

Converged Transport Architecture: Improving Scale and Efficiency in Service Provider Backbone Networks



What You Will Learn

Massive growth in backbone network traffic and increasingly volatile traffic patterns are causing significant economic challenges for service providers.

Current architectural approaches rely on network layers that are built and operated independently. This model contributes to scaling, provisioning, and operational inefficiencies that will soon push service providers past the point where revenue can keep pace with needed investments.

Technology advances within a single network layer, although crucial, will not eliminate these inefficiencies. An innovative new approach is needed that softens and optimizes inter-layer boundaries to create a highly converged transport architecture.

This converged transport architecture must tightly integrate packet, optical time-division multiplexing (TDM), and dense wavelength-division multiplexing (DWDM) layers across hardware, control planes, management planes, and administrative domains to create a network that is much more efficient, less complex, and easier to scale and operate.

Overview

In backbone networks, Internet traffic is growing at a rate that challenges business models based on existing network architectures. Both a service provider's business model and network architecture must adapt to the increasing and unpredictable demand for bandwidth, which arises from the proliferation of devices, increase in subscribers, demand for video and multimedia, and new technologies such as cloud services. In addition, there is rapid growth in non-Internet services, such as mobile backhaul, wholesale and transit, wavelength switching, and private-line business services, which remain very popular with businesses. All these services share the same underlying optical transport infrastructure with the foundation for Internet and IP services. Thus they rely on the backbone network's available capacity to operate. To keep up with increasing demand, service providers have little choice but to continue adding backbone capacity to support this growth and to create a reserve that can absorb the spikes in traffic that will inevitably occur.

Adding backbone capacity is costly. Until recently, service providers chose this option, because they were confident about getting a return on their investments. By 2015, however, adding capacity will carry more risk than reward. A breaking point will be reached where the costs of expanding backbone networks will exceed the revenue potential present in the traffic carried.¹ There are many factors moving service providers toward this breaking point, some technical, some architectural, and some organizational. They all result from the manner in which backbone architectures have traditionally been built. Continued reliance on conventional linear scaling solutions, adding components to increase capacity, as a way to address nonlinear issues will eventually result in decreasing returns to scale. To pull back from this breaking point, a new method for growing backbone networks must be found. This method must address multiple dimensions at once and eliminate architectural boundaries, isolated administrative domains, and the inefficiencies affecting the packet, optical TDM, and DWDM layers.

Existing solutions focus on solving a single dimension of a multidimensional problem. Alternative architectures for the IP core, such as the lean core and hollow core, focus on reducing capital costs. Separately, new multilayer control-plane protocols focus on increasing efficiencies between the transport and IP layers. If implemented independently, neither an alternative core architecture nor a new protocol will yield justifiable returns on investment. Therefore, to thrive, service providers must find ways to increase revenue and reduce capital and operational costs from their backbone networks, while increasing bandwidth, flexibility, and efficiency.

Alternative Backbone Architectures

Alternative backbone architectures attempt to eliminate the need for a traditional full IP core architecture, with the goal of decreasing costs. New developments in platforms and protocols offer alternatives to full IP core architectures, but sacrifices are made in flexibility and efficiency in favor of cost cutting. The two predominant alternative architectures, hollow core and lean core offer initial savings, risking the long-term investment protection inherent in full IP core architectures.

Hollow core architectures attempt to eliminate the costs associated with core backbone routers by replacing them with a transport switching function (most often assumed to be an optical transport (OTN) switching layer) that offers a lower overall cost per bit for a given interface speed.

In the hollow core model, these switches create a dense mesh of circuits between each of the edge and peering nodes. As with the full IP core, the packet, optical TDM, and DWDM layers are managed separately, and there is little control-plane integration. In many cases even this level of integration will not be possible because the interface between the packet and TDM layers will be via Ethernet VLANs multiplexed onto circuits by specialized transponders that preclude G.709-based inter-layer communication. This lack of control-plane integration prevents topology information from being shared between the packet, optical TDM, and DWDM layers. Routers know only the routing topology, and the DWDM and TDM switches know only the optical topology.

