

## Converge IP and DWDM Layers in the Core Network

The Cisco® IP over DWDM (IPoDWDM) solution for IP Next-Generation Networks (NGNs) enables the convergence of service providers' IP and DWDM core networks, increasing service flexibility and reliability while lowering operating expenses (OpEx) and capital expenditures (CapEx). The Cisco IPoDWDM solution efficiently manages traffic growth through the Cisco CRS-1 Carrier Routing System and the Cisco ONS 15454 Multiservice Transport Platform (MSTP).

### Introduction

To meet the demands for converged packet infrastructures, service providers are increasingly moving their networks toward an IP NGN-based architecture. This shift is driven by the need to reduce OpEx and CapEx while offering new revenue-generating services over a single, converged infrastructure. IP NGNs can enable convergence at the application, service, and network layers. This paper focuses on the network layer and specifically on the convergence of two important technologies within service providers' core networks: IP and DWDM.

### Core Network Infrastructure Evolution

One recent core network trend is the consolidation of multiple Layer 2/3 networks into a single IP/Multiprotocol Label Switching (IP/MPLS) infrastructure. In spite of this Layer 2/3 convergence, however, the underlying transport layer (Layer 1) of many service provider core networks has continued to use SONET/SDH, an infrastructure that was introduced in the early 1990s to support traditional time-division multiplexing (TDM)-based data and voice services. This SONET/SDH infrastructure was implemented to support three critical functions: grooming; protection and restoration; and thorough operational support (such as alarming and performance monitoring).

When introduced, SONET/SDH allowed the efficient multiplexing of lower-speed TDM circuits such as T1/E1 and T3/E3 to higher-speed OC-3 and OC-12 trunks for transport across service providers' core networks. Because of their synchronous nature, the trunk connections could be switched independently when needed without the need for complex demultiplexing as demanded by TDM-based data and voice services. The infrastructure also allowed the growth to faster aggregate links such as OC-48 and higher while also supporting the emerging IP networks of the day. This "grooming" function allowed the bandwidth efficiency and automation that service providers needed to support multiple switched TDM services in the core of their networks by using SONET/SDH add/drop multiplexers (ADMs) and cross-connect equipment. The SONET/SDH infrastructure, which typically uses a ring-based topology within the core, also allowed for protection and fast restoration (50 ms) during a failure on one part of the ring to maximize the availability of the overall network. With the advent of SONET/SDH standardization, a separate, standardized, message-based channel<sup>1</sup> was used for alarms, control, monitoring, and administration of the links from a centralized location.

---

<sup>1</sup> DCC – Data Communications Channel, a 192-kbps message-based channel for OAM&P between SONET/SDH equipment

In the latter part of the 1990s, DWDM emerged as a way to significantly increase the efficiency of the installed fiber plant by allowing transmission of multiple wavelengths over a single physical fiber. This function introduced another level of multiplexing and demultiplexing at the optical level to support greatly increased bandwidth at the core of the network, which followed the dramatic rise of IP-based networks fueled by the explosion of the Web. The SONET/SDH layer, which now handled increasing amounts of IP traffic, was mapped into wavelengths at the DWDM transport layer to be carried across the core long-haul<sup>2</sup> networks spanning regions and countries in many cases. This has remained largely the case in many service provider networks globally today.

### **The IP Explosion**

The volumes of IP traffic on these core networks have, however, continued to increase steadily to the point where the primary use of these core long-haul networks today is to carry massive amounts of transient IP traffic, significantly outpacing the traffic volumes of traditional voice and data services. Over the next 5 years alone, global monthly IP traffic is expected to rise to 26 *exabytes*,<sup>3</sup> accelerated by the application convergence of all video, voice, and data traffic to IP, resulting in a compound annual growth rate (CAGR) in excess of 56 percent globally. The convergence of traditional applications such as broadcast television, video on demand, and voice to new distribution models over IP as well as the explosion of new applications such as music and video podcasting and peer-to-peer (P2P) file sharing will only continue to fuel this tremendous growth of core IP traffic.

### **Core Network Infrastructure Challenges**

Despite the trend toward IP convergence, multiple equipment layers to support core long-haul networks continue to exist, creating OpEx and CapEx concerns for service providers as well as the challenges of profitability and return on investment. Furthermore, as customers demand increasingly stringent service-level agreements (SLAs), service providers must maintain higher levels of reliability while still having the flexibility or “speed to service” to accommodate change based on service demands or traffic growth characteristics within the network core. To meet these requirements, service providers must consolidate their core networks and move toward more efficient ways to handle the increased IP traffic loads – yet at the same time they are confronted with problems at multiple levels to achieve this objective.

