



WHITE PAPER

Challenges and Opportunities of Transitioning to 400G

State of the Union: 400 Gigabit Ethernet

Challenges and Opportunities of Transitioning to 400G

State of the Union: 400 Gigabit Ethernet

Coming soon to a network near you: 400 Gigabit Ethernet (400GbE). The standards are in place, thanks to the work of the IEEE, which approved the relevant 802.3bs standard in December 2017. The specifications are impressive, offering a data rate of 400 Gbps at Layer 2 with transmission ranges up to 100 meters with multimode fiber, or up to 10km over single-mode fiber.

Spirent was a leader in developing the 400GbE standard, and has [offered 400GbE solutions since 2014](#), aiding in the early testing of prototype equipment. In addition, several Spirent engineers participated in the 802.3bs process itself, with Spirent an active participant in IEEE standards working groups, involved in High Speed Ethernet contributions for over a decade.

With the standards already in place (IEEE 802.3bs also addresses 200GbE specifications), network operators are clamoring over the opportunity to add high-density network processing into their fiber and equipment racks—and for enabling massive throughput through a single switch.

Standing in the way are reasonable concerns, both on the original equipment manufacturer (OEM) and operator side. For OEMs, there are worries about interoperability, transceiver reliability, basic connectivity, error rates and power/cooling. Operators share those same concerns about equipment they are previewing and evaluating, as well as about traffic patterns and the best ways to deploy 400GbE in their networks and data centers.

We are now seeing operator demand driving OEMs—and the industry as a whole—to create a **vibrant ecosystem of performant, reliable, affordable and interoperable 400GbE**.

The Rush to 400GbE

It's been only a few years since the approval of the previous pair of Gigabit Ethernet standards, for 40GbE and 100GbE. In today's world, that's an eternity, especially since the rapid rise of two separate technologies: Internet of Things (IoT) and cloud computing in all its forms (such as SaaS, IaaS and PaaS).

While the amount of data generated or consumed by most IoT devices is small, the aggregation of all that traffic—whether it's a few bytes of telemetry, or streams of high-resolution video—is tremendous. That data is forging new paths across the Internet and enterprise networks to servers coping with an unprecedented number of connections. Those servers may be resident in a traditional data center, a colocation facility or a cloud service provider.

Cloud providers are experiencing their own boom, and not only because of IoT traffic. Many hyperscalers, like Amazon, Google, and Facebook, not only have their consumer-facing traffic, but also host business services including IaaS and PaaS offerings from Amazon Web Services, Google Cloud Platform, Microsoft Azure, as well as Alibaba, Baidu, IBM, NTT, Oracle, SAP and Tencent.

According to [Synergy Research Group](#) (January 2021), there were 111 new hyperscale data centres opened in the last eight quarters, with 52 of those coming onstream in 2020 despite Covid-19 causing a few logistical issues. When building out a new data center, a hyperscale lab, a large ISP or even an enterprise, are going to want the greatest density, in order to cram the most computing and networking potential into the smallest footprint—and that means 400GbE. It offers four times the density of 100GbE, not only in fiber, but also in transceivers, switches and racks. Even if the initial acquisition cost is higher, the long-term payoff for 400GbE is significant—*as long as 400GbE works as anticipated*.

For Data Center operators and Cloud providers, a significant amount of the traffic remains purely within the facility, as real and virtual servers, *n*-tier applications, clusters, application management, security and other applications operate. Even a data center with little traffic flowing internally or externally may have tremendous needs for raw bandwidth capacity and bandwidth aggregation.

Thus, it's our experience that nearly every network owner or operator building a new facility—or modernizing an existing one—is looking at 400GbE and beyond. It simply makes little sense to go with 100GbE, unless there truly are issues with cost or with equipment testing, certification and availability on the timeline needing for the project. Bear in mind that devices that support 400GbE also interoperate with lower capacity devices—200GbE, 100GbE, 50GbE and even 25GbE — which makes 400GbE ideal for accommodating incremental modernization projects.

