White Paper

Cisco nLight™ Technology: A Multi-Layer Control Plane Architecture for IP and Optical Convergence

Bandwidth demands on service provider networks continue to grow at exponential rates, driven by packet-based multimedia services such as video streaming, videoconferencing, and online gaming. With cloud networking, content and resources are shifting in real time, causing complex and dynamic traffic patterns. Exacerbating this problem, "over-the-top" media services create enormous challenges for service provider business models.

To help overcome these challenges, IP and optical networks require better integration. Today's network architectures are based largely on isolated IP and optical network layers, designed and operated mostly independently. This makes it difficult to increase service velocity, adapt to dynamic cloud topologies, enhance resiliency, and decrease total cost of ownership (TCO). Simple service requests can take months to fulfill. That simply is not good enough for the next-generation Internet of mobile, video, and cloud services.

This document examines a new multi-layer control plane architectural approach, which increases agility and programmability for IP and optical networks. As an example of the Cisco® Open Networking Environment (ONE) programmability, Cisco nLight™ technology defines a multi-layer control plane that allows service providers to integrate previously independent network layers and share intelligent information. By employing this approach, service providers can reduce network capital expenditures (CapEx) and operational costs (OpEx), while meeting or improving service-level agreements (SLAs) for mobile, video, and cloud services.

Introduction

Modern service providers face a host of challenges, many stemming from the complex, dynamic traffic patterns of mobile, video, and cloud services and increasing demand for bandwidth with high expectations about the quality of service. The expectations include network connectivity, more bandwidth and a "better than best effort" approach. The costs for meeting the expectations of the customer challenge many service provider revenue models. They must balance the cost of infrastructure upgrades against the return on investment. To remain competitive and profitable, service providers must be capable of offering new and improved services while gaining increased efficiencies from the network to lower costs.

The standard network architectures that exist today make adapting to evolving demands costly and difficult. Most existing networks are divided across layer boundaries–IP (Layer 3) and optical transport (Layer 1)–which are designed, deployed, and operated almost entirely independently. Bringing these layers together within a single, unified network environment will increase efficiencies in the network, improve time to revenue, and decrease network TCO. To accomplish this, service providers need a multi-layer control plane that allows relevant information to be exchanged across layers, allowing service providers to automate many functions across both IP-packet and optical transport domains. Such a multi-layer control plane must provide sufficient scale to support service provider networks, allow for interlayer communication without overburdening network elements, respect organizational boundaries, and respect the organizational knowledge base that service providers have developed for each network layer.

Service providers recognize that a multi-layer control plane offers many advantages. Currently, the industry has developed two primary models for multi-layer control planes: the peer model and the overlay model. However, both multi-layer models have inherent flaws. Figure 1 illustrates the peer and overlay models, listing their advantages and their flaws.

