Fibre Channel are not expressly designed to support the added loss of TAPs, with pre-engineering and the use of high-performance cabling systems it is possible to deploy TAPs and retain useful channel topologies.

As shown below, the evolution of higher-speed applications includes reduced loss budgets—underscoring the need for low-loss components and engineering guidelines.

| db link loss for transmission | | | | |
|---|---|---|---|---|
| Year | Application | Data rate | Standard | Loss budget (dB) |
| 1982 | Ethernet | 10 Mbps | IEEE 802.3 | 12.5 |
| 1991 | Fast Ethernet | 100 Mbps | IEEE 802.3 | 11.0 |
| 1998 | Short wavelength fast Ethernet | 10/100 Mbps | TIA/EIA-785 | 4.0 |
| 2000 | 1G Ethernet | 1,000 Mbps | IEEE 802.3z | 3.56 |
| 2004 | 8 FC and 10 G Ethernet | 10,000 Mbps | IEEE 802.3ae | 2.60 |
| 2010 | 16 GFC and 40 G Ethernet | 40,000 Mbps | IEEE 802.3ba | 1.9 |
| 2010 | 100 G Ethernet | 100,000 Mbps | IEEE 802.3ba | 1.5 |
| 2015 | 32 GFC | 28.800 Mbps | INCITS BSR 512-2015 | 1.86 OM4 |

## When designing a traffic access point (TAP) solution for a particular application, many factors need to be taken into consideration, including:

1 Loss created by the selected TAP splitter

2 Length and number of connections within the main and two monitor channels

3 Intended application (for example, 8G Fibre Channel or 10G Ethernet)

1 Loss created by the selected TAP splitter

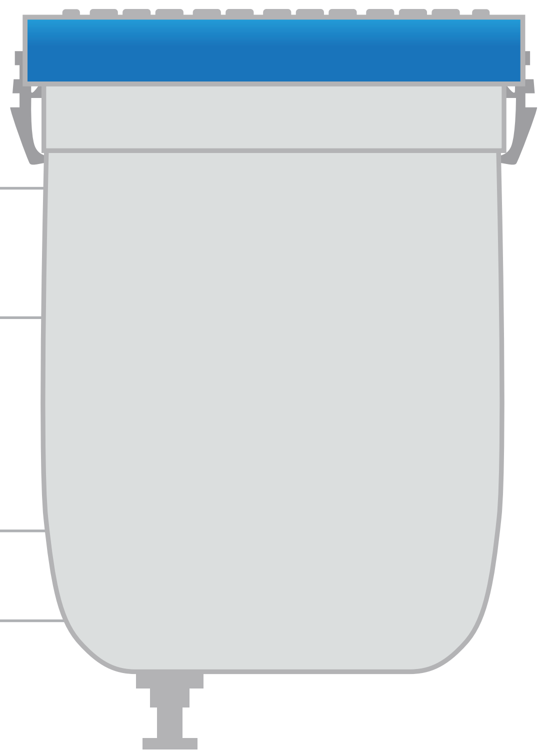

# 9 | Designing for fiber TAPs

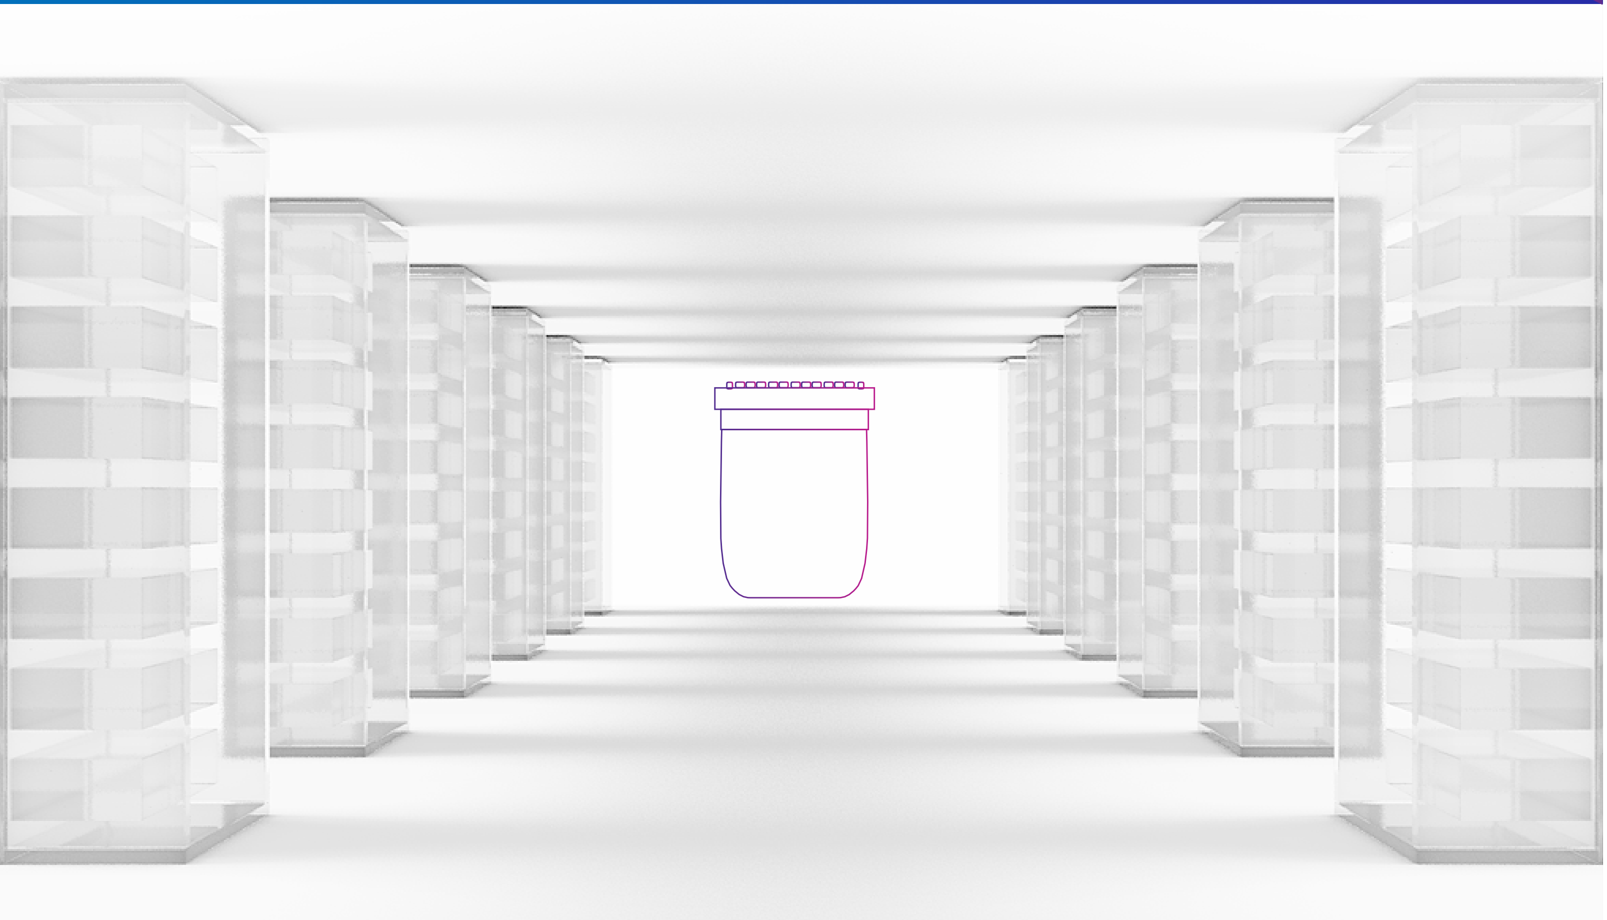

Using TAPs in high-speed fiber links can be complicated—
especially in a do-it-yourself retrofit application.
Instead of trial and error, today's best practice is to design and deploy an engineered solution
in the data center. Designing TAPs into the data center from the start enables the addition of monitoring
capability when it is needed in the future, while proving the operational links to be reliable and solid on day one.

## RESOURCES

**Design guide:**

SYSTIMAX® InstaPATCH® 360 traffic access point (TAP) solution design guide

**Calculator:**

Fiber performance (link loss) calculator

# DATA CENTERS

## Conclusions and authors

# Conclusions and authors

## A road map and resources for building an efficient data center

Even as the data center continues to be the lifeblood of an enterprise, it is undergoing tremendous change as cloud and multitenant options challenge the traditional model. To continue to effectively enable business applications and deliver a wide range of services to an enterprise's customers and employees, data centers must evolve and adapt to today's dynamic business and technology environment.

In this eBook, CommScope experts have explored many of the best practices that are the foundation of the data center model. Beyond the sheer diversity of applications, systems and technologies involved, data centers are also in a constant state of change—bringing new benefits, savings and potential to life.

Because of the number of solutions required and the agility with which they must be deployed, data center networks all over the world run on CommScope. As an industry leader with decades of expertise and ongoing innovation, CommScope designs and builds the solutions that power the data center in all its forms—always with an eye on collaborative development, competitive cost structures and the relentless growth in demand for capacity.

We invite you to contact a CommScope representative to see how we can help you make sure your data center is equipped for what's ahead.

Contact a CommScope expert today.
>>>

# Conclusions and authors

Return back to full Data Center eBook

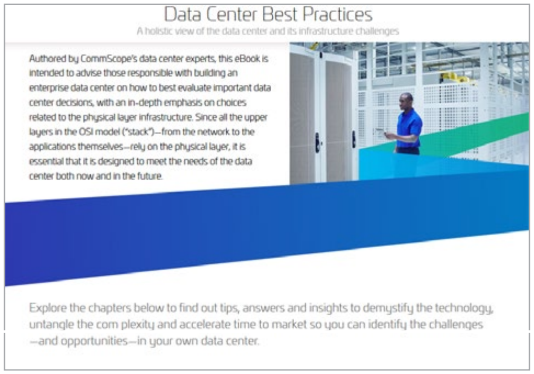
Data Center Best Practices

A holistic view of the data center and its infrastructure challenges

Authored by CommScope's data center experts, this eBook is intended to advise those responsible with building an enterprise data center on how to best evaluate important data center decisions, with an in-depth emphasis on choices related to the physical layer infrastructure. Since all the upper layers in the OSI model ("stack")—from the network to the applications themselves—rely on the physical layer, it is essential that it is designed to meet the needs of the data center both now and in the future.

Explore the chapters below to find out tips, answers and insights to demystify the technology, untangle the com plexity and accelerate time to market so you can identify the challenges —and opportunities—in your own data center.

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

COMMSCOPE®

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.



CO-110101.4-EN (03/22)