Lean core architectures are an adaptation of full IP architectures in which backbone routers have reduced network processing unit (NPU) functionality or memory. With reduced NPU memory, the router can only learn internal routes, which forces operators to use Multiprotocol Label Switching (MPLS) to forward subscriber traffic instead of IP. Another option for cutting the cost of NPUs would be reducing their functionality to a minimum so that they are only capable of providing transport functionality and forwarding via MPLS.

¹“Service Provider Capex, Revenue, and Capex by Equipment Type,” Infonetics, 2012; “Core and Edge Routers 5-Year Forecast,” Dell’Oro Group, January 2012; and “Evolving Backbone Networks Using and MPLS SuperCore, Juniper, June 2011.

Although nominally less costly to deploy because of their less-sophisticated NPUs, the lean core architecture and its associated label switch routers (LSRs) introduce problems that may ultimately erode potential savings. All IP services will be shifted to the edge or provider edge routers outside the backbone, effectively turning the backbone into an inner-core. Implementing a lean core architecture can also be a complex and disruptive undertaking that requires operators to re-architect an existing network.


Finally, lean core architectures do not address the other factors affecting backbone network costs and inefficiencies. This model still manages packet, optical TDM, and DWDM layers separately, and control-plane integration remains limited as in full IP core architectures. As with conventional architectures, topology information is effectively isolated within the packet, optical TDM, and DWDM layers, limiting operational efficiency.

Multi-Layer Control Plane: Generalized MPLS

Although new architectural approaches such as hollow core and lean core attempt to address the initial cost factors for evolving a network, they ignore the challenges associated with isolated packet, optical TDM, and DWDM layer control planes. This isolation leads to inefficiencies and difficulties in provisioning, monitoring, troubleshooting, service velocity, and more, when laborious coordination is required between network layers and disparate administrative organizations. To address layer isolation, the industry is developing a consolidated, multilayer control plane that would work across layers, based on MPLS and known as Generalized MPLS (GMPLS).

GMPLS defines two models of operation, peering and overlay. In the peering model (Figure 1), operators create a single routing domain that encompasses both the IP and transport layers. Layer 3 and Layer 1 devices share topology information, providing Layer 3 routers with visibility into Layer 1 transport paths, wavelengths, loads, risk groups, and so on. This communication enables Layer 3 routers to calculate best paths, request circuits that meet their requirements, and move or recover and restore circuits based on path conditions.

Figure 1. Peering Model




However, the peering model creates three significant problems:

- It forms a single routing domain that does not respect the boundaries that exist, to varying degrees, between the packet-layer and DWDM-layer administrative organizations.
- The packet-layer and DWDM-layer devices must scale to handle much larger routing domains. This not only adds to the memory requirements of every platform but also increases path computation loads across the entire network.
- Software testing and upgrades must span the packet and DWDM layers, slowing certification cycles and complicating deployments.

In the overlay model (Figure 2), the packet-layer and DWDM-layer routing domains remain isolated, eliminating the problems with scale and administrative domain overlap associated with the peering model.

Figure 2. Overlay Model



Instead of acting as a peer, the packet layer behaves as a client of the transport layer, which behaves as a server. The layers do not share topological information, but the packet layer can request that circuits be set up or torn down with a User-Network Interface (UNI).

The difficulty with the overlay model is that the packet and transport layers share very little information. The UNI interface only allows circuits to be set up or torn down. This approach provides too little, rather than too much, information sharing, and although it permits small improvements in the automation of circuit setup, it cannot address significant problems associated with circuit routing.

As currently architected, GMPLS offers either too much or too little information sharing between layers. If service providers are willing to consolidate administrative domains and absorb the scaling costs associated with the peering model, then much can be gained in terms of circuit provisioning, monitoring, and troubleshooting. However, few service providers are currently willing to take these risks.

If the overlay model is used, only a small amount of provisioning automation is gained, and a significant amount of functionality must still be performed manually, which bypasses the opportunity to reduce operational complexity and costs.

Better Solutions Are Needed

Service providers need to change the way their backbone networks are architected, if they are to remain profitable while evolving to accommodate massive growth in IP/MPLS traffic, unpredictability arising from cloud services, and the migration of SONET and SDH services onto packet. However, the solutions described above, lean core, hollow core, and GMPLS, are too one-dimensional and address only parts of the complexity and efficiency problems found among today's backbone networks. Service providers need a new, more comprehensive approach.

