LIGHTWAVE

EXECUTIVE SUMMARY

Ethernet in Data Center Networks

Kent Lusted, Principal Engineer, Intel; Co-chair, Ethernet Alliance High Speed Networking Subcommittee
Daniel Gonzalez, Business Development Manager, Anritsu Co.

JANUARY 19, 2022

KEY TAKEAWAYS

- Growing demand for information has created an explosion in data center traffic.
- Data center architectures are evolving to support increasing Ethernet transfer rates.
- Data center operators must optimize Ethernet media types for speed, power, reach, and latency.
- DPUs and IPUs are influencing data center design.
- Data centers are transforming into edge computing networks.
- To address requirements related to power and Ethernet speeds, data centers are turning to optical transceivers and high-speed breakout cables.
- Networking equipment manufacturers use Anritsu solutions to measure signal integrity.
- As distributed computing becomes more common, service levels and 400G network interoperability are priorities for data center operators.
- With multi-access edge computing and network virtualization, data center providers can maintain different SLAs for different applications.

in partnership with

Anritsu
Advancing beyond

Ethernet in Data Center Networks

OVERVIEW

Looking ahead, data center networking is expected to see continued evolution and innovation. This year, a growing number of data center operators will move beyond 100 Gigabit Ethernet toward 400G. Meanwhile, the IEEE 802.3 is beginning work on standards for what comes next via the new Task Force P802.3df. As data center network operators move to 400 Gigabit Ethernet and beyond, they will face new challenges such as signal integrity, network interoperability, and maintaining service level agreements (SLAs) for different applications. To address these issues, both data center operators and networking equipment manufacturers are turning to Anritsu test and measurement equipment.

CONTEXT

Kent Lusted and Daniel Gonzalez discussed Ethernet usage trends in data center networks, as well as the technologies helping operators meet growing bandwidth demands and verify network performance at high speeds.

KEY TAKEAWAYS

Growing demand for information has created an explosion in data center traffic.

Between 2010 and 2020, data center traffic experienced exponential growth. That growth further accelerated during the global pandemic. People working, learning, and playing at home increased the number of applications running in data centers and connectivity became even more critical. The growing numbers of users, access methods, and services have created a bandwidth explosion.

Figure 1: Voracious Demand for Data

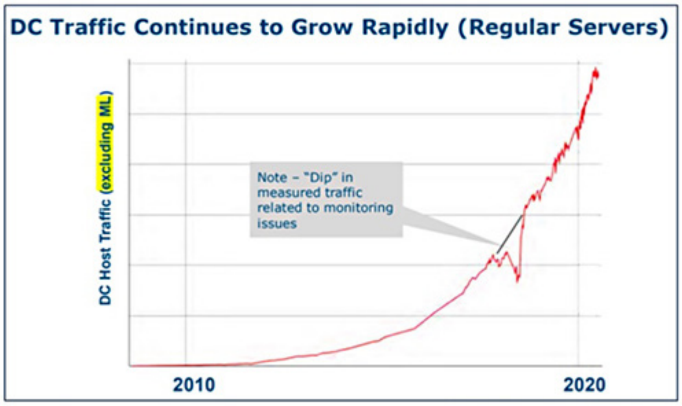

Courtesy: Cedric Lam, Google

Data center architectures are evolving to support increasing Ethernet transfer rates.

A cloud scale data center consists of a data center interconnect (DCI) that connects the data center building to others in the area, as well as routers, leaf and spine switches, top of rack (TOR) switches, middle of row (MOR) switches, and servers.

Historically, DCIs have supported 100 Gigabit Internet. Today, they are more likely to be 400 Gigabit Ethernet. For the server and compute element, data transfer rates in the past were 10 Gigabit Ethernet. Today, 25 Gigabit Ethernet is the norm and speeds in excess of 100 Gigabit Ethernet are expected. In the data centers of today and tomorrow, reach distances will decrease as a consequence of rate choices and the associated technology implications.

Figure 2: Next-Generation Data Center Networking

Courtesy: John D'Ambrosia, Futurewel

Ethernet in Data Center Networks

Data center operators must optimize Ethernet media types for speed, power, reach, and latency.

Data centers support a variety of media types which must be optimized for speed, power, reach, and latency. Higher Ethernet rates are forcing the industry to reassess long-held assumptions about:

1. **Power.** Data centers already have access to the maximum amount of power. As a result, data center designers must consider how to use available power more efficiently.
2. **Switch, router, transceiver, NIC architecture, and physical design.** These must be built using a holistic approach.
3. **Application latency.** Users engage in many different online activities ranging from search to ecommerce, social media, storage, and more. All of these applications have different user latency expectations.

