SDN-Enabled All-Optical Circuit Switching

White paper

Edition 2017

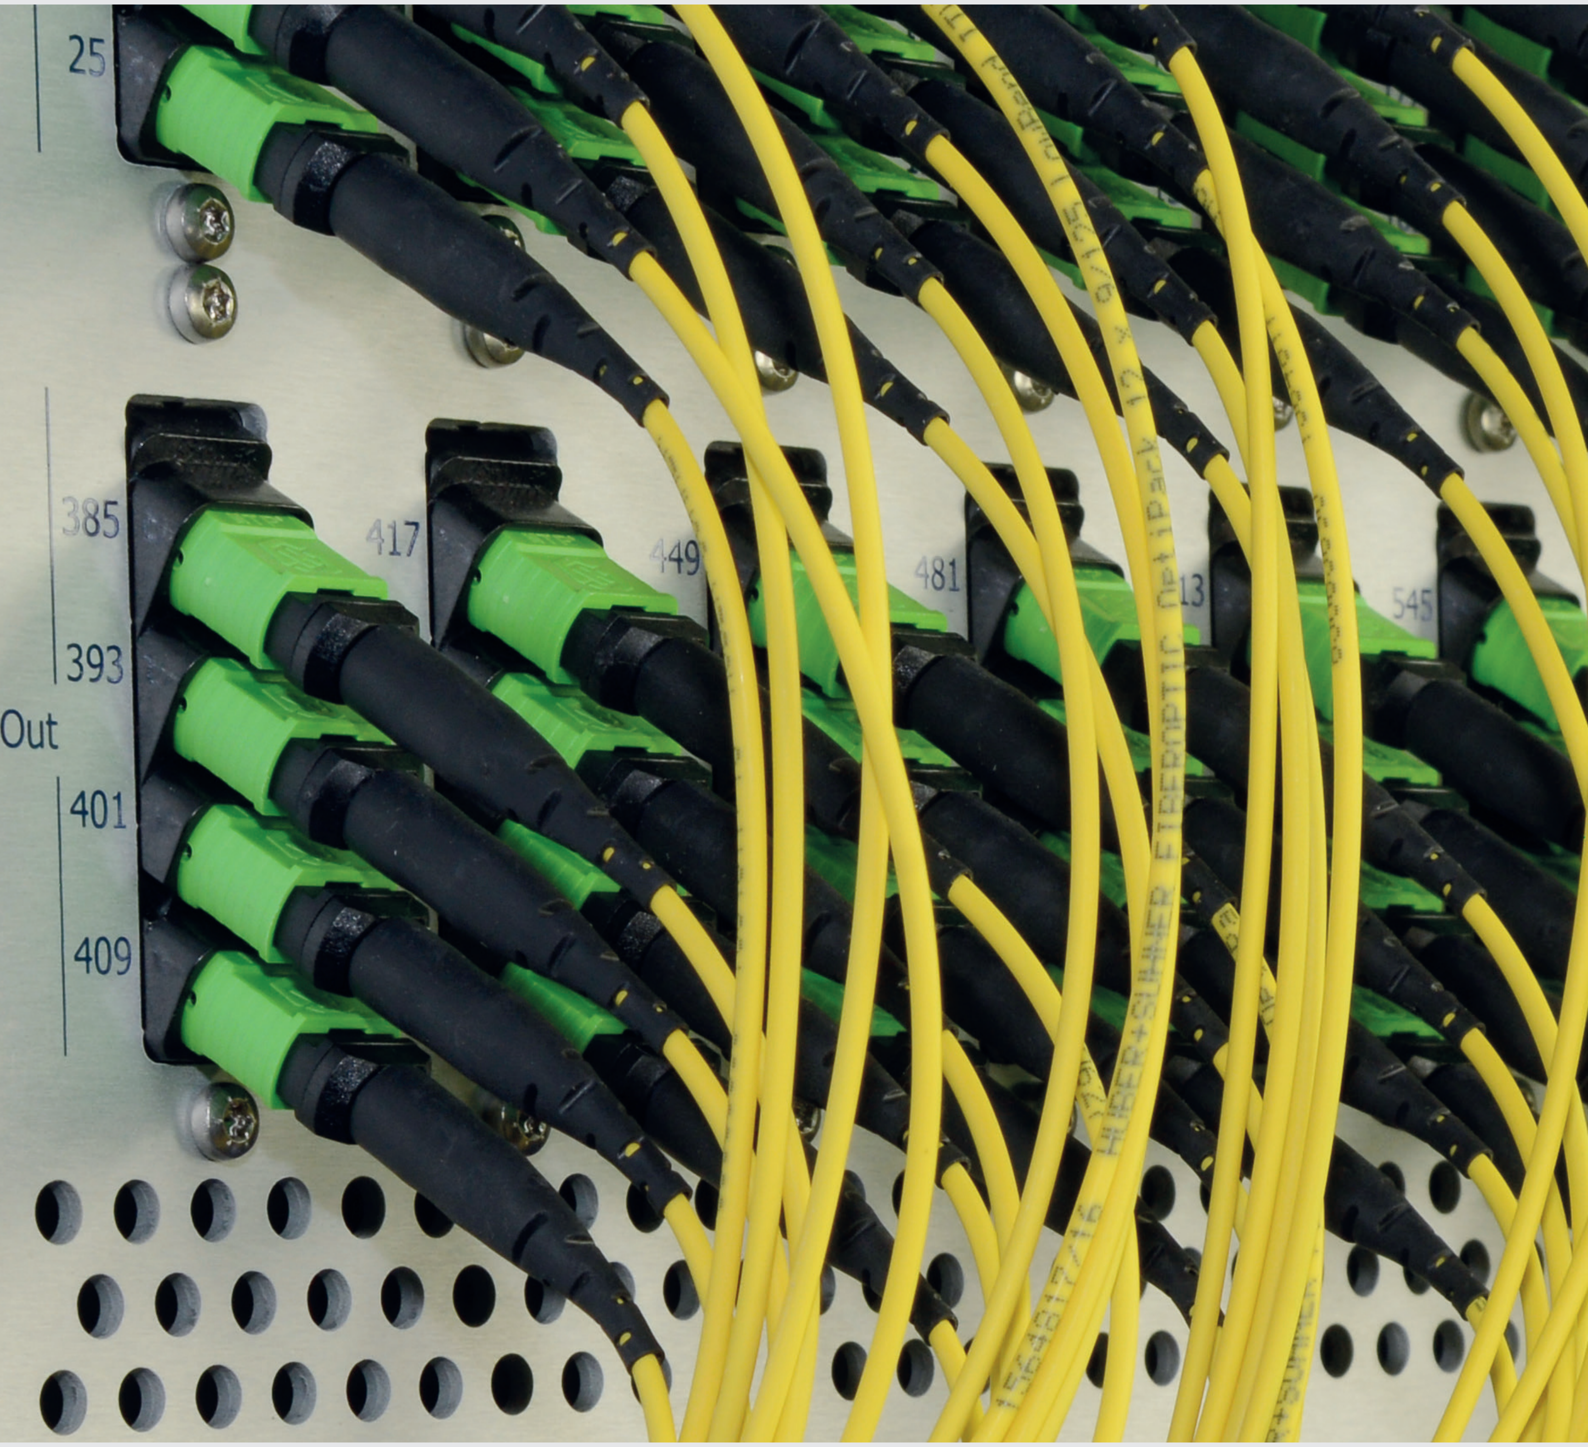

HUBER+SUHNER

Excellence in Connectivity Solutions

An Answer to Data Center Bandwidth Challenges

Executive Summary

Current data center network architectures cannot scale to meet the demands created by the proliferation of bandwidth-hungry mobile data applications and the associated rapid increase of internal (East-West) traffic volumes. The trend toward multi-tenant and disaggregated virtual data centers creates further challenges in provisioning Infrastructure as a Service (IaaS). Operators need more capacity in the same footprint together with the flexibility to allocate resources dynamically when, and to where, they are needed most. While optical interconnects are widely deployed to provide high speed transmission between processing elements in the data center, switching is still performed electronically, with fiber layer connectivity remaining mainly static.

Recent developments in all-optical circuit switching in combination with Software Defined Network (SDN) paradigms create compelling solutions to bring the fiber layer under software control. These solutions can provide dynamic, low-loss connectivity on demand between many housands of endpoints with speed-of-light data latency. SDN integrates the management of fiber-layer (Layer 0/1) optical circuit switches and conventional Layer 2/3 packet switches and routers under a common control plane to facilitate abstraction and virtualization of network resources. By augmenting network designs with SDN-enabled dynamic fiber cross-connects, operators can create scalable solutions that respond dynamically to changing business needs while reducing operational costs by automating service provisioning and bypassing bottlenecks.