### **Multiple Transport Layer Elements**

Some network inefficiencies result from the way core transport networks are built out today to support the IP layer over the SONET/SDH layer, supported by an underlying DWDM infrastructure. Consider the paths of two types of traffic entering and exiting a typical service provider point of presence (POP). The first scenario is IP traffic that needs a Layer 3 lookup at the POP and therefore is riding a wavelength that will terminate on a router. The second is called “pass-through” (or transient) traffic, which stays in the transport domain and bypasses the router to travel on to an adjacent POP in the service provider’s core network.

#### **Router-Terminated Traffic**

The IP traffic comes into the POP today typically through 10-Gbps SONET/SDH OC-192/STM-64 circuits, which are composed of colored wavelengths multiplexed through DWDM on to a physical fiber. This fiber is fed into a DWDM demultiplexer, which splits out the individual colored wavelengths. These individual wavelengths that are to be terminated on the router are then fed

---

<sup>2</sup> Long haul networks are typically characterized by reach of distances up to 620 miles (1000 km)

<sup>3</sup> 1 exabyte (EB) = 1 x 10<sup>18</sup> bytes. Source: Cisco Estimates, Ovum, Bernstein, and Public Company Data

into transponders, which convert them from optical (colored) to electrical and then to a standard short-reach wavelength (“grey light”). This optical-to-electrical-to-optical (OEO) conversion is used because historically short-reach optics have been used for connectivity inside the POP environment. The grey light is then typically fed into a short-reach interface on a SONET/SDH cross-connect,<sup>4</sup> which recovers the SONET/SDH clocking, performs any grooming necessary, checks for errors, and monitors for loss of signal (LOS) so that it can perform SONET/SDH-level restoration if needed. However, in most cases today, no grooming is actually needed because the full 10 Gbps is being connected to the router (rather than 2.5 Gbps or lower speed links in the past). Therefore, from a connectivity perspective, the cross-connect is serving essentially as a patch panel. The SONET/SDH cross-connect then feeds the 10 Gbps to the router, which performs performance monitoring at Layer 1 through Layer 3, monitors for LOS so it can perform MPLS Fast Reroute (FRR) restoration, and performs a Layer 3 and above lookup to route the packet to its destination. On the aggregation side the core router is typically aggregating multiple lower-speed links and grooming the IP traffic into well-used 10-Gbps links to present back into the core transport network.

#### Pass-Through Traffic

As traffic patterns in the core have become more distributed, the amount of traffic passing through a given POP purely at the transport layer (as opposed to terminating on a IP router) has tended to increase, and can sometimes be as high as 70 to 80 percent of the overall traffic that the POP handles. In this case the incoming DWDM link goes through a similar method of interconnections through the DWDM demultiplexer and transponders to the SONET/SDH cross-connect through short-reach optics. It checks for errors and monitors for LOS so that it can perform SONET/SDH restoration. Again the grooming function that would have occurred here previously is no longer required because typically full 10-Gbps links are being passed through the POP. Hence the cross-connect is again serving as a patch panel from a connectivity perspective. A similar process of interconnections occurs for outgoing traffic from the POP.

These OEO conversions and the associated electrical processing result in an additional cost in terms of space, because many racks of shelves may be required in a service provider POP, as well as additional power and cooling that is necessary because of the active electronics components that they contain. Furthermore, in this core network scenario the SONET/SDH functions are redundant because of the capabilities that have been integrated into the router.

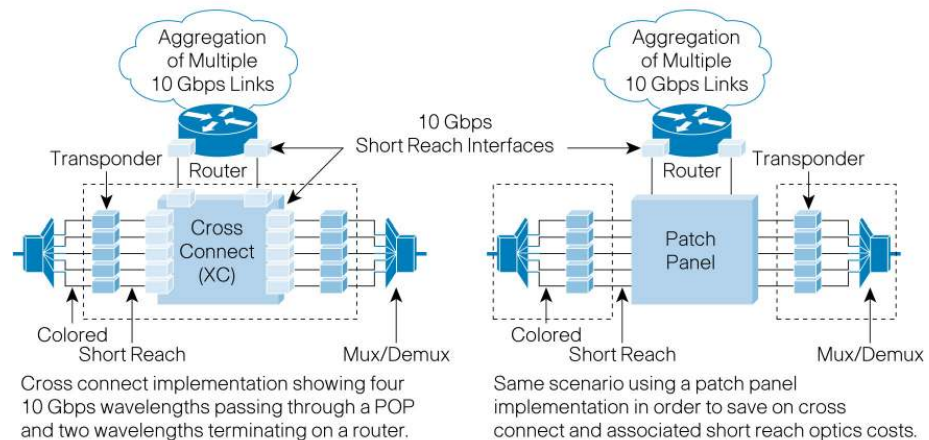
- **Grooming:** Because most traffic has moved to IP, the router now performs the grooming function by aggregating IP traffic and presenting it to the core transport layer within well-used 10-Gbps links.
- **Operational support:** The router and its associated interfaces can measure errors at Layers 1 through 3, collect performance statistics, generate appropriate alarms, etc.
- **Protection and restoration:** Using MPLS FRR, the router can provide 50-ms protection or better and do so much more efficiently than the traditional SONET/SDH protection schemes (such as BLSR<sup>5</sup>), which waste up to 50 percent of the bandwidth for protection purposes.

