Alcatel·Lucent

STRATEGIC WHITE PAPER

Entering the Terabit Era

Implementing the intelligent Optical Transport Network to face the exaflood phenomenon and support sustainable IP traffic growth

Explosive bandwidth demand and rising consumer expectations for rapid access to services constitute the exaflood phenomenon — a stressor on network infrastructures and on service providers' budgets that demands huge capital investments and ultimately results in higher operating costs. Today's service providers are seeking a way of reducing the cost of transporting high volumes of traffic while extracting value from innovative services. The next-generation Optical Transport Network (OTN) infrastructure supports emerging ultra-high rate services. A converged framework that fuses photonic and electronic networking, OTN integrates Dense Wavelength Division Multiplexing (DWDM), Reconfigurable Optical Add-Drop Multiplexer (ROADM) networking, and Optical Transport Hierarchy (OTH) sub-wavelength circuit networking. OTN provides a consistent, scalable, reliable, dynamically configurable optical layer that is cost-effective and responsive to ever-growing, unpredictable IP traffic demands. This paper explores the foundations of OTN, introduces the advantages of a converged OTN/WDM platform and describes its benefits as a key building block of the Alcatel-Lucent High Leverage Network™ Converged Backbone Transformation (CBT) Solution, enabling scalable and efficient IP bandwidth management at the lowest cost per transported bit.

Table of contents

1. The exaflood phenomenon

Initially, a 50 percent annual increase in Internet traffic might seem to be a boon to service providers, especially when combined with dramatic growth in multimedia applications and business and consumer IP services. The continued explosion in bandwidth demand, accompanied by consumer expectations of faster access to services, is referred to as the exaflood phenomenon. This situation puts tremendous stress on the existing network infrastructure, forcing service providers to invest heavily in expansion. However, more infrastructure invariably means higher operating costs, creating an apparently irreversible upward spiral.

The overall capacity of transport networks is being increased, with line rates moving from 10 Gb/s to 40 Gb/s and 100 Gb/s.

In contrast, the revenues associated with new IP services are not increasing at the same rate as the boom in traffic, decreasing service providers' ratios of revenue per transported bit. Service providers are faced with a momentous challenge as they seek new ways to support the cost-power-capacity crunch by lowering the cost per transported bit.

2. Converged backbone transformation: seamless IP-optical integration

Service providers must meet the challenge of reducing the cost of transporting high volumes of traffic while simultaneously extracting value from innovative services. A new architectural approach is needed — one that optimally combines optical transport and routing technologies to achieve scalable and cost-effective IP transport for any new service demand. This is the essence of the Alcatel-Lucent High Leverage Network Converged Backbone Transformation (CBT) Solution, in which the next-generation Optical Transport Network (OTN) infrastructure plays a key role (see Figure 1).

Designed to flexibly support high rate services from 1 Gb/s to 100 Gb/s and beyond, OTN and IP-over-OTN solutions provide a dynamically configurable optical layer that is responsive to IP networking demands. These solutions also enable multilayer optimization with superior service resiliency. With an OTN architecture in place, service providers can scale their IP backbone infrastructures and respond to huge traffic demands without incurring the extra costs associated with expensive routing capacity upgrades.

Alcatel-Lucent infuses OTN with additional intelligence for a solution that enables efficient bandwidth management at the lowest cost per bit.

Figure 1. Alcatel-Lucent CBT Solution

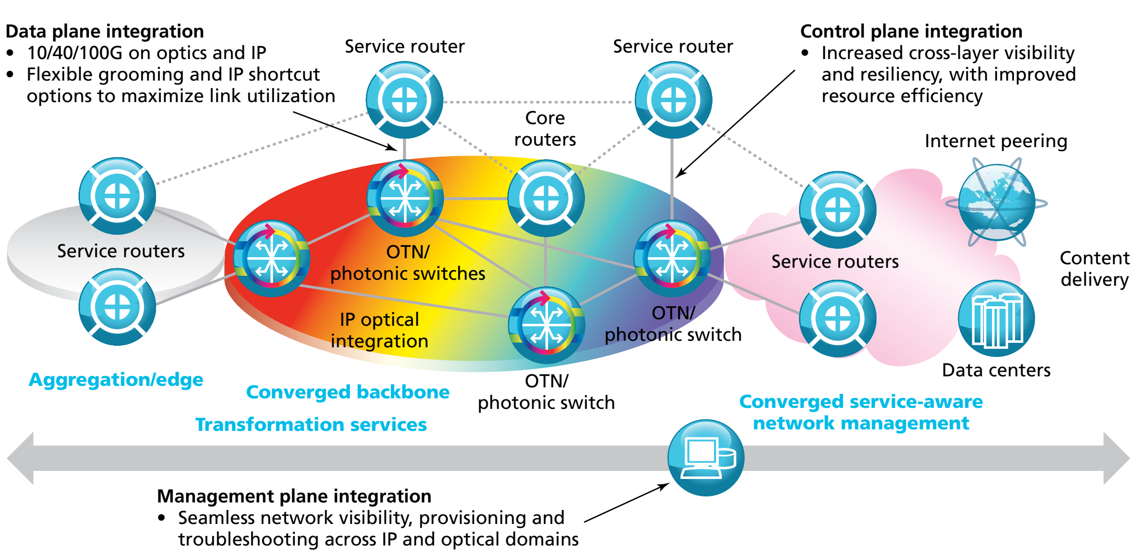

Integral to the Alcatel-Lucent CBT Solution is the Alcatel-Lucent 1830 Photonic Service Switch (PSS), featuring a new class of optical-core switching platform with terabit capacity and OTN support in addition to best-in-class Wavelength Division Multiplexing (WDM) functionalities.

New challenges require new approaches, and Alcatel-Lucent is committed to fully exploring innovative core-network transformation solutions that will help dramatically improve efficiency by transporting at the lowest cost per bit.

3. Rethinking network architecture

Today's backbone transport networks are typically based on a WDM layer that provides raw point-to-point capacity with 2.5 Gb/s, 10 Gb/s or 40 Gb/s line rates over long distances. Optical cross-connects in meshed topologies provide reliable SDH/SONET networking for managing bandwidth at the sub-wavelength level. These cross-connects create an intelligent, flexible transport network layer with highly efficient bandwidth management capabilities. Traffic flows between different network locations can be managed optimally and simply, with reliability and high quality, enabling efficient levels of connectivity in the IP backbone infrastructure.

