

Structured Cabling Design for Large IT/Service Provider Data Centers

Author:

Dave Kozischek, Marketing Applications Manager, Data Center and In-Building Networks

Introduction

“Structured cabling” is defined as building or campus telecommunications cabling infrastructure that consists of a number of standardized smaller elements (hence structured) called subsystems. For it to be effective, structured cabling is organized in such a way that individual fibers are easy to locate, moves, adds, and changes are easily managed, and there’s ample airflow around cabling.

Perhaps no environment requires effective structured cabling more than the data center. With no tolerance for downtime or network failure, the data center’s owners and operators are among the main consumers of training resources devoted to structured cabling. The reason is clear: even as fewer traditional data centers are being built in favor of outsourcing to the cloud – i.e., some type of IT service provider – there are still physical structures enabling the cloud, and these structures need to be cabled.

Fortunately, what constitutes effective structured cabling isn’t open to interpretation, rather, it’s clearly explained in the ANSI/TIA-942-B standard titled “Telecommunications Infrastructure Standard for Data Centers.” In this white paper, we’ll explore the standard and break down key considerations for making the most of structured cabling in the data center – no matter its size.

Consider the different types of data centers in operation today:

In-house data center: Also known as enterprise data centers, these facilities are privately owned by large companies. The company designs, builds, and operates its own facility – and can still provide a service for profit such as cloud services or music streaming.

Wholesale data center: Owned by IT service providers, also known as cloud providers, these data centers are in the business of selling space. Instead of building their own facilities, enterprises buy space and deploy their data center infrastructure within the wholesale facility.

Colocation data center: These facilities are like wholesale data centers, but enterprises just rent a rack, cabinet, or cage. The IT service provider is the one running the infrastructure.

Dedicated and managed hosting data centers: IT service providers operate and rent server capacity in these data centers, but each enterprise customer controls its own dedicated server.

Shared hosting data center: In these facilities, enterprise customers buy space on an IT service provider’s servers. These servers are shared among enterprise customers.

Today in the industry, a significant shift is underway in how these different types of data centers invest in their infrastructure. LightCounting and Forbes report* that cloud/IT service provider spending is up while enterprise IT spending is down, as shown in Figure 1.

Further evidence of this shift is reflected in Dell Oro’s graph of server investments, the lion’s share of which are shipping for installation in cloud-type facilities. See Figure 2.




Figure 1: Growth in Cloud/IT Service Provider Spending




Figure 2: Growth in Cloud/IT Service Provider Server Shipments

As enterprises increasingly decide to outsource some or all of their infrastructure to IT service providers, the result is not at all surprising: fewer data centers overall and hypersized facilities in their place. See Figure 3.




Figure 3: Shift from Enterprise IT to IT Service Provider Growth

The structured cabling requirements of these resulting hyperscale, multitenant data centers may differ from what has been installed in the past in the smaller single-tenant, enterprise-owned facilities – but TIA-942 provides guidance.

TIA-942 always recommends a star architecture, with different areas for cross-connecting and interconnecting cable. The standard defines five different cross-connect/interconnect areas: the main distribution area (MDA), intermediate distribution area (IDA), horizontal distribution area (HDA), zone distribution area (ZDA), and equipment distribution area (EDA).

These areas represent the full network from racks and cabinets to the main area where routers, switches, and other components are located. TIA-942 also provides guidance on redundancy definitions, and they rank those into four tiers, called ratings. Rated-1 is the lowest tier with the least redundancy. Rated-4 provides the most redundancy in a data center’s structured cabling and is typically deployed in large IT/Service provider data centers. The other basics covered by this standard include zone architectures and guidelines for energy efficiency. See Table 1 for a snapshot of the standard’s topics.

	Key Areas	Insight
	Architecture	Recommends a star topology architecture
	Cross-Connect vs. Interconnect	MDA, IDA, HDA, ZDA, EDA
	Redundancy Definitions	Rated 1-4
	Zone Architectures	Reduced topologies and consolidated points
	Energy Efficiency	Examples of routing cables and airflow contention

Table 1: Topics Covered by ANSI/TIA-942-B, Telecommunications Infrastructure Standard for Data Centers

When it comes to structured cabling, the standard addresses backbone and horizontal cabling as shown in Figure 4. Each of the distribution areas, or squares, is an area where there is a patch panel.

How much fiber is needed in each of those areas is a function of network speeds, network architectures, oversubscription, and switch configuration. Let’s look at a few examples under each of these considerations to illustrate how they affect a data center’s fiber count.