For these reasons, service providers have already started using manual patching in place of the cross-connect to save costs.

---

<sup>4</sup> The term cross-connect is used to refer to any such device that has an electrical backplane and performs OEO conversion, such as a broadband digital access and cross-connect system (DACs).

<sup>5</sup> BLSR: bidirectional line switched ring, a SONET transport network configuration in which network nodes are connected in a ring, and traffic can be re-routed in the other direction around the ring in the event of a cable cut or degradation of optical signal, thereby routing around the point of failure.

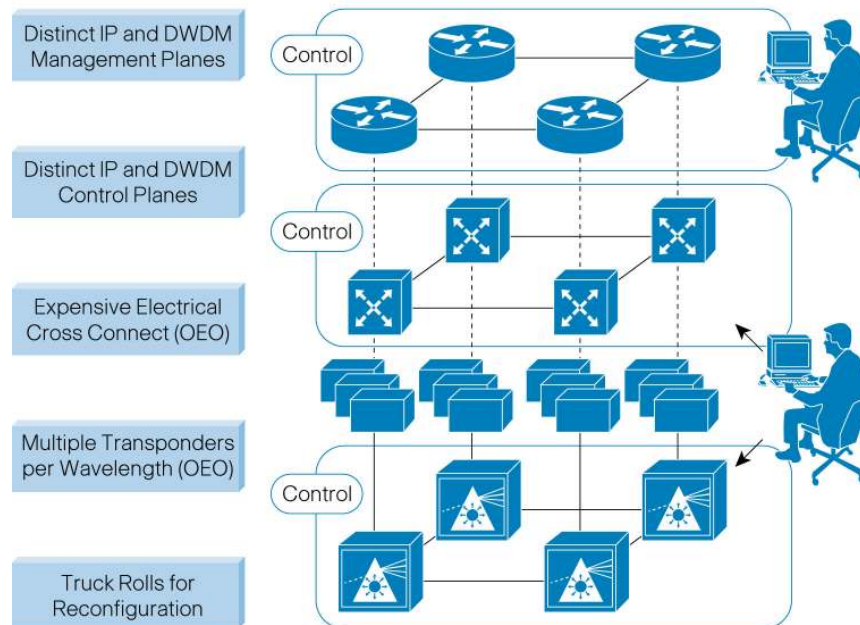
**Figure 1.** IP-over-DWDM Interconnect Model in Today's Network

### Manual Patching

Manual patching removes the cross-connects from the network but leaves the transponders at each hop along the path. Even for traffic that is transient or passing through a POP location, an OEO conversion must still be performed on every wavelength because of the limited ability of traditional DWDM systems to operate in ring and mesh configurations. This scenario still requires demultiplexing and multiplexing of the DWDM signal into individual wavelengths and converting it to a grey light through transponders. These grey signals must then be manually patched through a patch panel to their final destination within the POP before being multiplexed back into the DWDM system. Manual patching has become a common occurrence in today's DWDM networks because routers frequently present fully used 10-Gbps links that map directly into DWDM wavelengths without requiring grooming at the subwavelength level. However, where manual patching is used on intermediate POPs on transient connections, an operations person may need to manually patch the wavelength through to its destination at each of these intermediate sites when an add or change is required within the core network. This process is both cumbersome and error-prone, and it drastically slows down any network change, to the order of weeks for completion on large national and international networks.

