A SPIRENT IMPACT REPORT

When Worlds Converge

The Future of High-Speed Ethernet across Data Center, Telecom, and Enterprise Networking

spirent™

A SPIRENT IMPACT REPORT

Converging on A New Era of HIGH-SPEED Ethernet

New opportunity, unprecedented demand, and urgent innovation are defining Ethernet's most consequential transformation in decades.

We are pleased to introduce our first annual impact report tracking the latest trends, demand, and developments shaping the data center networking and high-speed Ethernet (HSE) market.

With more than two decades of Ethernet testing experience with global stakeholders, our breakthrough work spans deals with public cloud hyperscalers, data center providers, service providers, enterprises, government and military, network equipment providers, and chipset OEMs globally.

Our insights and predictions are based on valuable perspectives from the deals, strategies, and innovations shaping the high-speed Ethernet market. This includes the 340 HSE customer engagements Spirent supported around the world in 2023, with more than 25% of those being 400G and 800G engagements.

AI Advancements Mean More Ethernet, Everywhere

As all eyes focus on the power and promise of AI, there has perhaps never been more pressure to move faster, push the boundaries of speed, and relentlessly pursue every competitive edge available in this market.

Traditional north-south traffic data center strategies are quickly giving way to east-west approaches aimed at serving a drastically expanding market base. Stakeholders are taking new paths and traversing unproven terrain in an attempt to squeeze more capacity out of existing technologies for increased performance. Today, the rate of change is outpacing the ability of standards to keep up. Enterprises are not waiting on future developments to make progress and telecom operators are throwing out traditional playbooks to meet customers where they are in ambitious deployment cycles.

In Spirent's inaugural report covering this dynamic market, we proudly present a comprehensive look at key drivers, market impacts, and our predictions for what comes next.

Aniket Khosla, *Vice President of Wireline Product Management, Spirent*

What's Inside

High-Speed Ethernet Market Momentum

High-Speed Ethernet Market Update

Port speeds are continuing to evolve, and adoption remains strong across all levels of speed, with even 10G and 25G still experiencing steady demand among SMEs and 5G cell site deployments. Ahead of traditional demand curves, markets are already looking to 1.6T Ethernet to pursue AI-driven opportunities as soon as next year. (Our companion report, "Turning Ideas into Action for 1.6T Ethernet" details the current state of this dynamically evolving ecosystem.)

10/25G	100G	400G	800G	1.6TE
Steady demand driven by long-tail of SEM's and by 5G cell site deployments	Continued demand driven by service provider edge networks and medium sized enterprises	Waves of growth initially by hyperscalers refresh cycles then large enterprises, tier-2 and tier-3 cloud service providers & large Telco service providers (IP Core upgrades)	Growth driven initially by hyperscalers support for AI	Early demand driven by hyperscalers growing support for AI
>90M ports est. to ship between 2024-2026	>65M ports est. to ship between 2024-2026	>40M ports est. to ship between 2024-2026	>42M ports est. to ship between 2024-2026	>8.5M ports est. to ship between 2025-2026

70+ MILLION HSE ports shipped in 2023

with volume expected to explode to

240+ MILLION ports shipped between 2024 and 2026.

Source: Dell'Oro Group

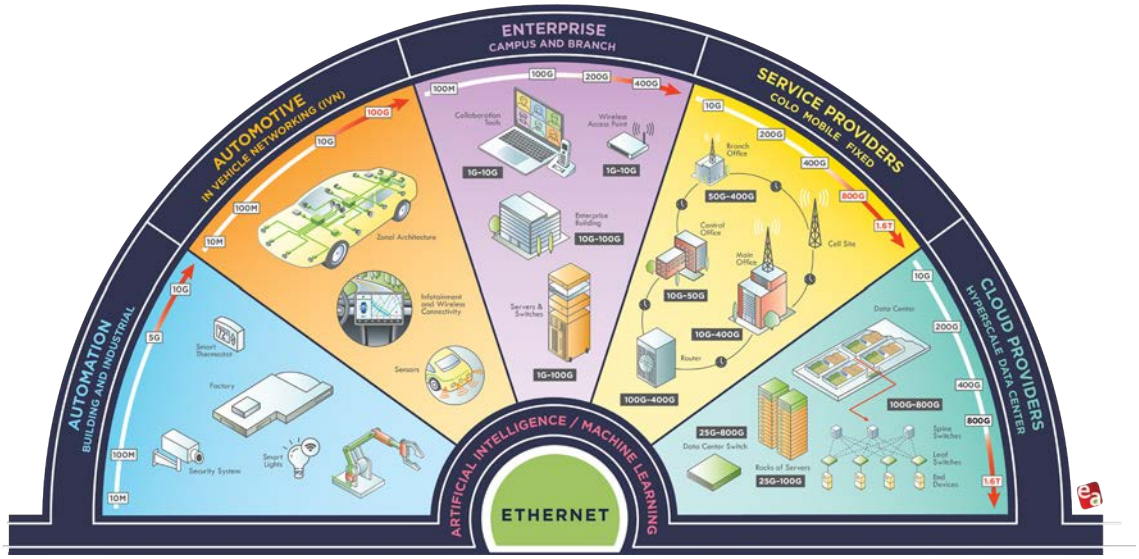

Source: Ethernet Alliance

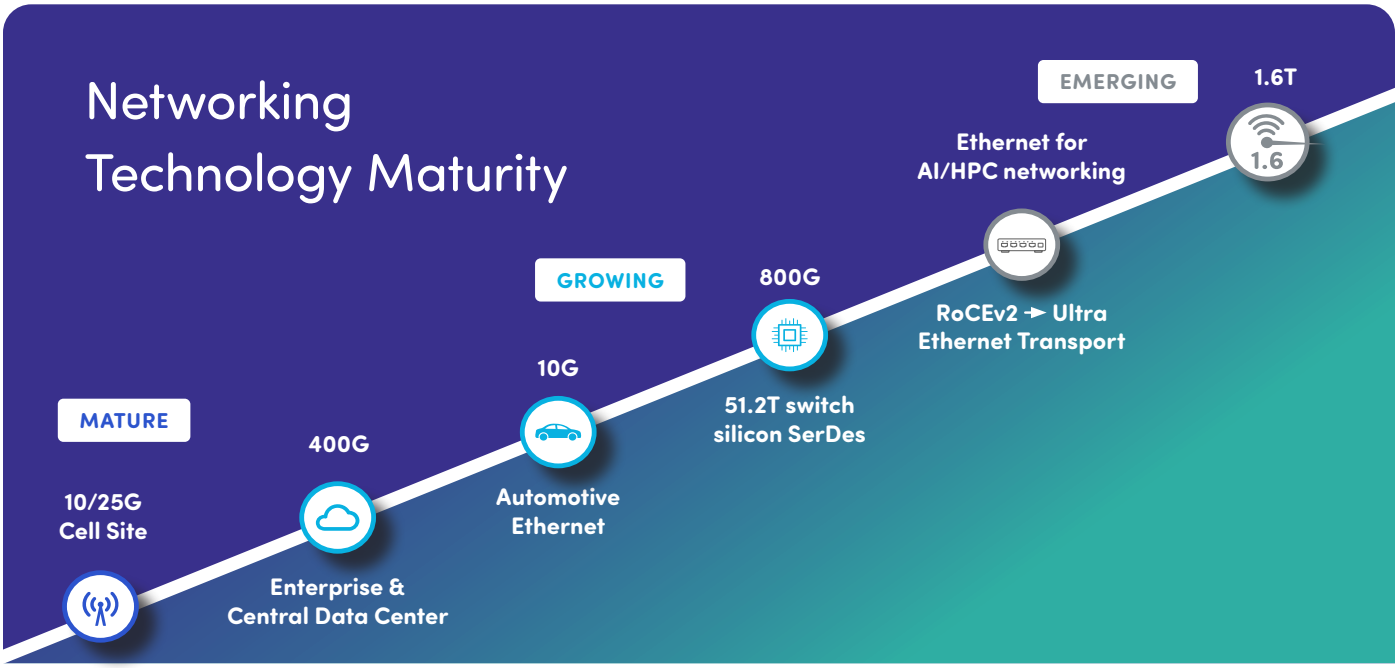

1.6T Ethernet

Projected exponential growth in data center traffic is driving 1.6T Ethernet research. Although IEEE 802.3dj finalizes in 2026, 2024 baseline features enable early development of silicon, optical components, and transceivers.

Ethernet for AI and HPC

Growing demand for standardized Ethernet-based IP networks for AI and HPC is initially focused on RDMA over Converged Ethernet (RoCEv2) and expected to evolve towards the new Ultra Ethernet Transport (UET) standard.

800G

Initial deployments will be led by hyperscalers to support AI in data centers with availability of 51.2Tbps switches. 800G offers higher bandwidth, lower latency, enhanced energy efficiency, and more connections, future-proofing data center interconnects for years.

400G

Adoption is shifting from cloud hyperscalers to large enterprises, tier-two and tier-three cloud providers, and large telcos for IP core upgrades. Large enterprises view 400G as future-proof, offering a simpler path to 800G and higher transmission rates at lower costs.

