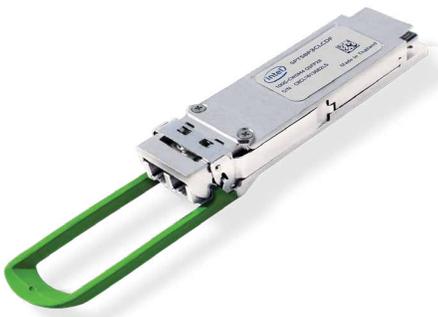




IT@Intel: Affordably Increase Network Bandwidth at 100 Gbps and Beyond

Intel® Silicon Photonics technology enables Intel's data centers to meet growth in demand for network bandwidth at a fraction of the cost of conventional optics technologies



Intel® Silicon Photonics helps enable faster data transfer over longer distances, while utilizing the efficiencies of Intel's high-volume silicon manufacturing to lower costs.

Executive Summary

Intel's demand for network bandwidth grows 30 to 40 percent a year. As a result, a key component of Intel IT's data center strategy is network innovation. We are increasing deployment of 100 gigabits per second (Gbps) connectivity, using Intel® Silicon Photonics.

When Intel IT compared Intel Silicon Photonics to conventional optics-based technologies, we found that it provides the following advantages:

- Uses less power (only 3.5 watts)¹
- Is less expensive²

Our adoption of Intel Silicon Photonics helps improve network utilization, reduces costs per port, and enhances overall data center efficiency. We have reduced the costs associated with all network components (physical cabling, active switch equipment, and optics). This has enabled us to lower the overall cost of transitioning to 100 Gbps. What's more, we are poised to break new ground with even faster network technology. We are ready to meet the data explosion head-on and satisfy Intel's demand for data processing for the foreseeable future.

Authors

Sanjay Rungta
Sr. Principal Engineer
Intel IT

Mohammad Ali
Network Engineer
Intel IT

Brad Horner
Network Engineer
Intel IT

¹ [intel.com/content/www/us/en/architecture-and-technology/silicon-photonics/optical-transceiver-100g-cwdm4-qsf28-brief.html](https://www.intel.com/content/www/us/en/architecture-and-technology/silicon-photonics/optical-transceiver-100g-cwdm4-qsf28-brief.html)

² Based on internal Intel IT measurements.

Contributors

Mario Villalta, Industry Engagement Manager, Intel IT

Acronyms

- DAC** direct-attach copper
- DOME** Design, Office, Manufacturing, and Enterprise
- Gbps** gigabits per second
- ROI** return on investment
- SMF** single-mode fiber
- ToR** top of rack

Background

Intel's data centers are the heart of a thriving, complex business. Intel IT operates 17 data centers with 56 data center modules. These data centers are home to over 305,600 servers that underpin the computing needs of more than 110,800 employees³. Intel's business is becoming increasingly data-driven, relying on machine learning, artificial intelligence, big data analytics, and automation. As data explodes, we are experiencing a 30 to 40 percent in demand for network capacity every year.

Intel IT is dedicated to supporting the needs of Intel's critical business functions—Design, Office, Manufacturing, and Enterprise (DOME). But we also are committed to operating our data centers as efficiently as possible. To meet these two goals, we have engaged in a multiyear evolution of our data center strategy—including the makeup of our data center networks.

Over the last 10 years, we have steadily upgraded our network ports from 1 gigabit per second (Gbps) to 10 Gbps to 40 Gbps. We are now migrating to multiple-100 Gbps inter-switch links to keep up with traffic growth (see Table 1). But “fast” is not the only thing we are concerned about. We also look at return on investment (ROI) to help ensure we are implementing the most cost-efficient solutions possible.

Table 1. 100 GbE Port Deployment Rate

| Year | Number of 100 GbE Ports | Annual Growth Rate (Increase) |
|--------------------------|-------------------------|-------------------------------|
| 2017 | 520 | |
| 2018 | 1,619 | 211% ▲ |
| 2019 | 3,758 | 132% ▲ |
| First six months of 2020 | 8,378 | 123% ▲ |

³ Number of data centers and servers as of February 2020. To define “data center,” Intel uses IDC’s data center size classification: “any room greater than 100 square feet that houses servers and other infrastructure components.”