Figure 4: Backbone and Horizontal Cabling Distribution Areas

Table 2 shows how network speed influences fiber count as a data center moves from 10 to 100G. On the left is the physical architecture with four racks or cabinets each with a switch on top and a switch at the end of the row. Next is the logical architecture in TIA-942's recommended star configuration for cabling, and finally on the right is the network speed. 10G only takes 2 fibers to support; 40G can operate over 2 or 8 fibers; and 100G takes 2, 8, or even 20 fibers depending on the transceiver. So you see that, depending on the network speed, as few as 2 fibers or as many as 20 fibers are needed for just one port. Takeaway: network speeds do affect fiber count. Check road maps (IEEE for Ethernet and, on the storage side, ANSI for Fibre Channel) for detailed information on per-port fiber counts.







Physical	Logical	Speed
		10G
		40G
		100G

Table 2: Network Speed Influences Fiber Count

Now let's look at how the network's logical architecture affects a data center's fiber count. In the example provided in Table 3, each architecture's speed will be constant at 40G with 8 fibers connecting each switch. Point-to-point architecture is the simplest – both logically being a star and physically cabled as a star with 8 fibers to each cabinet. A full mesh architecture connects each switch to every other switch, totaling 32 fibers for the same five switches. That logical mesh is “cabled” physically at the cross-connect, and it takes 32 fibers to do that. The final architecture in this example is the spine and leaf, in which every spine switch (Switches 1 and 2) has to connect to every leaf switch (Switches 3-5). In the same physical configuration with the same five switches, the spine-and-leaf logical architecture requires 16 fibers. So, depending on the data center's architecture, it can take an operator 8, 16, or 32 fibers for every cabinet. Takeaway: architecture redundancy increases fiber count.







Physical	Logical	Speed
	<p>Point to Point</p> 	<p>40G 8 Fiber</p>
	<p>Full Mesh</p> 	<p>40G 8 Fiber</p>
	<p>Spine and Leaf</p> 	<p>40G 8 Fiber</p>

Table 3: Network Architecture Affects Fiber Count

Next, let's consider how oversubscription impacts fiber count. Simply put, oversubscription is the ratio of circuits coming in vs. going out of a switch. In the example shown in Figure 4, the star architecture is used physically and logically with a constant network speed of 10G. The variable shown is the oversubscription rate. The example shows a 4:1 oversubscription with 24 10G circuits coming in and six of them going out; in the middle, 24 10G circuits come in and 12 go out for a 2:1 rate; and at the bottom is 1:1 with 24 10G circuits both entering and exiting each switch. Depending on the oversubscription rate, with all other variables remaining constant, the required per-switch fiber count can be 12, 24, or 48 fibers. Takeaway: the lower the oversubscription ratio, the higher the fiber count. Ultimately, the oversubscription rate is a function of network ingress/egress traffic needs – meaning the fiber count is driven by this requirement as well.

Physical	Logical	Speed	Over
<p>24 Servers</p> <p>Switch 5, Switch 4, Switch 3, Switch 2, Switch 1</p> <p>12 Fibers, 12 Fibers, 12 Fibers, 12 Fibers</p>	<p>Switch 1</p> <p>Switch 2, Switch 3, Switch 4, Switch 5</p>	10G	4:1 6 out Switch 1 24 in
<p>24 Servers</p> <p>Switch 5, Switch 4, Switch 3, Switch 2, Switch 1</p> <p>24 Fibers, 24 Fibers, 24 Fibers, 24 Fibers</p>	<p>Switch 1</p> <p>Switch 2, Switch 3, Switch 4, Switch 5</p>	10G	2:1 12 out Switch 1 24 in
<p>24 Servers</p> <p>Switch 5, Switch 4, Switch 3, Switch 2, Switch 1</p> <p>48 Fibers, 48 Fibers, 48 Fibers, 48 Fibers</p>	<p>Switch 1</p> <p>Switch 2, Switch 3, Switch 4, Switch 5</p>	10G	1:1 24 out Switch 1 24 in

Table 4: Network Oversubscription Impacts Fiber Count

Finally, a look at how the network's switch configuration drives fiber count. Using constant architectures and running 10G to all of the servers on the racks, we reconfigure what happens on the right side with the switch. In Table 5, all of the circuits going down are 10G; two of the 40G ports are quad small-form-factor pluggable 40G optical transceivers (QSFP), i.e., 8-fiber MTP® connection; they break out into four 10G to total 16 more ports – yielding two 40G ports going up, or $2 \times 8 = 16$. In the middle, we see the same switch with all four of the 40G ports going back up to the core – equating to $8 \times 4 = 32$ fibers. The final scenario shows an equal distribution of 10G going down as going up. 40G ports break out into 10G for $16 \times 10G$ ports, adding more 10G to make it even totals 64 fibers. Takeaway: just deciding how to configure the switch affects the fiber count in these scenarios from 16, 32, or 64 fibers.