### Continued Investment in Cross-Connects

In response to this added OpEx burden, cross-connect vendors have upgraded their systems to switch 10-Gbps signals, based on an extended multiplexing hierarchy as defined by the G.709 standard. Although this seems to solve the automated switching problem, additional short-reach interfaces as well as transponders (or alternatively, integrated DWDM interfaces on cross-connects) are required to provide the interconnection for pass-through traffic as more wavelengths are deployed to accommodate increasing traffic. Each wavelength still must go through a costly OEO conversion within the cross-connect device. Because these interfaces are expensive, service providers are not likely to predeploy many of them – and the result is necessary equipment upgrades and labor at multiple intermediate sites as traffic grows. Exacerbating this problem is that as the volume of IP traffic in the core increases from needing 10- to 40-Gbps links, these cross-connect platforms and their associated electrical switch fabrics will have to be upgraded further to accommodate this transition, a scenario that translates to additional CapEx and OpEx for the service provider. Even where the cross-connects have additional bandwidth capacity available to accommodate this growth, they typically maintain this additional idle bandwidth capacity until it is actually needed.

**Figure 2.** Today's Core IP and Transport Network Infrastructure

### Distinct Control Planes

As transport and IP networks have evolved – affected by differing requirements as well as evolving over differing periods of time – they have developed somewhat dissimilar and incompatible control mechanisms in order to direct the flow of traffic between their endpoints. Transport networks today primarily rely on fixed paths that are configured either physically through patch panels or manually through pre-provisioned circuits within cross-connects to connect one input port to an output port. In some cases, small portions of bandwidth are allocated for control-channel communication for pre-provisioned circuits to provide a semi-automated path setup between intermediate points. But the network topology is not updated in real time on the management systems, resulting in obsolete databases and potentially long lead times for implementation of a network change. Large networks can take weeks, because new wavelengths may need to be provisioned between distinct locations, and these networks often require operator-to-operator verbal communication. IP/MPLS networks have evolved with a much more automated model inherent with most packet-based network technology and hence these networks heavily rely on the exchange of control information through protocols such as Border Gateway Protocol (BGP) and Label Distribution Protocol (LDP) between routers to determine the optimum path between endpoints. The use of these control protocols in IP/MPLS networks provides the intelligence to self-optimize and self-route traffic based on varying conditions in dynamic ways that transport-level networks lack today. And given the dynamic nature of IP network traffic today, having these distinct and incompatible control mechanism between the transport and IP layers while each is still heavily reliant on the other creates a situation of increasing OpEx for the service provider as these networks continue to grow.

### Multiple Management Layers


Service providers also typically rely on separate and multiple network management platforms, protocols, and applications for operations, administration, maintenance, and provisioning (OAM&P) of the two networks. Transport networks typically use management protocols such as Transport Layer One (TL-1) and Common Object Request Broker Architecture (CORBA), which have strong ties to SONET/SDH standards, whereas IP networks rely on Simple Network Management

Protocol (SNMP) and the newer Extensible Markup Language (XML) mechanisms to manage devices. Few management applications provide the full suite of protocols and mechanisms to manage and operate both device types across the IP and transport networks. Furthermore, in most service provider networks today, organizational boundaries exist in that the operational and management organizations for the two networks are usually distinct and independent, and hence must rely on good person-to-person communication and coordination between them – across remote sites in some cases to ensure appropriate changes are made when needed across the two networks. This process thus can be error-prone and is usually very time-consuming.

### Cisco IP-over-DWDM Strategy for Core Convergence

The Cisco IP-over-DWDM strategy addresses all these challenges facing service providers today to reduce CapEx and OpEx, increase overall reliability, and improve the speed to service – while still allowing them to increase the overall scale and performance of the network to meet growing traffic demands of IP through this decade and beyond. This scenario is achieved by three critical points of integration between the IP and the DWDM network, as illustrated in Figure 3.