A Converged Transport Architecture

To overcome their challenges, service providers need a way to address all the dimensions at once, including scale, flexibility, and cost, without adopting solutions that sacrifice one in favor of the other.


Scale must be added without giving up flexibility. Flexibility must be added without harming scale or expanding routing domains beyond organizational tolerances. Costs must be lowered, not simply shifted from core to edge or from capital to operational budgets.

To improve scale, increase flexibility, and lower costs concurrently, service providers need a new architecture. This architecture must deliver advances across multiple dimensions, including packet and DWDM layers, control and management planes, administrative systems, network management systems (NMSs) and operating support systems (OSSs). Then it must converge them into a whole solution, greater than the sum of its parts. This is a converged transport architecture.

To achieve this convergence, service providers must focus not only on advances in packet routers and DWDM reconfigurable optical add-drop multiplexers (ROADMs) but also on eliminating boundaries between the packet and DWDM layers. Eliminating these boundaries through control-plane and NMS/OSS innovations can dramatically improve the efficient sharing of information within the network and across organizational boundaries, resulting in a more dynamic, efficient backbone network.

This convergence will be evolutionary, building on what is already in place, rather than replacing it. Layers will be consolidated and functions added, in contrast to the approaches that cut functionality in favor of savings.

Figure 3. Converged Transport Architecture



This new converged transport architecture (Figure 3) is based on the tight integration of:

- Converged transport router platforms that deliver massive scale, DWDM integration, OTN switching, and advances in transport-class reliability, while still allowing flexibility for mixed core architectures.
- An agile DWDM layer that enables dynamic 100Gigabit and higher wavelengths while eliminating the impediments to flexibility that affect much of the existing ROADM installed base.
- nLight, a multi-layer packet and optical path optimization and signaling capability that improves on GMPLS with crucial topology sharing and signaling functionality.
- Cisco Prime™ network management facilitates top-to-bottom converged operations, administration, management, and provisioning (OAM&P), while allowing separate administrative domains.

These technologies can be applied in part or in full, incrementally or all at once, to create an architecture that will improve, rather than impede, a service provider's scale, efficiency, administration, and profitability.

Converged Transport Routers

A converged transport router (CTR) is the service-delivery heart of the converged transport architecture. It delivers functionality that spans Layer0 to Layer3 and enables the simultaneous delivery of dynamic IP services, label-switched path (LSP)-based packet transport, and OTN-based circuit switching.

The CTR is not a new routing platform perse. Rather, it is a new class of routers that delivers advances and innovations in several or all of the following areas:

- Massive system scale: CTRs must deliver terabits of bandwidth per rack, with nonblocking scale, to several tens of terabits per system initially. Then delivery must increase to hundreds of terabits per system, over its lifespan, to handle the 29 percent CAGR traffic growth predicted by the Cisco® Virtual Networking Index (VNI). CTR system configurations must be flexible to meet varied service provider deployment scenarios, including single-chassis, back-to-back, and multichassis configurations, with simplified, cost-optimized interchassis connections and form factors optimized for both point of presence (PoP) and transport-class deployments.

- Agnostic switching fabrics: An agnostic and non-blocking switch fabric must be equally suited to packet, optical TDM, and even future traffic types, tying CTR line cards and chassis together.
- Advanced NPUs: CTR NPUs must provide unprecedented forwarding performance and flexibility, while lowering the cost curve through massive application-specific integration circuit (ASIC) integration, which reduces the number of ASICs from 14 in second-generation backbone routers to as few as two. Using fewer ASICs per NPU improves power consumption and improves deployment flexibility.
- Universal ports: Universal ports must be matched to CTR NPUs to break the model normally associated with line cards. Interfaces will no longer be dedicated to a single modulation or bandwidth.
- Coherent DWDM transponders: Integrated IPoDWDM transponders are one of the most useful and compelling technologies for decreasing transport costs. Coherent optical technologies enable IPoDWDM transponders to achieve higher speeds (100Gbps, 400Gbps, and beyond) over longer distances. Coherent DWDM transponders eliminate the need for external transponders and associated short-reach interconnecting optics, and perhaps even more important, they provide the physical means to accomplish packet and DWDM layer control-plane integration.
- Transport hardening: CTRs must provide granular and flexible isolation for all routing, control, and management functions to minimize the effect of any failures. Isolation must also be provided for the expanding and extensible set of core services, such as proximity routing, large-scale network address translation (NAT), mobility, and video monitoring.

