

Structured cabling in the data center can help enterprise business thrive coming out of the pandemic


Flexible, standards-based cabling systems provide cost effectiveness and flexibility.

BY RON TELLAS and **BETSY CONROY**,
Communications Cable and Connectivity Association

The ramping up of the digital economy fueled by emerging technologies and ever-increasing data has been placing huge demands on data center infrastructure for decades, with studies indicating a doubling of data center space and capacity between 2010 and 2019. While data center growth was already expected to continue increasing over the next decade, the COVID-19 pandemic is now changing the way people live and work and the way the world does business, further driving data center expansion. Data centers are needing to quickly add more capacity to support remote IT capabilities, more e-commerce, increased video streaming and gaming, and a greater need for applications like telemedicine, distance learning and online collaboration. And they're focusing heavily on low-latency connectivity to do so.

Hyperscale and large cloud data centers tend to be early adopters that shape the industry, with their practices ultimately becoming the standard for data center design

and deployment. Current connectivity trends within these spaces are supporting the need to quickly and cost-effectively ramp up capacity in response to emerging technologies and the demand for high-speed, low-latency performance in the evolving digital economy and COVID-19 world. As enterprise businesses re-evaluate the public cloud due to a variety of challenges and begin to realize the benefits of cost-effective hyperconverged infrastructure



The COVID-19 pandemic has changed the way people live and work, and the way the world does business—further driving data center expansion.

technologies and techniques that are becoming more mainstream, allowing them to mimic the hyperscalers, they need to adopt the right architecture, deployment topology and components. As they do so, it makes sense to shift away from point-to-point cabling and back to the timeless, flexible and standards-based best practice of structured cabling.

New norm spotlights low-latency connectivity

While emerging Internet of Things/Industrial Internet of Things (IIoT) and 5G-enabled technologies like self-driving cars, virtual and augmented reality, artificial intelligence (AI), machine-to-machine (M2M) communication and advanced data analytics were already fueling data center expansion, the COVID-19 pandemic has shed more light on the need to reduce latency (i.e. the time it takes to transmit and process data). Low-latency connectivity has become the umbilical to the outside world and digital readiness, as evidenced by significantly increased internet usage with extreme spikes in video streaming, e-commerce and online gaming that had Netflix and YouTube cutting back video quality, Amazon Web Services (AWS) ramping up capacity and Microsoft Azure working around the clock to deploy new servers (all while socially distancing 6 feet apart).

At the same time, enterprise businesses in both the public and private sectors have had to quickly reinvent themselves and embrace a fully digital approach to a remote workforce, maintain business continuity and provide some semblance of “business as usual” for customers—everything from telemedicine, distance learning and online ordering and delivery, to video conferencing, online collaboration and

virtual events. As a result, enterprise data centers have also had to quickly respond to the need for increased bandwidth and minimal downtime through virtual private networks and remote accessing capabilities. In fact, reports indicate that VPN usage in the U.S. is up by more than 100% since the onset of the pandemic in March 2019.

“In response to the recent pandemic, there has been an increased demand for data center capacity, network requirements and data security,” says Wendy Stewart, vice president of sales operations at Databank, a provider of enterprise-class data center, cloud and interconnection services with data center facilities throughout the U.S. “The drive for more data center capacity is partly a result of companies needing additional infrastructure to support the ongoing demand for remote life—work, school and play. Not only are companies addressing their current needs, they are planning for a future that consists of more remote and virtual living, which will require a greater focus on uptime, scalability, geo-diversity and security.”

The impact that the pandemic has had on data centers of all types and sizes is undeniable, creating massive pressure on existing IT infrastructure with data center managers scrambling to expand capacity while ensuring high-speed, low-latency bandwidth for maximum uptime. Latency is caused by a number of factors that include the distances data must travel, the number of network hops between switches and available bandwidth. While latency was previously just an annoyance—that 10 seconds it takes to load a web page, the occasional jitter in an online video game or the buffering of a Netflix movie—it simply can no longer be tolerated for emerging technologies and those that are now increasingly

prominent in the COVID-19 era. Consider the following innovations and their requirements.