Physical	Logical	Switch Configuration
		
		
		

Table 5: Network Switch Configuration Drives Fiber Count

Note that this switching configuration only addresses the Ethernet side of these servers. The fiber count would continue to climb if the servers also had a Fibre Channel network and/or ports for InfiniBand high-speed computing.

Furthermore, we've looked at how the four variables can independently increase the number of fibers needed in data centers, so imagine the impact that mixed variables can have in driving fiber counts up even higher. Changing the network's operating speed affects the fiber count, sure, but change the speed and the architecture? Or change the speed and the oversubscription rate? Fiber counts that were already relatively high go up even more.

What remains is the question of how to cable this type of data center. Typically today's increasingly large data centers extend to separate locations much like an enterprise campus as shown in Figure 5.

Indoor cable is typically used within each building, connected by indoor/outdoor cable and transitional optical splice enclosures. See Table 6.




Figure 5: Large IT/Service Provider Data Center

	Key Areas	Insight
	Meet Me Room	<ul style="list-style-type: none"> • Demarcation • Cross-Connect
	Main Distribution Area	<ul style="list-style-type: none"> • Racks/Cabinets • Cross-Connect
	Indoor Cabling	Plenum Rated
	Indoor/Outdoor Cabling	<ul style="list-style-type: none"> • Plenum/Riser • Armored Cable
	Optical Splice Enclosure (OSE)	Transition from Indoor to Outdoor Cables

Table 6: Data Center Cabling Areas

When it comes to deployment methods, there are three to consider:

Preterminated cable: Typically deployed for indoor plenum-rated cabling, these trunks are factory-terminated on both ends with 8- or 12-fiber MTP® connectors. They are ideal for MDA to HDA or EDA installations involving raceway or race floor where the entire fiber count is being deployed in one run at a single location at each end of the link. See Figure 6.

Pigtailed cable: These semi-preconnectorized assemblies are factory-terminated on one end with MTP connectors for easy high-fiber-count deployment while remaining unterminated on the other end to fit through small conduit or allow for onsite length changes. Often used in building-to-building installations, pigtailed cable is ideal for situations when conduit is too small for pulling grips and/or the cable pathway can't be determined before ordering. See Figure 7.

Bulk cable: This deployment option requires field connectorization on both ends, typically with MTP splice-on connectors. Bulk cable is best for deployments requiring center-pull installation and/or extremely high fiber counts (such as 1,728 fibers and up). See Figure 8.




Figure 6: Preterminated Cable




Figure 7: Pigtailed Cable




Figure 8: Bulk Cable

Table 7 provides an overview of the three deployment methods and their corresponding fiber counts.




Category	Method	Environment	Connector	Counts	Trunk Type	Fiber Type
 1 2	Preterminated Cables Pigtail Cable	Indoor 	MTP® to MTP Connector	<ul style="list-style-type: none"> • 144 • 192 • 216 • 288 • 432 • 576 	<ul style="list-style-type: none"> • Non-Armored 	<ul style="list-style-type: none"> • Multimode • Single-Mode
 1 2	Preterminated Cables Pigtail Cable	Indoor/Outdoor 	MTP Connector to Fiber	<ul style="list-style-type: none"> • 144 • 216 • 288 • 432 • 576 • 864 	<ul style="list-style-type: none"> • Armored • Non-Armored 	<ul style="list-style-type: none"> • Multimode • Single-Mode
 3	Bulk Cable	All	Fiber to Fiber	144 to 1,728	<ul style="list-style-type: none"> • Armored • Non-Armored 	<ul style="list-style-type: none"> • Multimode • Single-Mode

Table 7: Deployment Methods and Cabling Choices

Putting all of this information to practice, the following example illustrates how a four-way spine is cabled and the resulting fiber count. Table 8 comes from Cisco’s Massively Scalable Data Center White Paper, showing the Nexus 7000 switches. Based on manufacturer recommendations, there are 48 leaf switches with 32 ports down to servers and 32 ports going back up into the fabric. In this example, we use two Cisco 3064 switches at the top of each of the 24 cabinets to create an “A” fabric and a “B” fabric. Figure 9 shows how these recommendations translate to a logical architecture.





Figure 9: Cisco Four-Way Spine Architecture

Cisco Four-Way Spine Configuration	Device Count
Nexus 7009/7010 Spine Switches	4
N7K-F248XP-25 Blades per 7009 Chassis	7
N7K-F248XP-25 Blades per 7010 Chassis	8
Ports Used for Leaf Switches per 7009 Chassis	336
Ports Used for Leaf Switches per 7010 Chassis	384
Nexus 3064 Leaf Switches	48
Nexus 3064 Ports Facing Fabric	32
Nexus 3064 Ports Facing Servers	32
Key Design Parameters: <ul style="list-style-type: none"> • All 10G Ethernet, no 40G • Spine switches: Cisco Nexus 7000 series with 48-port blades • Leaf switches: Cisco Nexus 3064 • 32 ports facing fabric, 32 ports facing servers 	

Table 8: Cisco Four-Way Spine

As shown in Figure 10 and Table 9, Cisco’s spine-and-leaf architecture guidance provides for a four-way spine with 48 leaf switches. Starting at the rack and working backward, 32 ports go out to every leaf switch which translates to 64 fibers required per switch. With two switches on each rack, 128 fibers are needed to support this architecture for every cabinet. This design called for 10G and a 1:1 oversubscription as previously covered, and we will proceed with this example using fiber counts that are divisible by 12.