Evolution of 800G Ethernet Standards & Beyond

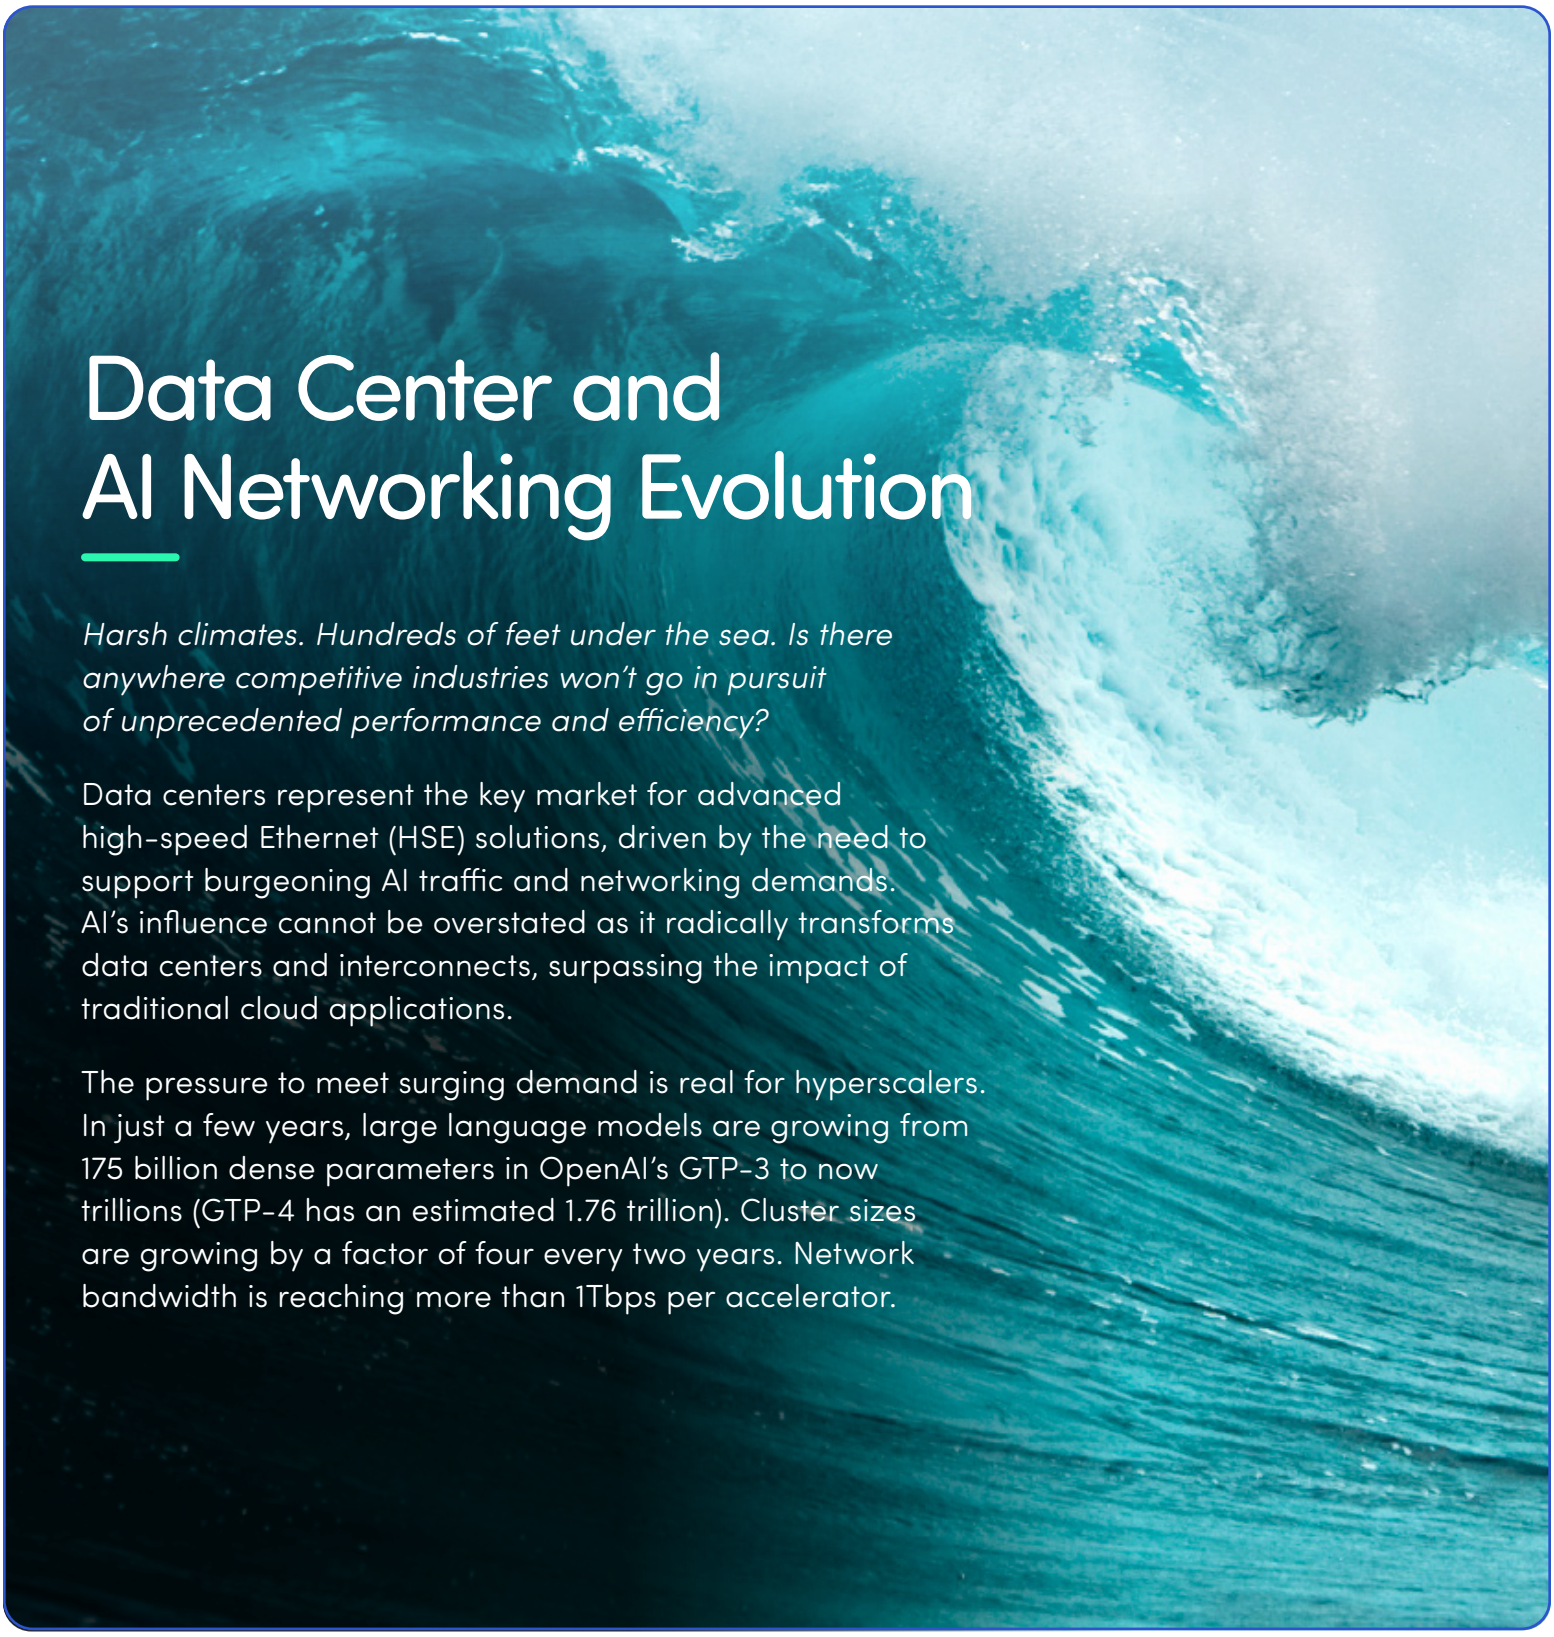

Data Center and AI Networking Evolution

Harsh climates. Hundreds of feet under the sea. Is there anywhere competitive industries won't go in pursuit of unprecedented performance and efficiency?

Data centers represent the key market for advanced high-speed Ethernet (HSE) solutions, driven by the need to support burgeoning AI traffic and networking demands. AI's influence cannot be overstated as it radically transforms data centers and interconnects, surpassing the impact of traditional cloud applications.

The pressure to meet surging demand is real for hyperscalers. In just a few years, large language models are growing from 175 billion dense parameters in OpenAI's GTP-3 to now trillions (GTP-4 has an estimated 1.76 trillion). Cluster sizes are growing by a factor of four every two years. Network bandwidth is reaching more than 1Tbps per accelerator.

Supporting AI is a massive undertaking with equally large costs. Players are now spending more on data centers than the entire 5G market over the past two years, with investments reaching $116 billion in 2024.

Dell'Oro has projected data center Capex will increase from $260 billion in 2023 to more than $500 billion in 2028.

The scale of growth is incredible, with analysts estimating over 2.2 million graphics processing units (GPUs) will be shipped during 2024 for data center AI processing with costs ranging from $15-$40K per GPU.

Effective and efficient power management of AI data centers is essential for meeting sustainability goals and maintaining profitability. McKinsey & Company predicts that data center power consumption will reach 35 gigawatts by 2030, enough to power 26.2 million homes. The Electric Power Research Institute sees data centers draining between 4.6 to 9.1% of U.S. electricity in 2030.* EPRI also finds that "At 2.9 watt-hours per ChatGPT request, AI queries are estimated to require 10 times the electricity of traditional Google queries."

The resulting outsized need for cooling is inspiring innovations such as building underwater data center facilities or locating them in nordic country ice, capping GPUs, using green coding, and even reverting to small language models.

** EPRI, Powering Intelligence: Analyzing Artificial Intelligence and Data Center Energy Consumption report (May 28, 2024)*

Meta is acquiring 350,000 GPUs to build out their clusters, with their new Llama3 training cluster hosting 24,576 GPUs. That equates to 16 GPUs per rack with 1,536 server racks per cluster. [source: Meta]

Google's AI-focused A3 supercomputer is incorporating around 26,000 Nvidia H100 Hopper GPUs. [source: HPC wire]

Amazon is scaling up to 20,000 Nvidia H100 GPUs in each UltraScale cluster. [source: Nvidia]

Oracle's OCI Supercluster scales up to 32,768 Nvidia A100 GPUs across 4,096 compute bare metal instances. [source: Oracle]

ByteDance is estimated to be receiving a total of 100,000 Nvidia A100 and H800 GPUs during 2023. [source: Data Center Dynamics]

Microsoft is estimated to already have over 600,000 GPUs with plans to scale to 1.8 million by the end of 2024. [source: Forbes]

Relentless advancement

80x bandwidth increase over 12 years

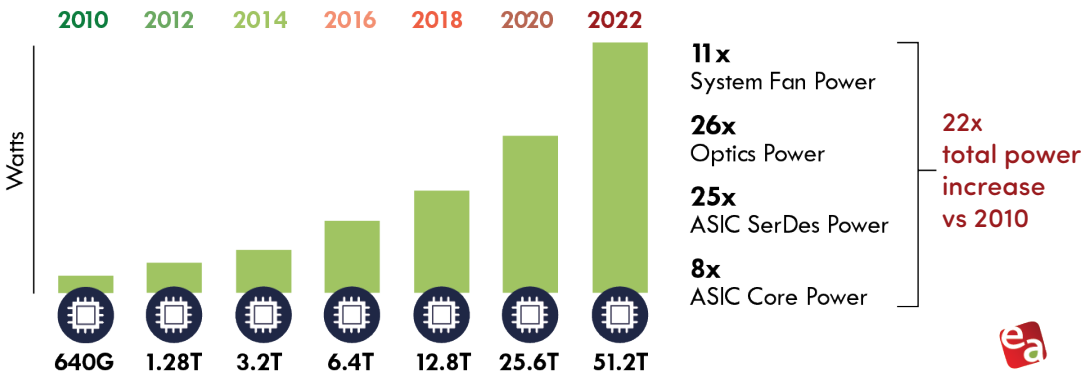

Sustainability / power consumption model, courtesy of Ethernet Alliance

AI's Impact on Data Center Networking Designs

GenAI is becoming a rapidly growing consumer of data center resources and workloads.