- Self-driving cars and M2M communications need millisecond signal transmission.
- Businesses can't afford latency causing online video conferences with customers to consistently cut out.
- Medical professionals and teachers need to be clearly seen and understood during remote learning and telemedicine sessions.
- The entertainment industry must ensure that virtual events, video streaming and online gaming provide an uninterrupted, seamless experience.

“With COVID-19, many of our data center customers are facing challenges related to either the need for rapid expansion or to the new reality of remote work,” says Mike Peterson, director of product management, interconnection at Flexential, a provider of data center, colocation, cloud, connectivity and data protection with nearly 30 data centers located throughout North America. “VPNs are becoming crowded with traffic, bandwidth is being consumed and outdated switched networks are being pushed in ways never thought of when they were designed. Companies are starting to understand that a traditional centralized approach leads to bottlenecks and latency issues. Running an optimized network architecture with network hubs near users, both internal and external, allows traffic decisions to be made closer to the edge and leads to increased performance and lower latency.”

Shifting away from the cloud

Large hyperscale and cloud data centers are adopting architectures,

Stats, standard set the tone for enterprise data centers

The authors of this article adeptly describe the role the global pandemic has played in changing the dynamics of data center planning and administration. Some cabling-market statistics released by market intelligence firm BSRIA support the authors' reporting, and add context to it. In late 2020, BSRIA reported in its Global Cabling Report that it expected the global cabling market to fall by 12%, to US \$5.8 billion from 2019's US \$6.6 billion. BSRIA said it anticipates that in 2022 consumption will return to the \$6.6 billion level. Among the information that BSRIA made public from its report were the variances from one market segment to another, from different geographic regions, and between copper and fiber cabling. "There were significant variances between countries, typically related

to the general economic impact of the pandemic," BSRIA said when it made its findings known in November 2020. "There have also been various government support policies for the country economies in general or for specific sectors, such as buildings. Sectors such as travel, hospitality and retail [have been] especially badly hit, while healthcare, data centers and in some countries, education, remain more buoyant." Copper cable and connectivity suffered more than fiber did, according to BSRIA. Copper sales declined 15% globally, with a fall of between 14% and 16% holding consistently across main regions of the globe. BSRIA attributes this decline to "a severe disruption in non-residential construction with the delay or cancellation of many projects." Fiber, on the other hand, declined 7%,

but with significant variations from region to region, "from a stable market in Asia Pacific to a drop of 14% in America," BSRIA stated. "In countries where fiber has held up, this was often supported by increased investment in data centers partly driven by the effects of the pandemic in driving demand for web services." As Carrie Goetz said, the modern work-from-anywhere approach "will give rise to more colo, on-premises and edge data centers closer to business locations where companies can maintain their own infrastructure." To that end, the Telecommunications Industry Association (TIA) soon will publish an addendum to its 942-B data center standard. The addendum will address infrastructure systems and design guidelines for edge data centers. —Patrick McLaughlin

topologies and components that provide low-latency, high-bandwidth connections while optimizing scalability, substantiating their critical role in supporting the new norm and increased internet usage, video streaming e-commerce and online gaming. Both Netflix and Hulu use AWS's public cloud for a large portion of their data center operations. These facilities have also played a pivotal role in enabling enterprise business continuity during the pandemic—thick cloud-based platforms like Zoom, Microsoft 360, Salesforce, DocuSign and others. Because these facilities are highly effective at enabling faster deployments, better performance, scalability and flexibility, they also provide the means for these large technology providers to

more quickly and cost-effectively roll out new technologies.

While enterprise businesses leverage the cloud for common business applications, as well as inexpensive storage and platforms for developing and stress-testing new digital innovations, the previously expected shift of comprehensive legacy IT to the public cloud has never fully come to fruition. Many enterprise businesses never felt confident enough to place sensitive or critical workloads in the cloud and have maintained a hybrid approach where internal systems remain privately hosted either on-premises or within colocation space. Many other enterprise businesses are now re-evaluating the public cloud approach, with several pulling apps back in house.

Recent surveys show that companies currently put only about 18% of their workloads in the cloud, and the majority of workloads are expected to remain privately hosted. With colocation data centers popping up in more places, placing them closer to users for better latency, it's expected that we will see even more enterprise businesses leveraging these spaces for their privately hosted assets.

