The benefits of optical circuit switches (OCS) in modern data center networks

Spine layer switch replacement

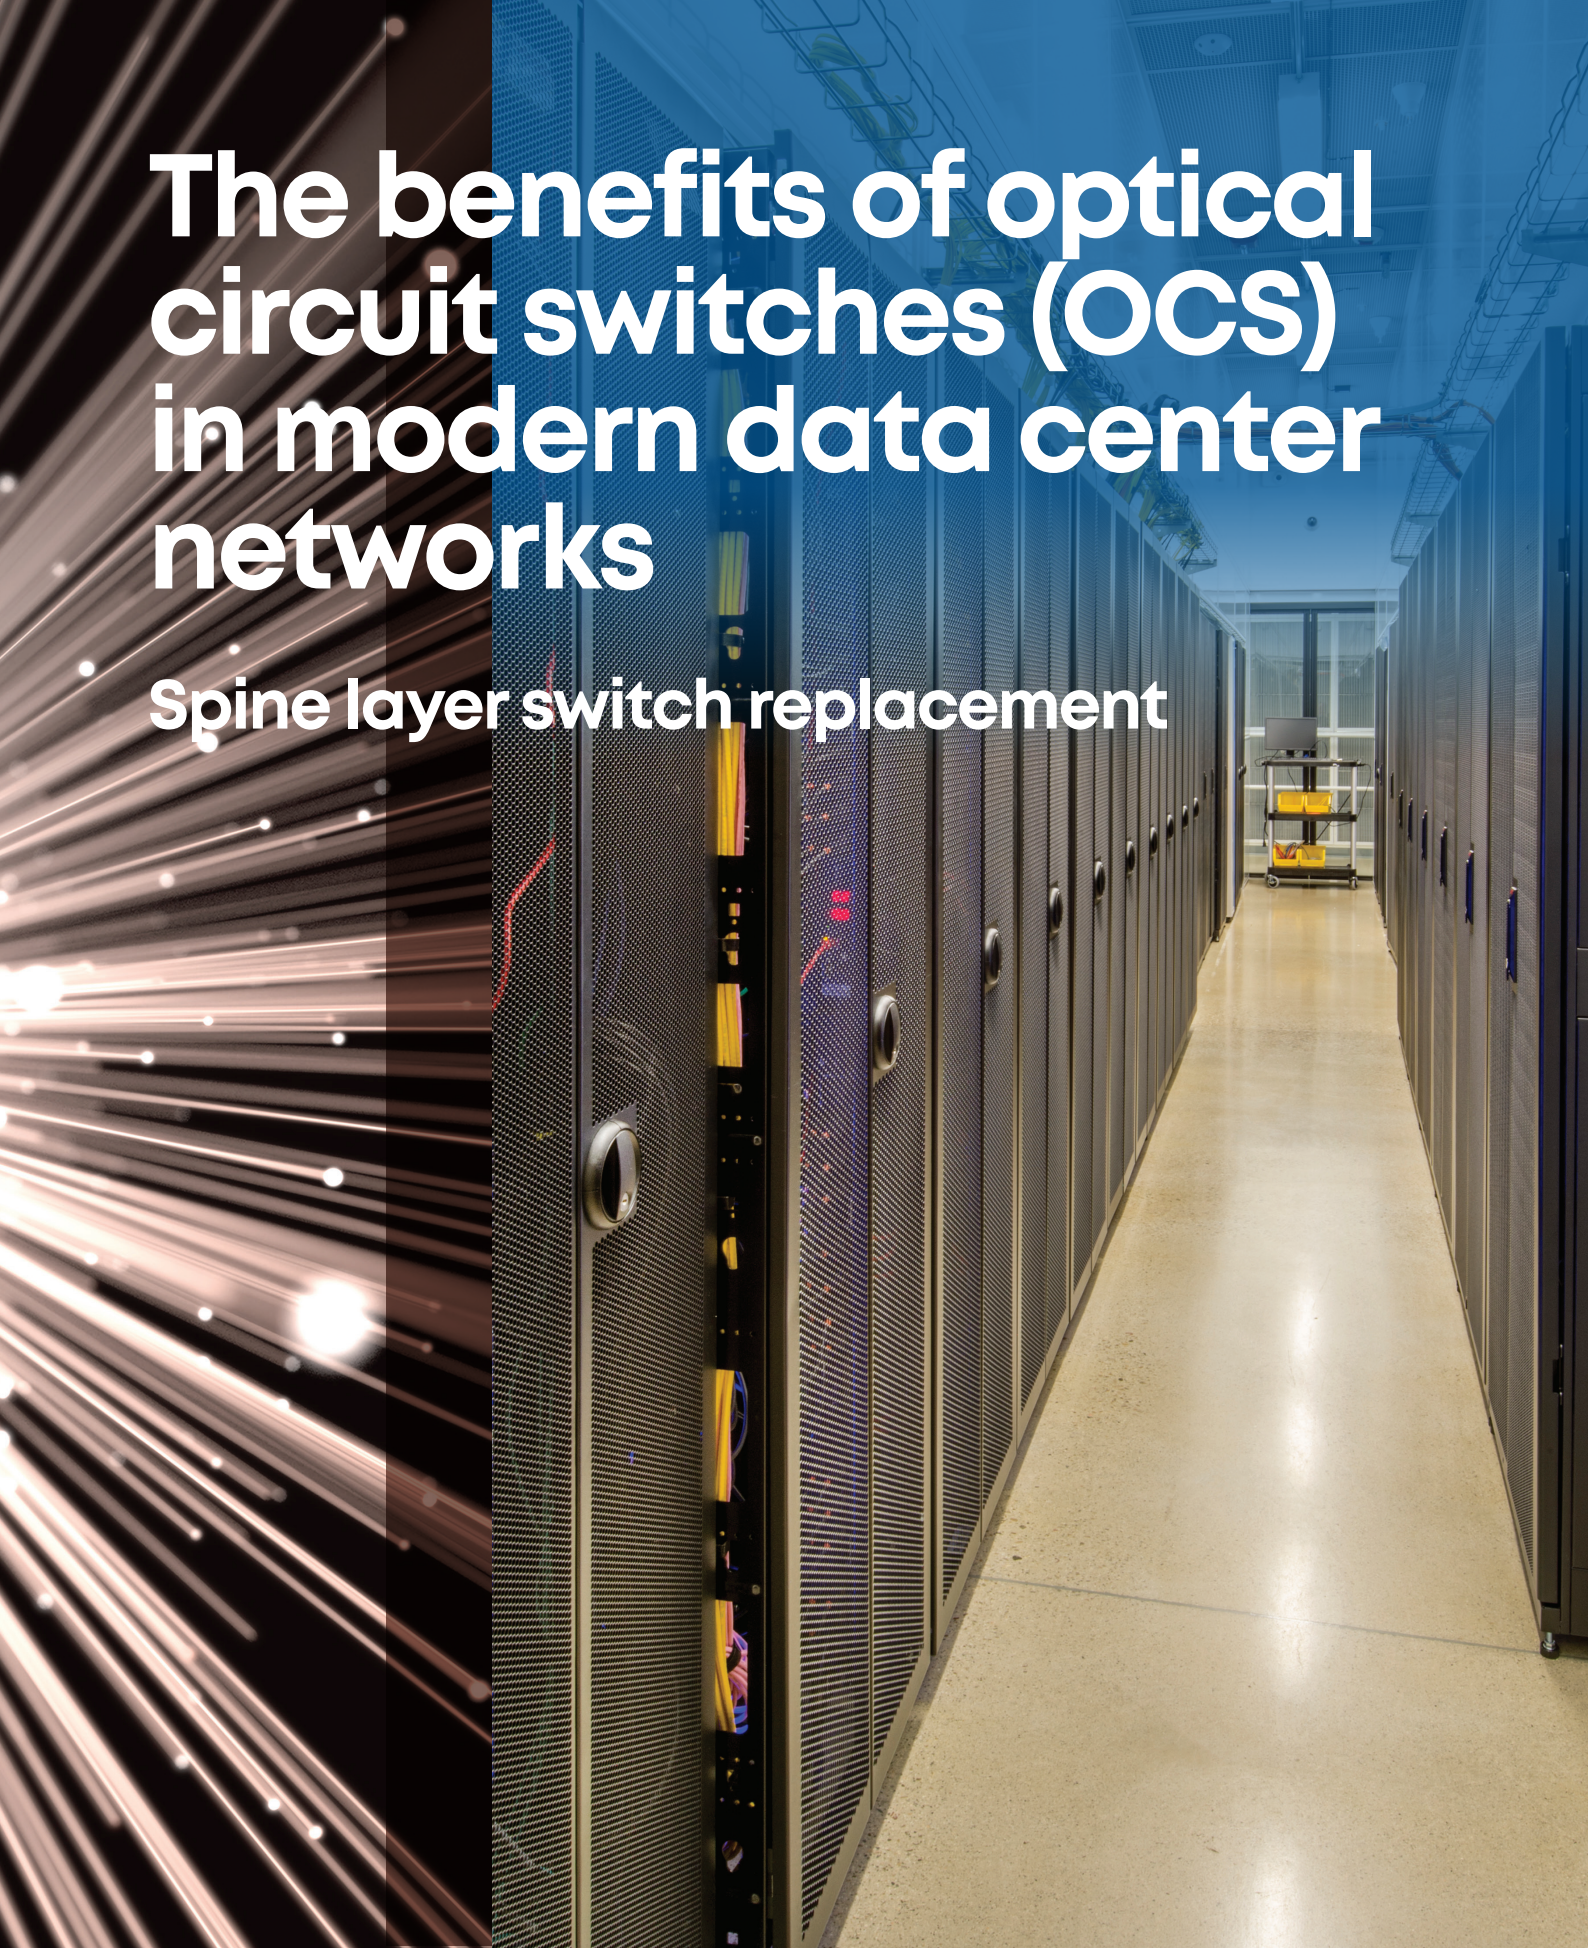

White Paper

HUBER+SUHNER

Introduction

Data center spine and leaf network architectures were developed to enable data centers to scale to meet the demands of ever-increasing data center traffic growth. The architecture has been widely deployed and is now the dominant network design in hyperscale data centers.

The spine-leaf architecture consists of two layers: the Spine (core backbone switches or spine blocks) and the Leaf (access switches or aggregation blocks). Spine switches connect all leaf switches, which in turn link directly to servers and network devices, facilitating high-capacity data transmission across the network. The interconnect ensures efficient data flow and supports enhanced network scalability without many of the bottlenecks associated with earlier architectures.

