White Paper

Making 100G OTN Economical:

OTN Switching & Packet-Optical Transport

Microsemi
Power Matters.™

Microsemi Corporate Headquarters
One Enterprise, Aliso Viejo,
CA 92656 USA

Within the USA: +1 (800) 713-4113
Outside the USA: +1 (949) 380-6100
Sales: +1 (949) 380-6136
Fax: +1 (949) 215-4996

E-mail: sales.support@microsemi.com

Abstract

Network operators worldwide are looking to 100G optical transport technologies to address unprecedented growth in bandwidth demands driven by Big Data and continued growth in mobile broadband and on-demand video. To unlock the full value of their 100G investments, network operators are beginning to turn to OTN switching as they seek to maximize network efficiency and resiliency while scaling their WDM infrastructure. Microsemi innovation in 100G OTN switching technology is critical to making this network vision a reality.

Contents

1 Introduction

Globally, mobile and internet traffic is growing at an unprecedented rate, driven by the rise of cloud services, in addition to mobile broadband and on-demand internet-based video services. Cisco predicts that from 2012 to 2017, global internet IP traffic will have tripled to 120 Exabytes per month and mobile data traffic will have grown 13 fold to greater than 11 Exabytes per month.[i] During the same period, Bell Labs predicts that traffic in Metro and Backbone Networks will grow by 560% and 360% respectively. [ii] This bandwidth growth is placing an incredible strain on the current generation of optical transport networks, which form the underpinnings of our modern IP networks. With the cost of deploying new fiber prohibitive, network operators have needed a way to substantially increase the capacity of their networks without laying new fiber.

Recent advances in optical transmission technologies, particularly coherent optical detection, enable the 10G wavelengths of today's typical DWDM transport network to be upgraded to 100G wavelengths without the need to trade-off optical reach, make costly fiber-upgrades, or add dispersion compensation equipment. This boon for network operators allows for a 10-fold increase in fiber capacity to 8Tb/s or more while leveraging existing fiber infrastructure. However, for a technology like this to make economic sense, the network must be capable of utilizing this new capacity efficiently; doing so is not without its challenges.

2 Challenges: Maximizing 100G Investments

In recent years, network operators have relied heavily on ROADM-based optical transport equipment, leveraging fixed point-to-point WDM connectivity to aggregate and transport client services throughout their Metro and Long-Haul networks on 10G wavelengths. These networks, if carefully engineered, can make reasonably efficient use of available optical spectrum. Transponders are used to map clients with bandwidths equivalent to the WDM line rate (i.e. 10GE or OC-192/STM-64 clients) onto 10G wavelengths, whereas muxponders are used aggregate sub-10G services (i.e. GE, OC-3/12/48, STM-1/4/16 clients) into 10G wavelengths for transport. With 10G services representing 50 percent or more of the ports feeding optical networks in 2013, the average service rate closely matches the WDM line rate, allowing for efficient use of available optical network bandwidth.[iii]

2.1 Poor Bandwidth Utilization

Despite the projected massive growth in network bandwidth demands, the overall distribution of client port rates feeding optical networks is not expected to change drastically. Even by 2017, 10Gbit/s or below end-to-end managed services will still represent greater than 95 percent of the ports feeding optical transport networks.[iv] A transition to 100G DWDM wavelengths creates significant rate mismatch discontinuity between client and WDM uplink ports, creating the potential for large amounts of inefficiently used optical spectrum.

Muxponders are a static solution, in which client ports are associated with a fixed WDM line port, and thereby a specific wavelength. This solution works well when all nodes in a network experience equal and unchanging service demands. However, in real-world network deployments, services evolve and node-to-node demands are unpredictable in both mix and distribution. It can become an operational challenge for network operators to avoid stranding optical bandwidth in their costly DWDM network.

For example, if an intermediate Central Office (CO) node in a WDM or ROADM network has a requirement for to add 30G worth of new 10G client services and all available client ports have been exhausted at the node, the network operator has two options:

1. Deploy back-to-back muxponders to de-multiplex up to 70G worth of pass-through optical traffic and aggregate it along with the new 30G of client services onto a 100G wavelength, or

