DATA CENTER Frontier Special Report

Future-Proofing the Distributed Data Center with High-Speed Fiber Networks

Written by Paul Gillin

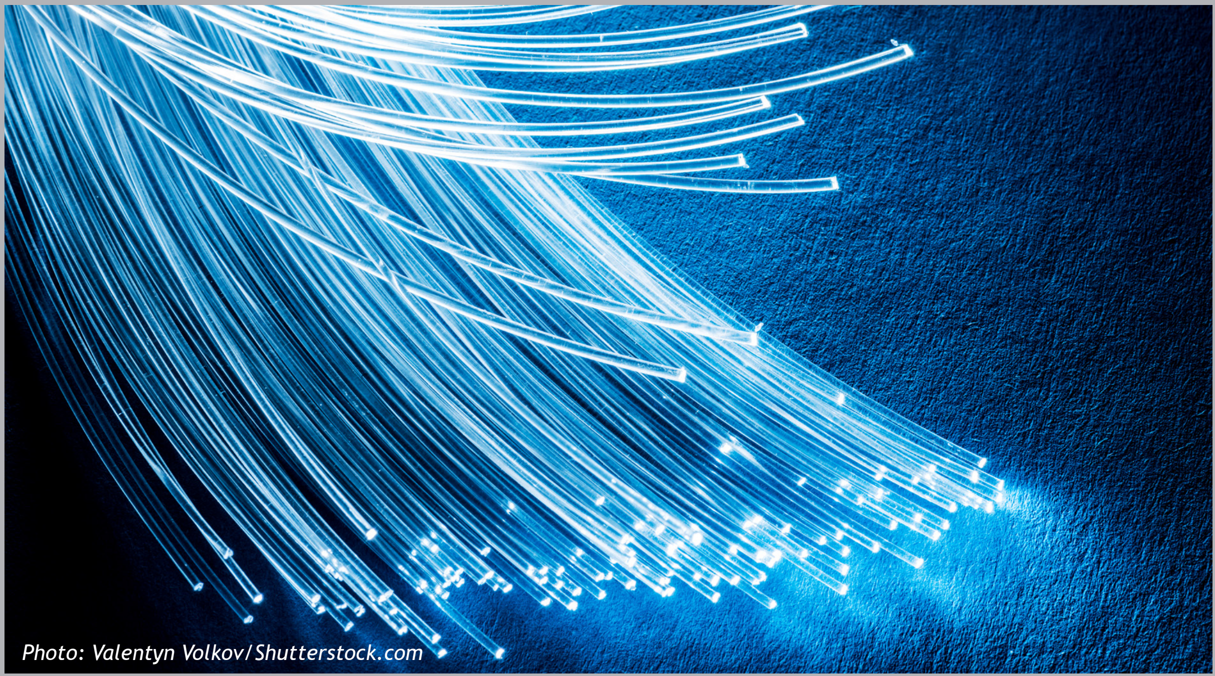

Photo: Valentyn Volkov/Shutterstock.com

Brought to you by

BELDEN

Contents

Introduction

The rapid growth of cloud services has created a popular impression that data centers are dinosaurs. In reality, the market is robust and growing. Gartner expects global data center spending to increase 5.3% in 2022 and to continue steady growth through at least 2024.

But future data centers will look quite different from their predecessors. Historically, data centers were large, centralized facilities intended to serve an entire company or a large geographic region. With the arrival of cloud infrastructure, that model is being replaced by a fabric of resources distributed across the globe to handle constantly shifting workloads according to the organization's need for availability, responsiveness, scalability, and security.

Growth in the use of intelligent devices—also called the internet of things—will drive that market to more than double to $1.06 trillion in 2030, according to Statista. This, in turn, is spurring the evolution of a new approach to infrastructure that moves computing closer to the points at which data is collected and disseminated, a topology called "edge computing."