Operators must be able to perform hitless software upgrades and patches, and NPUs must be architected, from the start, to span software changes with zero packet loss and zero effect on adjacent network platforms.

CTRs must also support Layer 0 through Layer 3 service management, using single or segmented views to accommodate organizational and administrative needs. The ability to partition network management allows service providers to partition or combine OAM&P functions as needed to accommodate administrative and organizational constructs, so they can more easily migrate to a converged transport architecture.

Service providers demand one essential capability for transport-class equipment: a very long service life.

The following factors all contribute to this capability: massive scale, low-cost multichassis, flexible form factors, agnostic switching fabric, advanced NPUs, universal ports, optical shelves that decouple optics from packet processing, and virtualization that allows for hitless upgrades and service expansion.

CTR is not the only element required for a converged transport architecture. However, it is one of the primary and most important elements, because it provides the foundation for massive scale and outstanding flexibility, and it ties the packet and agile DWDM layers into a powerful, flexible, and cohesive whole.

Agile DWDM

An agile DWDM layer forms the foundation for converged transport architectures, by delivering many features and innovations that promote significant provisioning and recovery flexibility. Provisioning is handled entirely by software, reducing manual operator intervention, which is one of the major contributors to high operational expenses and long provisioning and recovery delays. Agile DWDM innovations include: tunable lasers, colorless interfaces, tunable coherent receivers, omnidirectional interfaces, flexible spectrum, contention less ROADM routing, and rapid circuit restoration.


Cisco nLight

With the introduction of colorless and omnidirectional ROADMs, the DWDM layer becomes much more dynamic and flexible, allowing any interface or wavelength to be rerouted or recolored, using software, to meet current service requirements.

Given the inherent problems with the GMPLS peering and overlay models, conventional GMPLS is unable to provide flexibility that the packet layer can use efficiently. In addition, backbone routers are unable to request specific circuit routing based on service requirements or to react quickly to changes in the underlying DWDM network.

Cisco nLight extends GMPLS to create a more capable UNI that helps the transport layer share pathing, wavelength, and circuit information, while allowing the backbone routers to define specific conditions for the creation and routing of circuits (Figure 4).²

Figure 4. Cisco nLight Architecture and Information Flow



In effect, the transport layer becomes a database for the IP layer to use. The information shared by the transport layer provides enough detail about the underlying circuits and paths that the IP layer can make informed decisions about how a circuit should be routed - which helps ensure that the circuits meet requirements, such as latency, optical cost, common or disjoint circuit routing, common or disjoint paths, the inclusion or exclusion of specific shared risk link groups (SRLGs), or even specific paths. Consequently, circuits can be provisioned to meet specific requirements in seconds, rather than the hours or days that might otherwise be needed, as administrators make requests, reconcile requirements with databases of DWDM layer capabilities, manually create circuits, and then provide notification when circuits are ready. As a result, restoration takes only seconds.

The key economic advantage of Cisco nLight is the help it provides in maintaining compliance with strict service level agreements (SLAs), while using fewer network and operational resources, including fewer links, transponders, and manual configuration efforts.

²Crawford, Dan Crawford, Dan, Walid Wakim, and Calrence Filsfils. "Cisco nLight™ Technology: A Multi-Layer Control Plane Architecture for IP and Optical Convergence.

To meet SLAs, important backbone links must be protected through mechanisms in either the DWDM layer or the IP layer. Conventional DWDM-layer mechanisms rely on one of the following approaches:


- They use two transponders to create two isolated paths.
- They create a second 1+1 backup path.
- They (more rarely) use optical splitters to split the signal from one transponder down two optical paths, so the receiver can choose the stronger signal.