A key challenge for these networks, as previously indicated, is scalability. As Internet traffic increases with the growth of high-bandwidth applications, transport line rates increase from 10 Gb/s to 40 Gb/s and 100 Gb/s. While these rates clearly exceed the scalability of SDH/SONET as a sub-wavelength transport layer, the strain on the IP layer also grows: in the absence of a sub-wavelength transport layer, nearly all transit traffic crosses routers that directly interconnect to the WDM layer for raw wavelength capacity.

A new way of thinking about network architecture is needed. The most efficient solution is one that offers an optimal combination of optical transport and routing technologies, capable of transporting IP traffic at the most economical layer. The solution must be scalable and capable of leveraging sub-wavelength and wavelength layers within the OTN in line with the service mix and traffic distribution.

3.1 Drivers of a next-generation backbone network

Service providers' costs are growing proportionally with traffic but are misaligned with revenue growth and profitability. To compete, they must:

- Solve the cost-capacity challenge
- Maximize bandwidth monetization
- Ensure high availability and quality assurance
- Simplify operations and reduce total cost of ownership (TCO)

3.1.1 Solving the cost-capacity challenge by efficiently scaling the IP core

Service providers can support traffic growth while controlling costs using a highly efficient backbone network architecture, optimally combining transport and routing capabilities to manage traffic explosion at the most economical layer.

The architecture can leverage an OTN infrastructure that provides the most efficient sub-wavelength bandwidth management capabilities for traffic up to 40 Gb/s and 100 Gb/s, scales to multi-terabit capacity. The OTN infrastructure can optimize the transport of router traffic using multiple forwarding options (at the wavelength and sub-wavelength levels) according to service mix and traffic destination. As shown in Figure 2, traffic can be offloaded from core routers onto the optical layer, therefore expanding overall network efficiency and life cycle, while traffic flows can be individually managed for quality assurance and optimization of overall network utilization, performance and cost.

Figure 2. Optimizing IP transport with the Alcatel-Lucent CBT Solution

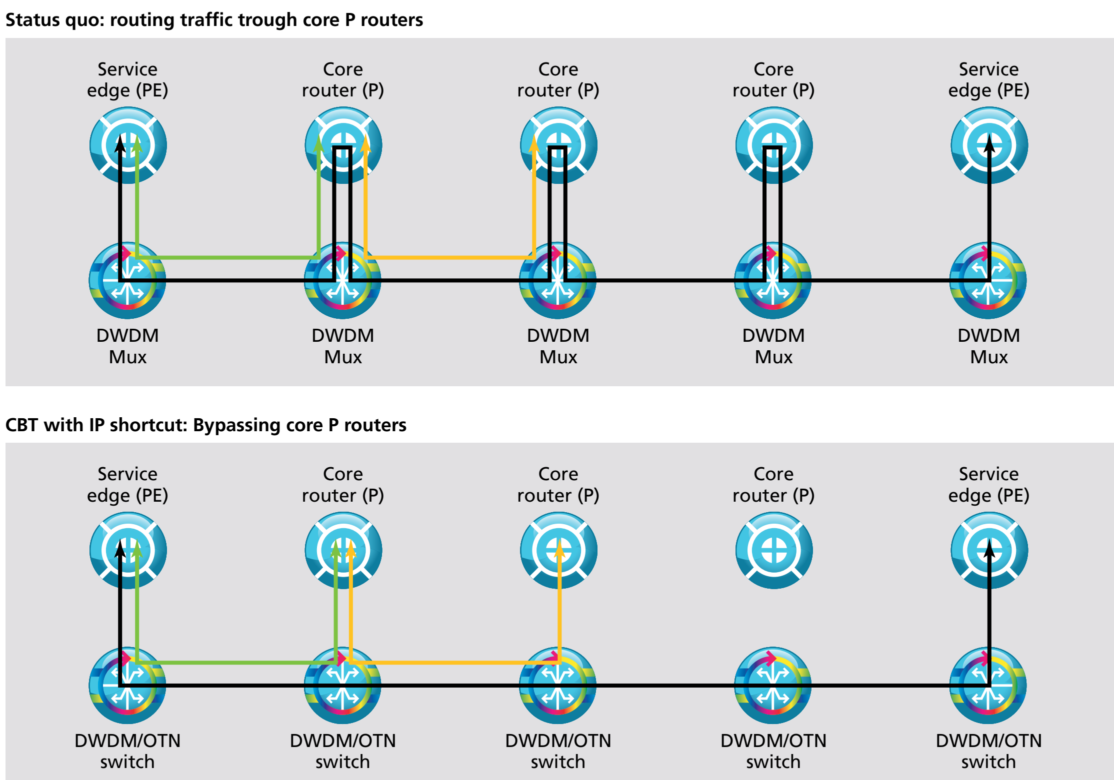

P = Provider, PE = Provider edge

This approach offers maximum granularity by enabling multiple traffic grooming options for the lowest cost per transported bit, eliminating unnecessary transit traffic in the IP layer and optimizing port efficiency. OTN integrates sub-wavelength and wavelength layers, enabling a highly scalable backbone infrastructure. With OTN, efficient IP traffic grooming options can be optimally and economically performed in the transport layer, avoiding overloading of the IP routing platforms.

3.1.2 Maximizing bandwidth monetization

Revenues per transported bit can be improved in an OTN with advanced automation and intelligence. Highly sophisticated restoration algorithms that make use of a standard Generalized Multi-Protocol Label Switching (GMPLS) control plane can minimize the number of wavelength and sub-wavelength bandwidth resources in the network that must be allocated for service resilience. This maximizes the reuse of available transport resources and frees network capacity for new, paying traffic.

3.1.3 Ensuring high availability and quality assurance

In an OTN, advanced control plane-based network protection and restoration capabilities enable Service Level Agreement (SLA) assurance and high availability. Consolidated transport-grade operations, administration and maintenance (OA&M) monitoring and fault management tools extend to new, higher bit rates. End-to-end SLA assurance and fault sectionalization capabilities are enhanced for multicarrier/multidomain networking scenarios.