Global Data Center Spending

9% increase in 2021 | **Continued steady growth through at least 2024**

Source: Gartner

Growth in the use of intelligent devices will drive the market to $1.06 Trillion in 2030

Source: Statista

This trend will grow as computing power becomes more distributed to meet the demands for faster response times, analytic workloads, and real-time processing for such purposes as remote equipment monitoring, predictive maintenance, real-time logistics, rapid health diagnostics, and securities trading.

Data growth creates complexity

The drive to digitally transform businesses to data-driven decision-making has organizations focusing intensely on how they can better capture and leverage data, even as volumes are growing at a 23% compound annual rate, according to International Data Corp. The more data that is generated, the more resources are needed to manage it.

Cloud platforms have given organizations a bounty of new choices about where and how to store and process that information. However, few are willing to put all their data processing eggs in a single cloud basket. On average, enterprises use 2.6 public and 2.7 private clouds, according to Flexera.

The result is that data processing facilities and resources are growing across the board.

In addition to the data center growth cited earlier, worldwide end-user spending on public cloud services is forecast to grow 20% in 2022 to $397.5 billion, according to Gartner. Colocation services are flourishing as well as organizations disperse processing power closer to the places where it is most needed and seek to take advantage of high-speed interconnect services. Grand View Research expects the data center colocation market will grow more than 13% annually to nearly $118 billion by 2028.

Even companies that make extensive use of public cloud services typically maintain some on-premises infrastructure to meet regulatory needs, provide backup and store highly sensitive data. Organizations are also increasingly shifting data and processing loads between owned infrastructure and multiple cloud providers as they seek the most robust services at the lowest prices.

The types of data organizations are capturing and using are also changing. Streaming data from sensors, video cameras, online transactions, and social media streams will make up more than 30% of all data generated in 2025, according to IDC. The new technologies that carry and consume these streams such as smart sensors, autonomous vehicles, high-speed 5G wireless networks, and robotics will place more intense demands on IT infrastructure. Latency will be a key factor in the successful deployment of many groundbreaking new applications.

Taken together, these factors will drive a shift away from centralized processing to multitier architectures, with data being managed both at the point of generation as well as further upstream in regional data centers and cloud servers. Edge computing will be a major factor in this transition. Gartner predicts that three-quarters of enterprise-generated data will be created and processed at the edge by 2025, up from just 10% in 2018. Small data centers and even individual servers in the field will gather this data, making decisions that require immediate action and uploading aggregated data to regional nexus points and the cloud.

Many edge and IoT devices will have special requirements to maintain uninterrupted operation in changeable and often adverse environmental conditions, including protections against the elements, power surges, and physical damage. Almost two-thirds of IoT adopters have reported that device failure rates in the field are much higher than in an office or data center, according to IoTForAll. The intelligence needed to maintain these distributed architectures must move into the network itself so that automated tools can leverage analytics and high-speed connectivity to detect and remediate problems in the field before they occur.

> An increasing number of business-to-business transactions are moving online.

E-commerce is also rapidly becoming a business staple. Online transactions were already growing rapidly before the onset of the COVID-19 pandemic, with more than 2 billion people purchasing goods or services online in 2020. The crisis accelerated that growth and e-commerce is now expected to make up nearly 22% of all retail sales by 2024, up from just 13.6% in 2019. An increasing number of business-to-business transactions are moving online as well reaching nearly $10 trillion last year.

Online customers are demanding. They expect immediate gratification and quickly navigate elsewhere if response times are slow. A difference of just four seconds in page load times increases bounce rates by 90%, according to Think with Google. The need to provide rapid response and frictionless transactions is pressuring companies to invest heavily in bandwidth, interconnection, advanced wireless services, and edge processing.

Evolving IT strategies

Taken together, these market dynamics demand that IT leaders evolve their data processing strategies with an eye toward more flexible architectures that move processing closer to the point of value. Future strategies are likely to include a combination of on-site infrastructure, public and private clouds, colocation services, edge data centers in the field or inside telecommunication carrier sites, and unattended smart devices.