Figure 1. Peer and Overlay Models

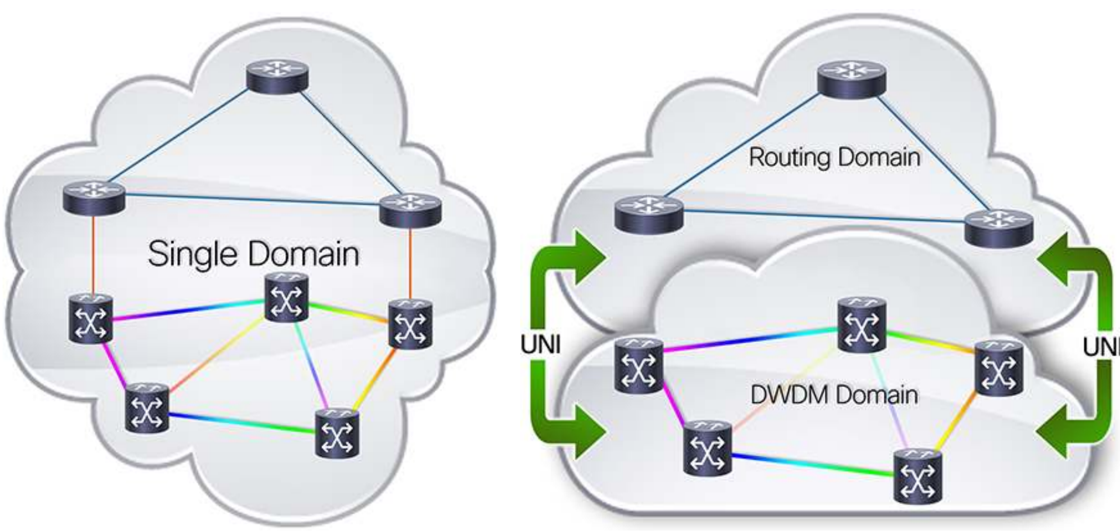

Peer Model	Overlay Model
This model treats the optical and routing layers as part of the same Interior Gateway Protocol (IGP) group. Optical network elements are treated the same as routing elements.	This model treats the optical layer as an autonomous administrative domain running its own control plane, while the routing layer runs its own separate control plane. Each layer is completely independent of the other, but the layers share a user-network interface (UNI) between them to allow for minimal communication, including turning up or tearing down a circuit.
• **Advantages**: Fully integrates the optical transport and IP network layers.	• **Advantages**: Provides appropriate levels of scale and respects the operational and network-layer boundaries within service provider organizations.
• **Flaws**: Overburdens the optical network elements with routing states and configurations, while overburdening the routing network elements with optical characteristics that do not necessarily add circuit-routing value. Scaling this model also becomes a major issue as the optical network grows, as this model does not respect operational boundaries or network-layer boundaries that exist within today's service provider environments.	• **Flaws**: Does not provide the necessary information to support an efficient system. Allows only the most basic crossover functions. Limits the ability to further integrate packet and transport organizational functions. Eliminates service-impacting misconfigurations

While the peer and overlay models offer multi-layer control plane solutions, they both fall short of real-world requirements for modern service providers. The peer model's inability to scale and overburdening of network elements with excessive information sharing create inefficiencies that outweigh benefits. The overlay model solves the bandwidth efficiency problems that are endemic to the peer model. However, in the attempt to optimize bandwidth utilization, the overlay model provides an inadequate amount of information that can be shared between layers. The overlay model also struggles with circuit protection, disjoint circuit routing, and efficient mapping of bundles (Figure 2).

Figure 2. Overlay Model Inefficiencies

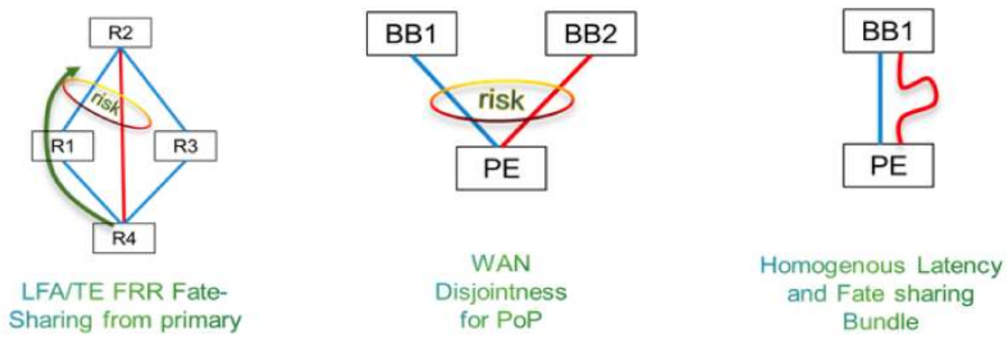

All three of these inefficiencies are directly related to the lack of awareness of the IP layer within the underlying dense wavelength division multiplexing (DWDM) layer. When activating a circuit, for example, the packet layer assumes the underlying network can meet a baseline level of protection. However, if a circuit is constructed along a path with a common Shared Risk Link Group (SRLG), a single failure will impact multiple circuits that may render the packet-layer protocol protection ineffective. To meet an SLA for a data center circuit, the service provider may need to ensure a circuit is completely disjointed from another circuit (or dedicated) to prevent a link failure to the data center. A similar problem can occur if the service provider wants to increase capacity by bundling links. In this case, a matching circuit on the same path is desired to deliver the same latency parameters. In either case, due to the limited information the overlay model shares between layers, the transport layer does not have enough information about what is happening in the packet layer.

The peer model provides too much information and the overlay model provides too little information. Therefore, a new model is required to effectively address multi-layer networks. This new approach to a multi-layer control plane must scale, respect organizational boundaries, respect the expertise organizations have developed around each network layer, and share enough information between layers to address the inefficiencies of the existing multi-layer network approaches. The solution: Cisco nLight technology.