3.1.4 Simplifying operations and reducing TCO

The OTN provides a range of operational benefits to service providers:

- Leverages intelligence and automation to optimize network operations across the optical and IP layers and accelerate time to service
- Reduces overhead costs by managing traffic efficiently and with maximum scalability
- Relies on consolidated transport network operational models to reuse existing operational skills
- Minimizes the number of transport entities necessary for each transported service
- Eases the pressures and costs associated with maintaining an always-on network by minimizing the need for on-site interventions during failures
- Utilizes the network optimally, ensuring maximum power efficiency by provisioning services automatically across the most economical resources in the network

4. Building OTN on a proven foundation

Operators have become highly dependent on — and attached to — the strengths of the previous transport network solutions deployed almost universally over the past two decades. Customers need proven transport functionalities such as deterministic performance, comprehensive OA&M tools, efficient grooming of multiple traffic flows, traffic monitoring, quality measurement and fault localization, and resilience against failures. These are all key elements of a reliable, easy-to-operate transport network infrastructure.

Standardized by the ITU-T, OTN capitalizes on past transport experience and extends capabilities to higher bit rates and multicarrier domains. OTN integrates these capabilities with WDM networking functions in a single consistent framework to provide a high-capacity foundation for next-generation OTNs. In essence, OTN builds on the graceful mix of photonic and electronic switching capabilities, which play complementary roles for maximum network efficiency:

- Photonic switching domain – Optical signals transiting a network node are switched at the wavelength level. This is best suited for cases in which the granularity of the transported service is close to the wavelength capacity. Photonic switching is used primarily to provision and restore wavelength services.
- Electronic switching domain – Each optical signal is terminated, and the entire signal — or its service-specific traffic contributions — can be individually switched. Electronic switching is used primarily to provision and restore sub-wavelength services that consume less than a wavelength of bandwidth.

OTN networking offers full management capabilities for the wavelength layer by introducing the Optical Channel (OCh). OTN integrates a new, expanded hierarchy for sub-wavelength bandwidth management: the Optical Transport Hierarchy (OTH), which offers a powerful networking tool capable of managing 1 Gb/s, 2.5 Gb/s, 10 Gb/s, 40 Gb/s and 100 Gb/s bandwidth pipes, including increments of multiple 1 Gb/s for maximum flexibility.

By combining photonic- and optoelectronic-based bandwidth management features and associated operational tools, OTN enables transparent transport for virtually any traffic type and rate across a scalable, reliable optical networking infrastructure, ensuring deterministic quality assurance.

With worldwide traffic rates set to jump each year, the ability to simply and efficiently manage ever-increasing traffic volumes is of huge benefit to operators. In this respect, OTN meets data-driven bandwidth needs and supports the emergence of new broadband services. Its multiservice-capable core infrastructure is based on lessons learned from SDH/SONET, incorporating additional high-bandwidth optical technology to meet the challenges of the telecommunications networking environment evolving toward full IP. OTN provides the gigabit-level bandwidth granularity required to scale and manage multi-terabit networks and offers service providers the tools for simplifying operations and improving network efficiency as traffic increases. OTN:

- Maximizes nodal switching capacity to the multi-terabit level, the gating factor for reconfigurable network capacity
- Minimizes the need for large numbers of fine-granularity pipes that complicate network planning, administration, survivability, management systems and control protocols
- Enables end-to-end quality assurance of client services while decoupling transport granularity from DWDM line system capacity
- Enhances SLA verification capabilities in support of multicarrier, multiservice environments

At the heart of OTN, OTH is a bandwidth management hierarchy in the electronic domain for very high rate signals up to 100 Gb/s. OTH leverages the familiar Time Division Multiplexing (TDM) technique, allocating any type of traffic flow into fixed-size bit frames (time slots) that are then switched by bufferless network nodes and transmitted over the digital signal. OTH provides unique, well-proven capabilities for transport networking, such as cost-effectiveness and minimized complexity, offering the best trade-offs among networking flexibility, hardware design simplicity and network efficiency. These OTH characteristics offer key advantages:

- *Transparency and multiservice* – Regardless of the nature of the transported service, in OTH any type of client traffic, including IP packets, is encapsulated in frames that are switched in bufferless network nodes and transported through logical pipes, at a constant bit rate and guaranteed quality with deterministic performance. This technique avoids the sophisticated traffic processing that is typical of packet technologies, such as classification, metering, queuing, Quality of Service (QoS) and congestion avoidance. Particularly advantageous in regional or backbone networks, where the traffic of multiple services from multiple access and metro locations is already well consolidated, this approach is a simple, efficient method for transparent, reliable transport over long distances.
- *Security* – In the electronic domain, OTH ensures that the networking technology used by the transport layer is fully independent from that used by the transported client service layer. This feature improves security against denial of service (DoS) attacks and guarantees full client traffic segregation, paving the way for an efficient shared network infrastructure among multiple service providers.

4.1 OTN hierarchy

The OTN architecture encompasses three hierarchical transport layers, as shown in Figure 3:

- *Optical Transport Section (OTS)* – Optical regeneration section layer devoted to the management of line optical amplifiers and related links. The OTS represents a managed multi-wavelength signal over a single optical span (for example, between line amplifiers).
- *Optical Multiplex Section (OMS)* – Optical multiplex section layer devoted to the management and multiplexing of wavelengths, and therefore to the management of multiplexers/demultiplexers. The OMS represents a managed multi-wavelength signal over multiple optical spans (for example, between DWDM equipment).
- *OCh* – Optical path layer devoted to the end-to-end management of wavelengths within the OTN. The OCh represents a single optical channel over multiple optical spans with flexible connectivity.

Figure 3. OTN hierarchy: integrated bandwidth management of photonic and electronic transport domains