Fiber Count Variables	Details
Cisco Spine-and-Leaf Rules	4 Spine Switches 48 Leaf Switches
Cisco Leaf Rules	32 Fabric 32 Server
Architecture	Spine-and-Leaf A + B Fabrics
Network Speed	10G
Oversubscription	1:1
Standard Fiber Counts	12-Fiber Divisible

Table 9: Cisco Spine-and-Leaf Fiber Count Variables

Figure 10: Cisco Spine-and-Leaf Architecture

When it comes to cabling this scenario, we have options. They may not all be good options, like the one depicted in Figure 11 that uses jumpers – over 3,000 of them. Better would be consolidating jumpers into 48 72-fiber cables as shown in Figure 12. Even better yet is Option 3: using high-fiber-count trunks, 576 fibers in each one, getting us down from 3,000 jumpers to six 576-fiber trunks. See Figure 13.




Figure 11: Cabling Option 1




Figure 12: Cabling Option 2




Figure 13: Cabling Option 3

To understand why some of these options are better than others, let's explore their relative ease of use (which translates to labor spend) both in their initial installation and during the moves, adds, and changes that are inevitable in a data center. Table 10 speaks for itself.

Would You Rather ...	Option 1 Qty 3,072: 2-Fiber Jumpers	Option 2 Qty 48: 72-Fiber Trunks	Option 3 Qty 6: 576-Fiber Trunks
Test and Clean	6,144 2-Fiber Duplex LC Connectors	576 12-Fiber MTP® Connectors	576 12-Fiber MTP Connectors
Document and Label	3,072 Jumpers and 6,144 Connectors	48 Trunks and 576 Connectors	Six Trunks and 576 Connectors
Pull and Install	3,072 Jumpers (Both Ends)	48 Trunks (Both Ends)	Six Trunks (Both Ends)
Purchase	3,072 Jumpers	48 (72 Fiber Trunks)	Six (576 Fiber Trunks)
Troubleshoot	3,072 Links, > 6,000 Connectors	48 Links, 576 Connectors	Six Links, 576 Connectors
Move, Add, or Change	One Jumper at a Time, Point-to-Point Configuration	Create Cross-Connect, Use Short Jumper	Create Cross-Connect, Use Short Jumper

Table 10: Deployment Differences from Jumpers to High-Fiber-Count Trunks

Furthering the case for high-fiber-count trunks are their impact on valuable data center real estate – the pathway for cabling. TIA-569 provides calculations to understand what percentage of tray/conduit/raceway is taken up by cabling along with a recommendation that the individual maximum fill ratio not exceed 25 percent. Though it may not be intuitive, it is a fact that a 50 percent fill ratio actually uses up an entire pathway, because the spaces between cables are part of the equation. With this in mind and referring to Figure 14, the first option using more than 3,000 jumpers isn't an option at all. However, the second cabling option (48 72-fiber trunks) does work in a 4 x 6-in tray but not quite as well in a 4 x 4-in tray. Both tray sizes can easily accommodate the six 576-fiber trunks option.




Figure 14: Fill Ratio Differences from Jumpers to High-Fiber-Count Trunks

Conclusion

High-fiber-count trunks can be the best fit in today's data centers. The days of trusty 12- and 24-fiber trunks to each rack are no more; now we're looking at data centers that are increasingly growing in scale and in the fiber counts required to support higher speeds, greater oversubscription rates, redundant architectures, and creative switch configurations. It's clear that large facilities are the new normal; enterprise IT customers will continue to shift away from small, single-tenant facility operators toward outsourcing all or part of their data center infrastructure. Fortunately, there are proven structured cabling methods and global manufacturers with many years of experience solving data center challenges – and the assurance of TIA-942 continuing to provide guidance.



CORNING

WESCO®

Corning Optical Communications reserves the right to improve, enhance, and modify the features and specifications of Corning Optical Communications products without prior notification. A complete listing of the trademarks of Corning Optical Communications is available at www.corning.com/opcomm/trademarks. All other trademarks are the properties of their respective owners. Corning Optical Communications is ISO 9001 certified. © 2017 Corning Optical Communications. All rights reserved. LAN-2184-AEN / September 2017