Cisco nLight Technology

Cisco nLight technology addresses the complexity and scaling problems inherent in the peer and overlay models. While the overlay model provides too little information between layers and the peer model provides too much information, the Cisco nLight Control Plane Protocol optimizes the information flow between layers. The Cisco nLight Control Plane is built on the client-server model. Information sharing flows from the "client" (the IP router interface) to the "server" (the DWDM network) and conversely (Figure 3). The Cisco nLight Control Plane supports Packet over DWDM, Packet over Packet transport, Packet over Optical Transport Network (OTN), and multiple layers of Packet over Packet transport over DWDM.

Cisco nLight technology innovations also cover intelligent nLight reconfigurable optical add-drop multiplexer (ROADM) features for zero-touch agility. Because the features are colorless, contentionless, and omnidirectional, Cisco nLight ROADMs remove the constraints of traditional DWDM networks. Cisco nLight ROADMs, together with the nLight Control Plane Protocol, creates a DWDM layer that is much more dynamic and flexible. DWDM networks using Cisco nLight technology can reroute any interface or wavelength in any direction across a meshed network, as well as allow the ROADM interface to change wavelengths when needed. Prior to the introduction of nLight ROADMs, the static nature of DWDM transport minimized any value associated with a multi-layer control plane, because any changes on the DWDM network required manual intervention.

Implementing a Multi-Layer Control Plane

Cisco nLight technology solves the complexity, scalability, and operational inefficiency problems of today's multi-layer approaches. The model addresses the shortcomings of the peer and overlay models while delivering advanced features that capitalize on the tighter coupling between packet and transport layers in the network (Figure 3).

Figure 3. Cisco nLight Provides Benefits of Peering with Intelligent Information Sharing

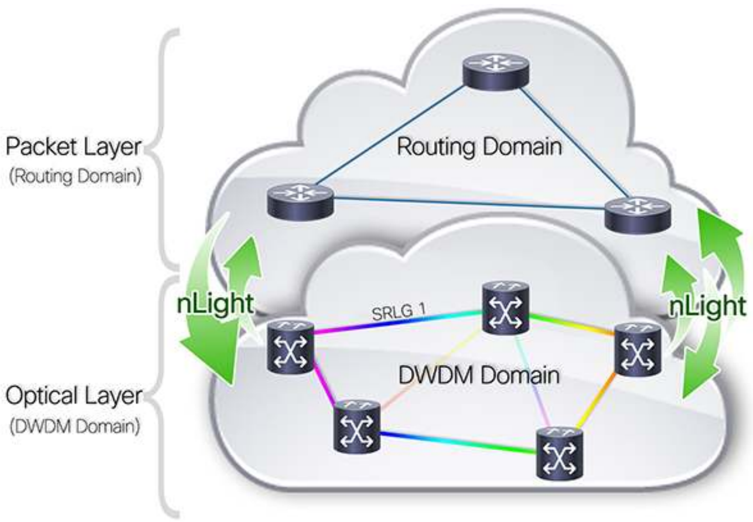

The Cisco nLight Control Plane focuses on Layer 3 and Layer 1 networks in which the Layer 3 network uses IP over DWDM (IPoDWDM) interfaces for converged efficiencies. However, the control plane is scalable and flexible enough to adapt to a number of different deployment options, including (but not limited to):

- **IP/MPLS over DWDM**, shown in Figure 3 with IPoDWDM
- **IP/MPLS service over IP/MPLS transport**, or a packet-over-packet network that maintains two separate IGPs to provide independent routing layers
- **OTN switching over DWDM**, in which the OTN switch acts as the client of the DWDM server
- **IP/MPLS service over IP/MPLS/OTN transport over DWDM**, or a "Russian nesting dolls" approach of using multiple layers that are truly independent, yet share relevant information

The Cisco nLight Control Plane scales with the network and has the ability to customize layer interaction so that service providers have the flexibility to unify layers or maintain layered boundaries as requirements dictate.