Intel® Silicon Photonics: An Affordable Path to 100 Gbps Networks

In today's ever-evolving data center, it is important to understand the myriad of networking speeds, as well as data center topology and various combinations of switching hardware, cables, and connectors. Each of these components impacts the cost of the network and the longevity of a networking solution. Intel's data center networks are complex systems with a wide range of functions, from silicon design to office and enterprise applications. A large portion of our data centers is used for chip design workloads. For these, we have found that a high-speed Ethernet-based solution is optimal when considering cost, rate of change, and business need. We carefully review all elements to choose the solution that is right for our data center workloads.

A Brief Overview of Intel's Network Topology

There are three key vectors in designing large-scale data centers:

- Ethernet switch
- Data center topology
- Physical media selection

Disciplined Change Enables Fast Adoption of Innovation

Intel IT runs Intel's data center services like a factory, affecting change in a disciplined manner and applying breakthrough technologies, solutions, and processes.

Our approach to data center strategy enables us to optimally meet Intel's business requirements. At the same time, we can provide our internet customers with effective data center infrastructure capabilities and innovative business services.

We tend to adopt higher-speed network technology almost as soon as it is available in the market. We started adoption of 40 Gbps in data centers in 2015 and adoption of 100 Gbps technology in 2017, to keep pace with network demand. Due to the scale of the data center, we focus on zero-touch provisioning and automation to keep the design consistent. Our focus is on day-zero and day-one automation to expediate the landing and support.

Early Adoption

Intel IT adopts higher-speed network technology almost as soon as it is available.



Through industry-wide discussions with fellow travelers, we have found that networks inside hyper-scale data centers are converging to a Clos topology. This is because the Clos topology is modular, scalable, and flexible. In this topology, multistage switches interconnect to support thousands of Ethernet ports, which in turn support servers within the data center.

Intel IT uses a five-stage Clos topology as shown in Figure 1. Connections between servers and the leaf are typically direct-attach copper (DAC) cable for cost effectiveness. In contrast, connections between leaf and spine and between spine and super/universal spine require fiber-optic single-mode fiber (SMF). This is because switches can be more than 300 meters apart. Multimode fiber can support a distance of only 70 meters, even at 25 Gbps. SMF also increases the longevity of the design, as it can support not only 100 Gbps but 400 Gbps and beyond.

Our Clos topology allows fixed-form-factor switches at most of the topology stages, including top of rack (ToR), leaf, and spine. We also use a chassis-less design (small form factor) instead of investing in a bigger chassis. Our evaluations show the cost of a port is 2x to 3x higher than a small chassis with numerous interconnections.⁴ With fixed-form-factor switches and a chassis-less design, we have found that we can adopt new technology faster compared to using other designs. Therefore, we can increase our use of 100 GbE with a lower total cost.

When we began to transition to 100 Gbps, we upgraded the fiber infrastructure to support SMF within our data centers. Initially, we used a product that required four strands of fiber (100G PSM4) for the transceiver. Later, we transitioned to a different product that required only two strands (100G CWDM4) for the transceiver, which lowered our fiber costs.

Optimizing the Optics Can Increase ROI

As seen in the Figure 1, there can be hundreds of transceivers required to interconnect the switches to build the high-speed data center fabric. To optimize cost efficiency, it was important to select the appropriate type of transceiver. We chose Intel® Silicon Photonics transceivers over conventional optics or discrete laser-based technology. Our evaluation showed that Intel Silicon Photonics transceivers offer the following benefits compared to other solutions:

- Less power consumption
- Support for higher density
- Optimum total cost equation

We standardized on the CWDM4 MSA QSFP28 model because it can span up to 2 kilometers on SMF. This model can also operate in a wide range of temperatures (0° to 70° Celsius) at low power (3.5 watts), which was important for variety of data center deployments.

⁴ Based on internal Intel IT measurements.