This White Paper highlights trends in software-defined data center network architectures and the new capabilities that are realized with SDN-enabled low-loss all-optical switching solutions.

Data Center Network

Data center operators face a host of challenges in addressing the explosive growth of network traffic. The steady rise of mobile, video and cloud-based virtual services, is rapidly outpacing the capabilities of current data center architectures. The resulting network congestion—and outright network overload—increases latency, degrades service and application performance; and ultimately accounts for an exponential escalation in operating expenditure.

To date, this problem has typically been addressed by deploying more packet routers and switches to ease congestion and ensure the availability of bandwidth when needed. But, this approach leads to an unsustainable cycle of increased cost and is challenging to scale. As an alternative, data center architects are designing new networks that combine all-optical circuit switching with L2/L3 packet switches and routers to bypass key congestion points and to support exploding capacity requirements. These networks improve scalability while reducing costs, rack space and power consumption.

The Potential of All-Optical Swichting

Operators are searching for bandwidth-on-demand capabilities that allow them to direct network capacity where and when it is needed, with performance guarantees to support latency, jitter and availability requirements for a range of diverse applications. Adoption of dynamic fiber cross-connects to bring the fiber layer under software control allows rapid provisioning, protection and reconfiguration of network resources on-demand and bring additional benefits to operators, including

- Eliminating manual patch errors and the potential for service interruption
- Maintaining current state of fiber layer connectivity in a software database
- Creating optical demarcation points in multi-tenant/ multi-service provider environments
- Facilitating bridge-&-roll during equipment commissioning, upgrade and replacement
- Providing physical isolation between virtualized network slices for enhanced security
- Enabling aggregation of optical taps for network monitoring

Since all-optical circuit switches do not require any optical-to-electrical conversions, they consume very little energy and add virtually no latency to the data path. Connections are fully transparent and format independent, which makes them ideal for use in future-proof data center network infrastructures where optical transmission rates continually advance.

While individual optical circuit switches can provide up to a few hundred ports of non-blocking connectivity, scaling beyond this level requires multi-stage cross-connect fabrics where the optical loss must be minimized to work with the restricted power budgets of low-cost datacenter optical transceivers.

However, by using today's best-in-class high-radix optical switch technology it is now possible to create 3-stage folded-Clos optical cross-connect fabrics that can scale incrementally to tens of thousands of ports but with optical losses of just a few decibels that fit within the power budgets of standard data center optical transceivers. Data latency is minimized since the transparent end-to-end path keeps all traffic in the optical domain. And because this optical switch technology can pre-provision dark fiber, there is no concatenation of set-up delay through the 3-stage fabric.

An Answer to Data Center Bandwidth Challenges

SDN-ENABLED ALL-OPTICAL SWITCHING

In today's Three-Tiered Data Center Network Architectures SDN-enabled all-optical switching is being introduced at a number of levels to optimize network resources in the data center, alongside existing packet routers and switches. In the applications outlined below, all-optical switching is not a replacement for packet switching and routing; it is added to the existing network to bring the capacity, flexibility and scalability needed at key congestion points to meet heavy and persistent traffic demands. Large data centers have clusters of tens of thousands of servers distributed across rooms, buildings and campuses. To provide reliable and scalable interconnections between all of these servers, data center networks are typically built up in tiers using a combination of packet routers and circuit switches at the core/metro, aggregation and access levels. This hierarchy enables efficient sharing of data and provides services across server farms while also connecting to outside networks. All-optical switching can be added to each tier of the data center to enhance performance and reduce costs. Figure 1 shows the basic tiered approach including potential locations for all-optical switching in each tier.

Low-Latency peering in the Core/Metro Tier

Peering allows data center operators to interconnect their networks to exchange traffic, typically using Layer 2/3 switches between the metro/core and peering arbitrators. In this function, transferring large blocks of traffic reliably and efficiently with low latency is critical.