The impact is so great that hyperscalers must either build a single new architecture that can handle current cloud and enterprise workloads plus AI (Google's approach) or completely rearchitect AI data centers.

Data center providers must evolve their architectures to cope with AI traffic growing tenfold every two years and accelerator GPU learning nodes scaling into the tens of thousands to manage the rapid increase in model parameters.

A new network called the "back-end network" has evolved with the sole purpose of handling data movement between GPUs.

While the traditional front-end Ethernet network must grow to ingest massive model training data sets, it is the back-end network that bears the brunt of escalating workloads from AI inferencing working on the new data sets. Those back-end data center training workloads require a multitude of GPU or other xPU hardware accelerators to scale up AI computing clusters.

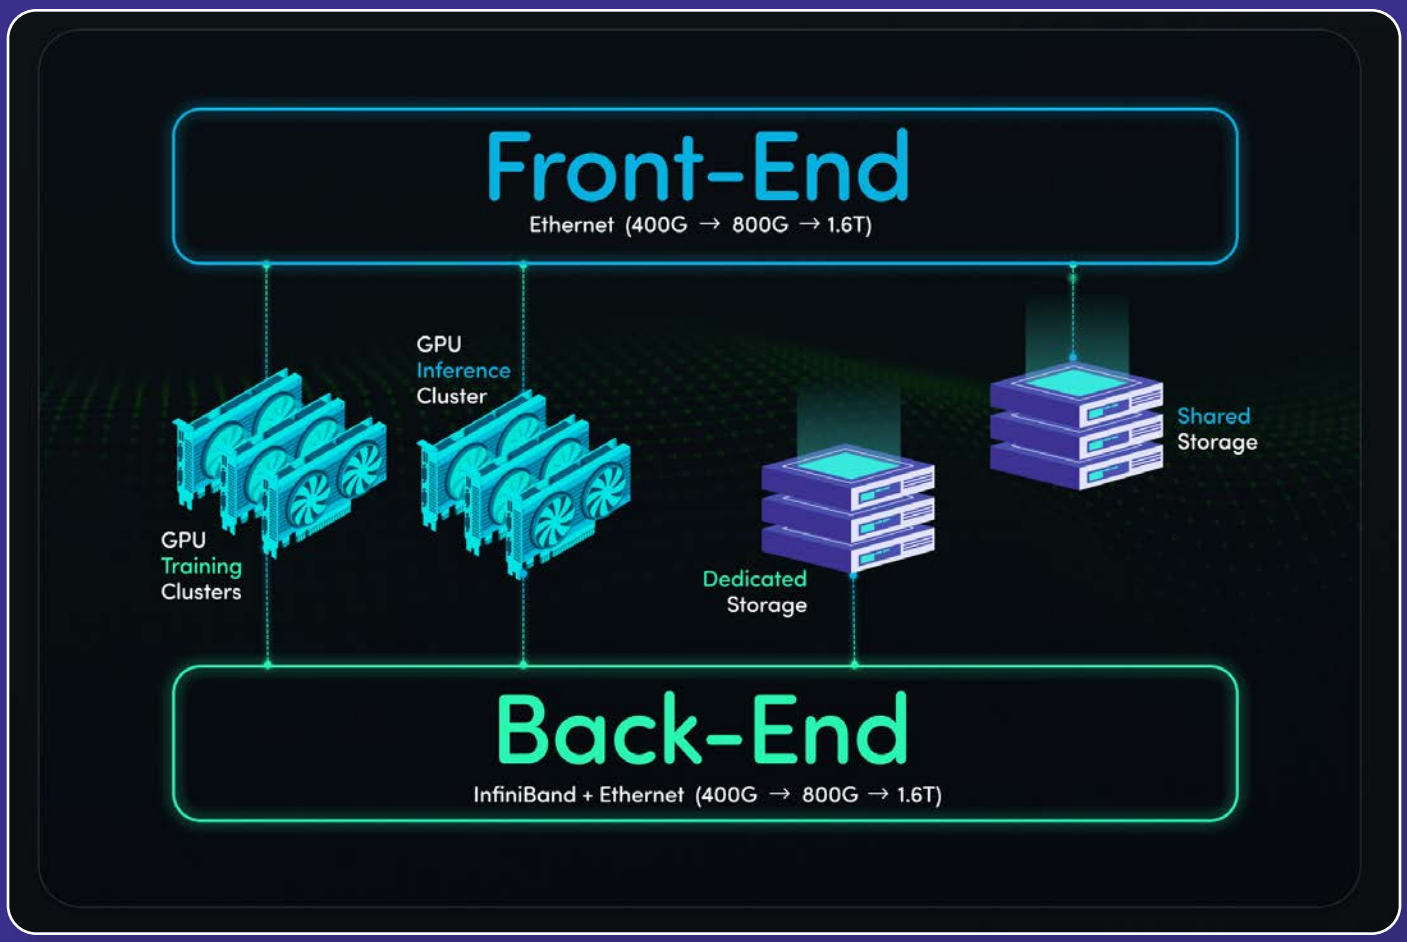

Connecting these accelerator nodes in large GPU clusters requires a back-end data center network fabric, which differs from the traditional front-end network used mostly to connect general purpose servers.

The back-end infrastructures demand a separate, scalable, routable network to interconnect thousands and even tens of thousands of xPUs to support AI training and inference.

To support AI, the back-end network fabric needs to provide:

EXTREMELY HIGH THROUGHPUT to handle compute-intensive and data-heavy workloads

EXTREMELY LOW LATENCY to quickly process workloads through many nodes

ZERO PACKET LOSS to lower latency

MASSIVE SCALABILITY for billions of parameters across thousands of nodes

Data Center Architecture Example

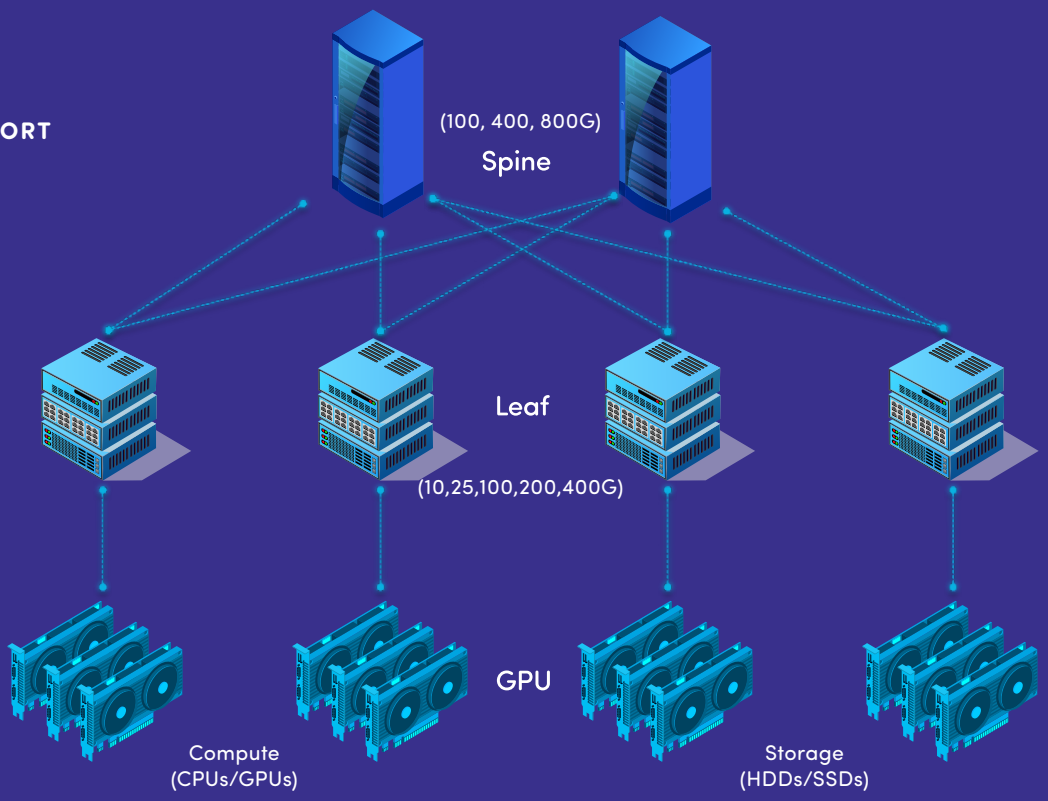

Spine-Leaf network architecture *is being adopted for back-end networking needs of east-west traffic prevalent in AI training and inference. This is flattening the traditional data center three-tier topology into a two-tier architecture offering higher bandwidth utilization, greater scalability, lower latency, predictable network delay and parallelism. The spine layer includes switches for routing and forwarding, which is the backbone of the network. The leaf layer includes switches for connecting compute nodes (GPUs, CPUs, etc.) and storage systems (SSDs, HDFS, etc.) The AI back-end network operates as a high-performance, scalable, and reliable network infrastructure optimized to handle the demanding requirements of AI workloads, ensuring efficient data processing, model training, and inference.*

These demands necessitate that new AI data center designs meet specific requirements, such as:

SENDING THE LARGE ELEPHANT FLOW TRAFFIC output of every GPU training result to all other GPUs in the cluster for harmonization

INCREASED NETWORKING BANDWIDTH to handle the large volumes of data exchanged

INCREDIBLY LOW AND DETERMINISTIC LATENCY and lossless packet delivery required for latency-sensitive front-end inferences, and to prevent back-end GPU timeouts and synchronization issues

DISTRIBUTED TRAINING WORKLOADS requiring parallelism techniques to synchronize east-west data across nodes

MULTIPLE AI TRAFFIC PATTERNS including AlltoAll, RingAllReduce, AllGather, and Broadcast, and more

> “AI network fabric spending is poised to reach $11.33 billion and grow by 27.1% CAGR through 2028.”
>
> **-GARTNER®**

Gartner®, Forecast Analysis: AI Network Fabric, Worldwide, By Naresh Singh, 29 April 2024. GARTNER is a registered trademark and service mark of Gartner, Inc. and/or its affiliates in the U.S. and internationally and is used herein with permission. All rights reserved.

It's impractical for data centers to meet these requirements by simply adding more racks. Instead, new data center architectures are essential.

During 2023, leading data center providers increased Capex by between 6% and 13% to build out AI infrastructures (xPU clusters) and xPU interconnect fabrics to meet these requirements.