Much of the original cloud appeal of shifting IT budgets from capital to operational expenditure, and having access to unlimited IT storage and compute resources, has been overshadowed by a variety of concerns and challenges, including the following.

Security—Companies are concerned about the increased risk of

cybersecurity and exposure of business-critical information due to the cloud-based model of accessing data from anywhere.

Compliance—Healthcare, financial and other industries regulated by stringent data privacy regulations are increasingly concerned about the replication and location of publicly stored data that can put them in violation.

Control—Lack of full control and visibility over data, such as where it is located and how it is managed, as well as less ability to customize, is causing frustration among many enterprise data center managers and IT professionals.

Cost—The shift to operational expenditure combined with a lack of visibility has often resulted in cloud overspending, and data transfer costs have added up well beyond what many enterprise businesses expected.

“Many companies that adopted the cloud-first mentality are realizing the security issues of having data in multiple locations with lack of control, which is why some sectors that deal with strict privacy data regulations like healthcare and finance have never fully embraced the public cloud, opting instead for a private or hybrid cloud solution,” says Carrie Goetz, RCDD/NTS, CNID, CDCP, AWS CCP, principal and chief technology officer of StrategITcom, who has decades of global experience designing, running and auditing data centers. “Others who are taking a good hard look at cost creep and anticipated growth are realizing that the public cloud may ultimately be too expensive, which is why companies like Dropbox and Uber have moved data assets back in-house to their own data center facilities where they will have more control and flexibility. Early in the pandemic, many enterprise companies leveraged the cloud

as the path of least resistance to get services up and rolling, but it was also a Bandaid for many. Moving forward, these companies are going to need to take a step back and determine what will meet their long-term strategy and support the new work-from-anywhere mentality. This will give rise to more colo, on-premises and edge data centers closer to business locations where companies can maintain their own infrastructure.”

While it used to be that the IT workload and computation needs of an enterprise business needed to reach a certain size before the public cloud ceased to be financially worthwhile, that is no longer the case due to the availability of cost-effective hyperconverged infrastructure technologies and techniques adopted by large hyperscale data centers (e.g., Microsoft, Google, Facebook) and cloud data centers that are now becoming more mainstream. Through advanced open-source protocols, white-box hardware and software-defined networking, enterprise businesses are now able to more closely mimic hyperscales and deploy highly virtualized server environments, which will further drive the shift away from the public cloud. These hyperconverged infrastructure environments significantly ease scalability and manageability for enterprise businesses to quickly and cost-effectively expand capacity and support digital innovation that will allow them to continue weathering the COVID-19 storm and even potentially thrive in it.

“Our primary business is colocation, and we have seen consistent growth from small- and medium-sized businesses in particular. While some customer applications are still going to the public cloud, for others it may not be a technical, performance or financial fit and these businesses

pivot to colocation space and our infrastructure,” says James Beer, senior vice president for eStructure, the largest Canadian-owned network and cloud-neutral data center provider with locations in Montreal, Vancouver and Calgary. “We are definitely not seeing a drop in colocation demand.”

Architecture and topology matter

In recent years, data centers have been trending toward the use of full-mesh, leaf-spine fabric architecture that reduces latency and supports data-intensive and time-sensitive applications in virtualized server environments where resources for a specific application are often distributed across multiple servers. Compared to a traditional three-tier architecture that creates a north-south traffic pattern through multiple switches, the leaf-spine approach optimizes east-west data center traffic for low-latency server-to-server communication by reducing the number of switches through which information must flow. This is accomplished by connecting every leaf switch to every other leaf and spine switch within the fabric.

From a topology standpoint, top-of-rack (ToR) deployments with short point-to-point connectivity took hold over the past decade as the primary means of supporting switch-to-server connections in the enterprise data center. However, many data center managers are now realizing that a ToR topology cannot effectively support a modern virtualized leaf-spine environment with the low-latency performance and scalability that the enterprise business demands. Only timeless, standards-based fiber-optic structured cabling offers the flexibility to support these environments.