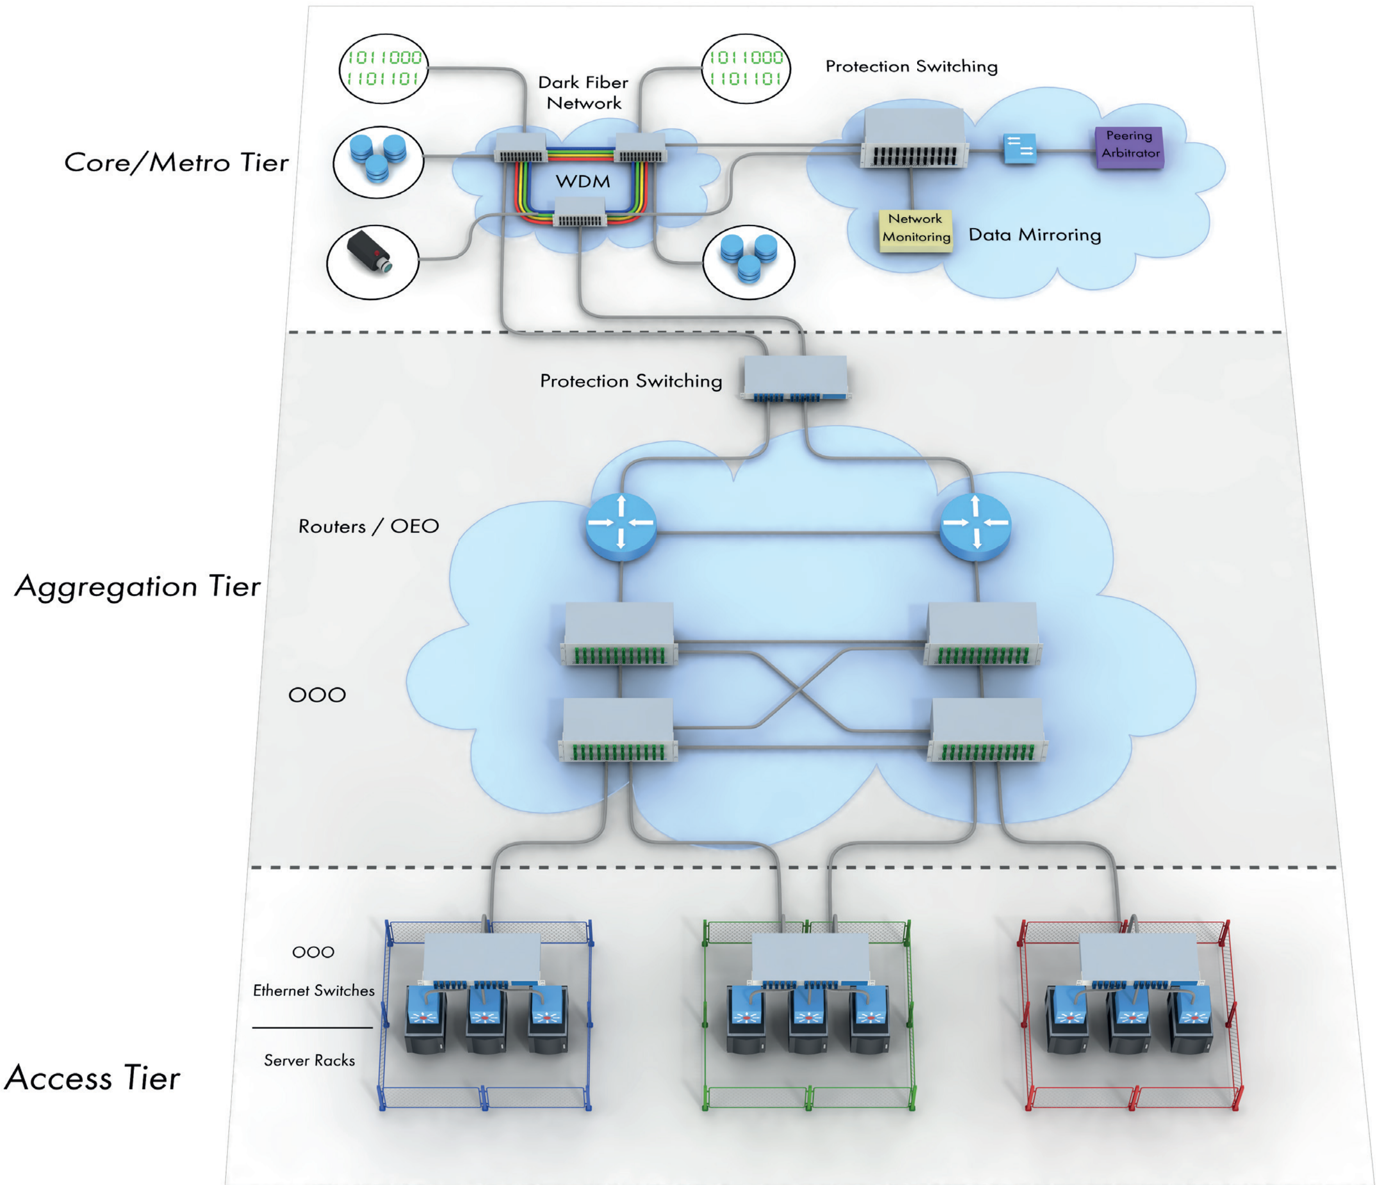

Figure 1: Locations for all-optical switching in the data center

An Answer to Data Center Bandwidth Challenges

All-optical switching in this layer provides both traffic provisioning and protection switching between the external network and the peering arbitrator. In addition, the optical switch can provide a number of value-added services including intrusion detection, lawful intercept and fault location. To perform data mirroring and network monitoring functions, optical taps that split off part of the signal can be integrated with the switch fabric.

Optical power monitors can also be integrated within the switch to monitor optical signals between the network and the peering switches. The integrated optical detectors monitor the optical power in all fibers carrying peering traffic to quickly detect and locate optical fiber and equipment faults. In the event of an equipment or fiber failure, the optical circuit switch can restore service by protection switching around the fault. The protection switching can be initiated by the higher-level data center control plane or performed automatically by the switch itself to minimize any service interruption.

Optical layer automated protection switching can be configured to perform 1:1, 1+1, 1:N or M:N protection of fibers and equipment. Furthermore, combining multiple functions into a single optical switch saves precious rack space and reduces operating expenditure while greatly enhancing the services offered and return on investment.

Router Bypass in the Aggregation Tier

Data center aggregation is often performed by stacks of packet routers and switches that connect server farms and provide access to the core/metro networks. Routers also host many other functions including connectivity, discovery, firewall, load balancing and intrusion detection.