As data centers optimize Ethernet media, they need to strike a careful balance between reach and cost. Each solution has a unique optimization point that addresses a specific data center challenge. For example, passive copper cables are relatively cost effective, but provide relatively short reach. Multi-mode solutions are more costly, but provide greater reach. With single mode fiber and coherent technologies, cost increases dramatically.

Figure 3: Reach vs. Cost Considerations

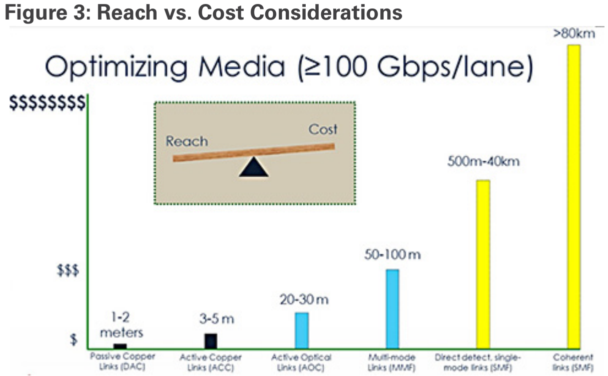

DPUs and IPUs are influencing data center design.

DPUs are hardware accelerators that offload networking and communication workloads from the CPU. With the exponential increase in network traffic to the network interface card in the server, increases in software-defined networking (SDN) have put a greater load on servers.

Infrastructure processing units (IPUs) accelerate and run the SDN and management software in hardware constructs away from server cores. Intel's IPU enables server cores to continue to run end customer applications, as well as provides system-level security, control, and isolation. The software framework offers a common look and feel for users, abstracting the IPU's features and capabilities. The Intel IPU is a hardware and software programmable solution.

Figure 4: IPUs Support the Data Center of the Future

> Ethernet is shaping today's and tomorrow's applications in the data center. At the same time, the applications running in data centers are influencing the development of Ethernet for today and tomorrow. It's a symbiotic relationship.
>
> *Kent Lusted, Intel and The Ethernet Alliance*

Ethernet in Data Center Networks

Data centers are transforming into edge computing networks.

Over the past three to five years, high-speed computing resources have shifted from centralized locations to a distributed model that uses high-speed, low-latency interconnections between resources. More recently, high-speed computing resources have pushed out as far as the network edge or as close to the user as possible. This trend has been driven primarily by increased demand for application support, Internet of Things (IoT) devices, and latency sensitive services.

As computing resources move closer to the edge, the latency key performance indicator (KPI) tightens. This KPI is application-service dependent. Latency affects the user experience for applications and must be considered when deploying Ethernet connects.

Figure 5: Evolution of the Data Center into an Edge Computing Network

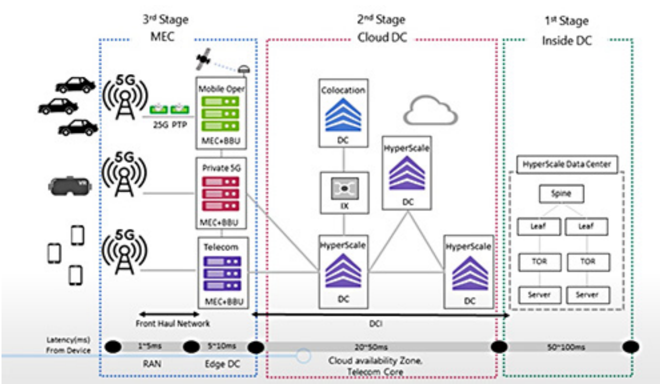

To address requirements related to power and Ethernet speeds, data centers are turning to optical transceivers and high-speed breakout cables.

Daniel Gonzalez made two observations about Ethernet trends within data centers:

- **Since data centers can't get more power, many are using optical transceivers to reduce power, while increasing the bitrate.** In hyperscale data centers, high-speed Ethernet optical interface advances have enabled providers to increase their leaf-spine connections up to 400G. As data center operators upgrade leaf-spine connections and deploy equipment from different vendors, they often face interoperability challenges. Not all 400G Ethernet optics are created equal and their performance on forward error correction (FEC) KPI thresholds varies.
- **High-speed breakout cables reduce costs, but they also present performance and distance tradeoffs.** High-speed breakout cables use the same pluggable interface as optics like quad small form-factor pluggables (QSFPs) or small form-factor pluggables (SFPs). However, they are fanout cables where one end supports the aggregate rate and the other end is a series of disaggregated interfaces. New breakout cables can support upward of 400G and reduce deployment costs. Data center providers must weigh the costs and benefits to identify the best fit based on application and deployment details.