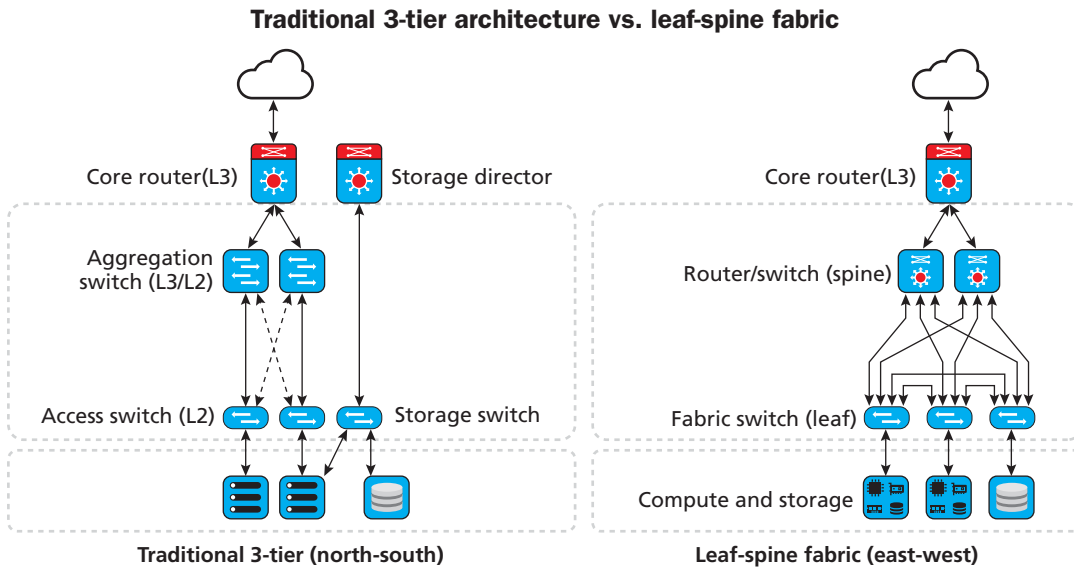
First of all, having a ToR leaf switch in each cabinet makes connecting

every leaf switch to every other leaf switch highly impractical. It also means higher latency due to the extra switch hop required when a server in

enterprise businesses will need to examine their spaces and determine what they can support. If they don't have the power capacity to fill a cab-

switch purchases and related maintenance and power.

In contrast, fiber structured cabling is far better suited to supporting virtualized environments, enabling low-latency communications between servers and providing the flexibility and ease of scalability needed to quickly and cost-effectively expand. With longer-distance structured cabling, larger leaf switches can be placed at the end of a row of cabinets and connected to multiple servers



A leaf-spine fabric connects every leaf switch to every other leaf and spine switch within the fabric. This optimizes east-west data center traffic flow for low-latency server-to-server communication.

one cabinet needs to “talk” to a server in another, as is common in a virtualized environment. A ToR topology also limits scalability because a single switch upgrade improves connection speed only to the servers located in the cabinet where the switch resides. Having a ToR leaf switch in each cabinet requires more power across the data center and demands higher port densities for spine switches, which can cause further scalability constraints. Furthermore, due to their inherent nature with small-form-factor SFP or QSFP modules or embedded transceivers, the assemblies used in ToR point-to-point connections, such as short-length direct attach cables (DACs) and active optical cables (AOCs), do not support multiple future generations of applications and will need to be replaced as speeds increase.

“As part of the reawakening and bringing IT resources back in house,

inet, ToR won't make a lot of sense and the use of centralized switches will be a better option,” says Goetz. “There's also the waste factor. Those that started out with DACs to support 10-Gbit/sec server links and are now ready to migrate to higher speeds will have to throw the DACs away and buy something new. Distance, speed and lifecycle costs must all be evaluated for the best long-term design strategy.”

A ToR configuration also means more switches to maintain, which can quickly become an operational burden. Port utilization is an additional cost concern. When using a ToR deployment, data centers may discover that they are not fully utilizing all switch ports due to power and cooling concerns that often limit the number of high-performance virtualized servers per cabinet. These unused switch ports across several cabinets can add up and ultimately equate to unnecessary

within the row via structured cabling interconnects or crossconnects. With all the servers in a row connected to the same end-of-row (EoR) leaf switch, there is no extra switch hop when two servers housed in different cabinets need to communicate.