Current data center aggregation architectures are optimized for handling the short, bursty traffic patterns that have dominated in the past. However, they are neither efficient nor cost effective for handling the larger persistent data flows between server clusters that are increasingly common with virtual cloud-based services. These persistent data flows are now common in routine operations like virtual machine migration, load balancing and data storage backup. A number of studies have shown that these so-called "elephant" flows are now dominating cluster-to-cluster traffic, with over 90% of network capacity being occupied with flows that last more than 10 seconds. Offloading elephant flows to dynamically-provisioned optical circuits provides a low latency, high speed path between endpoints and simultaneously relieves congestion at the packet layer.

Figure 3 shows how including all-optical circuit switches in the aggregation tier can create connections for persistent data flows that do not need the full services of a packet router or switch. This approach avoids the buffering and latency issues of going through routers because data traffic moves through an optical switch at the speed of light; no buffering is required. By playing to the strengths of each technology, a hybrid architecture combining packet routers and switches with all-optical

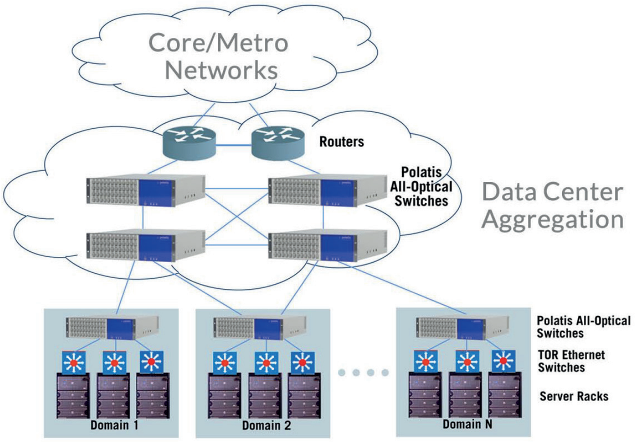

Figure 3: All-optical switches in the data center aggregation tier enabling new virtual cloud services.

switches allows the data center to handle both short bursty traffic and larger persistent data flows optimally. This is commonly referred to as router bypass.

Server Farm Aggregation Tier

The third data center application (shown in Figure 4) is server farm aggregation and it mirrors the aggregation tier application but on a smaller scale. Server farms are often individual racks that house multiple data servers interconnected with Top of Rack (TOR) Ethernet switches.

