Next Step For Your Network
40/100G Ethernet

Why are data centres migrating to 40/100G Ethernet?

Overview

With the ratification of new industry standards and increased demands on data centre throughput, 40/100G Ethernet will be an integral component of the next generation data centre. In fact, it is already an emerging influence on how organizations plan, build and operate their existing data centre architectures. The market proves this: manufacturers are already responding to the increased demand for Ethernet hardware, including cabling products, switches and transceivers.

The implementation of 40/100G Ethernet is dependent upon a variety of organizational factors, including existing infrastructure, budget, throughput demand and leadership priority. However, it is clear that the stage is set for the most dramatic change related to data center fiber optic infrastructures since their inception. In this paper, we will discuss the deep impact that this network speed transition has on data center cabling infrastructure, and the decisions that organizations will need to make to accommodate these changes.

> ***Like the physical universe, the digital universe is large – in 2020 containing nearly as many digital bits as there are stars in the universe. It is doubling in size every two years, and by 2021 the digital universe – the data we create and copy annually – will reach 44 zettabytes, or 44 trillion gigabytes.***

Data growth. The world revolves around digital data. We now rely on data to conduct business, engage in social activities and manage our lives. There is no sign of slowed growth in the production of, and demand for, more data – as well as faster access to it. According to the 2014 IDC Digital Universe Study sponsored by EMC [1]: "Like the physical universe, the digital universe is large – by 2020 containing nearly as many digital bits as there are stars in the universe. It is doubling in size every two years, and by 2020 the digital universe – the data we create and copy annually – will reach 44 zettabytes, or 44 trillion gigabytes."

Cloud. Among several other factors, the increase in cloud storage will drive the need for data throughput. "In 2013, less than 20% of the data in the digital universe [was] 'touched' by the cloud, either stored, perhaps temporarily, or processed in some way. By 2020, that percentage will double to 40% [1]."

The Internet of Things. Another factor contributing to the exponential growth of information is the advent of the Internet of Things (IoT). "Fed by sensors soon to number in the trillions, working with intelligent systems in the billions, and involving millions of applications, the Internet of Things will drive new consumer and business behavior that will demand increasingly intelligent industry solutions... [1]."

This exponential growth in information means processing speeds have to increase as well, so as not to slow access to data. High-performance cabling that can transfer data over 40/100G Ethernet will be a necessary addition to data centres looking to keep up with this digital data growth.

Virtualisation: the double-edged sword. Virtualisation can help data centers save on capital expenses, improve operational efficiency and create more agile infrastructures.

There are many types of virtualisation, from desktop to storage to server virtualisation. Server virtualisation, in particular, calls for fewer, more efficient servers, which, which translates to fewer server connections. Because there are fewer connections, is it important that these connections work properly.

Unfortunately, most data centres do not contain cabling infrastructure designed to meet the high-performance capabilities that virtualisation demands. This is particularly true for data centers built in the 1980s, before high-performance cabling even existed.

> ***When data transactions are interrupted due to network downtime, it translates to a very real loss of revenue. When Amazon.com went down in August 2013, the company lost £50,000 per minute. Considering how quickly lost revenue can add up, it makes sense that there is an extremely low tolerance for network downtime.***

Decreasing tolerance for downtime. When data transactions are interrupted due to network downtime, it translates to a very real loss of revenue. When Amazon.com went down in August 2013, the company lost over £50,000 per minute. Considering how quickly lost revenue can add up, it makes sense that there is an extremely low tolerance for network downtime.

The effect of downtime on revenue is even greater when considering end-user experience. According to one source, network downtime measured for user experience and business needs costs an average of £5,000 per minute!

Network administrators should have a contingency plan in place in the event of network failure. However, one of the most effective ways to mitigate this issue is to make sure the existing network is able meet the demands of increasing data throughput, including upgrading networks to be capable of handling 40/100G speeds.

Managing capital expenses. While migrating to 40/100G Ethernet creates an up-front capital expense, it saves data centres in the long run by future-proofing infrastructure. Not only will data centres be prepared for the increasing demands on data throughput, but the high-performance cabling infrastructure required of 40/100G Ethernet can grow with future hardware upgrades. This will reduce the need to tear out and replace cabling with each upgrade.

Factors in preparing for 40/100G Ethernet migration

Link distances and loss amounts

As data centre speeds increase, optical loss budgets decrease. Optical loss occurs over cabling distance and at mating points where connections are made. Since most data centre cabling runs are shorter distances (compared to long-haul campus runs), the inherent losses from distance in a data centre are somewhat negligible compared to the losses incurred from mating points.