Given AI training's intolerance to latency and packet loss, and with large volumes of traffic being exchanged across large elephant flows, data center architectures are evolving to enable high bandwidth east-west traffic between back-end network xPU clusters.

To date, high-speed networks for AI training have been remote direct memory access (RDMA) over proprietary, lossless InfiniBand, but there is growing focus on evolving Ethernet, an open standard with mass adoption, for this use. Compared to InfiniBand, Ethernet reduces cost and complexity and lacks scalability constraints. Progress on evolving Ethernet includes:

SUPPORT FOR RDMA over Converged Ethernet (RoCEv2), which enables direct memory access between devices over Ethernet to improve performance and reduce CPU utilization

EVOLUTION OF LOSSLESS ETHERNET, bringing advanced flow control, improved congestion handling, and advanced flow telemetry that improves the capabilities of modern switches

FORMATION OF ULTRA ETHERNET CONSORTIUM (UEC) and growing industry support for the UEC focus on architectures that optimize Ethernet for high performance AI and HPC networking. The UEC is collaboratively working on the Ultra Ethernet Transport (UET) specification to modernize RDMA (Remote Direct Memory Access) operation over Ethernet, optimized for AI and HPC workloads.

NVIDIA'S NEW SPECTRUM-X ETHERNET networking platform comprising its Spectrum-X range of Ethernet switches and BlueField–3 super-NICs. Spectrum-X offers Data Centers an Ethernet alternative to Nvidia's InfiniBand technology.

CISCO RECENTLY INTRODUCED THE NEXUS HYPERFABRIC AI cluster solution, co-developed with Nvidia, aiming to simplify AI infrastructure for enterprises using Ethernet networks. It consists of Cisco's 6000 series switches for spine and leaf that deliver 400G and 800G Ethernet fabric performance.

ARISTA UNVEILED ITS NEW ETHERLINK AI PLATFORMS consisting of 400G and 800G Ethernet spine and leaf switches supporting the emerging Ultra Ethernet Consortium's standards.

AI Data Center Networking Fabric Requires New Testing Approaches

Traditional AI data center interconnect fabric performance testing requires labs to operate real servers and GPUs, with test cases configured to generate AI workloads using real servers.

This is extremely expensive given the cost and long lead times to acquire GPUs, as well as the associated power consumption and real estate space required. For example, a basic test lab may comprise 80-160 GPUs across 10-20 servers at costs between $3 million and $6 million. More advanced labs could require thousands of GPUs, pushing costs to tens of millions dollars.

In addition, using traditional workloads to test AI fabric doesn't properly mimic GPU traffic loads and patterns.

Instead, new cost-efficient ways of stress testing AI data center networking are being used that emulate realistic GPU workload traffic.