Figure 1: Back-to-Back Muxponders

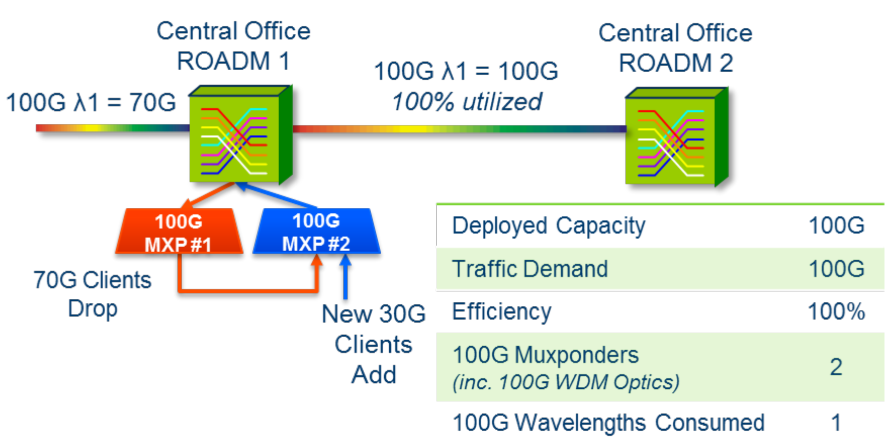

2. Deploy a new 100G wavelength that is only 30 percent utilized.

Figure 2: Underutilized 100G wavelengths

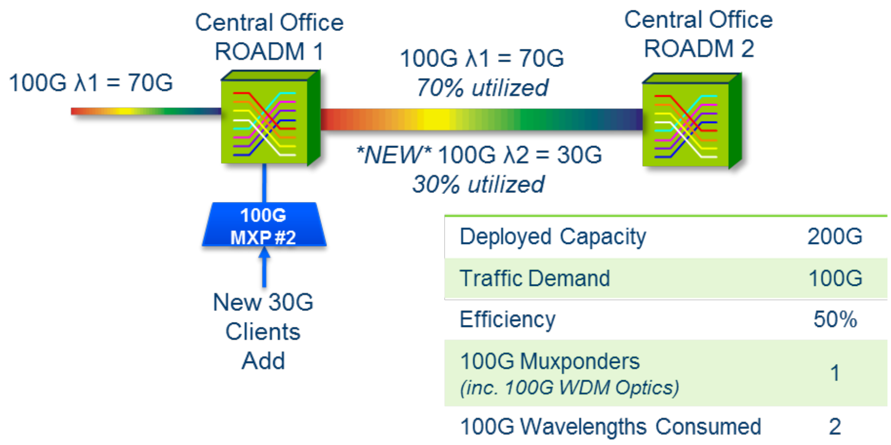

The first scenario results in better use of available 100G optical spectrum, but at higher equipment cost and introduces additional considerations, such as end-to-end link latency. This option is also very challenging to plan and manage given the dynamic nature of new cloud-based service demands on the network.

The second scenario, while easier to manage and deploy, requires that the operator light up a new 100G wavelength, consuming precious optical spectrum and potentially resulting in higher fiber infrastructure CAPEX spend if available fiber pairs have been exhausted at that node. This option also results in significant stranded network bandwidth, as the new wavelength will be only 30 percent utilized at the time of deployment.

Operationally, neither solution is ideal. Not only are clients fixed to a specific uplink wavelength, but since ROADM-based networks can only switch entire wavelengths, operators must send technicians into the field to manually perform any network reconfiguration at the service level, such as client re-routing or network re-optimization. Cloud services and data center interconnect can be particularly dynamic in their bandwidth requirements and thereby challenging to support with point-to-point transport infrastructure.

With both optical capacity and CAPEX at a premium, many operators will be looking for network architectures that address these limitations.

3 Making 100G Work: OTN Switching

Akin to the previous generation of Layer 1 optical cross-connect technology based on the SONET/SDH transport protocol, OTN switching is defined as the cross-connection of digital transport containers, known as Optical Transport Data Units (ODUks), that adhere to the G.709 Optical Transport Network standard multiplexing hierarchy. These transport containers vary in rate according to the client payloads that they carry, ranging from 100Gbit/s (ODU4) down to 1Gbit/s (ODU0).

OTN-switch-based networks differ from traditional muxponder-based infrastructure in that they provide:

1. Support for grooming of sub-wavelength client traffic, and
2. Physical separation of client and WDM uplink interfaces on a network element.

OTN switches typically support grooming and aggregation of traffic at the lowest available granularity specified by the G.709 standard, in this case 1Gbit/s, allowing intermediate nodes in a network to add or drop client traffic at granularities much less than the WDM line rate.