In traditional IP layer protection, network planners must determine the worst-case traffic loads through an interface during a failure scenario. Then their network design helps ensure that such capacity is always available, should a failure occur. This generally results in a situation where only 50 percent of link capacity is used, with the remainder reserved for use during link, node, or interface failures.

Cisco nLight builds on traditional IP layer protections to provide a more cost-effective and efficient approach. It uses IP-layer protections, such as fast reroute (less than 50-millisecond recovery) and proactive protection (boosting recovery to less than 15-millisecond link protection), while the DWDM layer provides dynamic path restoration in seconds.

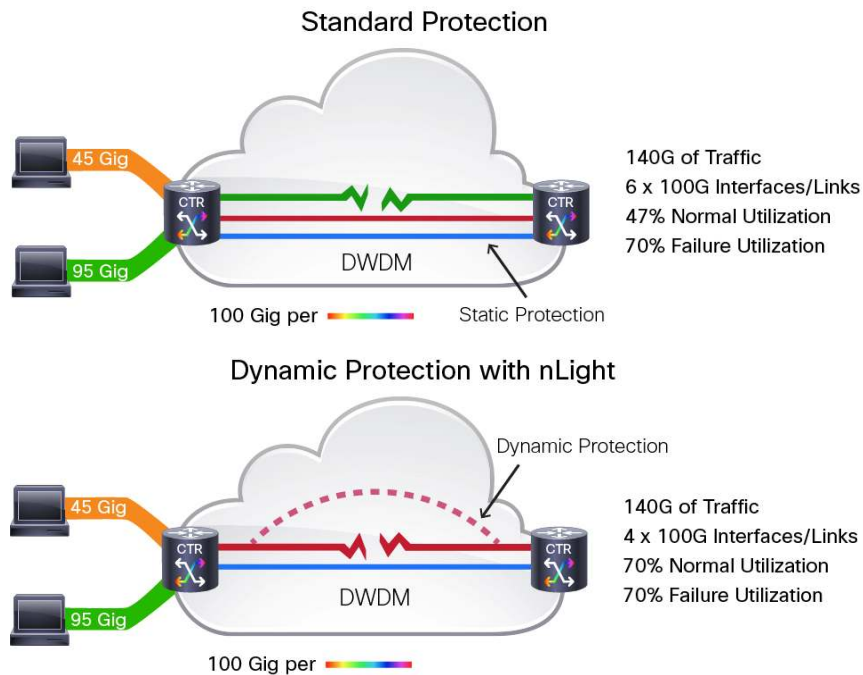
For example, in Figure 5, two links are configured between the CTRs, a primary link in red and a protection link in orange. When the path taken by the primary red wavelength is broken, IP protection switches the traffic to the orange protection wavelength in milliseconds. The DWDM layer then restores the primary wavelength (now purple) over another path and informs the CTRs, which then retune to purple before restoring traffic to the new primary link. The orange protection wavelength is rerouted as well to minimize (but not eliminate, in this case) shared paths. From an IP perspective, the topology does not change, and the inter-router links were protected without the use of redundant transponders or splitters.

Figure 5. Cisco nLight Dynamic Restoration



This combination of IP protection and DWDM restoration allows for a reduction in the number of interfaces and transponders required to protect a given traffic level. It also allows overall higher utilization over active links, as shown in Figure 6.

Figure 6. Cisco nLight Restoration Advantages



Even in this relatively simple scenario, the advantages are clear. The number of 100-Gigabit IP interfaces, 100-Gigabit DWDM transponders, and 100-Gigabit DWDM wavelengths drops by 33 percent, while link utilization increases from 47 percent to 70 percent. The capital savings are substantial, but service providers also realize operational savings due to the dynamic nature of the circuit restoration, which was accomplished rapidly without manual intervention.

Cisco Prime Portfolio

Converged transport architectures eliminate many of the barriers responsible for increasing costs and hampering flexibility in traditional backbone architectures. When packet and DWDM layers combine, control planes consolidate, and administrative domains blend, it is crucial that network management systems converge, as well, to provide a common platform for all OAM&P tasks.