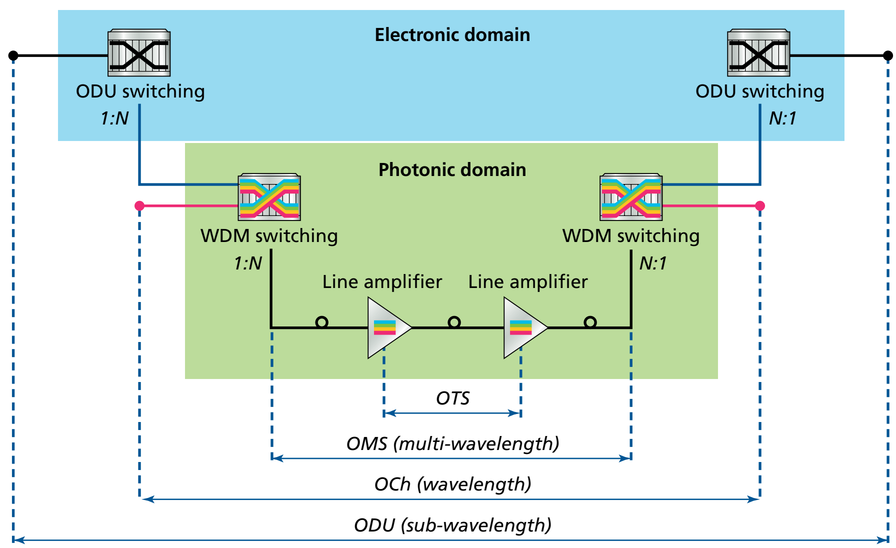

The OCh signal consists of an Optical Payload Unit (OPU), Optical Channel Data Unit (ODU) and Optical Channel Transport Unit (OTU). The OPU provides the functionality for mapping client signals — such as STM-64/OC-192, 40 Gigabit Ethernet (GE), 100GE and Fibre Channel (FC) — into the ODU. The ODU is a network-wide managed transport entity that functions as the primary bearer for client traffic and can transparently transport a wide range of client signals. In-band ODU operational tools support managed transport services in multi-operator optical networks in a client-independent way that is essential to the operation of these networks. Such tools support monitoring functions that enable SLA assurance for end customers, service providers and network operators. The tools provide for multiple levels of nested and overlapping connection monitoring as well as reliable performance using sophisticated fault localization capabilities.

Each ODU can carry a single client signal (Lower Order client container: LO ODU) with rates at 1.25 Gb/s, 2.5 Gb/s, 10 Gb/s, 40 Gb/s and 100 Gb/s or it can carry multiple (multiplexed) client signals (Higher Order server container: HO ODU) with rates at 2.5 Gb/s, 10 Gb/s, 40 Gb/s and 100 Gb/s.

In addition, a flexible-sized ODU container (ODUflex) has been standardized as capable of supporting the transport of variable-rate packet streams — for example, virtual local area networks (VLANs) within a 100GE port — to provide maximized bandwidth utilization and scalable transport of client services at the lowest cost per bit.

Robust protection schemes with deterministic performance are available in the three hierarchical transport layers to ensure both network and client survivability under 50 ms, along with GMPLS control plane-based restoration that supports flexible options for network resilience against multiple failures.

5. Efficiently scaling the IP backbone

Today, most backbone traffic travels through multiple core routers to and from its destination. With rising bandwidth demands, pressure on core routers increases in kind: the routers' throughput and port count increase in proportion to the overall transmitted traffic. Some operators are adding new routers and ports; others are upgrading their core routers to multichassis configurations and are adding corresponding WDM wavelength transport in the optical layer just fast enough to manage traffic growth. All this comes at a significant cost, power and carbon footprint — clearly not an economically efficient model when addressing the exaflood phenomenon.

5.1 Opportunities for greater efficiency

Analysis of traffic patterns across typical IP networks suggests that an overwhelming percentage of traffic passes through a relatively small number of locations for Internet peering, data centers and content distribution. Such transit traffic can represent as much as 75 percent of the overall traffic. This trend is likely to accelerate as video sharing, social networking and data center applications continue to evolve, placing additional stress on backbone networks. In addition, these real-time video and multimedia applications demand ever more stringent latency and delay capabilities from the network. Meeting such requirements is becoming a critical factor in determining the success of operators engaged in delivering high volumes of these services.

While certainly a challenge for service providers, these issues present a significant opportunity for optimizing the core network architecture. Offloading traffic from the IP core and handling it at the optical transport layer drastically reduces costs by avoiding router traffic congestion and capacity over-dimensioning. Economic benefits are further enhanced with intelligent control plane integration, yielding increased resiliency.

Capable of supporting terabit networking, OTN serves as the convergence layer for transporting a wide range of services whose bit rates do not allow efficient usage of the entire bandwidth associated with a single wavelength (lambda). Effective transport of such line rates involves supporting sub-lambda multiplexing, leading to resource efficiencies and cost savings. Moreover, by leveraging bufferless capabilities, an OTN-based backbone infrastructure ensures the low-cost transport of large portions of real-time traffic with deterministic quality and performance.

5.1.1 Integrated ODUk client/line switching

Next-generation photonic networks with Reconfigurable Add-Drop Multiplexers (ROADMs) allow for OPEX savings, automation and resilience at the wavelength level. Such networks are complemented with an electronic ODU layer that adds service-oriented sub-wavelength flexibility. Port and subport grooming introduces a switching and grooming layer below the 100G wavelength and complements the ROADM for all services that run below bit rates of 100 Gb/s.

DWDM transponders usually support one type of client interface to fill one wavelength. Transport of a service mix of GE, STM-16/OC-48, 10GE and 40GE cannot be achieved with a single transponder, therefore resulting in inefficient bandwidth utilization. In contrast, ODUk client/line switching supports full flexibility in associating client ports with DWDM lambdas while handling each service with its own quality assurance and even protecting clients, ports and line-side services (see Figure 4).

Figure 4. DWDM OCh switching with or without embedded ODUk switching

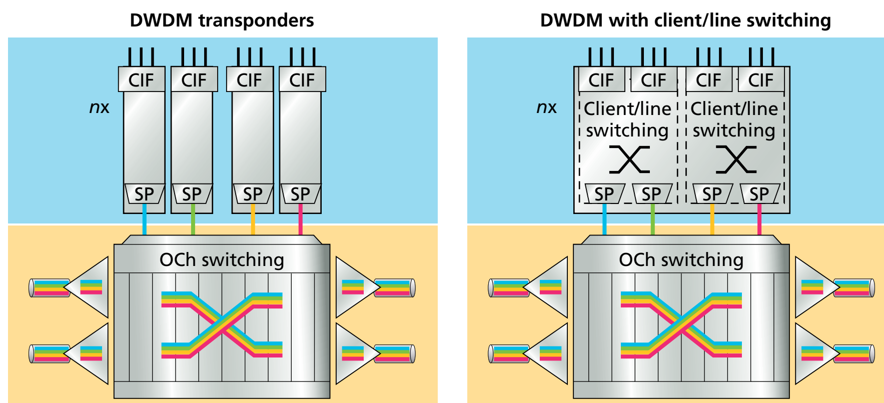

Dedicated muxponders, or switchponders, support the required client/line switching functionality. They offer multiplexing of lower-rate black-and-white client signals into a higher-rate carrier within the system. The multiplexed signal connects to the central fabric, with the following benefits:

- Flexible client-to-line assignment through an independent fabric
- Sub-lambda grooming between line ports
- Ultra-fast electrical protection and restoration of client signals
- Restoration with a per-service SLA
- Architecture with no single point of failure

The ODUk layer allows for service differentiation of transparent packet and TDM streams. Business-customer packet and TDM services can be delivered with the same network but with physical traffic segregation separate from the Layer 3 infrastructure. Moreover, as client service protection is not supported on DWDM transponders, ODUk switches support Sub-Network Connection/Network Connection (SNC/NC) protection schemes that offer interoperability with SDH/SONET networks.