Traditional spine and leaf datacenter architectures use Optical-to-Electrical-to-Optical (OEO) switches in both the spine and leaf layers. While effective, traditional OEO switches are now struggling to keep up with the exponentially increasing server-to-server traffic generated by Artificial Intelligence (AI), Machine Learning (ML), and other compute-heavy emerging applications. The need for more agile and scalable datacenter infrastructure is becoming increasingly apparent.

Network architects have devised a promising solution to the infrastructure problem using Optical Circuit Switches (OCS) in the spine layer as replacements for the OEO switches. As data volumes and traffic speeds increase, the need for multiplexing with OEO switching diminishes, creating opportunities to exploit the full benefits of all-optical switching with OCSs. Replacing the OEO switches with OCS enhances the basic scalability of the spine and leaf architecture to meet the increasing exponential growth of shifting traffic patterns while dramatically reducing power usage and overall costs.

This white paper explores the challenges faced by current OEO-based networks, introduces OCS as a solution, and discusses the benefits of adopting OCS in modern data center architectures. For organizations looking to stay ahead in the fast-paced digital era, OCS, coupled with sophisticated Software Defined Networking (SDN) and orchestration, represents a crucial step towards achieving optimal network performance and efficiency. Data centers that embrace this technology will be well-positioned to meet the challenges of tomorrow, delivering superior performance and sustainability.