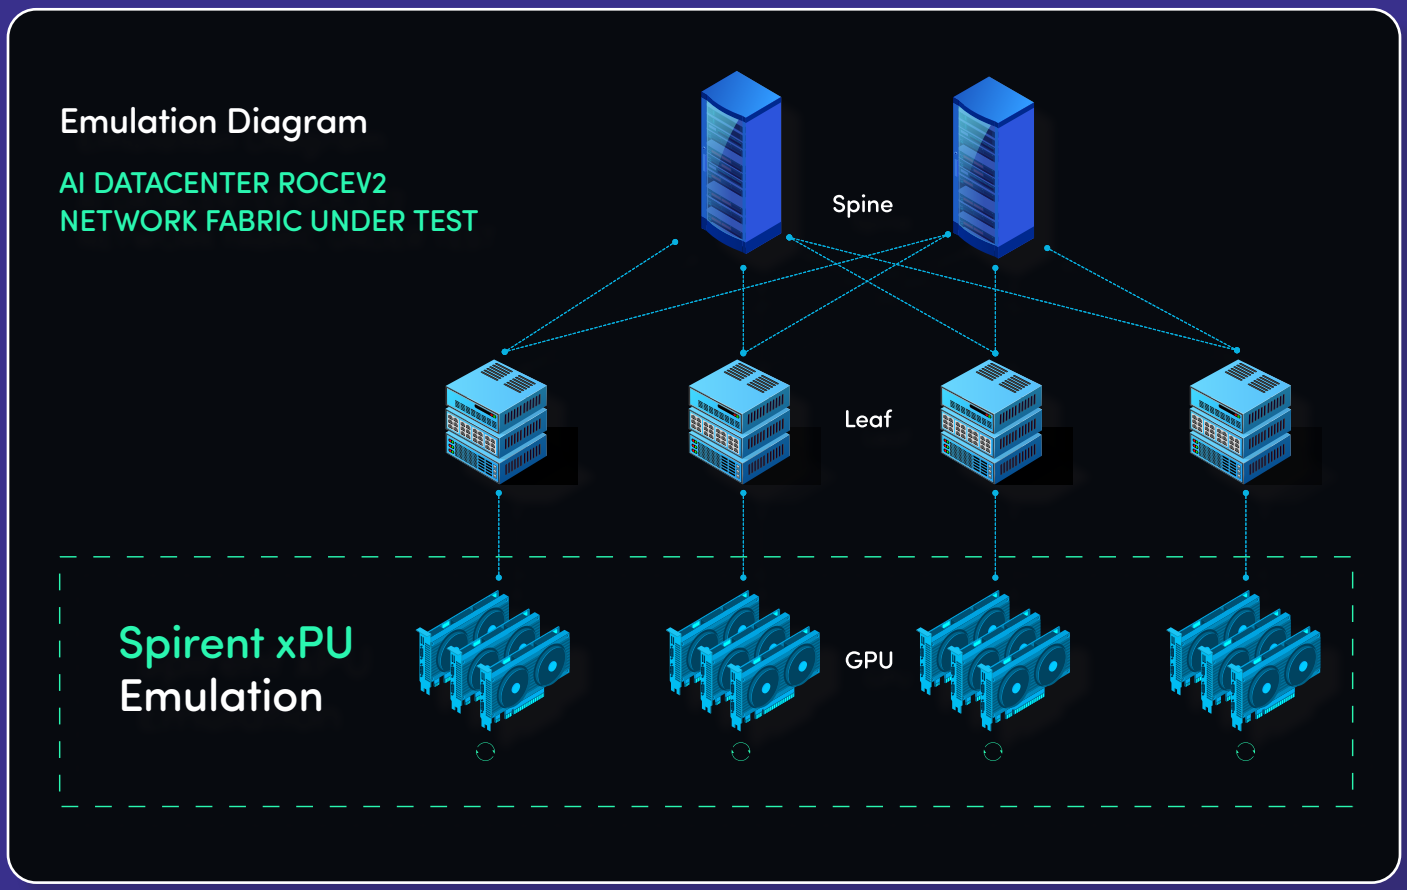

SPIRENT IMPACT REPORT

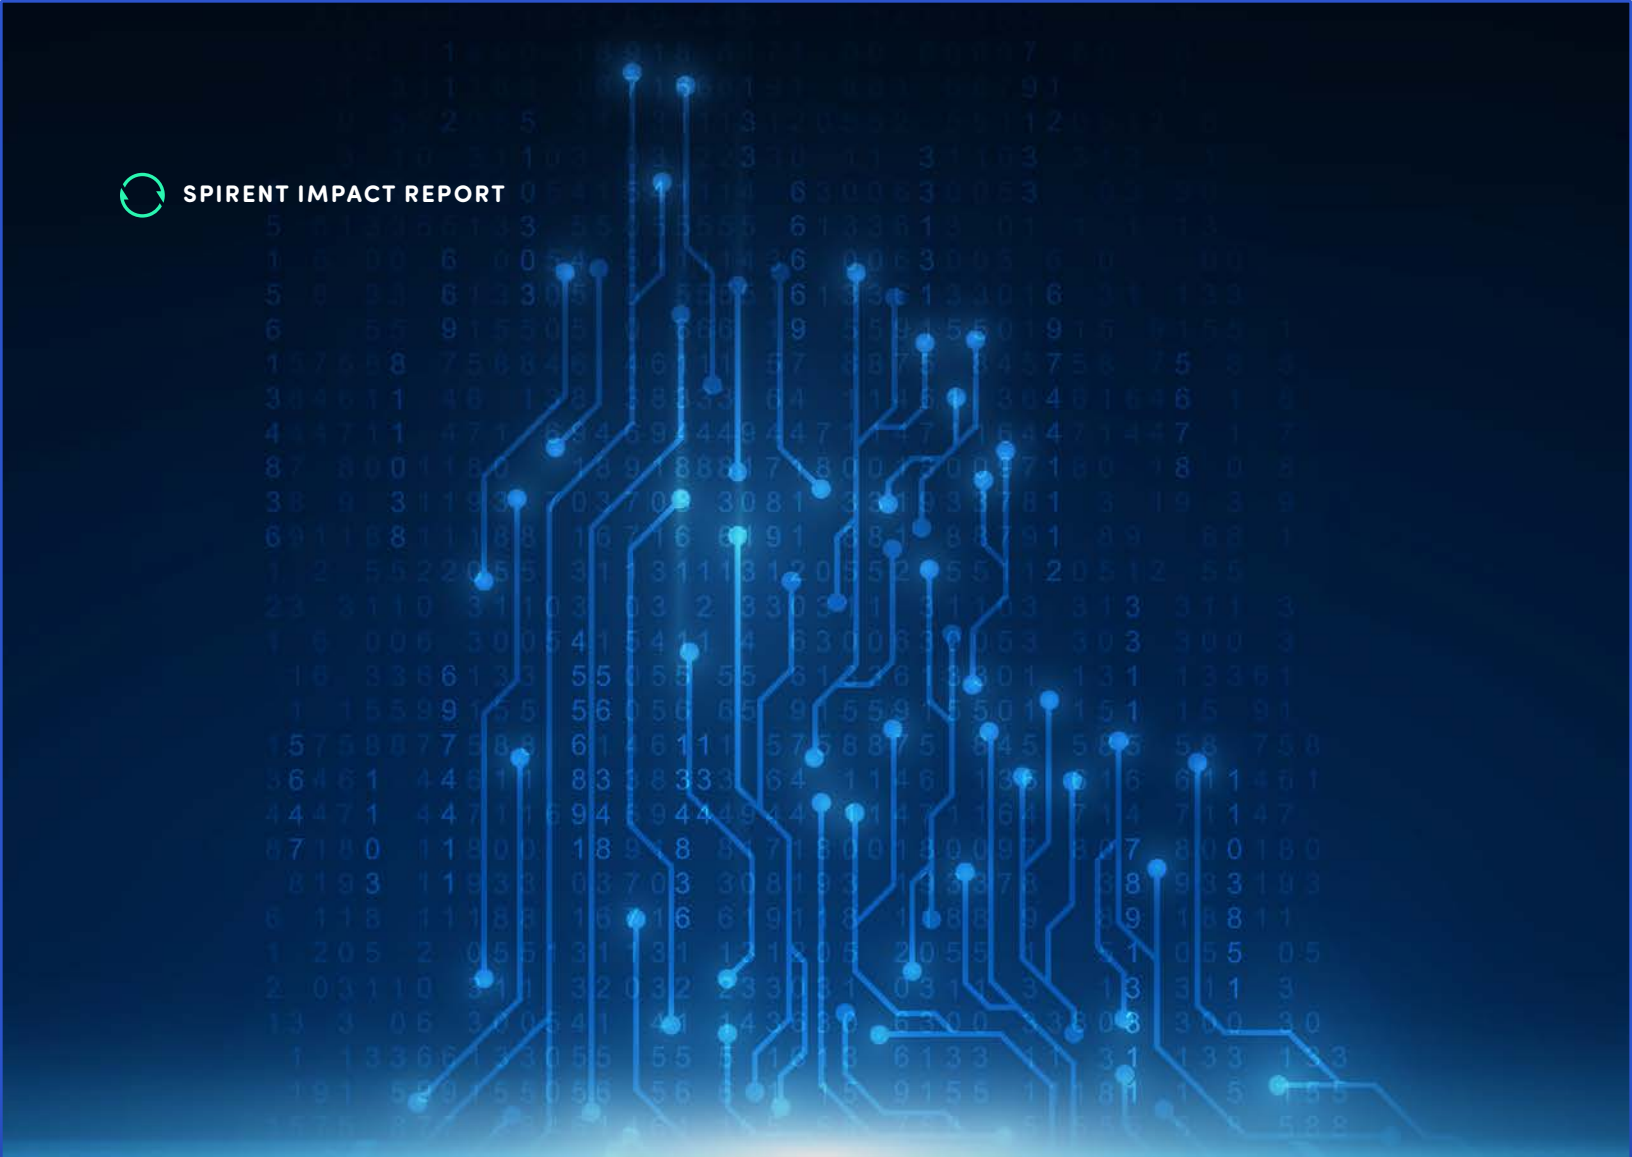

Data Center Networking Market Insights

Based on hundreds of high-speed Ethernet engagements and insights, Spirent is tracking the following high-speed Ethernet trends and market impacts.

Evolving data center high-speed Ethernet speeds

400G in the data center

- Large cloud hyperscalers are moving towards the latter stages of 100G to 400G refresh cycles and are beginning to focus on 800G.
- Tier-two and tier-three cloud service providers and large enterprises have begun to accelerate 400G refreshes, bypassing 200G.
- 400G is considered more future-proof due to its longer lifespan, simpler upgrade path to 800G, higher transmission rates, and lower cost per unit of capacity.

800G in the data center

- 800G's early momentum, driven by AI impacts on data center interconnect fabrics, will accelerate as GPU-hosting servers evolve from supporting 8 to 16 GPUs and beyond.
- Testing timelines suggest hyperscalers will start refresh cycles in late 2024 following 51.2Tbps switch availability. The rest of the market will follow in 2025 and beyond.
- Early testing engagements suggest hyperscaler 800G deployments are shifting from QSFP-DD to OSFP-XD pluggable form factors.
- Early 800G Ethernet Technology Consortium (ETC) standard and standardized interoperability issues were resolved with approval of the IEEE 802.3df standard.

1.6T future evolution

- Work continues towards the future 1.6T Ethernet standard through IEEE 802.3dj, which is anticipated to be finalized in 2026.
- A baseline set of features is expected to be completed this year by the 802.3dj task force, allowing the chipset ecosystem to start developing silicon.
- There is a debate around the potential adoption of OSFP-XD (pluggable optical module) using 16 x 100Gb/s lanes versus 8 x 200Gb/s, as initially demanded by the industry. This could speed time to market while the 200Gb/s technology challenges are resolved.

Dell'Oro has projected that nearly all AI back-end network ports will operate at speeds of 800G and above by 2027.

Ethernet Switches in AI Back-End Networks

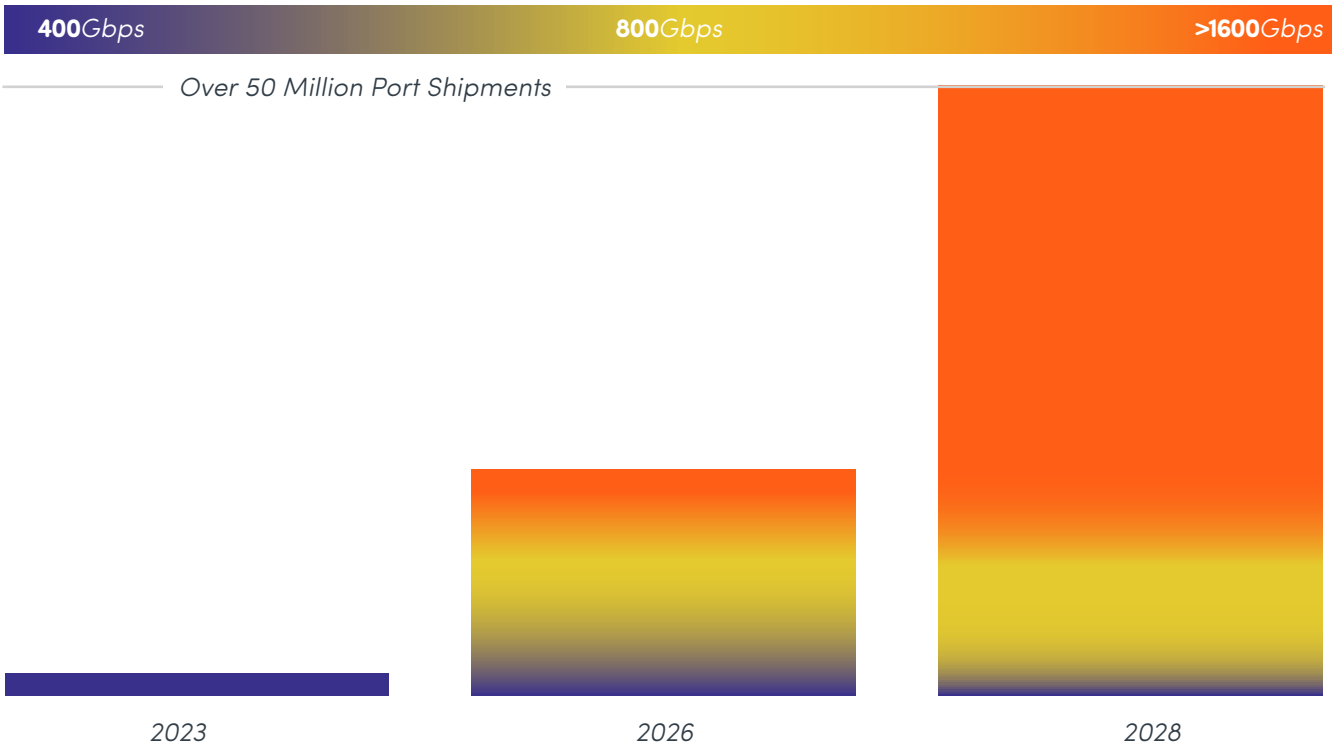

Source: Dell'Oro Group AI Networks Report

Dell'Oro forecasts Ethernet switch port shipments deployed in AI back-end networks to surpass 50 M by 2028.

Half of the Ethernet switch port shipments in AI back-end networks to be 800Gb/s by 2025 and 1600Gb/s in 2028, showing a very fast migration to the highest speeds available in the market.

Power and cooling innovations

Linear-drive pluggable optics (LPOs) are growing in interest to reduce escalating data center power consumption, cost, and latency. With linear-drive optics, the switch drives the module optics directly, thus eliminating digital signal processors (DSPs) and cutting the module's power consumption by half. However, adoption may be impacted by AI workloads with 200Gb/s lanes that may require DSPs to align the electrical and optical lanes. In addition, until the LPO-MSA and OIF releases specifications that will enable compatibility among suppliers, market adoption will remain slow.

While LPO's offer a tantalizing possibility of reduced power it's still too early in technical maturity to predict wide scale market adoption.

SPIRENT INNOVATION SPOTLIGHTS

Industry's first large-scale 800G test

Spirent supported H3C to complete the industry's first large-scale high-density 800G Ethernet test with up to 64 800G ports. The reliability and high performance of H3C S9827, a powerful series of switches for data centers and cloud computing networks, were validated with Spirent's award-winning high-density B2 800G Appliance. The H3C S9827 series is a new generation of 800G data center switches designed to handle large-scale workloads from AI-generated content (AIGC). The 800G switches use co-packaged optics (CPO) that integrate silicon and photonics.

Spirent test results showed a total switching capacity of up to 51.2T, with all 64 ports achieving 100% line speed forwarding under different traffic. Each port transmission rate reached 800 Gbps and the CPO technology fully met the high throughput demands for intelligent computing networking. These results make the H3C S9827 suitable for AIGC clusters and other high-performance data center core switching and related applications.

KEY TAKEAWAY:

Spirent's groundbreaking validation of H3C's 800G Ethernet switches underscores unparalleled reliability and performance, setting a new standard for high-density, AI-driven data center networks.

SPIRENT INNOVATION SPOTLIGHTS

Industry's first AI traffic emulation platform for data center networking using Ethernet

Spirent was honored with three "Best of Show" awards by Interop Tokyo 2024, including the "Grand Prize" for a ground-breaking AI traffic emulation solution capable of emulating realistic Artificial Intelligence (AI) workloads over Ethernet.

AI models demand unprecedented bandwidth, driving data centers to transition from 400G to 800G, and eventually to 1.6T speeds. This rapid evolution is crucial for meeting the massive data demands. However, AI data centers require a different approach compared to traditional ones. Recognizing this, Spirent has developed new testing methodologies to emulate AI workloads accurately.

For successful AI data center deployment, Spirent advocates a "trust but verify" approach with thorough stress testing of the network fabric to identify and mitigate potential bottlenecks. This proactive testing ensures that GPUs, which are expensive and critical to AI operations, do not sit idle due to network issues. Traditional testers cannot mimic AI traffic patterns. Spirent solution can emulate the high-bandwidth, low-latency demands of AI, offering a more realistic assessment of network performance before deployment.

KEY TAKEAWAY:

By emulating realistic AI workloads, Spirent helps customers gain confidence in their network infrastructure, maximizing efficiency and minimizing the risk of costly disruptions.

KEY TAKEAWAY:

Spirent test solutions help ensure the conformance, performance, interoperability, flexibility, and scalability of 800G deployments and AI workloads. Automated testing reduces network text complexity and simplifies and accelerates the path to 1.6T.