5.2 Beyond IPoWDM: efficient IP traffic grooming options

With OTN, the core routers can be efficiently offloaded using a variety of IP traffic grooming options: lambda-, port-, and subport-level. This capability minimizes the cost of transport relative to the service provider's traffic mix and approach to network integration because multiple grooming options beyond the lambda level deliver the greatest efficiency and operational flexibility. This is central to carrying out core network transformation and realizing the vision of intelligent OTN. Offloading must be done with full assurance of service quality, the introduction of operational simplicity, and full control plane integration. With all these criteria met, providers can transition efficiently and cost-effectively to a next-generation OTN.

Typically, a wavelength in the optical layer at 10 Gb/s or 40 Gb/s — or even 100 Gb/s in the most recent network deployments — provides far greater capacity than the overall amount of traffic bandwidth actually required between client peer nodes in the IP layer. As a result, many network connections do not require a full lambda. In addition, depending on the service mix, client interfaces likely carry various traffic flows, each representing a logical channel between two different peer nodes and requiring individual traffic forwarding to different destinations in the optical layer.

An intelligent OTN infrastructure enables multiple bandwidth management options with different granularities for maximizing the efficiency of transport networks and lowering the total cost of reliable transport according to the service mix and traffic distribution within the network. IP traffic from router ports and subports, such as VLANs within the same port, can be mapped to the most optimal transport granularity — a wavelength (OCh), fixed-rate virtual container (ODU) or variable-rate virtual container (ODUflex) — and then individually forwarded and managed across the optical layer with the highest reliability and quality assurance.

These highly granular flow-management options allow service providers to focus their routing resources on "high-touch" services that require more sophisticated treatment while moving transit traffic efficiently and seamlessly across lower-cost optical infrastructures (see Figure 5).

Figure 5. Efficient IP traffic grooming with OTN

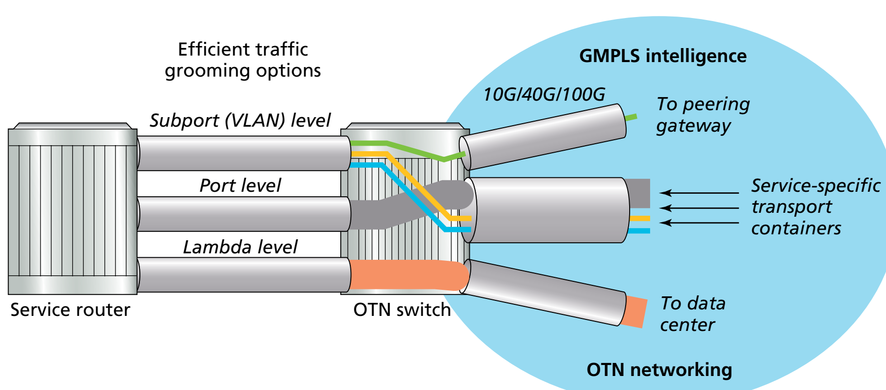

5.2.1 Subport-level grooming

Many incumbent providers and large operators offer both leased lines and a significant number of private Layer 2/Layer 3 data services, generating significant volumes of any-to-any traffic. When these combinations make up a large portion of the overall traffic matrix, the finer granularity offered by subport-level grooming becomes essential to maximizing network resources and maintaining the lowest cost per transported bit.

Subport-level grooming is a highly granular, flexible means of addressing growth, allowing maximum grooming flexibility by enabling VLANs or pseudowires within a router port to be logically mapped to a transport container such as a fixed-rate ODU or variable-rate ODUflex. For example, a 3 Gb/s VLAN flow as a portion of a 10GE port is mapped to an ODUflex, which is in turn groomed with other ODUs into a 40 Gb/s OCh wavelength. Different VLANs from the same router port are mapped to different virtual containers and are forwarded across the optical layer to their specific destination at the highest quality and lowest cost per bit.

Subport-level grooming enables finer granularity by supporting channelized or virtual interfaces such as VLANs, ensuring that IP router ports need not consume the full capacity of an optical transport port or wavelength. Without such granularity, providers would be forced to overbuild their IP and transport infrastructures — not an ideal financial scenario. Most important, subport-level grooming suits large service providers' existing organizational structures and operational models, optimizing their transport infrastructure without requiring alterations to network operations.

5.2.2 ODUflex: superior sub-level grooming

ODUflex is a new technology that enables traffic grooming between optical transport equipment and routers in a manner that efficiently addresses incremental bandwidth growth, in steps as granular as 1 Gb/s. Service providers no longer have to allocate a full ODU container to each connection: they can increase capacity in increments. ODUflex facilitates virtual channeling for 10 Gb/s, 40 Gb/s and 100 Gb/s Ethernet interfaces.

5.2.3 Port-level grooming

In addition to business services, most large operators offer residential broadband and Internet services that inherently add hub-and-spoke type traffic to any-to-any connectivity within the network. In these cases, a less granular grooming flexibility is required to better utilize resources, achieve the lowest cost per transported bit, and avoid overbuilding in both the IP and optical infrastructure domains.

With the port-level grooming option, a full port is logically mapped to a transport virtual container — typically a fixed-rate ODU pipe of the proper size. For example, a full 10GE port is entirely mapped to an ODU2, which is in turn groomed with other ODUs into a 40 Gb/s or 100 Gb/s OCh wavelength. Different ports from the same service router or a different router can be mapped to different virtual containers, groomed over a single wavelength, and forwarded across the optical layer to their specific destination at the highest quality and lowest cost per bit.

5.2.4 Lambda-level grooming

Lambda-level grooming, also known as IP over WDM (IPoWDM), offers the lowest level of grooming flexibility. For networks with very high volumes of hub-and-spoke traffic, it can be the optimal architecture because each physical port is equivalent to a single wavelength. In the lambda-level grooming model, a router port is mapped into a whole wavelength and the transport domain is purely photonic. The OCh wavelength signal is optically switched across the OTN optical layer using photonic switching capabilities and is delivered to its destination. With such a model, transponders can be hosted in the optical equipment or directly in the routing equipment. However, the latter option may prove impractical for existing organizational structures and operational models of large service providers, and for guaranteeing high optical performance, especially in long-haul applications.