**Figure 3.** Three Points of Integration Between IP and DWDM Networks

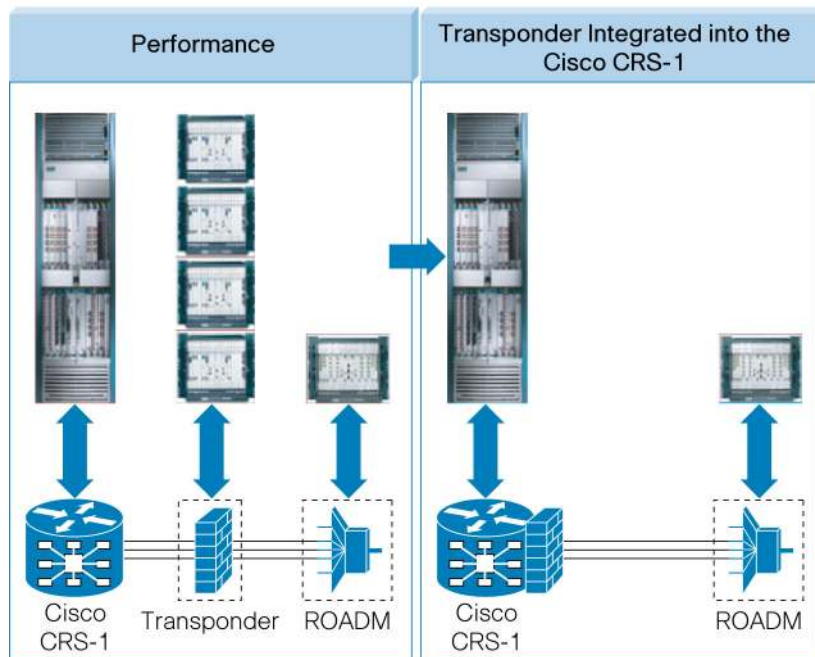


#### Element Integration

Element integration refers to the capability to take multiple, separate elements that operate in the network and collapse them into a single device without losing any of the desired functions for continued operation. Today, Cisco offers a physical line interface module (PLIM) for the Cisco CRS-1 Carrier Routing System with integrated transponder functionality at both 40-Gbps (OC-768c/STM-256c) and 10-Gbps (10 Gigabit Ethernet) speeds. The Cisco CRS-1 1-Port OC-768c/STM-256c Tunable WDMPOS Interface Module and the Cisco CRS-1 4-Port 10GE Tunable WDMPHY Interface Module enable the Cisco CRS-1 to interconnect directly with existing DWDM systems, eliminating the need to invest in external transponder shelves in typical service provider POP configurations. The result is significant savings in both CapEx and OpEx for the service provider by eliminating 50 percent of the short-reach optics required for interconnections to the DWDM equipment. Resiliency is hence inherently enhanced by this reduction in active devices within the network. Furthermore, the Cisco CRS-1 now has visibility into the DWDM layer and can proactively monitor the optical paths end to end for transmission-level performance. This means that even if wavelength paths were to gradually degrade because of environmental, component, or other factors imposed on the physical fiber plant, by monitoring the Forward Error Correction (FEC) code the router can identify when the corrected errors justify preemptive action at the IP/MPLS layer. It is hence able to trigger FRR to back up paths prior to reaching the limits of the error correction capability of FEC (which would result in a complete LOS) – and thereby achieve virtually hitless switchover before the traffic is interrupted. This mechanism is comparable to, or better than, traditional SONET/SDH-based failover mechanisms. The failure response time is immensely

improved, as well as the overall end-to-end resiliency of the network integration – a result that is not achievable in more discrete configurations with separate transponders today.

**Figure 4.** Elimination of Transponder Shelves with the Cisco CRS-1 Integrated DWDM PLIMs



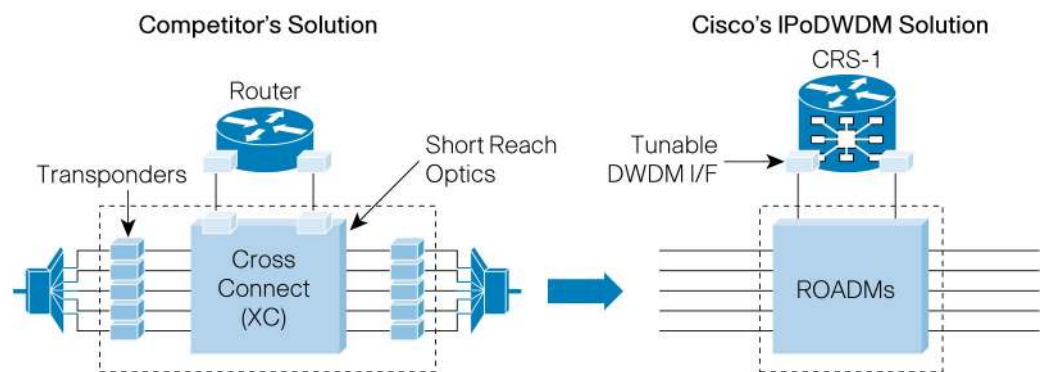
The Cisco CRS-1 1-Port OC-768c/STM-256c (40-Gbps) Tunable WDMPOS Interface Module uses a modulation scheme that is compatible with existing 10-Gbps DWDM systems and amplifiers that many service providers have deployed today. As a result, service providers can now increase their core trunk link speeds fourfold from 10 Gbps to 40 Gbps while still preserving a large portion of their original transport investment in 10-Gbps DWDM systems. Furthermore, the PLIM interfaces are fully tunable across the ITU C-Band to maximize flexibility with existing systems. An enhanced version of the interface supports DPSK modulation. It features integrated Tunable Dispersion Control (TDC) modules with odd and even ITU channels to accommodate 50-GHz spacing applications. Support for Enhanced Forward Error Correction (EFEC) allows the DWDM signal to reach distances of over 2000 km (1242 miles) without requiring optical regeneration (depending on fiber quality) – an industry first compared to other routers available on the market. The Cisco CRS-1 1-Port OC-768c/STM-256c Tunable WDMPOS Interface Module is also equipped with typical SONET/SDH-like OAM&P through G.709 framing to fully support the management functions of previous SONET/SDH-type interfaces while interfacing directly to the DWDM layer.