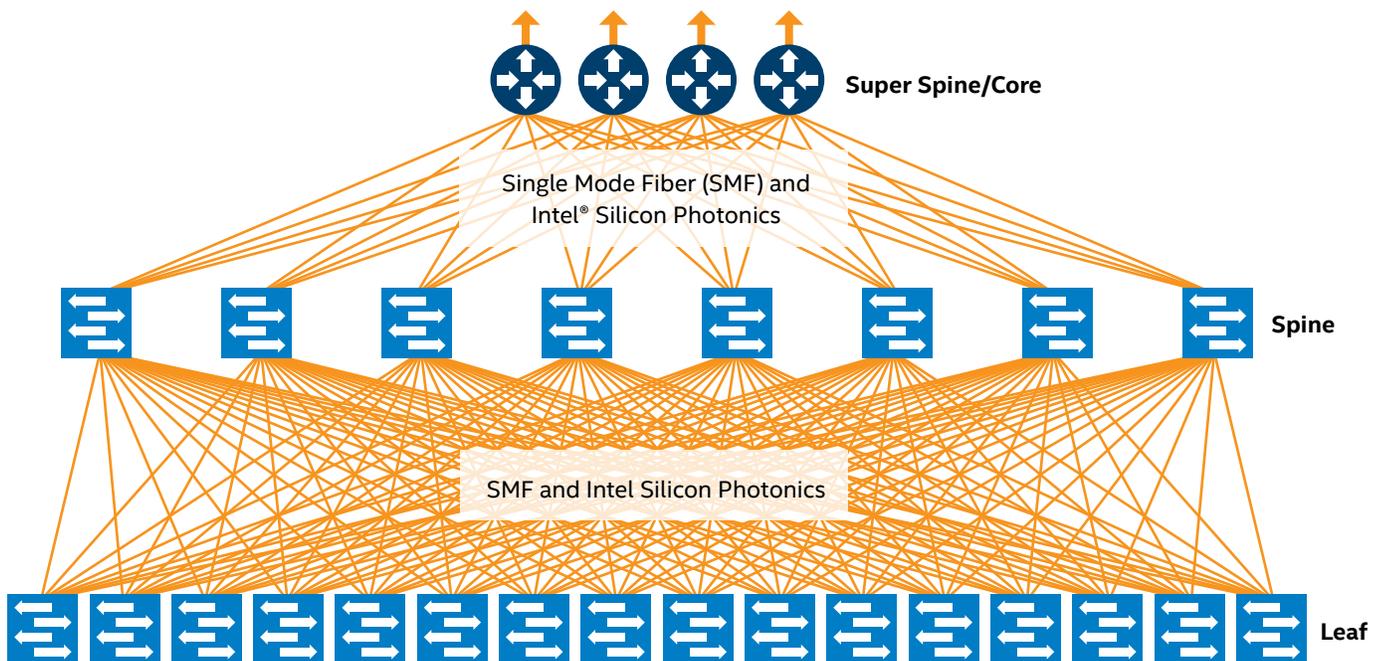


Figure 1. The optical connections (orange lines) can be hundreds of meters apart. We are migrating to SMF as we deploy more 100 Gbps connectivity.

Results

Over the last 18 months, we have deployed more than 8,000 100 GbE ports across Intel and 3,000 Intel Silicon Photonics transceivers. The deployment of Intel Silicon Photonics has enabled us to meet the growth in demand for network bandwidth without significantly increasing costs.

Conclusion

Adoption of Intel Silicon Photonics is a key element of our high-performance computing data center strategy. It will enable us to maximize network utilization, cost reduction per port, and overall data center efficiency. When deploying 100 Gbps technology, we have found it is imperative to evaluate all the associated costs: physical infrastructure, switching hardware, and optics/transceivers. By optimizing each of these, we have been able to cost-effectively transition to 100 Gbps, using Intel Silicon Photonics. We will soon begin exploring 200 Gbps and 400 Gbps technology—positioning our network to meet continuing data explosion and processing demand for years to come.