5.3 Optimal filling of network resources

As traffic increases, service providers face the challenge of deciding how and when to migrate to higher bit rates in both IP and optical layers: from the router's 10GE interface to an intermediate 40GE interface and finally to the 100GE interface and from a 10 Gb/s link to a 40 Gb/s link or directly to a 100 Gb/s wavelength. Such approaches basically translate into the following options:

a. Using n x 10GE link aggregation
b. Migrating to 40GE as an intermediate step
c. Migrating to 100GE
d. Optimized migration using router port virtualization

Link aggregation combines multiple parallel links into one large logical link although its use typically decreases actual throughput efficiency to about 60 percent. In addition, in a 100 Gb/s lambda-based DWDM network, when the first 10GE port is connected to a DWDM transponder, the entire 100 Gb/s wavelength is switched through the network although it is filled to only 10 percent of its capacity.

A ROADM-based network switches and restores the entire 100 Gb/s wavelength and provides no access to more efficient grooming and dropping of traffic within this wavelength. 40GE interfaces are relatively less efficient when compared to available 100GE options but several operators see value in deploying 40 Gb/s wavelengths as a DWDM line rate for 4 x 10GE transport.

5.3.1 Virtualizing router ports using OTN grooming

Router port virtualization offers an efficient capacity handover to the OTN. Multiple 10GE ports can be combined into one 100GE interface, and the traffic can be logically separated into separate virtual streams. For example, OTN subport grooming can combine three 100GE ports, each with a filling rate of 30 percent, into a single 100 Gb/s wavelength for optimized network efficiency.

Along with the support of IP offloading options, OTN subport grooming offers efficient usage of low-filled router ports with port virtualization while both port and sub-port grooming provide optimal wavelength filling in a DWDM network for maximized efficiency and bandwidth monetization. OTN grooming helps solving the cost-capacity crunch by efficiently scaling IP metro and core networks.

5.4 Service differentiation

Core networks do not exclusively transport best-effort IP traffic. There is a need to transport traffic from multiple Internet service providers (ISPs) and business customers, transparent wavelengths for “carrier’s carrier” applications, and legacy TDM traffic. This traffic mix requires a grooming layer separate from packet grooming.

All these different services can be groomed into 100 Gb/s lambdas to maximize bandwidth monetization and can be protected and restored per service with high availability and quality assurance based on their SLAs.

6. Taking quality assurance and operational simplicity to the next level

An essential requirement for a transport infrastructure is its ability to supervise the quality of the transported services throughout the entire network. This is key to ensuring that SLAs are met and to the prompt enforcement of automated recovery actions in case of quality degradation.

Although existing transport technologies do support such features, new needs are emerging and posing additional challenges as service providers’ optical networks expand in both size and capacity. A service provider (or “lead” operator) often needs to interconnect to another network operator to deliver end-to-end connectivity to its end customers. In such a multicarrier networking environment, new capabilities are necessary. For example, lead service providers must be able to protect and supervise their customer services end-to-end — including the portion of network capacity that belongs to other operators.

OTN expands the capabilities of transport networking to address those potential limitations imposed by existing transport technologies that add to operational complexity and increase network costs.

6.1 Enabling transparent capacity services for the carrier’s carrier

A service provider often needs a capacity service from another carrier (the carrier’s carrier) to offer end-to-end connectivity to its customers. In some cases, the capacity service can be a GE connection, but if service provider runs SDH/SONET, the capacity service of the carrier’s carrier must be an SDH/SONET connection.

As shown in Figure 6, if Provider B — the carrier’s carrier — uses SDH/SONET as the transport technology, it cannot transparently carry Provider A’s SDH/SONET payload and overhead because these would terminate in Provider B’s network upon multiplexing or cross-connection. As a consequence, Provider A’s network would be unable to support end-to-end resilience — for example, with the use of SDH/SONET ring protection within its network. In such a case, Provider B could possibly deploy a passive all-optical solution. However, this would lack the supervision capabilities necessary to ensure quality.

With OTN, the transparent transport of SDH/SONET capacity services is no longer an issue. As Provider B deploys an OTN infrastructure, Provider A’s entire SDH/SONET frame is mapped into an OCh container, which is transparently carried through Provider B’s network and provides the necessary networking, supervision and resilience capabilities for a capacity service.

This approach allows a carrier to transparently carry the transport infrastructure of other service providers, enabling the delivery of transparent capacity services with simplified end-to-end operations.

Figure 6. Transparent transport of other carriers' infrastructures with OTN

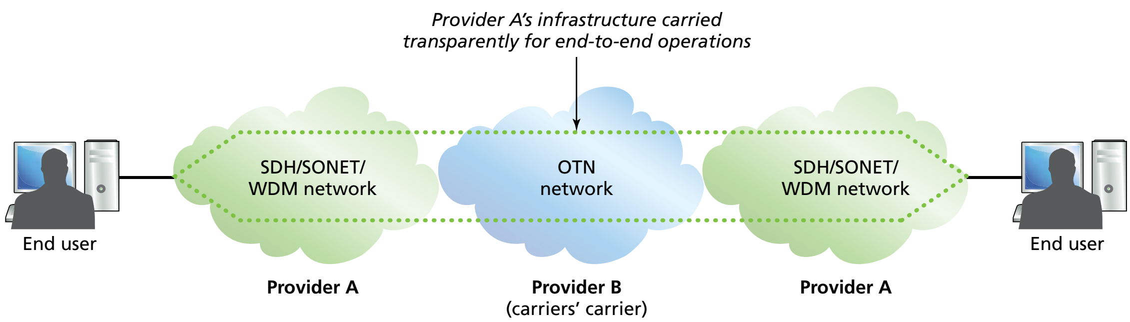

6.2 Enhancing quality control in a multicarrier environment

Supporting automatic SLA verification and fault localization in multicarrier scenarios is complicated. For example, existing SDH/SONET automatic supervision mechanisms such as tandem connection monitoring (TCM) allow one network operator within a group to monitor the end-to-end connection, or they allow multiple network operators to separately monitor their own networks. Because it is not possible to concurrently automate both types of monitoring, SLA verification often requires tight intercarrier cooperation and offline processes that add significantly to operational inefficiencies and costs.