Figure 6: Direct Connection vs. Breakout Cables

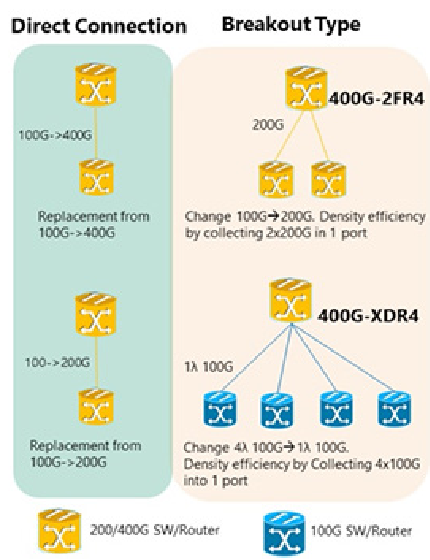

Ethernet in Data Center Networks

Networking equipment manufacturers use Anritsu solutions to measure signal integrity.

To test the latest versions of 100, 200, 400 and even 800 Gigabit optical interfaces, manufacturers use Anritsu instruments like the MP1900A Signal Quality Analyzer. The MP1900A Signal Quality Analyzer is a benchtop, high-speed bit error rate tester that supports the latest Ethernet rates up to 800 Gigabits per second.

Pluggable optical host interfaces have multiple lanes of synchronized or pulse amplitude modulation or PAM4 signals. Each one carries forward error correction (FEC) patterns that are generated in the pluggable optic and are converted to an optical waveform.

Optical waveforms are evaluated to determine if the signal integrity remains intact over the physical medium which could be optical, coax, or cable, depending on whether it is an active optic cable or a direct attach cable.

Figure 7: Anritsu Test and Measurement Solutions for 400G/800G Optics

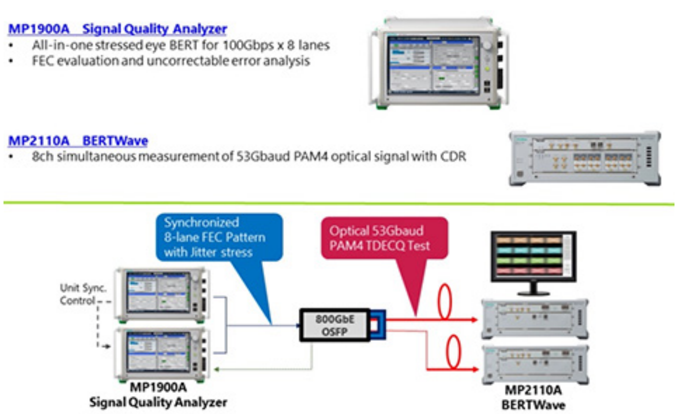

As distributed computing becomes more common, service levels and 400G network interoperability are priorities for data center operators.

Advances in Ethernet data transport speeds are improving network latency and throughput. As a result, data centers are shifting from centralized models to distributed and interconnected models.

Sharing high-speed computing resources over a converged network offers a variety of benefits. Rather than creating dedicated, redundant disaster recovery sites, it is now possible to use high-speed resources to distribute computing across multiple locations. Those locations can be interconnected to create pooled resources for computing-intensive applications.

Coherent, pluggable 400Gbase-ZR optical modules can transport 400G Ethernet over individual wavelengths over various optical network devices. The challenge, however, is ensuring 400G network interoperability with multi-vendor pluggable modules.

In addition, maintaining high service levels can be difficult in a distributed network with multiple demarcation points. Many data center operators use Anritsu test and measurement equipment to verify whether the network meets their key performance indicators. To ensure that the quality of the link remains consistent, KPIs can be measured frequently and across various demarcation points which may be managed by different providers.

Figure 8: Ethernet DCI KPIs Across Multiple Demarcation Points

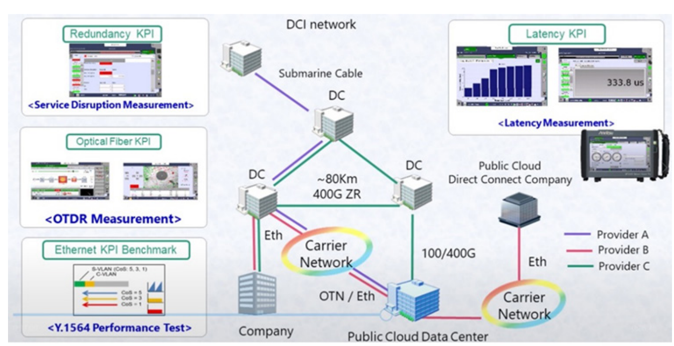

Ethernet in Data Center Networks

With multi-access edge computing and network virtualization, data center providers can maintain different SLAs for different applications.