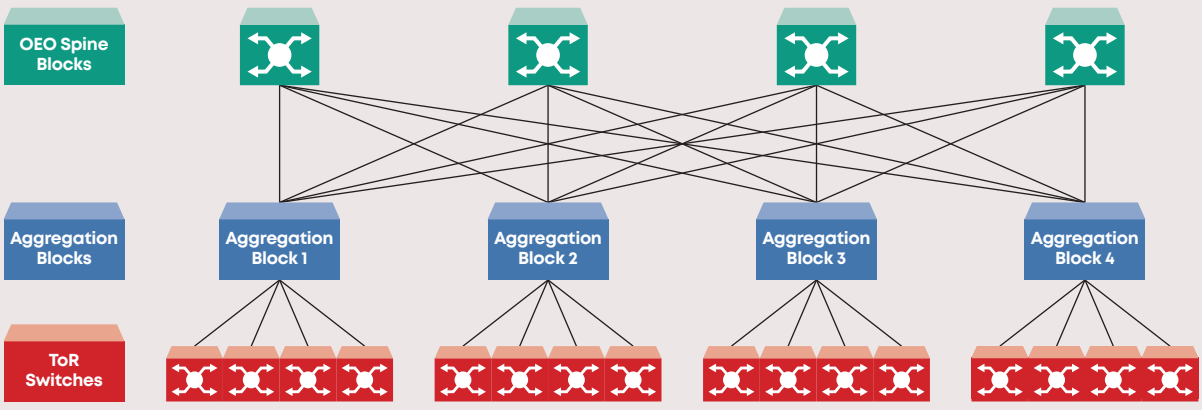

Traditional data center network architecture

Introduction

Challenges in current data center networks

The modern data center network is a complex ecosystem that requires high performance, scalability, and reliability to support a myriad of applications and services. Despite the technological advancements in network infrastructure, several key challenges persist with the use of traditional OEO switches.

Bandwidth limitations
OEO switches convert optical signals to electrical signals for retiming and switching, and then reconvert them back to optical for retransmission. These conversions limit signal bandwidth to specific transmission speeds and require replacing equipment when upgrading to higher traffic rates and more efficient transmission formats. As data traffic continues to grow at an unprecedented rate, these switches become bottlenecks, restricting the performance and scalability of the network. For instance, large-scale data transfers for Machine Learning, high-definition video streaming, and real-time data analytics demand bandwidth capa-cities that traditional OEO switches struggle to support efficiently.

Power consumption
The conversion process in OEO switches consumes substantial power. This high power consumption not only increases operational costs but also contributes to greater environmental impact due to higher energy usage and the need for heat dissipation. In an era where sustainability and energy efficiency are paramount, the high power requirements of OEO switches pose a significant challenge. Data centers are under constant pressure to reduce their carbon footprint, and inefficient power usage by network equipment directly impacts those efforts.

Latency issues
OEO switches introduce latency due to the signal retiming process, limiting the ability to support real-time applications like self-driving cars and remote surgery. As applications increasingly demand lower latency to ensure optimal user experiences, the limitations of OEO switches become more pronounced.

Scalability constraints
The need for more ports and higher throughput in data centers is constantly increasing. OEO switches face scalability issues, making it challenging to meet the growing demands without significant infrastructure overhauls, which are costly and complex. As businesses expand and data centers scale up, the limitations of OEO switches in handling increased traffic and port density become a critical bottleneck.

Solution

Optical Circuit Switches (OCS)
Optical Circuit Switches (OCS) offer a compelling alternative to traditional OEO switches by directly routing optical signals without the need for conversion. This direct optical switching can address many of the limitations faced by OEO switches.

Spine layer packet switches often shift large aggregate flows with little statistical (de)muxing necessary – especially where a lot of aggregation is going on in the leaf or aggregation blocks. With clever cloud server placement, courtesy of predictive analytics, and clever packet flow routing, courtesy of a fine-grained SDN controller, statistical (de)muxing can be minimized so that all of the traffic on a given port at the top of an aggregation block needs to go to only one other aggregation block port, in which case the optical circuit switches are sufficient. All the finer-grained packet switching that is still necessary is done in the Top of Rack switches and aggregation blocks.

Transitioning from packet switching to optical circuit switching in the spine layer

Data rate independence
Today's data centers are heterogeneous, operating at multiple data rates and bandwidths. The fixed data rate of the transceivers associated with OEO switches introduces significant restrictions that OCS overcome. OCS are agnostic as they are able to connect at multiple rates, introducing future-proofing and enabling faster scale up and scale out.

High bandwidth capacity
As optical circuit switching is data rate independent it can handle significantly higher data rates, leveraging the full capacity of optical fibers. This allows for more efficient use of network resources and supports the growing bandwidth demands of modern data centers. For example, in a scenario where large volumes of data need to be transferred quickly between storage systems and processing units, OCS can provide the necessary bandwidth to avoid congestion and ensure smooth data flow.

Power efficiency
By eliminating the need for optical-electrical-optical conversions, OCS drastically reduce power consumption. This leads to lower operational costs and a smaller carbon footprint, aligning with sustainability goals. Data centers that deploy OCS can achieve significant energy savings, contributing to both cost reduction and environmental sustainability.