As connections in the data centre increase to improve manageability, performance suffers. This is because added connections contribute to increased dB loss. Therefore, a balance must be maintained between manageability and performance.

Choosing the right cabling product can combat the issue of balancing manageability versus performance. Cabling products with low optical loss rates will ensure that a structured cabling environment is running at its peak. When comparing dB loss rates of cabling products, look for "maximum" instead of "typical" loss rates. While typical loss rates can allude to the performance capabilities of a product, they are not helpful when determining loss budgets.

> ***Cabling products with low optical loss rates will ensure that a structured cabling environment is running at its peak. When comparing dB loss rates of cabling products, look for "maximum" instead of "typical" loss rates.***

Ethernet Transmission Speeds And Loss Amounts Over The Years

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
1998	Short Wavelength Fast	10/100 Mbps	TIA/EIA-785	4
2000	Ethernet 1G Ethernet	1,000 Mbps	IEEE 802.3z	3.56
2004	10G Ethernet	10,000 Mbps	IEEE 802.3ae	2.6
2010	40G SR4 Ethernet	40,000 Mbps	IEEE 802.3ba	1.9
2015	100G SR4 Ethernet	100,000 Mbps	IEEE 802.3bm	1.9

> ***TIA-942 offers a roadmap for data centre cabling infrastructure based on the concept of a structured cabling environment.***

Cabling infrastructure design

Due to the exponential port growth experienced by data centres during the last two decades, cabling infrastructure is often reduced to a cluttered tangle commonly referred to as "spaghetti cabling." Spaghetti cabling leads to decreased efficiency, increased dB loss and more cable management challenges.

The Telecommunications Infrastructure Standard for Data Centres, or TIA-942, was developed to address various data centre infrastructure design topics, including the problem of spaghetti cabling. Among other aspects of data centre planning and design, TIA-942 focuses on the physical layout of cabling infrastructure.

TIA-942 offers a roadmap for data centre cabling infrastructure based on the concept of a structured cabling environment. By creating logical segments of connectivity, a structured cabling system can grow and move as data centre needs change and throughput demands increase. Therefore, implementing a structured cabling system in accordance to TIA-942 standards is the ideal way to prepare for migration to 40/100G speeds.

The heart of a structured cabling system is the main distribution area, or MDA. All equipment links back to the MDA. Other terms used to define this area include: main cross connect, main distribution frame (MDF), and central patching location (CPL). The principle of a structured cabling system is to avoid running cables from active port to active port (often referred to as "point-to-point"). Instead, all active ports are connected to one area (the MDA), where the patching is done. This is also where moves, adds and changes (MACs) take place.

TIA-942 calls for the use of interconnect points, which are typically in the form of patch panels (also referred to as fiber enclosures). Patch panels allow for patch cables to be used in the front of the racks or cabinets where the equipment is housed. The patch cable would then connect to a fiber optic trunk, and then to another patch panel in the MDA.

There are several advantages to implementing a structured cabling system. First, using fiber optic trunks significantly reduces the amount of cabling bulk both underfloor and in overhead conveyance. Implementing a structured cabling system also reduces airflow congestion, which reduces power usage.

Another distinct advantage to a structured cabling system is that it allows for modularity, meaning connector changes can be made without having to remove horizontal or distribution cabling. For example, a chassis-based switch with 100BFX ports is connected to a patch panel using SC fiber optic patch cable. To upgrade the chassis and install new blades with LC ports, you no longer have to replace the entire channel as you would with a point-to-point system. Instead, the module within the patch panel is replaced. Underfloor and overhead cabling would remain undisturbed.

However, it should be noted that this method does add insertion loss to the channel because it adds more mating points. To offset insertion loss created by additional mating points, high-performance fiber optic cables should be used for implementation.

Conclusion and next steps for data centres

Data centers are experiencing the most significant change in cabling infrastructure since the introduction of fiber optic cabling. No longer is it a question of if data centres will migrate to 40/100G Ethernet, but when. Installing a high-performance, fiber optic structured cabling infrastructure is essential to a successful migration.
We covered why migration to 40/100G Ethernet is imminent, as well as the decisions data centre managers will need to make to prepare for implementation. There are several next steps you can take to prepare for this drastic change:

1. Determine your current and future data center needs, including throughput demand, data production rates and business-driven objectives. In what ways does your current data center infrastructure support and/or fail those needs?
2. Use this information to determine when your data center should migrate to 40/100G Ethernet.
3. Map out your current data centre infrastructure.
4. Use this map to create a plan for the hardware and cabling infrastructure upgrades necessary for migration.
5. Create a plan for migration, including internal communication strategy, budget, timeline, and roles and responsibilities of those involved.