Sharing Information Between Layers

Each layer in the network contains a host of information necessary for transporting data from point A to point B. The DWDM network contains information regarding the end-to-end transport network, including how a circuit gets from point A to point B, over what direction, using which resources and which combinations of other circuits. At the same time, the network layer devices maintain information about the underlying traffic requirements, including how much latency can be tolerated, what level of SLA the traffic demands, and how the traffic will be protected.

To realize efficiencies, the Cisco nLight Control Plane Protocol must facilitate information sharing between these two network layers. Unlike the peer approach described previously, information sharing in the Cisco architecture does not overburden any of the layers with layer-specific information that does not add value. The Cisco nLight approach allows each network layer to share abstracted information that improves overall network efficiency. The DWDM layer acts as a server and the IP layer acts as the client, sharing information such as circuit IDs, path-level latency, optical cost, Layer 0 shared risk link groups, and the optical path.

With the ability to share only relevant information between layers, the approach can optimize and automate many aspects of service delivery that previously required extensive manual effort (and therefore, extensive time and costs). Consider the process of turning up a new circuit. Today's circuit turn-up procedure is almost entirely manual. It requires network engineers to have access to a detailed network database, and requires extensive manual coordination between packet and optical organizations. This manual process entails the following:

- The packet team requests a new circuit from the DWDM transport team.
- The DWDM team verifies the available path based on the requested service type.
- The DWDM team verifies that the performance and resources are available to ensure the path is a valid candidate for the requested service.
- The DWDM and packet teams coordinate circuit turn-up.

Even this description does not capture all of the complexity involved in this process. For example, ingress and egress circuits may be different, effectively doubling the turn-up effort required. In addition, the DWDM team typically does not have the automated tools to easily verify the parameters described above, and must rely on manually updated databases. As a result, the entire process can take weeks or even months.

Now, consider the circuit turn-up process using Cisco nLight technology. Here, the Cisco nLight Control Plane automates many of the manual processes using interlayer signaling (Figure 4). When turning up a new circuit or service, the packet-layer request can include parameters such as:

- Matching path
- Disjoint path
- Specific latency or latency bound
- Avoiding or including specific SRLG(s)
- Lowest optical cost circuit

As a result, the Cisco nLight technology process entails just two steps, which can be accomplished in minutes:

- The client IP router signals a new circuit request, along with the requirements for the new circuit, to the server DWDM layer.
- The DWDM layer signals a valid wavelength to the packet layer, or signals a path error message if the request cannot be filled, and updates its internal network database.

Figure 4. Cisco nLight technology Provides the Benefits of Peering with Intelligent Information-Sharing

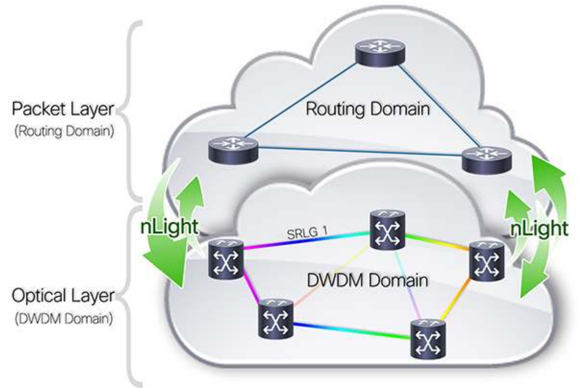

Diagram A illustrates an integrated packet and optical network using Cisco nLight technology information-sharing and signaling.

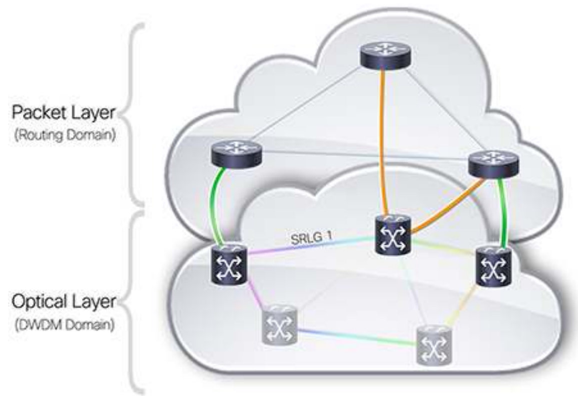

Diagram B shows a request for a low-latency circuit, illustrated as the orange line.