The Cisco CRS-1 4-Port 10GE Tunable WDMPHY Interface Module has capabilities similar to those of the 1-Port OC-768c/STM-256c Tunable WDMPOS Interface Module, such as EFEC for extended reach and SONET/SDH-like OAM&P through G.709 framing, but is cost-effective at a 10-Gbps LANPHY price point through the innovative use of a WDMPHY-based technology. Thus service providers can take advantage of the benefits of 10-Gbps-like core IP links that map into 10-Gbps DWDM wavelengths with the SONET/SDH-like OAM&P to expand their network as traffic needs increase.

Element integration also includes the incorporation of photonic switching into optical multiplexing equipment, which delivers a reconfigurable optical ADM (ROADM) capability directly on the Cisco

ONS 15454 MSTP,<sup>6</sup> eliminating the need for costly and complex OEO conversions where the traffic simply needs to pass through a site without having to terminate on a router for IP processing. In cases where termination is necessary, the ROADM hands off the optical wavelength, keeping it in the optical domain without the need to perform an electrical conversion in order to hand off the traffic to the router, where the electrical conversion is used only for IP processing. ROADM also provides automatic gain and transient control, eliminating the operational expense of sending technicians into the field to manually tune the systems whenever a wavelength needs to be added or dropped at a remote site. Multi-degree ROADMs (2 through 8 degrees of freedom) allow wavelengths to remain in the optical domain while being passed from one ring or network segment to another, further eliminating the need for OEO conversions and leveraging the ability of core routers to initiate DWDM-compatible wavelengths.


**Figure 5.** Cisco IP-over-DWDM Solution with All-Optical or Photonic Transmission



Keeping traffic purely in the optical domain as much as possible has the added advantage of “future-proofing” a service provider’s core transport network. Pure optical transmission is inherently more tolerant to bit-rate variations where moves to higher rates and new protocols may still be required in the future, and hence more robust because photonic processing is intrinsically insensitive to protocol changes, unlike typical electrical processing elements.

Both of these major element-integration components (the Cisco CRS-1 and the ONS 15454 MSTP) enable service providers to greatly reduce their CapEx and OpEx in ways described above, while simultaneously improving the overall resiliency of the network. Figure 6 highlights the CapEx savings differential of using ROADMs on MSTP together with integrated DWDM PLIMS on the Cisco CRS-1 across an actual service provider core network, excluding the cost of the common components. The service provider saves up to 66 percent on the IP-over-DWDM solution compared with traditional patch-panel or cross-connect approaches.

<sup>6</sup> Further information about the Cisco ONS 15454 MSTP with ROADM technology is available at: <http://www.cisco.com/en/US/products/hw/optical/ps2006/index.html>.