For example, if the decision is made to provide end-to-end monitoring, manual processes are required for fault localization. If the decision is made to provide per-operator monitoring, the lead operator cannot determine overall signal quality. The lead operator would therefore need to rely on customer complaints or work to ensure the tight coupling of management systems across carrier boundaries. This lack of direct end-to-end monitoring for service assurance has long been a barrier to further lowering operational costs.

By offering a full set of quality supervision, connection monitoring and SLA verification capabilities for the easy and definitive support of multicarrier scenarios, OTN helps carriers augment SLA assurance levels for their end customers while dramatically lowering OPEX. As shown in Figure 7, ODU containers provide nested and overlapping supervision capabilities for every stakeholder in the transport domain: client user, lead service provider and network operators.

For example, the customer can own the OCh endpoints and their monitoring capabilities while the lead service provider owns the OCh leased circuits for which the network operators provide the OCh connections. Two fixed levels of connection monitoring capabilities and six variable levels of nested and overlapping connection monitoring are defined for this purpose.

Figure 7. Nested OTN monitoring levels: enabling efficient SLA assurance in a multidomain environment

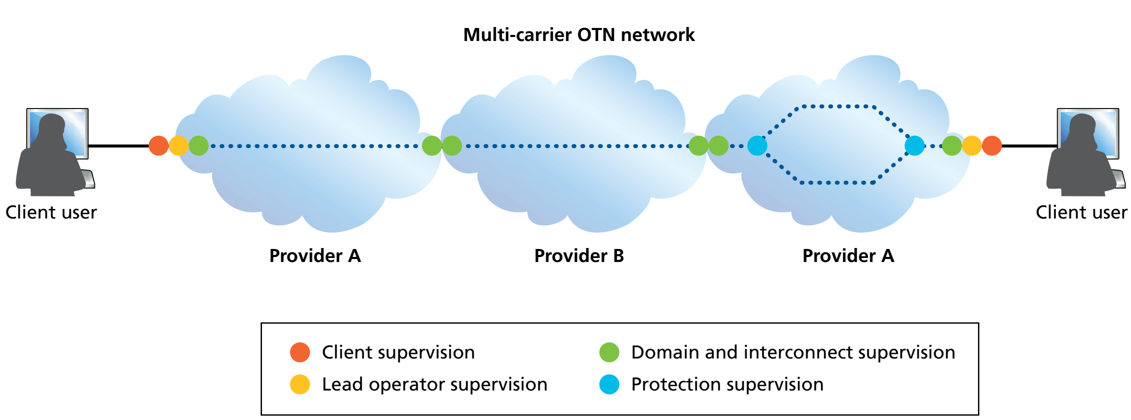

6.3 Decreasing network recovery time

Many of today's data-network reliability problems are solved in the transport layer as the first line of defense against severe network faults such as fiber cuts. Recovery from these optical layer outages must be rapid given the ever-growing bandwidth and number of users per fiber. OTN protection mechanisms provide fast transport layer recovery for deterministic performance. In general, providing transport layer recovery as close as possible to the physical media layer tends to be the most efficient solution — minimizing the need for spare capacity across all affected upper layers as well as reducing the number of transport entities involved.

6.4 Control plane intelligence for enhanced reliability and operations

Control plane intelligence is a key component of evolving OTNs in many critical aspects, enabling efficient bandwidth use and monetization. It also expands service quality assurance from access to core and provides graceful scaling, improved resiliency and simplified, automated operations. To reduce expenditures and respond to increasing bandwidth demands, OTN leverages the experience of the Automatic Switched Optical Network (ASON)/GMPLS control plane intelligence frequently deployed in today's optical networks.

GMPLS control plane-enabled OTN solutions simplify network operations by delegating several key operational processes to the control plane for automation, yielding a self-running network. Automated processes include the discovery of network topology, resources and services, as well as end-to-end optical connection routing for optimal resource utilization, flow-through service provisioning and mesh restoration.

By supporting GMPLS control plane intelligence, an OTN infrastructure enables the following benefits:

- Less need for manual and time-intensive operational procedures
- Higher network monetization with optimized resource utilization
- Better SLA assurance and network reliability provided by control plane-enabled restoration schemes, including protection and restoration combined (PRC)
- Higher interoperability across multivendor, multilayer and multicarrier networks
- Efficient tailoring of optical transport capacity according to IP layer demands
- Rapid turnup of revenue-generating services

GMPLS allows for dynamic bandwidth provisioning at the most efficient level and offers resiliency to multiple failures as well as intelligent restoration mechanisms. When a failure has been detected in a given network path, the protection schemes switch the traffic to the protected path. The GMPLS control plane then recalculates the new protected pipe to ensure that additional support will be available should a new failure arise. This process can be reiterated as many times as necessary according to the availability of network resources, guaranteeing an always-on transport infrastructure that is able to survive multiple failures.

GMPLS restoration enables SLA assurance with minimum resource allocation, freeing up bandwidth that can then be used for additional paying traffic. Another key advantage of GMPLS is operational simplicity. The network is self-healing, removing the need for manual interventions and allowing operators to relax their schedules for on-site interventions while still ensuring the highest service quality.

6.5 IP and optical cross-layer control plane

For maximum operational control of network resources, rapid service provisioning and increased resiliency, GMPLS user network interface (UNI) protocol enhancements enable a rich information exchange between the IP and optical-layer control planes. The standard UNI, with Alcatel-Lucent extensions, provides carrier-class resilience and facilitates cross-layer information sharing between the optical and the IP networks, including:

- Cross-layer control and automation
- Dynamic OTN bandwidth allocation on IP demand
- Fast provisioning
- Non-disruptive elastic bandwidth modification
- Cross-layer fault localization and reporting
- Coordinated multilayer restoration

The cross-layer control plane solution leverages photonic and electronic switching to optimize network costs. The solution reduces CAPEX by forwarding and protecting bits at the most economic layer, freeing expensive capacity. OPEX is also optimized by increasing service availability, providing synergies and consistent operations and services across layers, and guaranteeing the highest network power efficiency.

7. Entering the terabit era with the Alcatel-Lucent 1830 PSS

The Alcatel-Lucent 1830 PSS portfolio adds best-in-class multi-terabit electronic ODU switching functionality to its best-in-class DWDM networking capabilities — 100G coherent, Tunable ROADM [T-ROADM], long optical reach, photonic OA&M, design tools, and more — as proven by more than 70 customer deployments.