While traditional online transaction processing workloads are likely to remain in the data center or private cloud, data may also traverse the network for such uses as analytics, reporting, and sharing with business partners.

The location and type of computing that is used will be driven by the characteristics of the workload. For example, applications like autonomous vehicles and streaming video delivery demand sub-five millisecond response times. In such cases, the processing is best distributed across multiple tiers with stream processing at the edge, a mid-tier control plane managing multiple devices and cloud servers aggregating and analyzing data at a high level.

This environment may include a combination of owned infrastructure, co-location services, carrier services, and public and private cloud. Operations and maintenance may be provided by dedicated staff and a network of service providers and contractors.

Other applications, such as real-time ad delivery and securities trading, require high-speed interconnection of the type provided by co-location services. Networks will need to be segmented to allocate dedicated bandwidth to latency-sensitive processes and multitiered compute fabrics will be deployed based upon required response times, workload characteristics, security needs, and other factors.

While traditional online transaction processing workloads are likely to remain in the data center or private cloud, data may also traverse the network for such uses as analytics, reporting, and sharing with business partners. For companies that operate internationally, this will create new demands on network infrastructure as well as processing considerations driven by data sovereignty regulations and cost.

Future-proofing the data center

With the pace of change only likely to quicken in the future, IT leaders should begin thinking now about how they can future-proof their infrastructure. Here are some of the factors they should take into account.

Resiliency

Distributed architectures inherently have more points of failure than centralized ones. Enterprises will want to build backups into every point in the value chain so that workloads can be shifted to alternative regions or between on-premises and cloud destinations. Cloud services present some attractive new opportunities in this area.

Workload portability will become an increasingly critical issue as processing becomes more distributed. By taking a standards-based approach to infrastructure deployment and making extensive use of virtualization and software container technology, IT organizations can gain the maximum flexibility to choose where data is processed and managed. Use open-source technologies where possible and exercise caution about becoming dependent on platforms that are tied too closely to individual cloud providers.

At a time when ransomware attacks are growing dramatically, continuous backup and disaster recovery are essential. There are many cloud-based data protection services that provide for continuous backup and rapid data restoration across on-premises and cloud infrastructure as well for software-as-a-service.

Security

The COVID-19 pandemic has demonstrated the weaknesses of perimeter-based security as virtual private networks were overloaded and employees flocked to cloud services often without adequate controls provided by the IT organization. Because hybrid work arrangements are likely to be with us for some time, CIOs need to rethink security practices across the organization.

Zero trust security uses principles like cloud-based identity, access management and network segmentation to limit attack surfaces and contain threats.

Distributed architectures are inherently less secure than those inside the data center because they create new attack vectors. The use of encryption, both for data in motion and at rest, is table stakes. Organizations are also increasingly examining "zero trust" options that redefine access around users and applications rather than physical devices. Zero trust security uses principles like cloud-based identity, access management and network segmentation to limit attack surfaces and contain threats.

Performance

As the volume of digital transactions grows, organizations will need to put robust e-commerce foundations in place. IT leaders will need to rely on a network of cloud regions, co-location providers, content delivery networks, and telcos to speed customer interactions. This will introduce security, reliability, and latency issues that must be accounted for in infrastructure planning.

Performance demands go beyond transactions. As the use of data analytics grows, so do data sets, reaching petabyte-scale for some machine learning applications. Moving data volumes of that size on a timely basis will require substantial increases in bandwidth for most organizations.

Cost

Cloud computing services have increased the latitude customers have in budgeting and cost allocation, but they have also introduced new expenses that aren't always apparent and that can add up over time. For example, egress fees, which charge for data taken out of public cloud storage, can be substantial when large data sets are involved. A 150 TB file transfer from a cloud platform costs an average of $1,000 in egress fees alone, and large enterprises typically move many times that amount on a typical day.

Software-as-a-service provides an abundance of choice, but it is up to their customers to watch for collateral costs created by unused or underused accounts as well as security risks like data exfiltration and configuration errors that leave data exposed. Ninety percent of respondents to a recent McAfee report said their organizations had experienced some security issues with cloud infrastructure, mostly attributable to user error.