**Figure 6.** CapEx Savings Differential of the IP-over-DWDM Solution

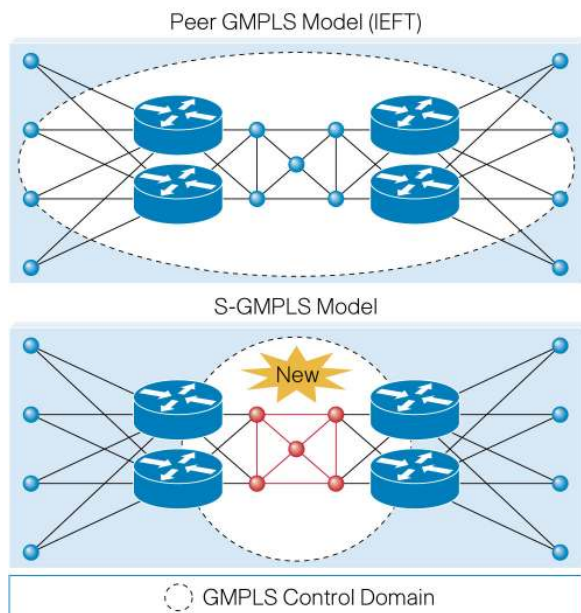
### Control Integration

The second critical integration component of the IP-over-DWDM strategy is that of control integration between the planes of the two networks, IP and DWDM. As mentioned earlier, disparate manual provisioning methods at the transport layer can incur high OpEx and slow service activation that could take weeks to implement, if not longer. Cisco has been leading a multiyear strategy since early 2000 through the IETF to extend the control plane intelligence offered by IP/MPLS to the transport layer within service provider networks. The vision is to allow Generalized MPLS (GMPLS)-based networks composed of optical devices such as the Cisco ONS 15454 as well as IP routers such as the Cisco CRS-1 to dynamically find and provision optimal paths based on user traffic requirements. In other words, a flow that potentially starts on an IP network and is transported by the optical network and then switched through a specific wavelength on a specific physical fiber by the intermediate optical nodes that are GMPLS-capable can be controlled by the overall intelligence of the network. The vision is all-encompassing because it provides a path for optical elements within the transport network today to become peers of the router elements in the IP network, and it provides the capability to autoprovision wavelengths directed by the IP control plane. The benefits to service providers are significant savings in OpEx across the network through the capability to enable fault correlation between the networks in real time while improving the overall speed to service from end to end.

However, both Cisco and the industry realize that the vision of GMPLS, and in particular the currently proposed full peer model of GMPLS, has some major challenges to overcome if it is to be successful in service provider core networks. One of the major challenges for the full peer model of GMPLS is the lack of separation between the routed and optical domains in terms of administrative boundaries, resulting in the exposure of control and topology of the transport network between the transport and data groups or between the service provider and its customers, an exposure that creates both security and operational risks. This model also requires that all the transport nodes be able to run the full GMPLS protocol suite (Open Shortest Path First [OSPF], Resource Reservation Protocol [RSVP], etc.) to interoperate. This is a significant burden on some of the existing transport infrastructure that was primarily designed for manual provisioning. Hence control integration focuses on a phased migration approach to enable the greater vision over a period of time yet allow many benefits to be realized without a full implementation.

To enable this phased migration, Cisco has introduced a new GMPLS model known as the Segmented GMPLS (S-GMPLS), a hybrid of the full peer and the typical overlay model of transport networks today (Figure 7). In the S-GMPLS model, only the border routers receive information from the optical devices and other routers. In fact, the border routers use logical router instances to shield and segment the topology information between the IP domain and the optical domain. The border router acts as a gatekeeper between the two, enabling a segmented administrative boundary and helping ensure management separation between the two networks. This scenario facilitates the deployment of GMPLS in service provider networks because segmented administration between the networks can be maintained where required.

**Figure 7.** Segmented GMPLS (S-GMPLS) Model



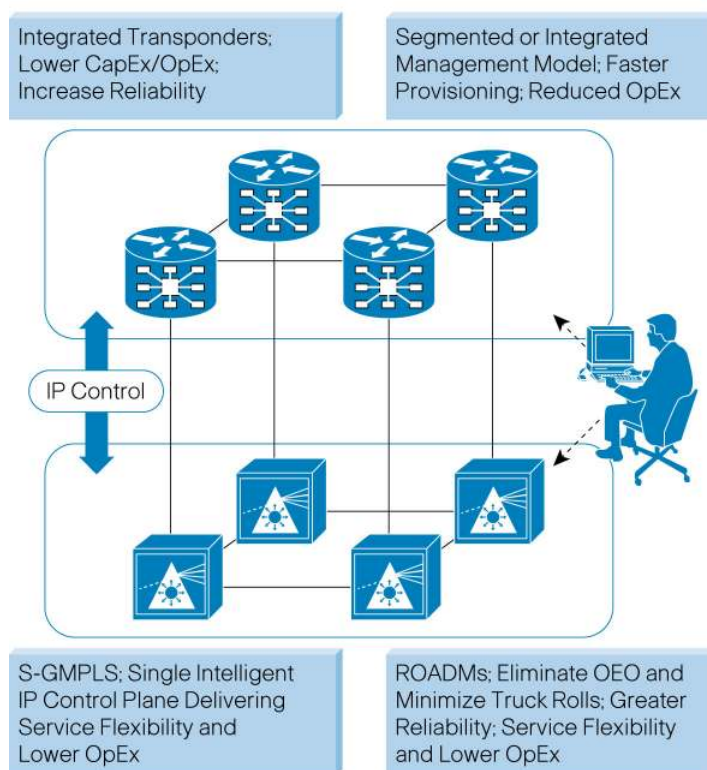
The initial phase of this migration is the enablement of remote provisioning of wavelengths through an integrated management system or any other third-party solution by using multi-degree ROADM technology on the Cisco ONS 15454 MSTP. This provisioning also enables complete autopatching of the intermediate nodes between two endpoints while keeping the transport signals purely in the optical domain, eliminating equipment upgrades to all these intermediate points and reducing OpEx for the service provider while improving speed to service. Subsequently S-GMPLS can be introduced to enable the intelligence in the transport network. The remote provisioning and autopatching capability of the ROADM technology combined with S-GMPLS can help service providers move toward dynamic service activation capability within their core networks.