Unlike a ToR switch, larger EoR switches do not limit scalability. The upgrade of an EoR switch speeds connection to all servers in the row, rather than just those in a single cabinet. Lower port requirements at spine switches also provide room for growth, and standards-based fiber structured cabling supports multiple generations of applications—multimode and singlemode fiber currently support up to 400 Gig, and standards bodies have their sights on 800 Gig and beyond.

It should also be noted that unlike point-to-point assemblies, standards-based fiber cables available from a variety of reputable cable

manufacturers are interoperable with much longer warranties and third-party verified performance to work with any vendor's equipment. Port utilization is also maximized with structured cabling and EoR switches because the switch ports are not confined to single cabinets. Switch ports on higher-density EoR leaf switches can be divided up, on demand, to any of the servers across several cabinets in a row. Only having to maintain one switch per row rather than a switch in every cabinet can also help reduce cost.

The second figure shows the difference in server-to-server communication between a ToR topology with point-to-point connections versus an EoR topology using structured cabling.

It is clear to see that when server A needs to talk to server B in a ToR topology, the signal must travel through three switches—from one ToR leaf switch, to the spine switch, and back to the second ToR leaf switch. Even if every ToR leaf switch was connected to every other ToR leaf switch, the signal would still need to travel through two switches. The EoR configuration with structured cabling requires just one switch hop for server A to talk to server B.

"ToR switch vendors aren't going to educate customers on structured cabling because they'll sell fewer switch ports, and there was originally the belief that this topology is the easiest solution. And everybody loves the easy button," adds Goetz. "But with

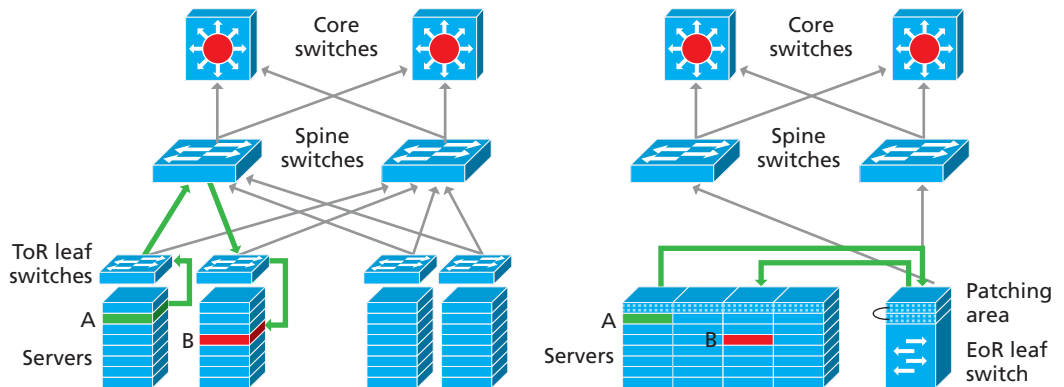
open-source white-box solutions and additional, and sometimes more-feature-rich, vendors, there are now more players—and there's a lot to be said for understanding the long-term ramifications of taking the easy route and re-evaluating all of the options based on future business needs."

The beauty of the crossconnect

Using fiber structured cabling throughout the data center enables the

New services can be quickly brought online at the crossconnect. Spine switches in the main distribution area (MDA) can be connected to the crossconnect via permanent fixed links and new leaf switches can be easily connected to unused spine switch ports at the crossconnect. This is especially ideal in a colocation data center when there is a need to quickly connect customer equipment to service provider equipment outside of the

Server-to-server communication, ToR point-to-point vs. EoR structured cabling



With server-to-server communication in a top-of-rack topology, the signal must travel through three switches. In an end-of-row topology with structured cabling, server-to-server communication requires just one switch hop.

use of distribution areas with traditional crossconnects for flexible, standards-based connections between equipment, including leaf switches to servers, leaf switches to spine switches, and servers to storage devices. The use of crossconnects eases expansion in virtualized server environments, as well as enabling the clustering of servers to more easily share compute and storage resources. The beauty of the crossconnect is that by using fiber panels that mirror the ports on connecting equipment, data center managers enable an "all-to-all" scenario where any equipment port can be connected to any other equipment port by simply repositioning fiber patch cords at the front of the fiber panels.

meet-me room and without having to access equipment, providing an additional level of security and assurance without interruption of data center operations. That is why the crossconnect is considered such a valuable asset in these environments.