Multi-access edge computing (MEC) refers to computing resources that are interconnected over a network, but distributed as close to the end users as the edge network allows. One key to making MEC work is splitting network resources into multiple virtual networks. While the physical computing resources remain the same, the Ethernet traffic is split into multiple logical networks.

As a result, it is possible to run virtually independent networks over the same hardware. These networks can each have its own SLA which is managed independently. For example, one network slice might support streaming real-time video, while another supports gaming or real-time analytics for autonomous vehicles.

The KPIs that are important in a MEC environment include latency, as well as metrics related to network synchronization over Ethernet transport such as time error measurement for Precision Time Protocol or PTP. These KPIs can determine whether each virtual network instance is meeting its service level agreement. Without metrics, it is impossible to guarantee the customer experience.

Figure 9: Ethernet MEC KPI for 5G Applications

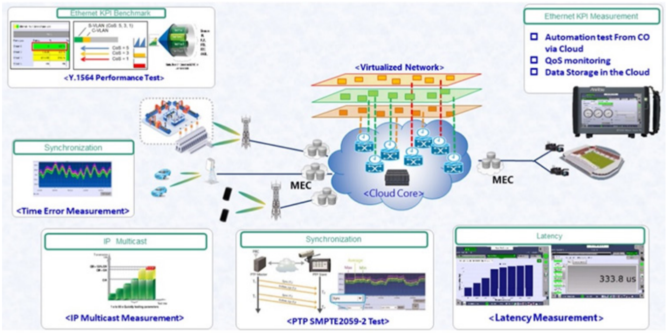

> Interoperability has become a very hot topic. While vendors follow standards, there are many ways to manipulate the registers when it comes to optics. And this goes for cables, as well. The importance of interoperability is knowing that you can cover a variety of different customer deployments and multi-vendor configurations."
>
> *Daniel Gonzalez, Anritsu*

ADDITIONAL INFORMATION

- The Ethernet Alliance is a global community of end users, system vendors, and component suppliers with a mission to promote existing and emerging IEEE 802 Ethernet standards and accelerate industry standards adoption by demonstrating multi-vendor interoperability. Ethernet Alliance expands the ecosystem that supports Ethernet development by facilitating interoperability testing, such as Plug Fests. It also provides interoperability assurance, collaborative interaction with industry organizations, global outreach across its membership, and thought leadership.
- IEEE 802.3 2020 Ethernet Bandwidth Assessment. Downloaded this report from the IEEE 802.3 website.
- Ethernet Alliance Technology Exploration Forum (TEF) 2021. Recordings and content from this event are accessible from the Ethernet Alliance website.

BIOGRAPHIES

Kent Lusted

Principal Engineer, Intel; Co-chair, Ethernet Alliance High Speed Networking Subcommittee

Kent Lusted is a Principal Engineer focused on Ethernet PHY Standards within Intel's Ethernet Products Group. He won an Intel Achievement Award in 2002 for his contributions towards delivering the world's first client and dual-port server Gigabit Ethernet controllers. Since 2012, He has been an active contributor and member of the IEEE 802.3 standards development leadership team and is currently the Vice-Chair of the IEEE 802.3ck 100G SERDES electrical interfaces Task Force. Kent is also co-Chair of the Ethernet Alliance Higher-Speed Networking Subcommittee that hosts interoperability events to verify new designs and validate the specifications used to create new products and solutions.

Daniel Gonzalez

Business Development Manager, Anritsu Co.

Daniel Gonzalez possesses over 20 years' experience in digital and optical transport testing, development, training, and execution spanning technologies including TDM, SONET, OTN, ATM, Carrier Ethernet, and Physical Layer Signal Integrity.

As a Business Development Manager for Anritsu Company, Daniel is responsible for providing technical support to sales, marketing and customers in North & South America. Daniel holds a B.S. in Telecommunications Management from DeVry University, is a member of OIF (Optical Internetworking Forum) Networking and Operations Working Group (WG), IEEE Communications Society (ComSoc) and Ethernet Alliance, and holds a Personal Certification for MEF (Metro Ethernet Forum) Carrier Ethernet 2.0. He also has several of his articles published in *Lightwave*, *ISE Magazine*, *Mission Critical*, and *Pipeline* publications.