Additionally, this sub-wavelength grooming is realized using a centralized electrical fabric, whereby any client port on any client card can be routed to any of the available WDM line ports in the system. In contrast to a muxponder-based solution, this physical separation of client and WDM line ports decouples the cost of deploying new client ports from that of deploying more expensive WDM line ports. Clients are no longer tied to specific wavelengths, which is an important capability to support an evolving network.

From an operator perspective, OTN switching in 100G networks offers the following compelling advantages:

1. Improved CAPEX efficiency
2. Lower OPEX and faster service velocity
3. Improved network resiliency
4. ”Pay as you grow” service flexibility

3.1 Improved CAPEX Efficiency

Using the example discussed previously, deploying 30G worth of new services can be as simple as adding 30G worth of client ports to the node. These clients can be aggregated along with the existing client traffic onto the appropriate 100G WDM uplink card, maximizing the utilization of the existing fiber infrastructure. Client signals in a shared wavelength need not be destined for a common node given that they can be redirected to another node en-route. The resulting infrastructure and equipment CAPEX spend can be up to 35 percent less than muxponder-based network architectures. [v] Large-scale network modeling projects have concluded that the use of OTN switching at intermediate nodes in a network to groom and aggregate sub-wavelength traffic can increase wavelength utilization by up to 250 percent.[vi]

In Metro networks in particular, transport network investments are made on the basis of lowering cost-per-bit more than any other consideration. With the cost of 10G optics continuing to fall, an inefficiently utilized 100G wavelength does not provide lower cost-bit than 10G alternatives. By substantially increasing utilization of 100G wavelengths, OTN switching is the key to making 100G economical in Metro networks.

3.2 Lower OPEX & Faster Service Velocity

Operationally, OTN switching greatly simplifies the deployment and management of services. The ability to dynamically assign client services to any of the available 100G wavelengths means that adding, removing or re-provisioning client services no longer requires manual operator intervention at intermediate nodes; instead this is accomplished remotely via OTN switch configurations. Not only does this result in significant OPEX savings, it also enables network operators to rapidly monetize high-value, differentiated service offerings.

3.3 Improved Network Resiliency

Sub-wavelength grooming and the ability to dynamically assign clients to WDM uplink ports also allows network operators to employ highly-meshed network architectures that are capable of much more rapid fault restoration than pure ROADM-based networks. In OTN switched networks, monitoring and protection occurs at the service layer, rather than the optical layer, so time-consuming and complex optical path computation and tuning processes necessary for pure-optical protection schemes are not required during a fail-over event. The result is a 10x or greater improvement in fault restoration times, enabling robust carrier-grade networks with significant improvements in QoS.[vii]

3.4 "Pay as You Grow" Service Flexibility

Lastly, the ability to upgrade client ports separately from WDM line ports allows network operators to reduce the capital outlay associated with new client services as compared to fixed point-to-point WDM architectures.

A simple example is a scenario where a network operator has a new requirement to support a high-value service from an emerging data center customer requiring 16G Fiber Channel clients. In a muxponder-based network, if the existing muxponder equipment does not support 16G Fiber Channel clients, the network operator must not only procure new 100G equipment, incurring the cost of both client and 100G coherent optics, but they must also deploy a new 100G wavelength to support the service. Conversely, in an OTN-switch-based network, the network operator can "pay as you grow," only needing to deploy a new client card that supports 16G Fiber Channel clients.

The new client can also be added to any line interface that has sufficient capacity and can be easily moved to a different line interface at any time. This translates into maximum service flexibility for the network operator, a factor that is becoming especially important given the emergence of transport services geared towards data center and cloud markets where on-demand, dynamic bandwidth is critical.

Network operators worldwide are embracing these benefits, deploying OTN switching platforms throughout China, Europe and North America. In fact, a recent survey conducted by Infonetics revealed that 86 percent of network providers interviewed have deployed, or are planning to deploy, OTN switching in their Metro and Core transport networks.[viii]

4 Packet Optical Transport

Metro networks are challenged to support and manage both traditional transport services, as well as emerging packet services. In Metro networks, Ethernet services are the largest and fastest growing service type, be it for 3G/4G mobile backhaul, enterprise access, or data center interconnectivity. To gain efficiencies in both equipment and operational costs, Metro network operators have been seeking a platform that can both merge Carrier Ethernet and existing transport services onto a common equipment infrastructure, and provide packet grooming and aggregation functionality to maximize the utilization of the underlying optical fiber network. Furthermore, as GMPLS, ASON, and now SDN enable true multi-layer convergence in the network control plane, the notion of a converged multi-layer platform is ever more practical. The result is a new class of metro network equipment called the Packet Optical Transport Platform.

