Mission Critical

Data center and mission-critical facility solutions

January / February 2020

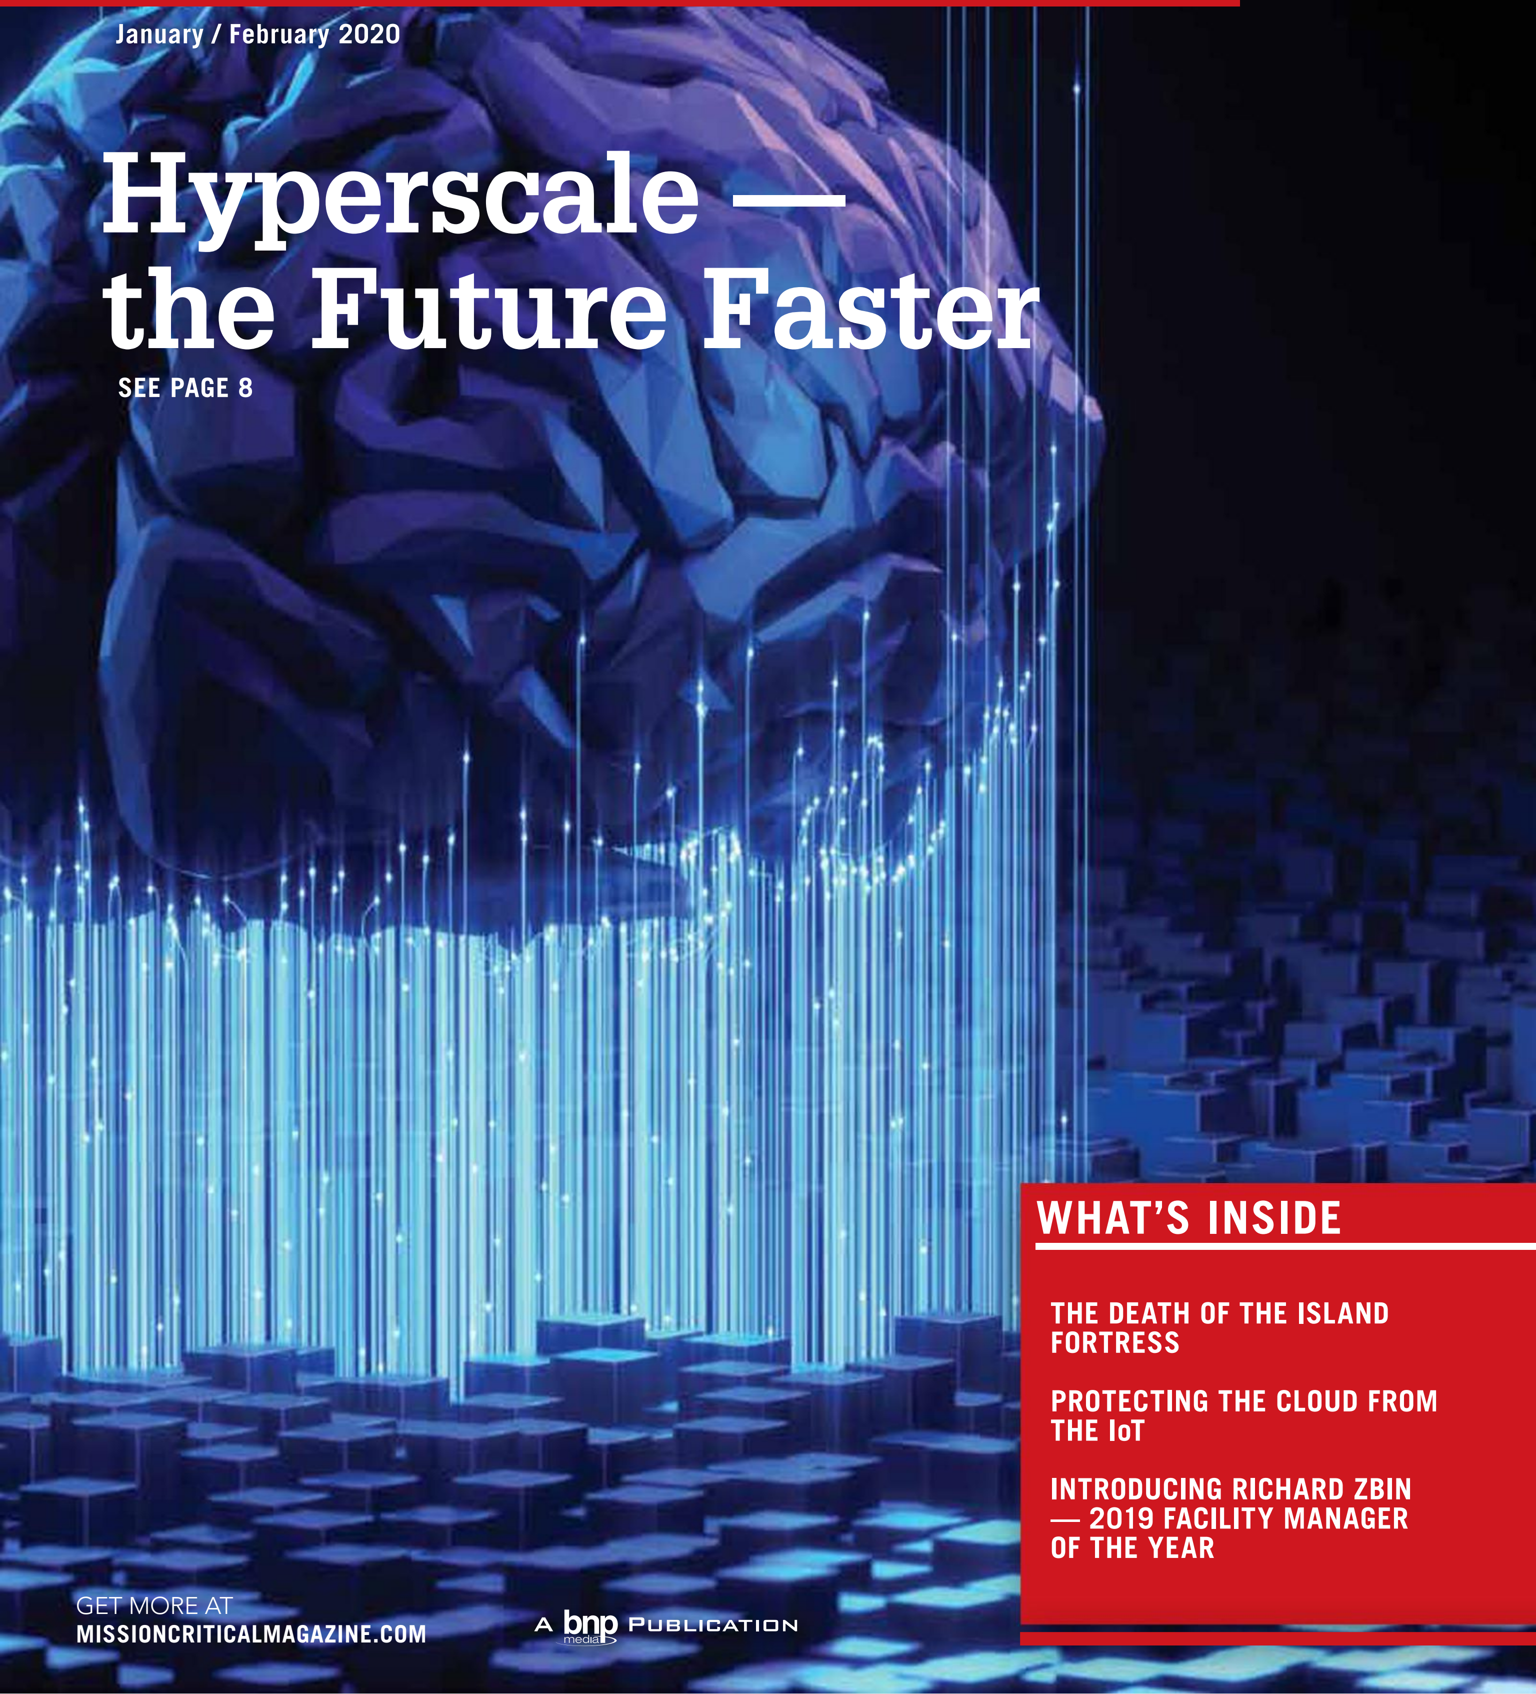

Hyperscale — the Future Faster

SEE PAGE 8

WHAT'S INSIDE

GET MORE AT
MISSIONCRITICALMAGAZINE.COM

A bnp media PUBLICATION

What Does a Future-Ready Data Center Look Like?

Scalability is the defining factor for success

Across the planet, end users and organizations are generating and consuming unprecedented volumes of data. But it can be challenging to comprehend just how fast data is growing worldwide and what the implications are for storing, managing, and accessing all of that data.

Many business, industrial, and consumer factors are driving this exponential data increase, including the Internet of Everything (IoE), AI, social media, and demand for streaming content delivery. According to International Data Corp. (IDC) predictions, by 2025, individuals will average close to 5,000 digital interactions daily versus their current 700 to 800, helping to fuel the expansion of the global datasphere from 33 zettabytes in 2018 to 175 in just seven years. Each of those zettabytes equals approximately 1 trillion gigabytes.

MASSIVE GROWTH, MAJOR CHALLENGES

This staggeringly fast pace of growth creates a daunting set of challenges for data management. As data storage needs outgrew traditional data centers, let alone endpoint devices, the cloud has increasingly picked up the slack. Many desktop computers no longer have storage, and, in the not-too-distant future, smartphones are likely to have increasing amounts of off-device storage, while everyone still expects instant on-demand access to any information they need.

Successful implementation of capabilities from 5G to AI to autonomous vehicles hinges on low or zero latency that is reliable, resilient, and sustainable. As global reliance on cloud services escalates, latency will become the key performance indicator.