Large numbers of server racks are arranged into clusters or server farm domains, and TOR switches are connected to the data center aggregation tier by optical transceivers. In the server farm environment, a premium is placed on protecting both data and data access Large server farm domains may contain thousands of servers spread across a variety of physi-

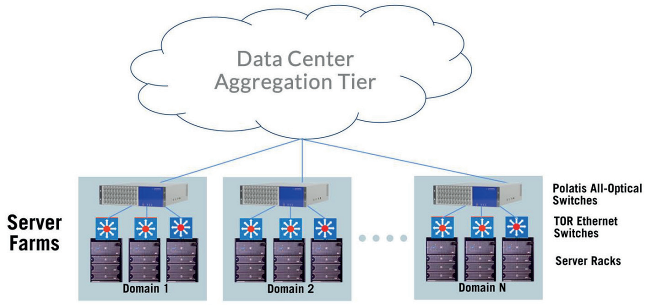

Figure 4: All-optical switches for server farm aggregation.

An Answer to Data Center Bandwidth Challenges

cal locations. Properly managing equipment cost, space, power and cooling while continuously improving processing efficiency is critical to the success of the data center operation.

As server farms grow larger, and more and more services become cloud based, the need for server farm aggregation grows as well. Introducing a reconfigurable optical layer can significantly improve throughput by improving the efficiency of communications between server farms and enabling disaggregation of compute, memory and storage resources. This reduces costs by increasing overall service delivery efficiency along with equipment utilization. In the event of a major disaster, the enhanced physical fiber layer switching also allows significantly faster recovery and restoration of services.

All-Optical Switching and Software defined networking

Software Defined Networking (SDN) is driving the adoption of Layer 0/1 optical circuit switching in the data center. A growing number of multi-tenant data center operators are looking to SDN to improve current operations and accelerate the development of new capabilities to support cloud-based virtual services.

In today's data centers, much of the networking equipment is managed and controlled with vendor-specific proprietary software. This limits innovation because the introduction of new capabilities is dependent on equipment vendor development cycles, meaning that it can take months to design, test and field new services. To add features, or to integrate new technologies, likewise requires proprietary system upgrades that must be done individually with each equipment supplier.

The power of SDN is that it allows operators to write high level software applications based on open standards that automate basic network operations and facilitate new service creation; which can be done in hours or days rather than weeks or months. The control flexibility of SDN greatly simplifies integration of new technologies like all-optical switching into the existing network infrastructure. SDN's ability to dynamically monitor and react to network changes allows operators to take greater advantage of the inherent capabilities of optical switching.

SDN offers a holistic way of managing the components of the network and a new way of orchestrating network operations to reduce costs and speed up service innovation and delivery. With SDN, operators can adjust network behavior in real-time and deploy new applications and network services in a matter of hours or days, rather than weeks or months.

The promise of SDN is much more rapid and dynamic deployment of customized services and applications. SDN gives data center operators more flexibility to reconfigure the network, allocate capacity where it is needed, and quickly react to unexpected traffic loads and other network dynamics. This makes for a much more efficient use of resources, improving utilization and ultimately lowering costs. In SDN architectures, network intelligence is centralized in virtualized SDN controllers that maintain a global view of the network. The network controller communicates in a vendor-independent way with SDN enabled switches or routers via standards-based protocols such as NETCONF, RESTCONF and OpenFlow.

Promoted by the Open Networking Forum (ONF), OpenFlow is an emerging industry standard for the communications interface between the control plane and network elements. Figure 5 shows the logical view of the SDN architecture defined by the ONF. In this example, the network data plane infrastructure, including L0/L1 optical circuit switches as well as L2/L3 packet routers and switches, is all managed by the SDN control plane via OpenFlow. Each network device has an embedded OpenFlow agent that allows direct interfacing with the SDN controller. Vendor-independence is one of the key advantages of SDN, and it helps operators avoid being locked into a particular vendor's proprietary solution.

Figure 6 shows how Polatis' SDN-enabled all-optical switching can optimize resources in a multi-vendor data center network. Here we show a hybrid packet-optical data center aggregation layer operating under an OpenDaylight SDN control plane using OpenFlow and NETCONF protocols to manage the Polatis optical circuit switches and other network elements. This arrangement is just one representation of a configuration-on-demand architecture that optimizes the utilization of data center resources. In this scenario the topology and properties of the Polatis optical switch layer are published to the SDN controller and orchestration layers via a plug-in residing in the service abstraction layer.