Functionally, a Packet Optical Transport Platform (also known as P-OTP or P-OTN) provides both Ethernet/MPLS and OTN switching on a single platform, reducing equipment cost and enabling packet grooming and aggregation operations to be performed before services are mapped into the transport infrastructure, increasing optical infrastructure efficiency. P-OTPs differ from OTN switched platforms in that they:

1. Add comprehensive packet services support, including centralized Ethernet/MPLS switching and aggregation, QoS, Carrier Ethernet OAM, and synchronization functionality; and
2. Support simultaneous centralized switching of native packet and OTN traffic efficiently onto any of the available 100G WDM line ports.
3. Integrating this functionality with existing transport features drives three new architectural requirements for P-OTPs:
4. WDM uplink cards must be able to *process* and *groom* both OTN *and* Packet traffic simultaneously from the central fabric. These are commonly referred to as Hybrid Uplink Line Cards. Recent advances by Microsemi in OTN processing silicon have enabled this.

Figure 3: P-OTP/P-OTN Hybrid Packet and OTN Uplink Card

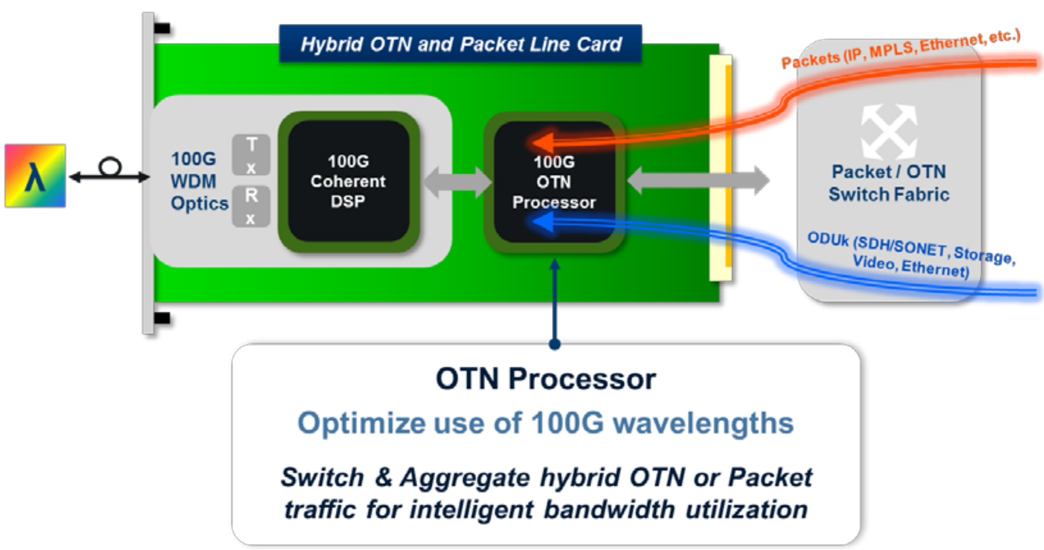

5. Line cards must support mapping of individual packet flows into flexible OTN transport containers, defined by ODUflex, enabling operators to most efficiently transport and switch sub-rate packet services throughout their network.
6. To facilitate the dynamic bandwidth nature of emerging on-demand services, such as data-center-driven packet traffic, uplink and client line cards must be capable of re-sizing containers hitlessly using ITU’s G.hao/G.7044 protocol.

Enabling a transport platform to aggregate, groom and process both packet and OTN traffic flows across all cards in a chassis is a powerful tool for making maximum use of costly 100G bandwidth. At the same time, packing Layer 2 switching, Layer 1 switching, and DWDM transport functions into a single transport platform clearly reduces overall equipment spend.

5 DIGI-120G: OTN Processing Innovation to Enable the 100G Era

Microsemi's DIGI-120G OTN processor is critical to enabling flexible, resilient and efficient 100G OTN transport networks, because it delivers the functionality necessary for OTN switches and P-OTP/P-OTN platforms. As the industry's first fully-channelized 120G OTN processor, DIGI-120G enables high-capacity fabric-connected line-cards for OTN switches. Additionally, as the industry's only OTN processor to support simultaneous *processing* and *grooming* of both packet and OTN client traffic, DIGI-120G is the only solution that allows OEMs to build the Hybrid WDM Line Cards that are an important building block for Packet Optical Transport platforms.