People costs are usually the largest single line item in a CIO's budget and administering networks of distributed data processing equipment can require a lot of people. With data center operators already facing desperate skills shortages, automation should be applied wherever possible. Virtualizing infrastructure lays the foundation for automation and is now applicable across servers, storage, networks, desktops, and data. Moving to software-defined infrastructure insulates IT organizations against equipment obsolescence and can significantly reduce staffing needs.

Network topology

The distributed, heterogeneous enterprise IT landscape in the future will require far more robust and automated networking capabilities than were used in the past. Legacy networks were designed for static workflows. However, today's business climate is dynamic and mobile, subject to sudden spikes in activity and often unpredictable traffic patterns. New technologies like robotics, virtual and augmented reality, video streaming, and real-time equipment monitoring will also require organizations to allocate plentiful bandwidth to meet latency demands.

IT leaders should look to fiber networks constructed from a mesh fabric to give them maximum flexibility and automation. A mesh network links devices together through other branches or nodes. They are optimized for routing data efficiently between devices and clients, enabling organizations to maintain consistent connections with a minimum of vulnerability to disruption or unexpected changes in traffic flows.

New Ethernet fabrics based upon industry-standard protocols such as Shortest Path Bridging (SPB) and Transparent Interconnection of Lots of Links (Trill) enable all links in a network to be active with multiple equal-cost paths. This improves performance to optimize routing while permitting the network to quickly repair itself and route around outages.

Perhaps more importantly, fabric networks lay the foundation for software-defined networking (SDN) and network functions virtualization (NFV). SDN uses virtualization to improve flexibility and management by abstracting the control plane from the data forwarding functions in network devices. Essentially, all devices can be treated as a single pool and capacity allocated dynamically depending upon the needs of applications. This enables centralized management, control, automation, and policy enforcement across both physical and virtual networks.

NFV applies to services commonly provided by dedicated hardware such as routers, firewalls, and load balancers. It enables those important functions to be deployed at a lower cost on commodity hardware and managed centrally. In most cases, there is no need for on-site troubleshooting, which reduces travel and personnel costs and enables faster setup and remediation.

Virtualization benefits

Among the benefits IT organizations can realize from virtualization are:

- The ability to rapidly reconfigure networks in minutes rather than days
- Create sub-segments for performance and security purposes
- Automate many functions that formerly required human intervention

Virtualized networks also support multiple physical topologies and can be constructed incrementally over time.

Fiber advantage

Fiber-optic cable and connectors are the best way to future-proof network infrastructure because of their many advantages.

- **High bandwidth and transmission speed**
  Fiber broadband connections can reach 940 Mb per second with low lag time and can achieve even greater throughput using compression.
- **Extreme growth capacity**
  Wavelength-division multiplexing splits optical signals into different colors of laser light that can share a single strand of fiber. This enables nearly limitless capacity growth.
- **Performance over long distances**
  Because light signals degrade more slowly than electrical signals carried over a wired transport medium, a single fiber link can span between 20 and 60 miles without significant loss in performance and much farther using regeneration and amplification.
- **Fault tolerance**
  Ring and mesh topologies using multiple fiber cables provide nearly 100% protection against downtime.
- **Maximum data security**
  Light pulses are physically more difficult to intercept than electrical signals and attempts to tap into fiber-optic cables are more likely to be noticed by system administrators.
- **Support for multiple protocols and applications**
  Fiber can be used with any major network protocol.
- **Reduced design complexity**
  Because of fiber's capacity and high signal quality over long distances, there is less need for repeaters and redundant cabling. Fiber networks are also easier to visualize for management purposes.

Fulfilling your cabling requirements with products from a single vendor provides consistent quality, simplicity of purchase and installation, minimal waste, and a single point of support. As IT leaders expand their infrastructure footprint in the new age of distributed processing, fiber ties it all together in a package that will support growing data and processing volumes for years to come.