Cisco Prime is designed to promote this convergence. Cisco Prime is a modular suite of applications that provide A-to-Z management for next-generation converged transport architectures. Its optimal network management approach helps consolidate the management of large and diverse networks through a centralized OSS, while accelerating time-to-value for new equipment and technology investments.

The integrated solution provided by Cisco Prime offers a foundation for continuously improving network operations, helping to lower CapEx and OpEx, deliver high-quality services to retain customers, and increase revenue streams. Through centralization and automation, Cisco Prime provides more consistent service delivery and management to simplify the design, provisioning, and management of carrier-class networks - resulting in the reduction of operating expenses and acceleration of time-to-revenue for new backbone-affected services and technologies.

Building Blocks for Converged Transport Architectures

The converged transport architecture is made possible by four key elements:

-
- A new class of backbone router, the converged transport router, which delivers unprecedented scale, reliability, DWDM integration, transport flexibility, and a very long lifespan.
 - An agile DWDM layer that allows the flexibility and adaptability normally associated only with the packet layer.
 - Extensions to GMPLS that tightly integrate and consolidate the packet and DWDM layers without introducing scalability limits or organizational concerns.
 - An NMS built for end-to-end and top-to-bottom management of converged transport architectures, offering full OAM&P capabilities in concert with existing management systems.

These elements, woven together, allow service providers to finally address all their scale, flexibility, and cost concerns at once, without sacrificing one for another.

Evolving Towards a Converged Transport Architecture

The advantages of converged transport architectures are clear, but perhaps the steps and paths needed to build one are less so. Few, if any, “greenfield” backbone networks are being built today, so starting from scratch is an unlikely option. Instead, service providers will need to evolve from their current layered backbone architectures toward converged transport architectures in steps.

These steps are not rigidly defined or sequential, and the path will conform to each service provider’s unique characteristics. But service providers can begin realizing benefits immediately by deploying crucial technologies, such as Cisco nLight, coherent IPoDWDM transponders, and packet and OTN integration where and when possible.

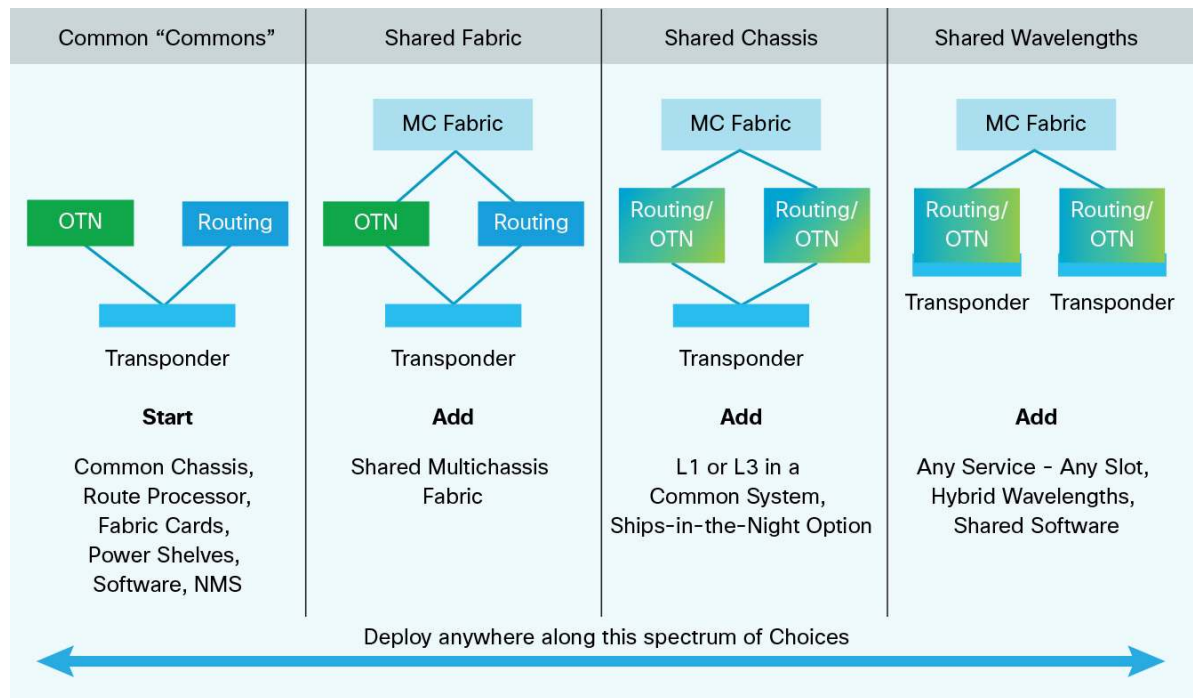
Cisco nLight is a critical element that enables consolidation of the packet and DWDM layers by enabling much-needed automation for provisioning and recovery. And, because Cisco nLight is platform-independent, it can be deployed on any routing platform based on Cisco IOS® XR which has IPoDWDM interfaces, including the installed base of Cisco Carrier Routing System (CRS) routers.