SPIRENT INNOVATION SPOTLIGHTS

Supporting the high-speed Ethernet ecosystem

Spirent was an active participant and contributor at the OFC 2024 event, showcasing its new B3 800G Appliance, a high-density native 800G OSFP and QSFP-DD test platform, and the first in its class to support IEEE 802.3df specifications.

Together with a dozen ecosystem partners, Spirent demonstrated its award-winning 800G test platforms for validating the reliability, performance, and interoperability of the latest high-speed networking solutions and devices across various interconnect technologies. (Read our blog post with key takeaways from this event.)

Telecom Service Provider Wireline Networking Market Insights

While hyperscaler cloud providers are migrating to 800G, service providers are evolving wireline networks, from edge network devices to edge, regional, and centralized data centers.

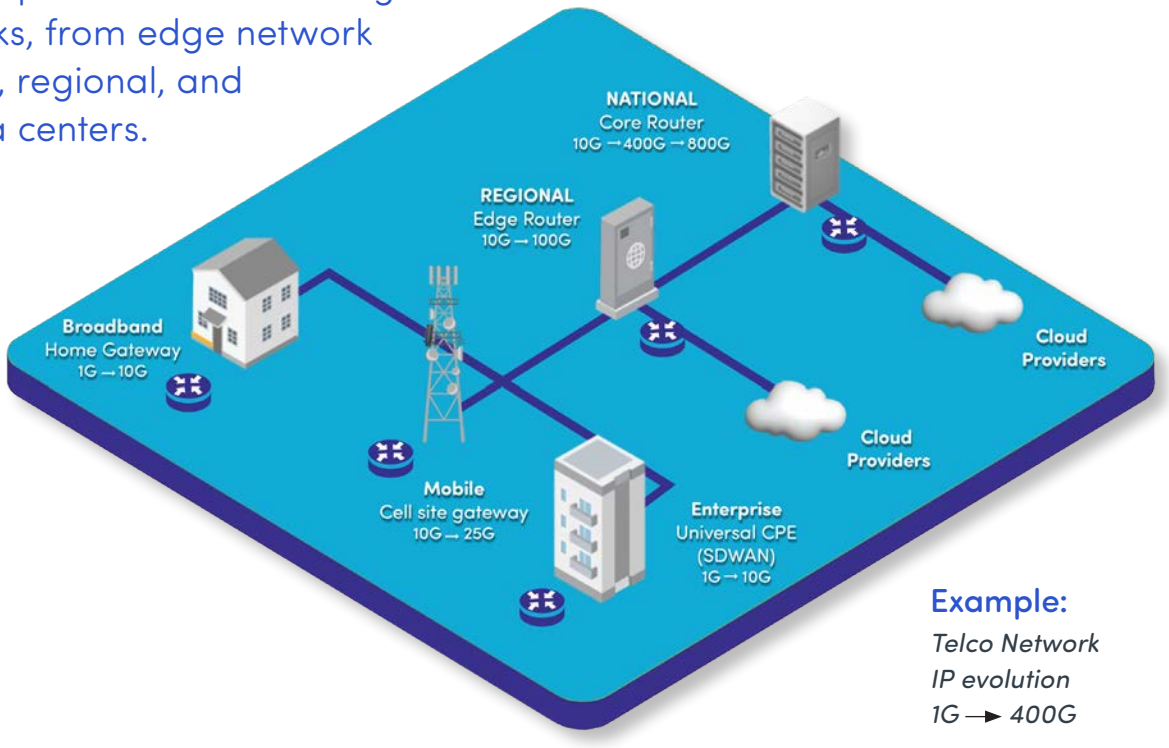

Example:
Telco Network IP evolution 1G → 400G

5G IP core: 400G momentum grows

To date, telecom service providers have initially focused on RAN and edge transport networks with the introduction of 5G Non-standalone (NSA) networks. Now, as they transition to 5G Standalone (SA), the focus has shifted to IP core deployments. The growth in 5G traffic is driving a wave of IP core network upgrades to 400G to reduce costs. 400G also offers higher speeds per port and decreased energy consumption, reducing the number of ports and rack space required.

The QSFP-DD form factor is currently becoming the most popular choice for module/connector for service providers evolving to 400G.

5G RAN: higher capacity needed

Growing 5G traffic and increasing 5G RAN deployments are triggering the need for more capacity and higher speeds at cell and aggregation sites. This is driving transport network upgrades from 1G to 10G and even 25G for cell site interconnects, and 100G at edge aggregation sites. For example, burgeoning virtualized and Open RAN networks are driving the demand for 25G capacity on eCPRI interfaces to support MIMO radios. This increases the need for 10/25G at the cell site to terminate radio links and 100G or higher connections towards the Centralized Unit (CU). Most virtual DU (Distributed Units) are also being deployed today with 100G network interface cards (NICs).

800G: building awareness

Service providers remain focused on the journey to 400G but realize they need to consider 800G in the next number of years as traffic increases with the higher bandwidth demands of AI and 5G's evolution to 5G Advanced and 6G. Nascent enhanced services such as extended reality and ultra-high-definition video are also generating early awareness in 800G.

This is leading many service providers to ask vendors to ensure 400G upgrade cycles will be 800G-ready to enable a seamless refresh in three to four years. Late followers may even bypass 400G and refresh from 100G directly to 800G.

AI inference: edge capacity will grow

As service providers endeavor to adopt AI/ML, including hosting and monetizing third-party applications, network transport infrastructure is in urgent need of preparation.

AI/ML processing will be highly distributed across telecom networks with inference primarily taking place across various edge locations, including far-edge at cell sites or hub sites, mid-edge at aggregation sites, and near-edge at core sites.

Inference is more relevant to edge locations due to the real-time low latency requirements of use cases and the fact that most edge locations only need to host smaller pre-trained models that require less storage and compute overheads versus centralized data center locations where training and learning processing will take place.

As inference use cases at the edge grow, a significant amount of AI traffic will traverse edge locations, prompting the need for early capacity upgrades in access and transport networks.

Early forecasts suggest edge locations could require substantial additional capacities with far-edge sites requiring 25-50G speed grade upgrades, mid-edge sites requiring 100-200G and near-edge sites requiring 400G with potentially a faster refresh cycle to 800G.

Network slicing: hard and soft slices

Taking advantage of new technologies coming into play in the next few years, telecom networks are evolving at the physical layer while incorporating network slicing and flexible transport layer routing.

Service providers are focusing on the transport infrastructure for network slices, especially for wholesale and enterprise private network offerings. Both dedicated, isolated hard slicing for mission-critical services, and flexible, on-demand soft slicing for shared services like VPNs are being tested and offered. A hybrid mix of capabilities is viewed as key for delivering a variety of use cases, service-level agreements, and operational efficiencies.

Hard slices are being tested using optical channel flexibility via ODUflex and emerging FlexE, which uses Ethernet to provide flexible use of optical transport bandwidth.

Today, soft slices are using MPLS with segment routing traffic engineering (SR-TE) technology and VLANs. As SR-TE evolves to SRv6 with Flex-Algo, service providers have the flexibility to define topology and path segment algorithms.

IPv6 and segment routing

The introduction of IPv6 and segment routing over IPv6 (SRv6) is creating a growing market trend currently favored in greenfield builds. IPv6 simplifies wireline network management by reducing routing hops and enhancing network reliability. It also enables fast rerouting of traffic to alternative paths, which is crucial for service providers.

MPLS remains favored in brownfield sites, however service providers are developing evolution plans that include hybrid support of MPLS and IPv6.

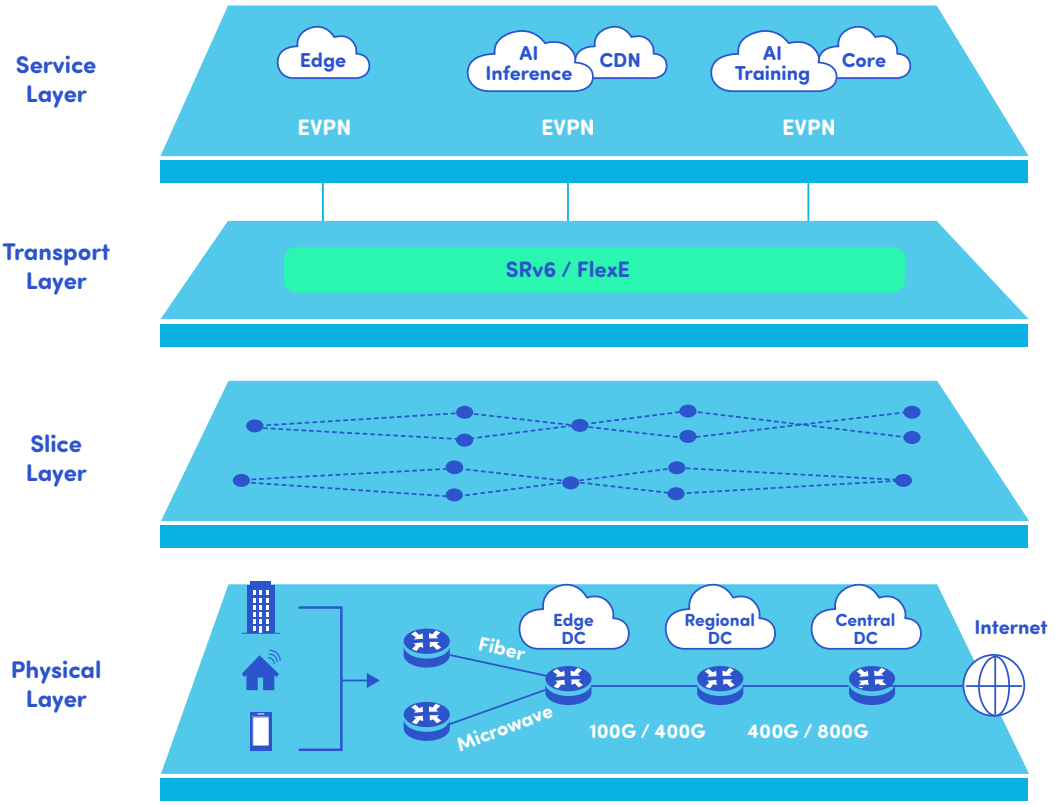

Future Telecom Network

SPIRENT ENGAGEMENT SPOTLIGHT

China Mobile uses Spirent's IP network digital twin platform

To empower the evolution of IPv6, cloud, 5G, and intelligent network capabilities, China Mobile developed and deployed segment routing over IPv6 (SRv6) in the backbone, metropolitan, cloud, and SD-WAN networks. With the rapid transition to IPv6, 5G, and cloud-native, it was essential for China Mobile to ensure the service quality, reliability, and stability of that network.