### Management Integration

The third integration aspect is that of management integration between the IP and DWDM networks. Service providers can significantly simplify operations by using an integrated management approach across both the IP and transport networks. This is achieved by a set of Cisco tools that help consolidate OAM&P across both the Cisco CRS-1 and the Cisco ONS 15454 MSTP. These tools combine troubleshooting and wavelength provisioning from the Cisco CRS-1 DWDM ports to peer router ports through the optical transport layer, reducing operational complexity and costs. And because current service providers' operational organizations are often segmented, these tools also provide for a segmented administration mode, allowing different

operations teams and users to have segmented access between the IP and the transport equipment that is being managed. Service providers can thus take advantage of the cost benefits of a unified management platform while still providing access to segmented operational groups where needed. The tools are also compatible with existing third-party management systems or service provider-developed operational support systems through the support of TL-1, CORBA, and SNMP for northbound interfaces. XML, direct command-line interface (CLI), and HTTP/S are provided for direct element communication. As a final component of the overall IPoDWDM management integration initiative, Cisco provides additional tools for the full management lifecycle, including tools to facilitate design of the optical transmission layer, to plan for known and unknown traffic projections, and to generate installation instructions for the final management of the elements end to end for both the transport and IP network elements within service provider networks.

**Figure 8.** Converged IP-over-DWDM Network Solution



### The Cisco IP over DWDM Solution for IP NGN

As service providers look to reduce their OpEx and CapEx, they should consider converging the IP and DWDM layers. The technology and the means currently exist to enable this convergence and deliver savings and increased revenue through accelerated speed to service and reliability in an increasingly competitive environment. These goals can be achieved without compromising the performance and scale that is needed to accommodate traffic growth for the next decade resulting from application convergence of all video, voice, and data traffic to IP. The new converged IPoDWDM core network solution is illustrated in Figure 8. To deliver on this promise of IPoDWDM convergence, service providers need a partner that understands their business and has a clearly defined vision and strategy, including an execution plan for that strategy, and all the critical components of the end-to-end solution – Cisco.

**Americas Headquarters**

Cisco Systems, Inc.  
170 West Tasman Drive  
San Jose, CA 95134-1706  
USA  
www.cisco.com  
Tel: 408 526-4000  
800 553-NETS (6387)  
Fax: 408 527-0883

**Asia Pacific Headquarters**

Cisco Systems, Inc.  
168 Robinson Road  
#28-01 Capital Tower  
Singapore 068912  
www.cisco.com  
Tel: +65 6317 7777  
Fax: +65 6317 7799

**Europe Headquarters**

Cisco Systems International BV  
Haarlerbergpark  
Haarlerbergweg 13-19  
1101 CH Amsterdam  
The Netherlands  
www-europe.cisco.com  
Tel: +31 0 800 020 0791  
Fax: +31 0 20 357 1100

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](http://www.cisco.com/go/offices).

©2007 Cisco Systems, Inc. All rights reserved. CCVP, the Cisco logo, and the Cisco Square Bridge logo are trademarks of Cisco Systems, Inc.; Changing the Way We Work, Live, Play, and Learn is a service mark of Cisco Systems, Inc.; and Access Registrar, Aironet, BPX, Catalyst, CCDA, CCDP, CCIE, CCIP, CCNA, CCNP, CCSP, Cisco, the Cisco Certified Internetwork Expert logo, Cisco IOS, Cisco Press, Cisco Systems, Cisco Systems Capital, the Cisco Systems logo, Cisco Unity, Enterprise/Solver, EtherChannel, EtherFast, EtherSwitch, Fast Step, Follow Me Browsing, FormShare, GigaDrive, HomeLink, Internet Quotient, IOS, iPhone, IP/TV, IQ Expertise, the IQ logo, IQ Net Readiness Scorecard, iQuick Study, LightStream, Linksys, MeetingPlace, MGX, Networking Academy, Network Registrar, Packet, PIX, ProConnect, ScriptShare, SMARTnet, StackWise, The Fastest Way to Increase Your Internet Quotient, and TransPath are registered trademarks of Cisco Systems, Inc. and/or its affiliates in the United States and certain other countries.

All other trademarks mentioned in this document or Website 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. (0705R)