Reduced latency
OCS provide near-zero latency by avoiding the signal conversion process. This is particularly beneficial for applications requiring real-time data processing and low-latency communication. For instance, in virtual reality applications where immediate feedback is critical, OCS can ensure that data is transmitted and received almost instantaneously, enhancing the user experience.

Scalability
OCS architectures are inherently more scalable, capable of supporting a larger number of ports and higher aggregate throughput. This makes them well-suited for the dynamic and growing needs of modern data centers. As businesses grow and their data requirements increase, OCS can easily scale to meet these demands without the need for frequent and disruptive infrastructure upgrades.

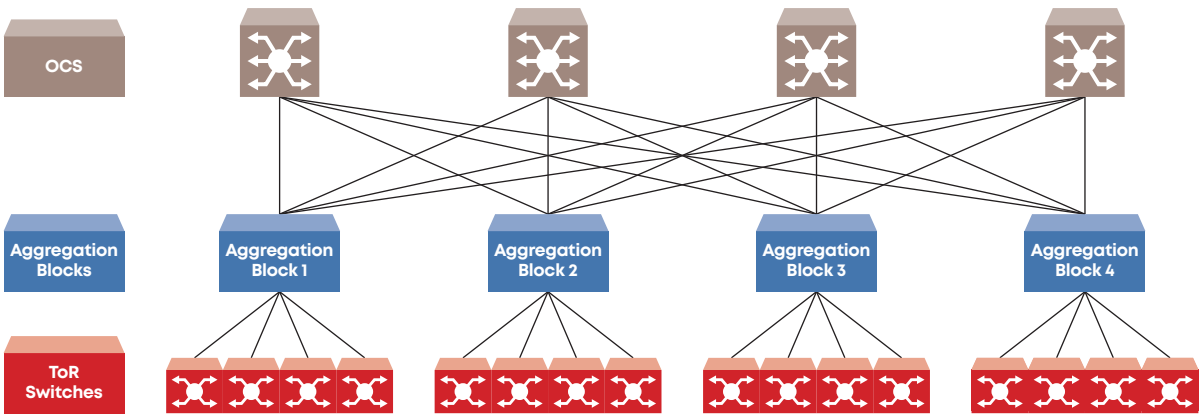

Network architecture – spine layer replaced with OCS

Using OCS in data center spine layers

Adopting OCS in data centers provides a multitude of benefits, enhancing overall network performance, efficiency, and sustainability.

Enhanced network performance
The higher bandwidth and lower latency of OCS improve overall network performance, enabling faster data transfer and more responsive applications. This enhanced performance is crucial for data-intensive tasks such as big data analytics, machine learning, and high-performance computing. For example, a data center supporting a large-scale AI project can benefit from the rapid data transfer capabilities of OCS, accelerating the training and deployment of AI models.

Cost savings
The power efficiency of OCS translates to significant cost savings in terms of both energy consumption and cooling requirements. Additionally, the scalability of OCS reduces the need for frequent and costly infrastructure upgrades. A study by one hyperscaler showed that a data center using OCS in the spine layer could reduce its energy consumption by up to 40% compared to one using traditional OEO switches.

Improved sustainability
The reduced power consumption and operational efficiency of OCS contribute to a more sustainable data center operation, supporting environmental responsibility initiatives. Data centers that prioritize sustainability can leverage OCS to contribute to achieving their green goals while maintaining high performance and efficiency.

Simplified network management
OCS can simplify network architectures by reducing the complexity associated with OEO conversions. This can lead to easier management, maintenance, and troubleshooting. Network administrators can benefit from a more streamlined infrastructure that is easier to monitor and maintain, reducing the risk of errors and downtime.

Future-proofing the network
The redistribution of data across storage devices or network paths is often necessary when adding new storage devices, upgrading network infrastructure, or after a system expansion, in order to optimize performance, increase redundancy or balance workload. By redistributing the data, operators can ensure that the data is evenly spread across the available resources, minimizing hotspots and enhancing overall system efficiency.

The process can be complex and may involve significant data movement, but it is crucial for maintaining optimal performance in a growing data center. As data traffic continues to grow, OCS provide a future-proof solution that can adapt to increasing demands without major overhauls. Their scalability and performance characteristics make them a strategic investment for long-term network evolution.