Challenge

The SRv6 network usually comprises dozens of core routers, hundreds of convergence routers, and numerous access switches. The cost and time associated with replicating these complex networks in a lab was cost-prohibitive and impractical.

Due to the nature of the technologies involved, the time and resources required to complete each complex testing cycle was increasing to unacceptable levels. For example, during the test preparation stage, thousands of status checks are performed to confirm whether the environment meets the test requirements. This is followed by running hundreds of tests to validate network functionality, performance, reliability, and interpreting results.

Requirements

China Mobile needed an automated and simplified test solution to realistically emulate and test the many SRv6 network scenarios, including representative traffic generation, network event scenarios, and network state reproduction. The solution also needed to support network scalability requirements as traffic increased, ensuring stability throughout network convergence, connectivity changes, and routing updates.

Solution

Spirent partnered with China Mobile and third-party network equipment vendors to develop a web-based IP network digital twin platform. Spirent's IP network digital twin provides the ability for 1:1 simulation/emulation that reproduces hundreds of network properties, protocols, events, and services and the ability to support hundreds of network devices and thousands of network connections in a single topology network testing environment.

The digital twin included Spirent's innovative TestCenter network traffic generator and protocol emulator, supporting SRv6, and companion network emulation hardware to stimulate real-world network conditions, including impairments and adversarial conditions.

Specifically, China Mobile benefited from actionable, real-time analytics with insights for multi-dimensional network information, and a "What You See Is What You Get" interactive simulation with step-by-step wizard interfaces to speed testing cycles and time-to-revenue. High-level integration for the network tester, network impairment emulator, SDN controller, and network device, along with quick and flexible test scenario customization was made possible via an expandable, common data interface.

Results

China Mobile maintained network stability and more effectively managed interactions between different networks and technologies, successfully achieving a:

- 90% reduction in Capex by replicating the real environment.
- 70% reduction in time and cost for network validation, improving efficiency and accelerating time-to-market
- 90% reduction of human errors

Enterprise Wireline Networking Marketing Insights

Spirent's high-Speed Ethernet engagements have seen enterprises growing their use of Ethernet networks for enterprise data centers and campus switching networks.

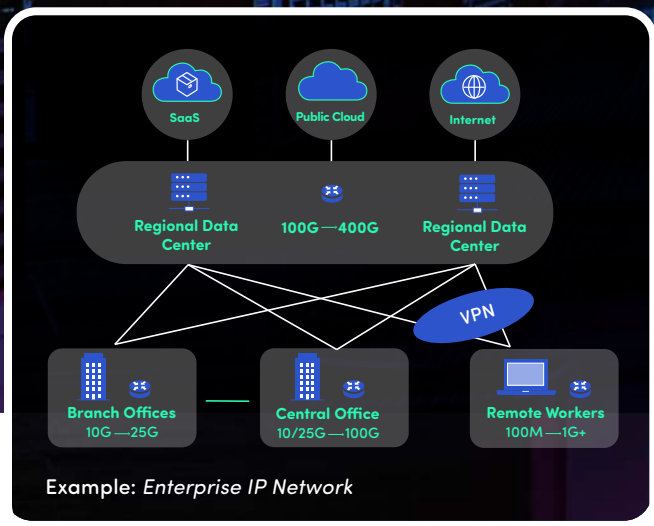

Example: *Enterprise IP Network*

Enterprise Ethernet networks

In enterprise data centers

Enterprise data centers are still utilizing 100G Ethernet networks. However, there is growing interest among large enterprise data centers to gradually transition to 200G and 400G Ethernet for switch spine and fabric. We expect to see continued 200/400G growth over the next few years due to supply availability, cost-per-bit value, demands for more bandwidth, and continued growth in data traffic, especially related to remote working and cloud transformation.

In campus networks

Campus networks typically use 10G for campus core and aggregation layers, but 25G and 100G are gaining in popularity for medium and large campus networks.

At the access layer, 1G remains dominant but demand is growing for faster speeds due to the increasing adoption of Wi-Fi 6 and 6E access points as networks upgrade from 802.11ac to 802.11ax.

We have also seen interest in a new market segment, Campus Network-as-a-Service (CNaaS), with new players focused on delivering cost-efficient networks, value-added services, and innovation, not just enabling technologies.

Enterprise network security equipment

The enterprise network security market is substantial, with extensive testing of next-generation firewalls and gateways currently underway.

We are seeing continued growth in remote work and cloud applications that are driving the security service edge (SSE) market and cloud-centric security solutions.

While 1G to 10G firewall interfaces continue to dominate the high-end enterprise firewall market, 100G next-generation firewall interfaces are the primary focus for enterprise data centers. There is also early interest in 400G interfaces for larger data centers as the number of connected remote applications increases.

SPIRENT ENGAGEMENT SPOTLIGHTS

Major banks ensure vendor-neutral performance and compliance

With millions of dollars per minute on the line, one of the largest banks in the world selected Spirent to ensure ultra-high network reliability. Spirent used high-speed Ethernet test equipment and expertise to develop a validation program for a new infrastructure that provides independent assurance that newly acquired 100G switches will perform to specifications when deployed in the production network.

A major U.S.-based bank partnered with Spirent to develop a vendor-neutral validation program for new high-speed Ethernet infrastructures. The program revealed that some vendors' switches could not deliver the full 100G line rate when multiple ports were in use. Comprehensive testing determined the maximum number of ports that could be enabled while still achieving full line rate, thus enabling the bank to maximize performance, lower latency, and improve reliability.

As the number of network devices increased, this international bank struggled to perform regulatory compliance testing across all device and vendor combinations. Spirent automated the test environment, deploying a Layer 1 switching fabric to allow test beds and associated equipment to be remotely powered, automatically configured, and used to perform tests. Now the bank turns up test beds 300x faster with fewer resources, has reduced energy consumption by 40%, and can run automated tests 24/7 to ensure regulatory compliance.

One of the largest banks in the U.S. needed to deploy new AI-powered services and commissioned an AI data center. Spirent TestCenter gave the bank the ability to test the performance characteristics of 400G switches with the opportunity to add on 800G, and AI traffic emulation to validate future target AI data center architectures such as RoCEv2.

A US-based financial exchange was in the process of a multi-year cloud migration consisting of a hybrid mix of public and private services. With Spirent TestCenter and TestCenter Virtual, the exchange is able to perform full line-rate testing of all links, including packet loss and round-trip latency measurements, to ensure a level playing field in terms of connection latency with their multiple trading entities.

KEY TAKEAWAY:

Spirent helped ensure banks hit ultra-high network performance and compliance targets with the latest 100G and 400G Ethernet test solutions.

SPIRENT ENGAGEMENT SPOTLIGHTS

Gaming providers ensure end-user experience

A leading provider of cloud-based gaming services was experiencing high latency issues when delivering games via a top cloud provider's platform and struggled to identify the root causes. The gaming services used Spirent experts and high-speed Ethernet test solutions to augment their internal testing capabilities, identify potential sources of the high latency, and help develop best practices to minimize these sources. As a result, the gaming provider significantly improved the latency of their offering over the production cloud platform.

One of the largest gaming providers in the world needed to ensure an optimal experience for end-users of their services. They wanted to ensure that any changes to their networks were thoroughly tested and vetted prior to delivering live services. The gaming provider partnered with Spirent to develop a lab-based test environment which faithfully recreates real-world protocols, devices, and traffic loads. By using full emulation of high-speed Ethernet devices and traffic, the gaming provider can comprehensively validate new Ethernet infrastructures to ensure high-quality end-user experiences.

KEY TAKEAWAY:

Spirent helped minimize gaming latency with comprehensive network testing and emulation.