Challenges of Transitioning to 400GbE

This isn't the first rodeo for most OEMs and network/data center operators. They've purchased 10GbE, 25GbE, 40GbE, 100GbE, and now 400GbE equipment. Many of their concerns are reasonable extrapolations of older challenges. Given the rapid growth of networks and data centers, and high demand for the greatest possible bandwidth and port density, there is a lot of pressure on the industry to address these questions quickly.

On the physical side, the industry is facing the following challenges:

Interoperability. While the IEEE standard is approved, experience shows that real-world multi-vendor interoperability requires addressing numerous details in product specifications, design and manufacturing. This will require not only attention to detail, but real-world tests in High-Speed Networking plugfests (like those at the University of New Hampshire, where Spirent participates and assists in the program), but in hyperscale test labs and prototype networks. In short: There's good progress, but the industry's not there yet.

Transceiver reliability. We are seeing issues where optical transceiver manufacturers are finding it difficult to stay within the power budget of the device, resulting in heat issues, which cause problems with fiber alignment, optical distortions and packet loss. This reliability issue stems, in part, from the old design of the Quad Small Form-factor Pluggable (QSFP) transceiver form factor use in 400GbE — which was designed for 40GbE. This is a real challenge at the physical layer. There are similar challenges for the newer modules used in 400GbE systems, such as the CFP8 and QSFP-DD form factors.

Optical coding. IEEE specification moved from a familiar type of coding used for 100GbE and below, called NRZ (binary non-return to zero), which used two levels of optical signaling, to PAM-4 (pulse amplitude modulation), with four levels of signaling. Encoding and decoding four levels requires significantly more sophistication, as well as more expensive hardware. PAM-4 is also more sensitive to noise on the optical line. PAM-4 is a big issue for even a single manufacturer to master which makes multi-vendor interoperability even more challenging.

Excessive link flaps. Whenever there are a series of errors or other failures on an optical connection, both transceivers need to perform AN/LT auto negotiation and link training before data can begin flowing again. When such faults happen often—several times per minute—it's referred to as link flap and can dramatically affect throughput. With mature optical technologies, link flap is rare, and is usually caused by a bad cable, configuration error or a defective transceiver. With new technologies such as 400GbE, link flaps can also happen due to heat issues, design issues with the transceiver modules, or problems with the switches themselves.

About Spirent

Spirent Communications (LSE: SPT) is a global leader with deep expertise and decades of experience in testing, assurance, analytics and security, serving developers, service providers, and enterprise networks. We help bring clarity to increasingly complex technological and business challenges. Spirent's customers have made a promise to their customers to deliver superior performance. Spirent assures that those promises are fulfilled.

For more information visit:
www.spirent.com

The Importance of Testing

With 400GbE, the industry—including OEMs and network/data center operators—are still addressing basic connectivity issues, managing everything from transceiver reliability to link flap, from excessive frame error to excessive packet loss. In some cases, links aren't coming up at 400GbE.

Testing is the fundamental building block to delivering reliable, high-performing systems and starts with validating the link's ability to pass up to line rate traffic as well as testing the functionality, performance, scalability and Quality of Experience (QoE) of the upper-layer engines that deliver services.

Spirent is ready to help, with solutions for OEMs, integrators, hyperscale data center, cloud providers and data center operators. Spirent 400GbE test platform creates and evaluates test beds at the scale needed to thoroughly test a 400GbE system, while work is well under way with technology partners to demonstrate interoperability for the next Ethernet rate, 800GbE.

By processing billions of real-time results to validate tests and identify problems, these High-Speed Ethernet test solutions provide engineers immediate feedback to debug problems and accelerate development. The solution also delivers more results with tight correlation, and more information to find those obscure bugs.

Spirent has the experience with 400GbE to handle the industry's toughest challenges and is already on its fourth generation of 400GbE products, winning multiple awards over the years, most recently at the Lightwave Innovation Reviews and Tokyo Interop.

With comprehensive test coverage, hands-on test experience, and deep technical expertise, Spirent solutions help answer questions faster and in a single test run, where multiple runs are necessary with other test tools, whether in cloud computing, streaming services, data centers or terabit routers. No matter the high-speed Ethernet scenario, Spirent is ready.

For information, visit www.spirent.com/HSE