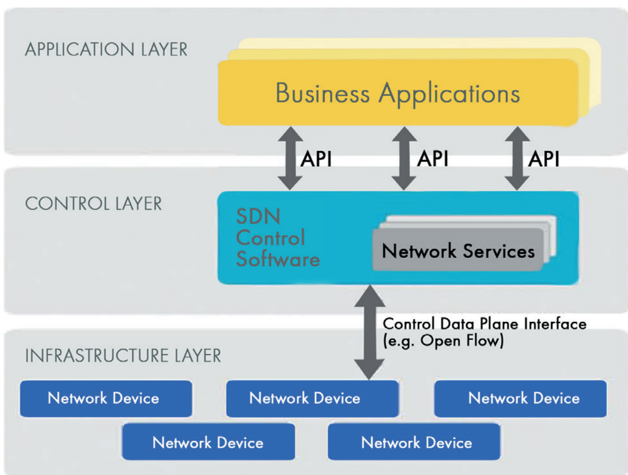

Figure 5: Logical view of a Software Defined Network (source: The Open Network Foundation).

An Answer to Data Center Bandwidth Challenges

In this arrangement, persistent East-West data flows can be identified by the Packet-Circuit FlowMapper in the SDN controller using data analytics calculated from monitoring flows between network endpoints. Once identified, the controller can dynamically reconfigure the network using OpenFlow and NETCONF to create a low latency optical circuit for these large data flows through the optical switch layer. Optical switching is ideally suited to managing persistent data flows that do not need any sub-wavelength grooming. Offloading these elephant flows to the optical circuit switch layer relieves congestion and reduces latency through L2/L3 packet-switched routers, thereby substantially improving the efficiency and throughput of the network infrastructure

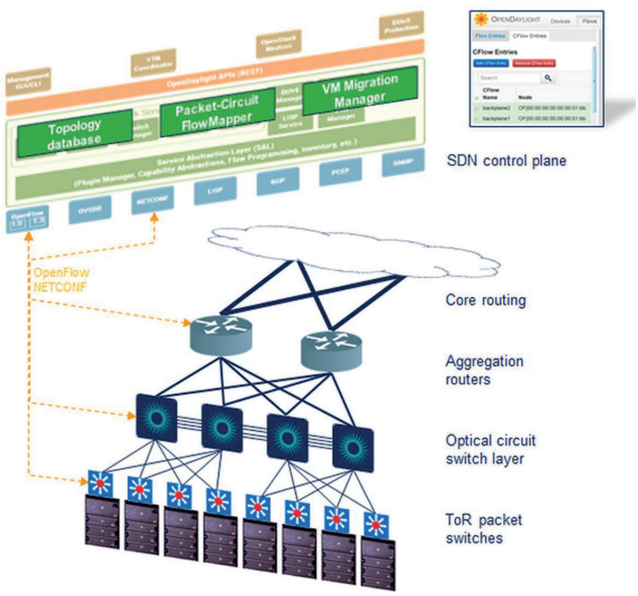

Figure 6: Data center operators can use SDN-enabled all-optical switching to optimize network resources.

Conclusions

SDN-enabled all-optical switching is poised for larger scale deployment in many data center networks. Dynamic fiber cross-connects provide operators with flexibility at the fiber layer to allow their networks to scale and respond dynamically to changing business needs, to automate more basic operations and to reduce the cost of running their networks.
The increasing adoption of SDN technologies also removes barriers to innovation at the transport layer and makes it easier for operators to improve competitiveness by integrating novel optical circuit switching solutions in a multi-vendor network infrastructure.

About POLATIS' SDN-Enabled All-Optical Swichtes

Polatis delivers the world's lowest loss all-optical switching solutions for remote fiber-layer provisioning, protection, monitoring, reconfiguration and test, with over 5 billion port-hours in service to date. Built on the dependable, field-proven DirectLight™ optical matrix switch technology, Polatis dynamic fiber cross-connects scale from 4x4 to 384x384 ports and make fully transparent connections with low loss and no back reflection that are entirely independent of the color, direction, power level or data rate of traffic on the fiber.

As the performance leader in optical circuit switching technology, Polatis believes that dynamic fiber cross-connects are a key element in enabling automation and virtualization of software-defined optical network infrastructure. Polatis is committed to providing its customers with open and programmable interfaces that make it easy to integrate our products with open-source SDN controllers. Polatis optical circuit switches are ideal for software-defined network applications where connection loss, stability and reproducibility are critical. Polatis dynamic fiber cross-connects have a full-featured embedded OpenFlow agent that is easily integrated with popular SDN controllers such as OpenDaylight. Polatis is also working with leading transport SDN vendors to support emerging standards and ensure that our customers benefit from the most dependable and cost-effective SDN solutions.