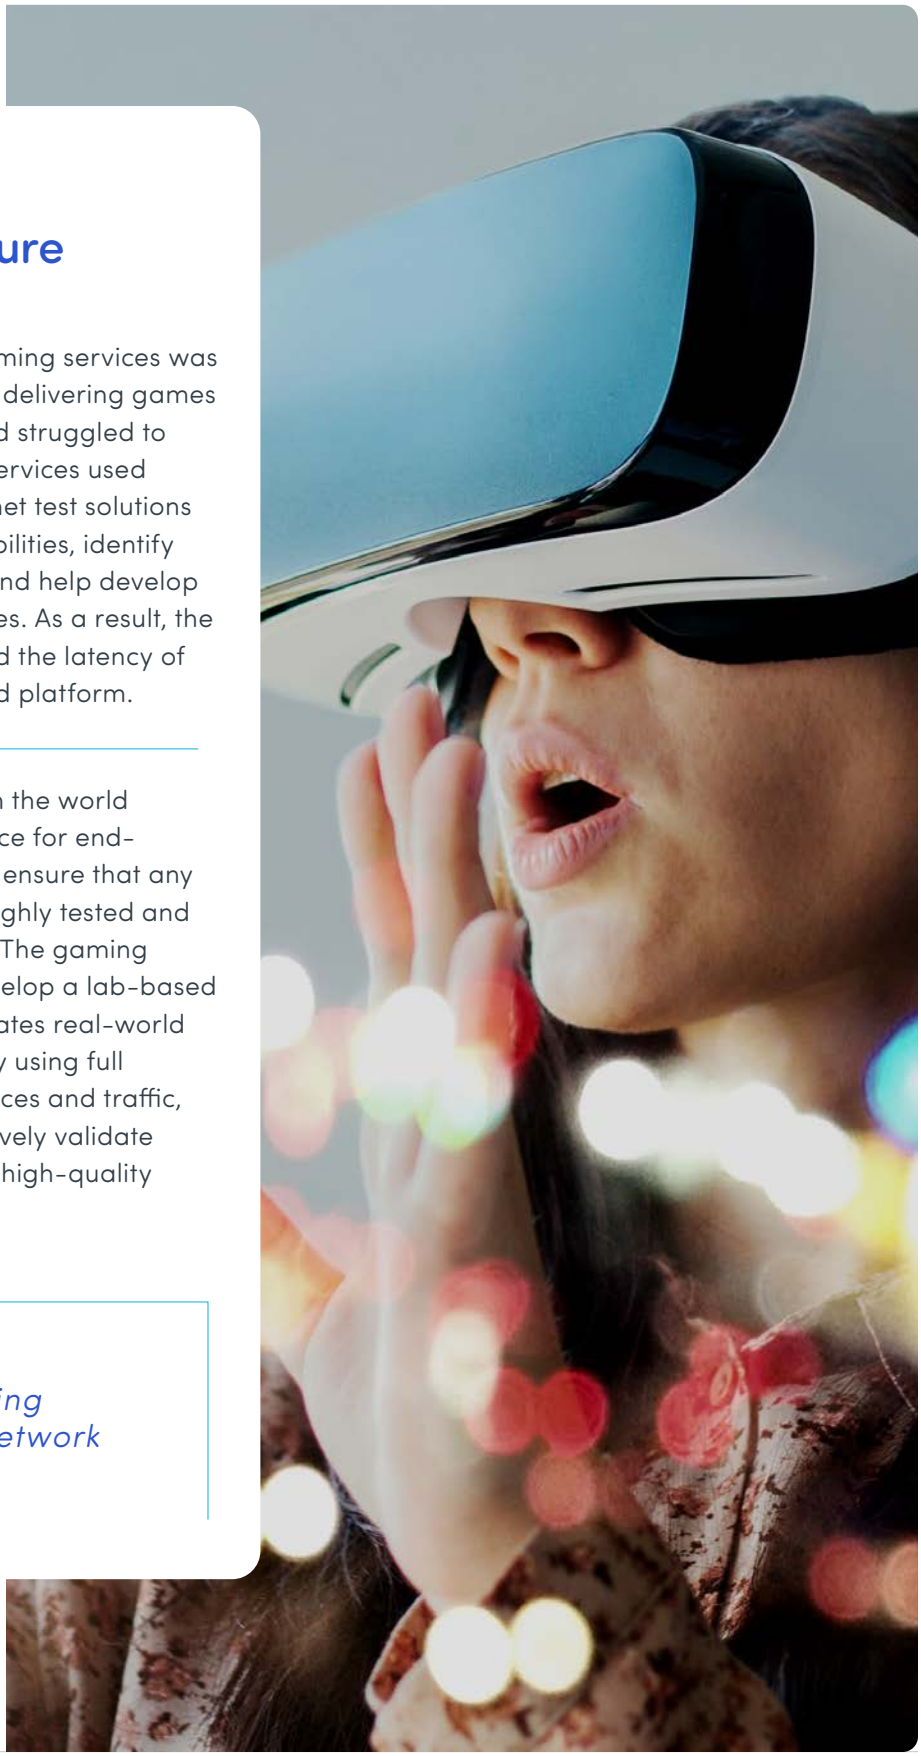

High-Speed Ethernet Market Evolution

The pace of innovation in the high-speed Ethernet market is accelerating faster than ever, with cycles being compressed and markets outpacing standards. AI's data deluge is pushing changes to Ethernet.

The stakes are high and the industry is moving quickly to meet them.

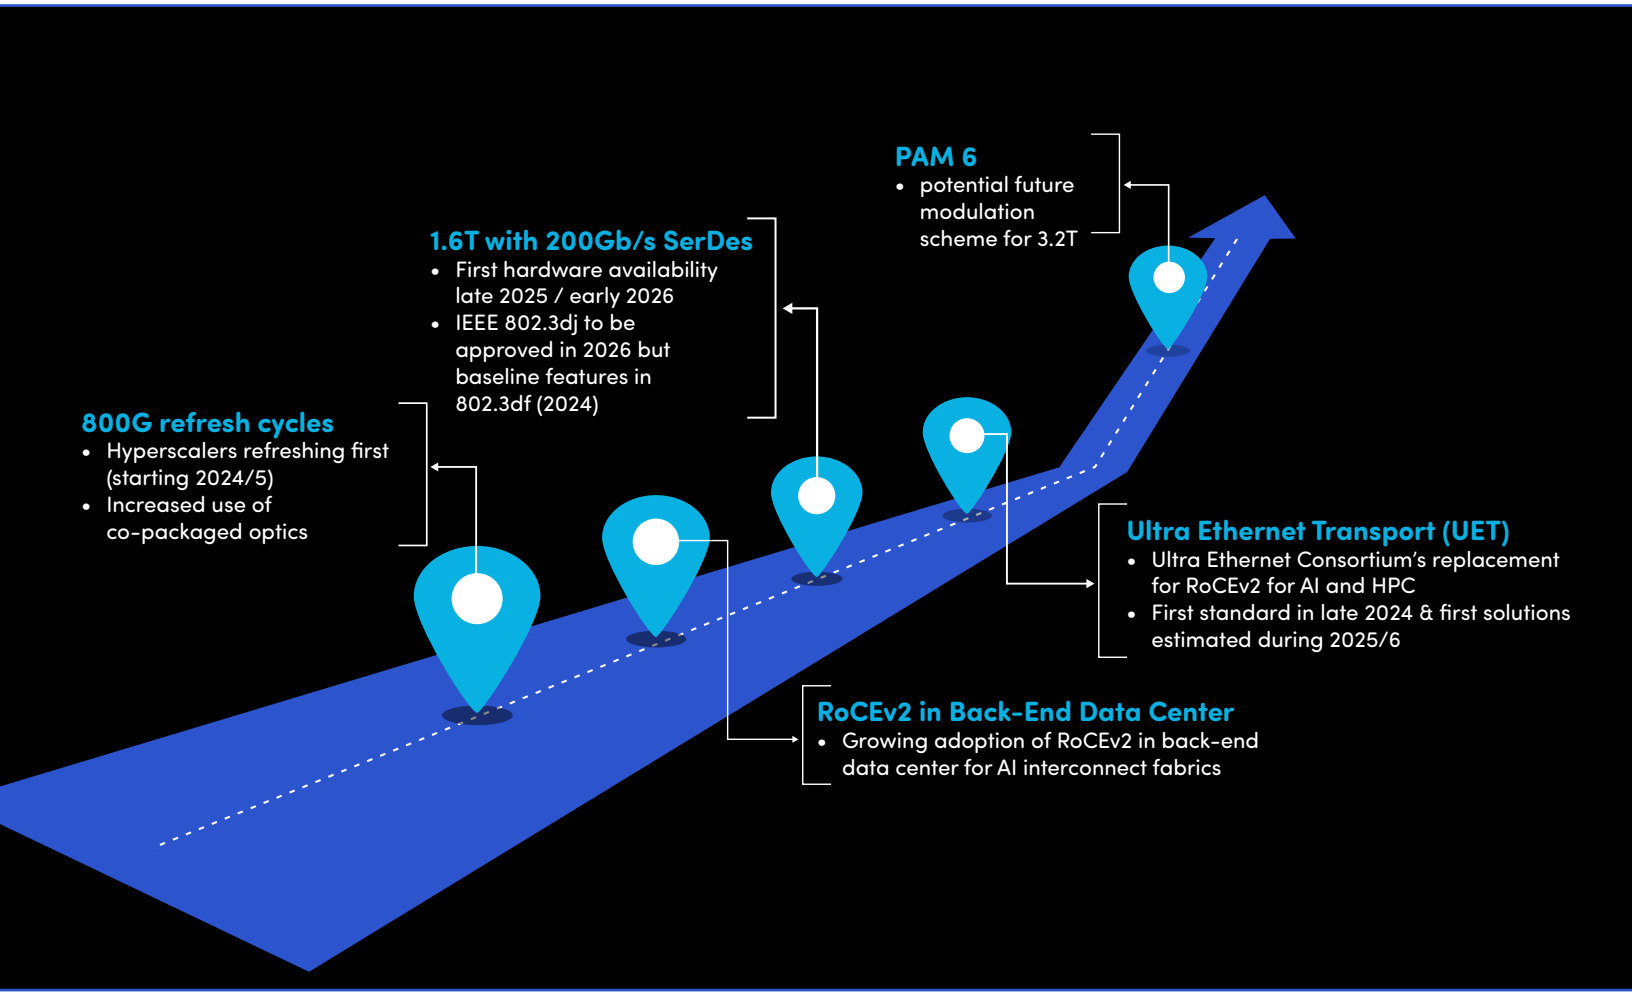

Innovations Enabling high-speed Ethernet and AI Networking

Spirent's Ethernet testing expertise spans decades, with a global reputation for supporting high-speed Ethernet innovation. Most recently, this has included speeding the release of next-gen high-speed Ethernet test capabilities with solutions for emulating large AI traffic loads at scale.

Our high-speed Ethernet solutions range from line rate testing to efficient large-scale traffic generation, impairments, and timing and synchronization. We incorporate realism in test scenarios and test automation for efficiency and cost reduction, while streamlining overall processes.

Spirent has taken complexity and cost out of AI Ethernet fabric pre-production testing by simulating xPUs and AI workloads so our customers don't have to build labs with massive numbers of expensive, physical GPUs and servers. These AI fabric tests ensure our customers are getting the best possible performance from AI data centers, without wasting resources.

Our ability to remain ahead of the market innovation curve differentiates us, with Ethernet solutions that range from 1G to 800G, including:

- The industry's first high-density test solution capable of emulating realistic AI workloads over Ethernet
- The Spirent B3 800G Appliance is the first in its class to support IEEE 802.3df specifications
- The high-density, native 400G QSFP-DD test platform enabling delivery of high-performing, interoperable 400G infrastructure and support emerging technologies like AI/M

Spirent stands ready to validate interoperability of the latest technologies and participates in leading standards bodies and alliances, including IEEE, Ethernet Alliance, Ultra Ethernet Consortium, Broadband Forum, Open Alliance, and Open Compute Project.

Spirent is dedicated to serving this dynamic market through innovation, collaboration, and vendor independence.

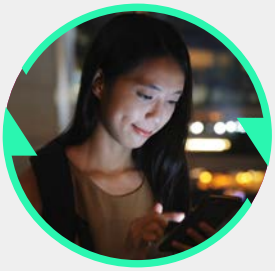

About Spirent

Spirent Communications plc. (LSE: SPT) provides innovative products, services, and managed solutions that address the test, assurance, and automation challenges of a new generation of technologies, including AI, 5G, edge computing, cloud, autonomous vehicles, and beyond.

For more information visit: www.spirent.com.

Learn More

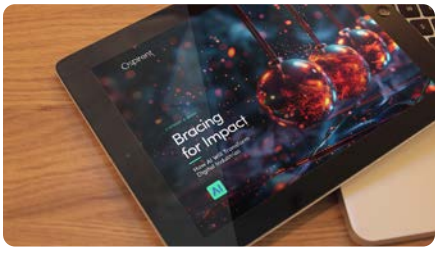

eBook
Bracing for Impact - How AI Will Transform Digital Industries
Artificial intelligence (AI) systems are expected to transform digital industries in ways still being discovered.

Infographic
Testing the AI Fabric

White Paper
Turning Ideas into Action for 1.6T Ethernet

Ethernet Alliance
Ethernet Roadmap

Americas 1-800-SPIRENT
+1-800-774-7368 | sales@spirent.com

Europe and the Middle East
+44 (0) 1293 767979 | emeainfo@spirent.com

Asia and the Pacific
+86-10-8518-2539 | salesasia@spirent.com

www.spirent.com

Rev A | 10/2024