Coherent DWDM transponders are instrumental in erasing the boundaries between the packet and DWDM layers, and virtual transponder features allow management of these resources to be shared or segmented, as needed, to meet organizational requirements. However, these transponders are not limited to the CTR. Like Cisco nLight, they are also applicable to the Cisco CRS.

A significant benefit of implementing Cisco nLight and coherent DWDM is that service providers can begin addressing existing organizational issues, since the deployment and management of these resources can be, and ultimately must be, shared.

Another important advantage of the converged transport architecture is deployment flexibility, as shown by the example in Figure 7. For a mixed routing and OTN environment, several deployment options are possible, with varying levels of integration. These options include shared commons, shared fabrics, shared line cards, and even shared wave lengths, with each level of integration achievable according to service provider targets. Service providers can converge as little or as much as they wish for a given backbone requirement.

Figure 7. Converged Transport Router Deployment Models



Mixed vendor deployments will remain standard for the majority of service providers. Consequently, Cisco is working to standardize the GMPLS extensions that enable Cisco nLight functionality. Coherent DWDM transponders implement standards-based protocols, including ITU G.709 OTN framing and alarms, so their usefulness is assured even with non-Cisco DWDM implementations. (Until Cisco nLight is standardized, however, some automation and recovery capabilities will be unavailable.)

Because Cisco Prime can integrate into existing OSS and BSS systems, allowing end-to-end and top-to-bottom OAM&P, service providers can improve their operational efficiencies, even without deploying agile DWDM ROADMs or CTRs.

The CTR, combined with agile DWDM, Cisco nLight, and Cisco Prime, provides the most comprehensive solution and the greatest level of transport convergence through features, such as agnostic fabric and integrated OTN. However, it is possible to evolve toward a converged transport architecture over time by implementing elements of the architecture whenever and however possible.

Conclusion

To successfully address the challenges posed by dramatic growth and increasing unpredictability, service providers must evolve their backbone networks while simultaneously eliminating costly inefficiencies.

These inefficiencies result primarily from the boundaries between the packet, optical TDM, and DWDM layers. The boundaries slow service, impede flexibility, and complicate dimensioning, monitoring, and troubleshooting efforts.

Therefore, service providers must eliminate boundaries that affect hardware, control planes, management, and other dimensions of a network. They must also address other factors, such as NPU complexity and the lack of flexibility in the DWDM layer. The resulting converged transport architecture will allow service providers to reduce

the inefficiencies that are sapping profitability, and they can build a durable foundation for delivering modern services efficiently and profitably.

For More Information

Visit <http://www.cisco.com/go/core>



Americas Headquarters
Cisco Systems, Inc.
San Jose, CA

Asia Pacific Headquarters
Cisco Systems (USA) Pte. Ltd.
Singapore

Europe Headquarters
Cisco Systems International BV Amsterdam,
The Netherlands

Cisco has more than 200 offices worldwide. Addresses, phone numbers, and fax numbers are listed on the Cisco Website at www.cisco.com/go/offices.

Cisco and the Cisco logo are trademarks or registered trademarks of Cisco and/or its affiliates in the U.S. and other countries. To view a list of Cisco trademarks, go to this URL: www.cisco.com/go/trademarks. Third party trademarks mentioned are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any other company. (1110R)