Related Content

If you liked this paper, you may also be interested in these related stories:

- Data Center Strategy Leading Intel's Business Transformation paper
- Building a Multi-Cloud-Ready Enterprise Network paper

For more information on Intel IT best practices, visit intel.com/IT.

Intel® Silicon Photonics—A Closer Look

Intel® Silicon Photonics is the result of more than a decade of research and development.

Production shipments for Intel's 100G optical transceivers commenced with PSM4 products for data center deployment in 2016. In 2017, the portfolio was expanded to include 100G CWDM4 transceivers. In 2018, Intel Silicon Photonics received the SEMI Americas Award for the First High Volume Integrated Silicon Photonics Transceiver.⁵ In that same year, Intel passed the production-run-rate milestone of more than a million units per year. Also in 2018, Intel expanded the product portfolio to include extended temperature versions and began sampling 400G products for next-generation data center applications.

⁵ Semiconductor Today, July 2018, "Intel wins SEMI's 2018 Award for the Americas for process and integration of silicon photonics transceiver." semiconductor-today.com/news_items/2018/jul/intel_110718.shtml

IT@Intel

We connect IT professionals with their IT peers inside Intel. Our IT department solves some of today's most demanding and complex technology issues. We want to share these lessons directly with our fellow IT professionals in an open peer-to-peer forum.

Our goal is simple: improve efficiency throughout the organization and enhance the business value of IT investments.

Follow us and join the conversation:

- [Twitter](#)
- [LinkedIn](#)
- [#IntelIT](#)
- [IT Peer Network](#)

Visit us today at intel.com/IT or contact your local Intel representative if you would like to learn more.

Intel technologies' features and benefits depend on system configuration and may require enabled hardware, software, or service activation. Performance varies depending on system configuration. No computer system can be absolutely secure or error free, and therefore Intel makes no warranty, representation, or guarantee regarding the performance of its products, the information contained in this paper, or the suitability of its products for any particular purpose. Check with your system manufacturer or retailer.

Cost reduction scenarios described are intended as examples of how a given Intel-based product, in the specified circumstances and configurations, may affect future costs and provide cost savings. Circumstances will vary. Intel does not guarantee any costs or cost reduction.

THE INFORMATION PROVIDED IN THIS PAPER IS INTENDED TO BE GENERAL IN NATURE AND IS NOT SPECIFIC GUIDANCE. RECOMMENDATIONS (INCLUDING POTENTIAL COST SAVINGS) ARE BASED UPON INTEL'S EXPERIENCE AND ARE ESTIMATES ONLY. INTEL DOES NOT GUARANTEE OR WARRANT OTHERS WILL OBTAIN SIMILAR RESULTS. INFORMATION IN THIS DOCUMENT IS PROVIDED IN CONNECTION WITH INTEL PRODUCTS AND SERVICES. NO LICENSE, EXPRESS OR IMPLIED, BY ESTOPPEL OR OTHERWISE, TO ANY INTELLECTUAL PROPERTY RIGHTS IS GRANTED BY THIS DOCUMENT. EXCEPT AS PROVIDED IN INTEL'S TERMS AND CONDITIONS OF SALE FOR SUCH PRODUCTS, INTEL ASSUMES NO LIABILITY WHATSOEVER AND INTEL DISCLAIMS ANY EXPRESS OR IMPLIED WARRANTY, RELATING TO SALE AND/ OR USE OF INTEL PRODUCTS AND SERVICES INCLUDING LIABILITY OR WARRANTIES RELATING TO FITNESS FOR A PARTICULAR PURPOSE, MERCHANTABILITY, OR INFRINGEMENT OF ANY PATENT, COPYRIGHT OR OTHER INTELLECTUAL PROPERTY RIGHT.

No license (express or implied, by estoppel or otherwise) to any intellectual property rights is granted by this document.

Intel, the Intel logo, and other Intel marks are trademarks of Intel Corporation or its subsidiaries in the U.S. and/or other countries.

Copyright © Intel Corporation

Other names and brands may be claimed as the property of others. 0720/WWES/KC/PDF 343725-001US