These new demands are forcing the cloud to expand its capabilities into data delivery in addition to storage, putting pressure on IT infrastructure to evolve rapidly and dramatically. As just one example, Australia recently rolled out a brand-new national broadband network that is already insufficient because data streaming grew so rapidly during the seven years it took to build it.

HYPERSCALE MOVES FRONT AND CENTER

As the traditional data center network struggles to deliver the latency required to support accelerated cloud computing, hyperscale data centers have become the preferred option. Despite what the name appears to imply, hyperscale is primarily about scalability and resource optimization rather than the physical size of the facility.

	Deployed	Planned	Future
MESH	40G OM4	100G OM4F	400G SMF
Leaf-Spine	40G OM4	100G OM4	400G SMF
Server – Leaf (TOR)	10G AOC	25G AOC	40G AOC

FIGURE 1: Anticipated hyperscale network architecture growth.

Hyperscale data centers deliver four key advantages:

- Modularity in terms of design, construction, and performance;
- Reliability, functionality, and a yield greater than the sum of its parts;
- The ability to automate; and
- Economies of scale, such as lowering the cost per port, whether for a new build of a 35-MW system or adding 1 MW to an existing data center.

Let's take a closer look at typical hyperscale cloud data center architecture. Assume a data center currently deploys more than 30,000 servers in a mesh-plus-pod or pod-plus-leaf-spine structure, with three pod connections making up a mesh network. A typical spine switch supports up to 432 40-GbE ports (36 x 12). At a 3:1 ratio, 108 ports will be occupied by the uplink network (mesh level), meaning the maximum top of rack (TOR) switches in the pod cannot exceed 324. In reality, the number of TOR switches should be closer to 200 or 300 maximum. Each TOR switch can provide 48 10-GbE ports for downlink and four 40-GbE ports for uplink. A 4-by-40-GbE network link is set up between the leaf and spine and has four MTP12 channels with OM4 fiber to support the network.

But with the bandwidth consumption growing exponentially, 10 GbE will be insufficient. Plans already call for using 25 GbE in new hyperscale data centers, which will increase speeds between leaf and spine to 100 GbE, using MPO12 with OM4 to support the link. To be truly future-ready, however, requires the main stream to be 40 GbE from server to leaf switch, as shown in Figure 1.

POWERING UP FOR THE FUTURE

Most hyperscale facilities are extremely large, with power draws ranging from 10 to 70 MW. These power requirements obviously impact site location, energy infrastructure access, and sustainability. Minimizing hyperscale energy costs — the primary expense for data centers of any size — is a top priority. Even in smaller data centers, switches generate significant amounts of heat, making the design and implementation of cost-effective heating and cooling systems critical to both sustainability and market competitiveness.

As a result, the ability to capitalize on automation and smart technology plays a key role. The more thermostats, lights, and other disparate energy-related components can be automated and centrally controlled, the better the savings potential.

Planning for a robust wired and wireless network is also a critical foundational element for successful data centers of the future. Although wireless may seem to epitomize the future, a hybrid model will be at the network core. Wireless service providers have to backhaul the network, and from a signal transfer bandwidth reliability perspective, fiber and optic cabling offer the strongest choices.

The more 5G expands, the more cable is necessary, putting power over Ethernet (PoE) systems into the limelight. Power is increasingly shifting from 110-V and 220- to 240-V AC to USB, which can originate or be delivered in a PoE system. PoE offers numerous advantages, including the fact that it is low voltage, is easy and inexpensive to install, and can carry data and energy power at the same time.

ADAPT OR DISAPPEAR

Two competing demands will define future data center architecture: storage and latency. As storage continues to move from devices to centralized facilities, edge computing is likely to explode. With data being captured from so many different sources, nodes are going to become as common as streetlights to ensure nothing gets lost.

This new reality creates both an opportunity and a threat to existing data centers, especially those that are 1 MW or below. The good news is that they have a valuable role to play in working together with edge computing and hyperscale network architecture to successfully meet storage and latency requirements — assuming they can innovatively adapt to the new reality. That means being designed or upgraded to deliver hyperscale speed, low costs, and scalability.

In fact, the future may actually see a wave of data center towers located in city centers — smaller data centers built within former office buildings.

RETHINKING THE DATA CENTER

Bottom line, scalability is going to define future data centers. First and foremost, this requires design-based approaches that take a holistic and long-term view of energy consumption, data storage, management, and delivery needs across the business, industrial, and consumer spectrum. Second, cost-effectively integrating and enhancing scalability depends on the ability to build it on products and technology platforms that can be mixed, matched, rearranged, and reused to meet constantly evolving needs.

The move to future-ready data centers is underway, but more work is needed to ensure the successful transition of the existing wireless network, hardware technology, and network architecture to manage the integration of all new use cases and diverse requirements. As transition ramps up quickly, data centers must ensure they are working with an innovative technology provider that understands their needs and can help them develop future-ready, next-generation data center solutions. ■

Lauren Straub is director of sales and technical services at Molex Connected Enterprise Solutions, the business unit focused on structured cabling products and IoT solutions.

John Berryman is group product manager at Molex Connected Enterprise Solutions. He is responsible for a global team of product managers focused on product planning and product marketing.