For more information, visit us at www.polatis.com
or call +1 844-POLATIS.

North American Headquarters
HUBER+SUHNER Polatis, Inc.
213 Burlington Road
Suite 123
Bedford, MA 01730
U.S.A.

+ 1 781 275 5080 phone
+ 1 844 765 2847 toll free
+ 1 781 275 5081 facsimile
info@polatis.com

European Headquarters
HUBER+SUHNER Polatis Ltd
332/2 Cambridge
Science Park Cambridge CB4 OWN
United Kingdom

+ 44 1223 424200 phone
+ 44 1223 472015 facsimile
info@polatis.com

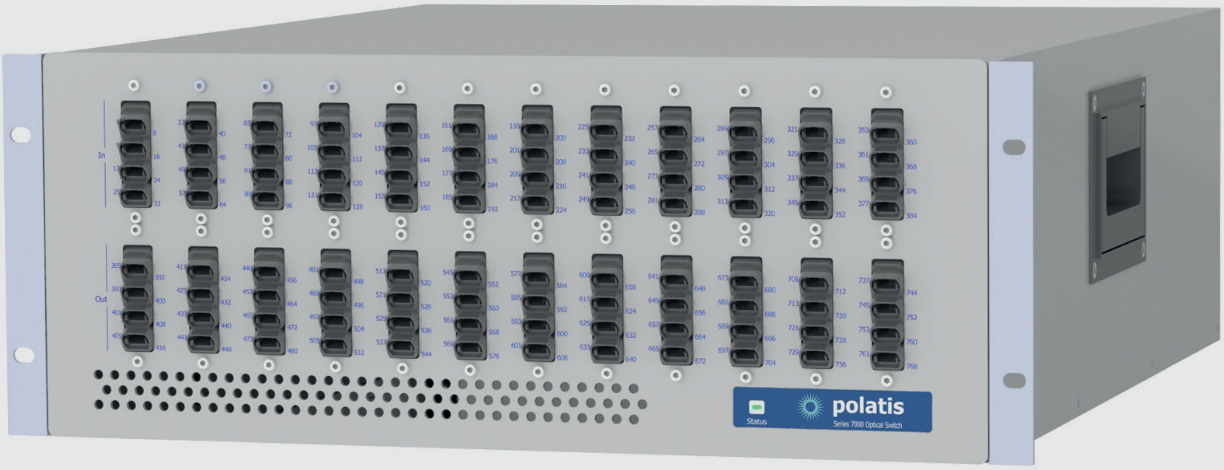
Status
polatis
Series 7000 Optical Switch

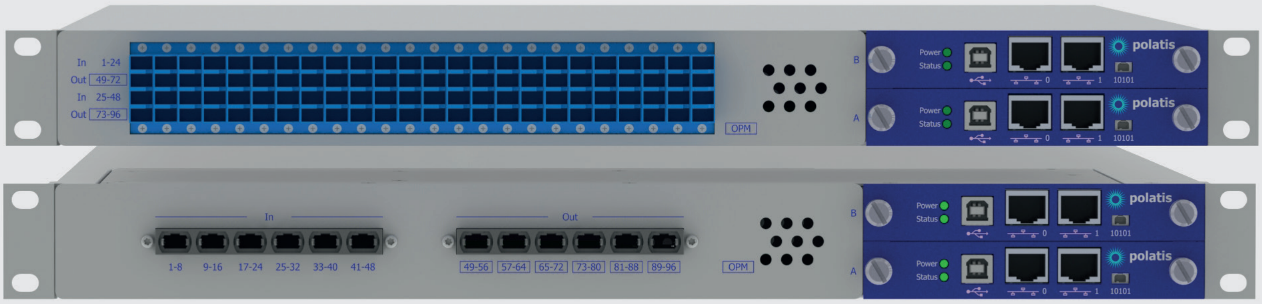
In 1-24
Out 49-72
In 25-48
Out 73-96
OPM
Power
Status
polatis
10101
In
Out
1-8
9-16
17-24
25-32
33-40
41-48
49-56
57-64
65-72
73-80
81-88
89-96

HUBER+SUHNER AG
Division Fiber Optic
Degersheimerstrasse 14
9100 Herisau
Switzerland
Tel. +41 71 353 4111
Fax +41 71 353 4444
hubersuhner.com

4291/4743/10.2017

HUBER+SUHNER is certified according to ISO 9001, ISO 14001, ISO/TS 16949 and IRIS.