Data centers can deploy OCS with confidence, knowing that their infrastructure will be able to handle future data growth and technological advancements. Furthermore, as OCS are signal agnostic, there is no need to upgrade them as transmission speeds increase and protocols change.

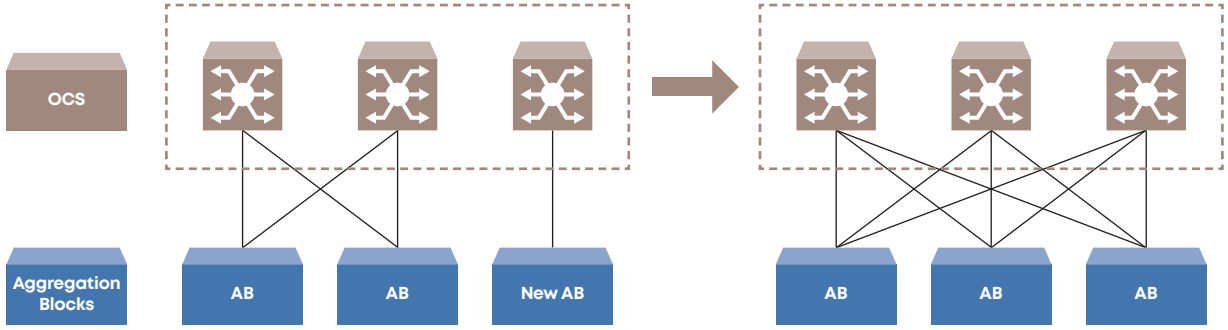

Automated network expansion to support new configurations and add capacity

Role of a sophisticated SDN/orchestrator

To fully realize the benefits of OCS in replacing spine layer packet switches, it is crucial to integrate a sophisticated Software-Defined Networking (SDN) controller or orchestrator. An SDN/orchestrator can dynamically manage and optimize the optical paths, ensuring efficient utilization of network resources.

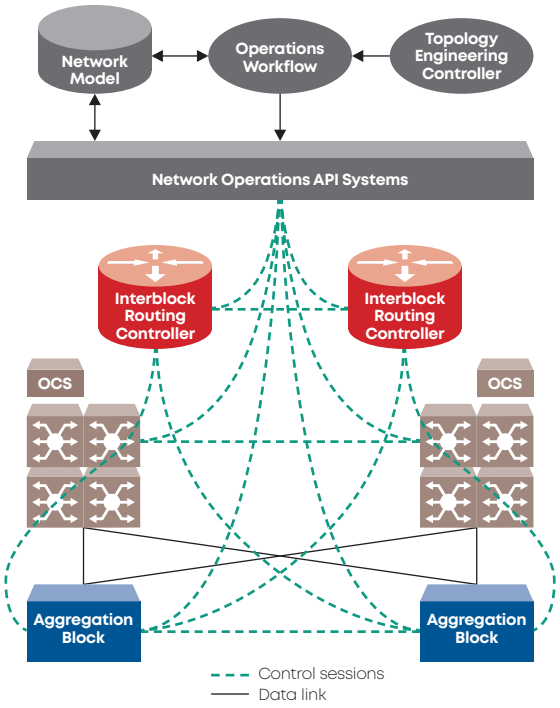

Example of a traffic management orchestrator

Dynamic path management
The SDN/orchestrator can dynamically allocate and manage optical paths based on real-time traffic demands, optimizing network performance and resource utilization. For example, during peak traffic periods, the SDN/orchestrator can reroute data flows to less congested paths, ensuring smooth and efficient network operation.

Automated provisioning
Automation of provisioning and configuration reduces manual intervention, lowering the risk of errors and improving operational efficiency. Data centers can benefit from faster deployment times and reduced human error, leading to more reliable and efficient network operations.

Traffic engineering
Advanced traffic engineering capabilities enable the SDN/orchestrator to balance loads across the network, preventing congestion and maximizing throughput. This is particularly useful in environments with variable traffic patterns, such as cloud data centers, where the SDN/orchestrator can ensure optimal resource utilization.

Enhanced visibility and control
The SDN/orchestrator provides comprehensive visibility into network operations, allowing for proactive monitoring and rapid troubleshooting. Network administrators can gain insights into network performance and health, enabling them to address issues before they impact service delivery.