Figure 4 - Microsemi's DIGI-120G: Enabling Hybrid P-OTP/P-OTN Packet and OTN Line Cards

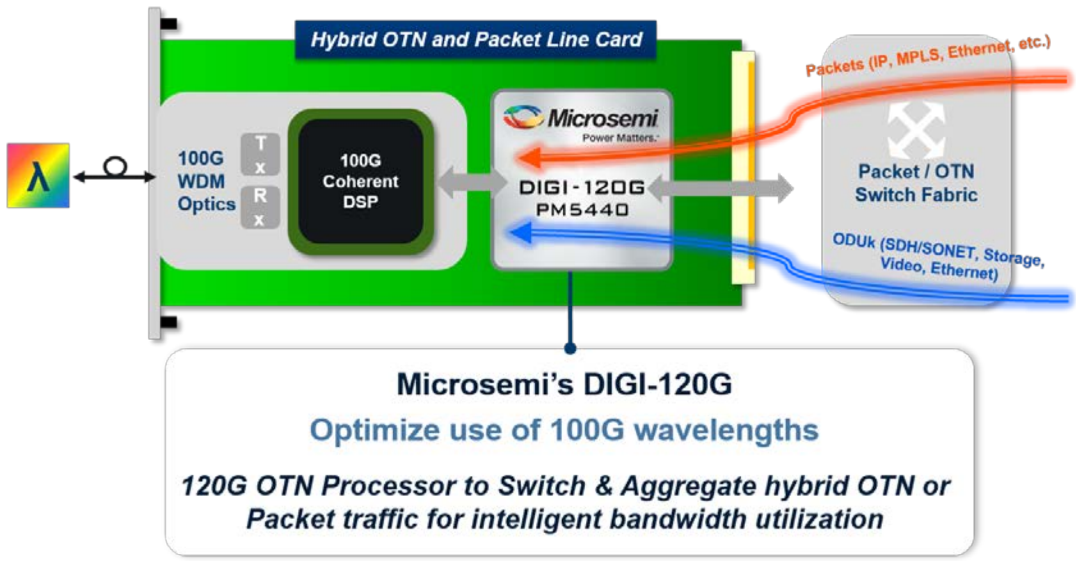

Furthermore, DIGI-120G is the industry's first OTN processor to enable dynamic and hitless packet traffic scaling, so Packet Optical Transport platforms can be optimized to handle transport of emerging high-value, packet-centric services.

With DIGI-120G, OEMs can develop a complete portfolio of equipment solutions for P-OTPs/P-OTN with a single hardware and software investment – from dense 10G/40G/100G Multi-Service Client Cards to 10G/40G/100G Hybrid WDM Line Cards.

6 Conclusion

Feedback from Network Operators worldwide is unanimous: 100G optical transport is a necessity to address the growing bandwidth problem in both Metro and Backbone networks driven by the 'Big Data' phenomenon. Network architectures employing OTN switching enable the most efficient use of underlying fiber infrastructure, and more flexible and resilient 100G networks, while at the same time offering significant reductions in CAPEX and OPEX. Emerging P-OTP/P-OTN optical transport equipment, marrying both OTN and packet switching functionality, are positioned to deliver even greater network and service efficiencies, while at the same time allowing network operators to embrace the dynamic packet-based services that characterize the fast-growing cloud and data center markets. Innovation in OTN switching silicon, led by Microsemi's DIGI-120G, is the key technology innovation that brings it all together, enabling Hybrid WDM line cards that deliver the efficiency and future-proof features to make mass deployment of 100G in optical networks a profitable reality.

---

[i] "Cisco Visual Networking Index: Forecast and Methodology, 2012–2017.", Cisco Systems Inc., May 2013

[ii] "Bell Labs Metro Network Traffic Growth: An Architecture Impact Study.", Alcatel-Lucent, 2013

[iii] "1G/10G/40G/100G Networking Ports (2014 Edition).", Infonetics, April 2014

[iv] Ibid.

[v] M. Bertolini et al., "Benefits of OTN Switching Introduction in 100Gb/s Optical Transport Networks.", OFC/NFOEC Technical Digest, paper NM2F.2, 2012

[vi] A. Deore et al., "Total Cost of Ownership of WDM and Switching Architectures for Next-Generation 100Gb/s Networks.", IEEE Communications Magazine, November 2012

[vii] "CD ROADM - Fact or Fiction", Fujitsu Network Communications, 2013

[viii] "OTN, MPLS, and Control Plane Strategies: Global Service Provider Survey," Infonetics Research, May 2013