1830 PSS-1 1830 PSS-4 1830 PSS-16 1830 PSS-32 1830 PSS-36 1830 PSS-64

Highlights of the Alcatel-Lucent 1830 PSS portfolio include:

- Single product integrating WDM and OTN switching modules
- Scalable product size variants from access (1830 PSS-1) to core (1830 PSS-64)
- End-to-end networking in the WDM and OTN layers across product variants
  - Allows for sub-wavelength ODU switching and any-client/any-line assignment to maximize wavelength filling factors and add flexible bandwidth management
- New Alcatel-Lucent 1830 PSS-36 and 1830 PSS-64 shelves
  - Scalable multi-terabit OTN switching options at multi-terabit capacity
  - Support for any mix of client traffic
- T-ROADM configurations
- Next-generation 40G/100G coherent optics
- GMPLS control plane intelligence, ensuring the greatest possible network efficiency
  - Cross-layer automation
  - Highly resilient transport
  - Dynamic bandwidth provisioning
- Common network management
- Common cards across the product portfolio

7.1 Alcatel-Lucent 1830 PSS-64 and 1830 PSS-36

The new Alcatel-Lucent 1830 PSS-64 and 1830 PSS-36 form factors offer integrated terabit OTN and DWDM capabilities for the next-generation intelligent optical core. These advanced optical switches manage traffic at the most economical layer. Extensive automation and GMPLS control plane intelligence minimize the need to reserve network capacity for resiliency. Instead, providers can operationalize more of their available bandwidth, generating additional revenue and maximizing profits.

The Alcatel-Lucent 1830 PSS-64 and PSS-36 feature a unique set of capabilities:

- Multi-terabit switching architecture
  - State-of-the-art 2 Tb/s capacity
  - 120 Gb/s per slot backplane with built-in readiness to upscale to 8 Tb/s
- Universal switch unit that can support any traffic mix
  - 100 percent TDM to 100 percent packet
- Single compact, highest-density chassis
  - Lowest power per transported bit
  - Very low power consumption (less than 2 W per Gb/s)
- Support of OTN
  - Multiple transport networking options, including OTH, Carrier Ethernet, SDH and SONET
  - Most efficient bandwidth management of low-rate signals (GE, STM-1/OC-3)
  - Up to 40 Gb/s or 100 Gb/s
- Flexible IP traffic grooming options at the OTN layer, including port-level and subport-level grooming
- GMPLS control plane
  - Enables fast restoration and protections schemes at ODU-layer service granularity
  - Ensures superior SLAs

7.2 Economically viable 100G deployments

The Alcatel-Lucent 100G polarization division multiplexing - quadrature phase-shift keying (PDM-QPSK) coherent solution is commercially available with significant market traction since its launch in June 2010. Since then, Alcatel-Lucent has deployed its optical 100G solution in dozens of networks worldwide.

The Alcatel-Lucent 1830 PSS solution supports single-carrier 40G and 100G coherent detection. This 100G solution is unique in the market, providing:

- Ultra-fast electro-optical analog-to-digital converter/digital signal processing (ADC/DSP) at 25+ Gbaud real-time processing, developed by Alcatel-Lucent Bell Labs
- High transmission performance as well as high tolerance to physical impairments
- Full compatibility with existing 10G and 40G channels
- High level of integration in the complementary metal oxide semiconductor (CMOS), for lower cost/power consumption
- Compliance with the ITU 50 GHz wavelength grid, fully tunable across the C band
- Elimination of the need for external polarization mode dispersion/chromatic dispersion (PMD/CD) compensators

High-capacity 100G coherent modulation is a key component in building scalable converged OTN and photonic networking solutions. 100G applications are expanding rapidly and have proved to be economically viable compared to 10G and 40G deployments. Typical applications include:

- Plug-and-play 100G line cards in existing 10G networks when capacity is exhausted
- Lower cost per bit and dispersion compensation module (DCM)-free networks for greenfield deployments
- 100GE interconnection for IP-optical integration
- IP offloading, achieving the lowest cost per transported bit

8. Conclusion

OTN supports efficient IP traffic growth and maximizes the value of the network without compromising service levels. Able to support massive bandwidth growth while maintaining transport expenses at the lowest cost per bit, OTN allows service providers to fully monetize their traffic and take economic advantage of rising demands.

Drawing on our established WDM and cross-connect technology leadership, Alcatel-Lucent has expanded our leading WDM portfolio with a rich set of integrated terabit OTN switching capabilities. The Alcatel-Lucent 1830 PSS portfolio offers service providers the foundation for a next-generation transport infrastructure: one that is capable of cost-effectively scaling the IP backbone and entering the terabit era.

9. Acronyms

ADC	analog-to-digital converter
ASON	Automatic Switched Optical Network
CAPEX	capital expenditures
CBT	Alcatel-Lucent Converged Backbone Transformation Solution
CD	chromatic dispersion
CIF	client interface
CMOS	complementary metal oxide semiconductor
DCM	dispersion compensation module
DoS	denial of service
DSP	digital signal processing
DWDM	Dense WDM
FC	Fibre Channel
GE	Gigabit Ethernet
GMPLS	Generalized MPLS
HLN	High Leverage Network
HO	Higher Order
IP	Internet Protocol
IPoWDM	IP over WDM
ISP	Internet service provider
ITU-T	International Telecommunication Union - Telecommunication Standard
LO	Lower Order
MPLS	Multi-Protocol Label Switching
NC	Network Connection
OA&M	operations, administration and maintenance
OCh	Optical Channel
ODU	Optical Channel Data Unit
OMS	Optical Multiplex Section
OPEX	operating expenditures
OPU	Optical Payload Unit
OTH	Optical Transport Hierarchy
OTN	Optical Transport Network
OTS	Optical Transport Section
OTU	Optical Channel Transport Unit
P	provider
PDM-QPSK	polarization division multiplexing - quadrature phase-shift keying
PE	provider edge
PMD	polarization mode dispersion
PRC	protection and restoration combined
QoS	Quality of Service
ROADM	Reconfigurable Add-Drop Multiplexer
SDH	Synchronous Digital Hierarchy
SLA	Service Level Agreement
SNC	Sub-Network Connection
SONET	Synchronous Optical Network
SP	service provider
T-ROADM	Tunable ROADM
TCM	tandem connection monitoring
TCO	total cost of ownership
TDM	Time Division Multiplexing
UNI	user network interface
VLAN	virtual local area network
WDM	Wavelength Division Multiplexing