Scalability and flexibility
The SDN/orchestrator can easily scale and adapt to changing network conditions, providing the flexibility needed to accommodate future growth and evolving application requirements. As data centers expand and new technologies emerge, the SDN/orchestrator with traffic engineering, based on predicting traffic flows and workloads, ensures that the network infrastructure can keep pace with these changes.

Conclusion

The transition from OEO switches to Optical Circuit Switches (OCS) in the spine layer of data center networks offers substantial benefits, including higher bandwidth, lower latency, greater power efficiency, and enhanced scalability. Integrating a sophisticated SDN/orchestrator is essential to fully harness these benefits, ensuring dynamic management and optimization of optical paths.

As the digital landscape continues to evolve, adopting OCS in conjunction with advanced SDN/orchestration solutions can ensure that data center networks remain agile, efficient, and capable of meeting the demands of the future while keeping one eye on their power consumption.

The HUBER+SUHNER advantage

POLATIS® has significant advantages over other all-optical (OOO) switching solutions for the spine layer, including:

- The industry's lowest optical loss and superior performance in stability.
- The broadest range of symmetric (NxN) switches, in matrix sizes from 16x16 to 384x384 ports, with modular scalability to connect thousands of fiber endpoints.
- High density switch matrices occupying very little rack space.
- Protocol and data rate agnostic so can switch signals of any type.
- Switching time <50ms for a single connection.
- Near-zero signal latency to support time-critical applications.
- True dark fiber switching, which requires no light to make and hold connections, enabling preprovisioning of future paths.
- Fully software-controlled for a seamless interface with leading orchestration solutions.
- Support for the broadest range of Software Defined Networking (SDN) interfaces including NETCONF and RESTCONF.
- Robust by design to be highly reliable for mission critical applications, with dual redundant, hot-swappable network interface cards and power supplies.
- Eco-friendly, low power consumption.

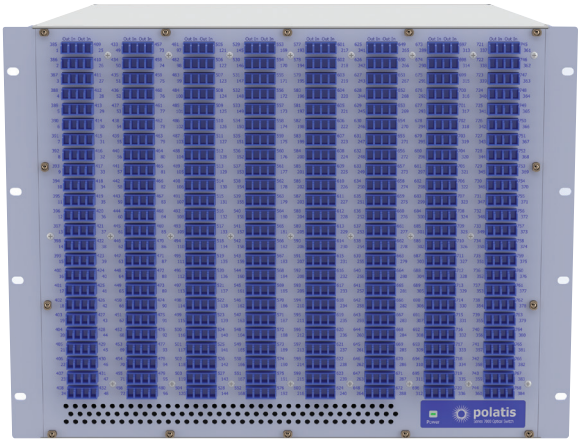

POLATIS® 7000 Series (384x384)

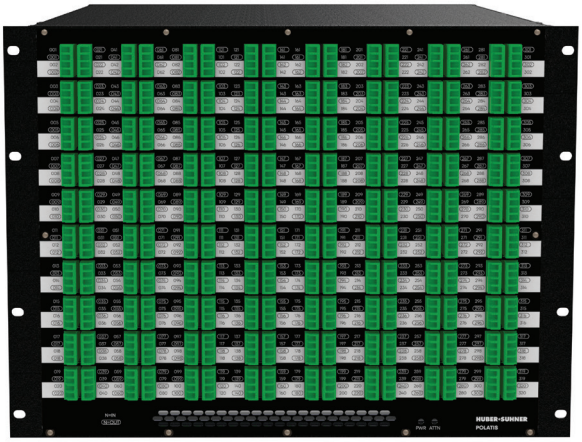

POLATIS® HS Series (320x320)

The HUBER+SUHNER advantage

The POLATIS® team has made a significant contribution to academic studies, projects and industry initiatives relating to new data center network architectures. Our network experts can provide advice and guidance on the best way to integrate optical circuit switches into your network architecture.

HUBER+SUHNER offers a broad range of products for data centers such as fiber cables, patch cords, fiber management, structured cabling solutions, POLATIS® optical circuit switches, transceivers, WDM components and more.

Worldwide sales and support are available to make sure data center systems continue to operate day in and day out.

HUBER+SUHNER
POLATIS® optical circuit switches
Americas: +1 781 275 5080
EMEA/Rest of World: +44 (0)1223 424200
info.polatis@hubersuhner.com
polatis.com
hubersuhner.com

4014/4115/06.2025