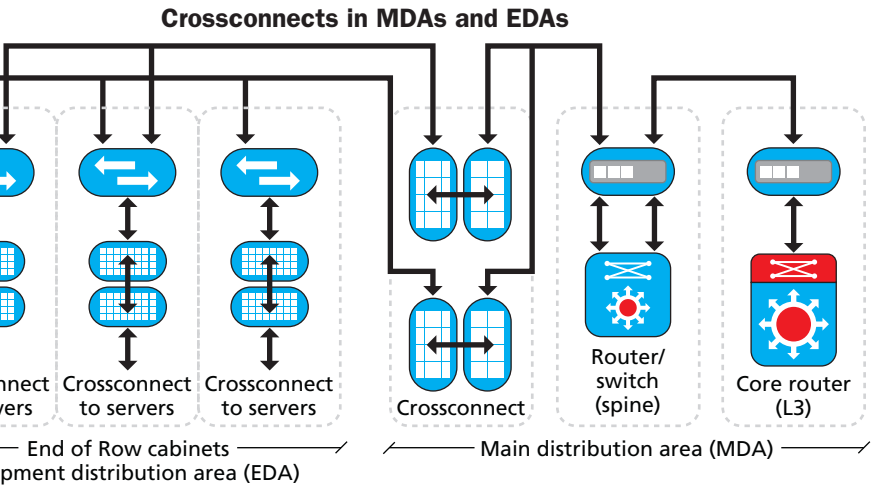
"Structured cabling is a convenient, cost-effective, efficient and secure method for connecting end users to their ISPs in the colo space," says Databank's Stewart. "No longer do you need five days to connect end users to ISPs. Connection can be completed within 24 hours as a result of the infrastructure already being in place."

Beers of eStructure agrees. "Our crossconnect volumes are increasing quarter over quarter. Clearly there

is a lot of industry network activity and traffic volume growth and customers needing to connect to more service providers and more locations. As a carrier-neutral colo provider, we have multiple site fiber entrances and meet-me rooms, and the crossconnects are a source

of revenue and the means to connect our tenants to their preferred carriers. These end customer crossconnect opportunities also allow us to attract additional carriers.”

The benefits of the crossconnect hold true within equipment distribution areas (EDAs) located in the enterprise customer colocation space or on-prem data center space. If new servers need to be added to a row, they can be easily connected to the EoR leaf switch at the crossconnect. The distances supported by fiber structured



Crossconnects located between end-of-row leaf switches and servers, and between leaf switches and spine switches, support flexibility and rapid expansion.

cabling, from 100 meters to 10 kilometers depending on the application, also means that these crossconnects can be located virtually anywhere in the data center.

“Once you deploy a cabinet or cage, a crossconnect allows customers to take advantage of adjacent services, partners, carriers and other ecosystems,” says Flexential’s Peterson. “Crossconnects run to a meet-me room when connecting to certain services or partners or they may run directly from a customer cage to another cage. More


and more, these crossconnects are becoming fiber as distances increase, and even more so as bandwidth needs increase. For us, the crossconnect is all about supporting our customer connectivity needs and their digital journey.”

As enterprise businesses and their supporting colocation providers leverage hyperconverged infrastructure technologies and techniques that originated in hyperscale and cloud data centers and are enabled by the timeless, flexible and standards-based best practice of structured cabling, they will have the agility to cost-effectively expand as they strive to reinvent themselves in the COVID-19 era with new digital online tools and capabilities. ◆

Ron Tellas is an instrumental member of the Data Center Committee for the Communications Cable and Connectivity Association (CCCA), which serves as a resource for well-researched, fact-based information and education on issues and technologies vital to the structured cabling industry. He also holds the title of technology and applications manager for Belden.

Betsy Conroy is an industry freelance writer and consultant to CCCA. She has 20 years of experience helping organizations deliver relevant, authoritative information on a variety of ICT-related technology topics.

The crossconnect



The crossconnect creates a convenient, highly flexible “any-to-any” scenario for virtualized environments where any equipment port can be connected to any other equipment port via fiber patch cords.