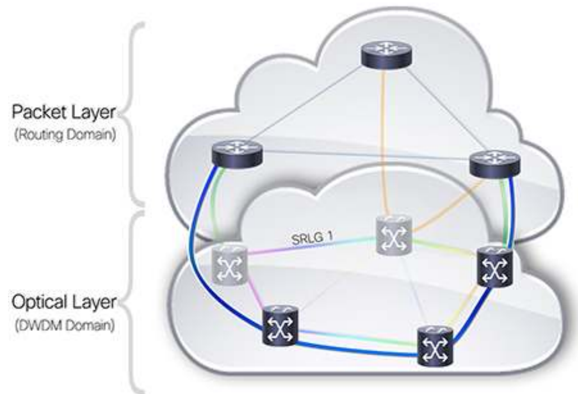

Diagram C illustrates a request for a disjoint path, depicted as the blue line.

<table>
<tr>
<td>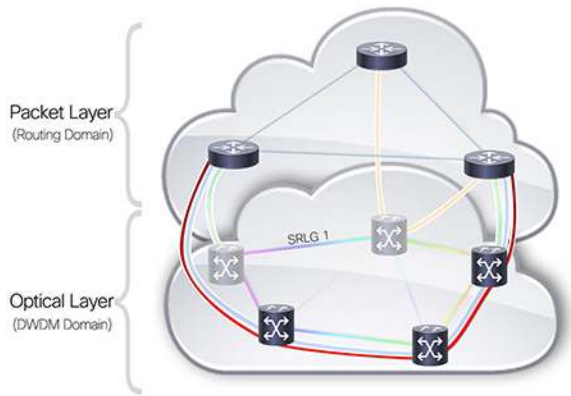
</td>
<td>In Diagram D, the red line shows a request for a matching path or link bundle.</td>
</tr>
<tr>
<td>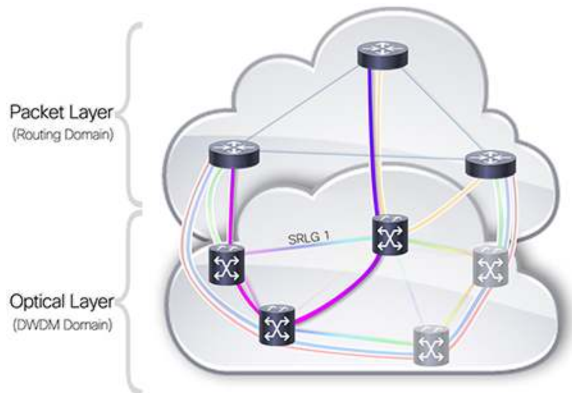
</td>
<td>Diagram E shows a request to avoid a specific Shared Link Risk Group (SRLG 1), shown as the purple line.</td>
</tr>
</table>

As illustrated, Cisco nLight technology simplifies circuit creation and optimizes operational efficiency by treating the network as a pathing database. All network information is automatically updated between layers and maintained by the network (rather than by manually updating databases), thus supplying up-to-date and accurate data between layers to optimize paths and adjust for topology changes. Should a circuit fail within the network, Cisco nLight technology is capable of restoring the circuit, without manual intervention.

Circuit Protection and Restoration

Today, service providers protect packet traffic using optical-layer protection and IP-layer protection. Optical protection uses either a “1+1” approach or optical fiber protection. Using 1+1 protection requires two transponders to send a single packet interface in two directions. Optical fiber protection splits the output of the transponder in two directions and switches at the far end based on power changes. Both of these protection approaches increase the cost of the network. The 1+1 approach requires twice the number of transponders and reserves half the true bandwidth capacity of the DWDM system in case of failure. The optical protection mechanism also cuts DWDM capacity in half, and is not reliable over long distances, so it is rarely used in service provider networks.

IP layer mechanisms require added network planning in the transport network to determine the worst-case traffic through an interface during failure scenarios (link, node, and interface), and require network engineers to design the optical transport network so as to assure that the capacity is available in case of a failure. Using the IP layer protection mechanism results in 50 percent interface utilization of the optical transport network (Figure 5).

Figure 5. Comparing 1 + 1, Fiber Protection, Layer 3 Protection, and Layer 3 with IPoDWDM for Packet Services

1 + 1 Protection	Fiber Protection	Layer 3 Protection
		
1 Layer 3 interface 1 DWDM interface No protection against Layer 3 interface failure No protection against router failure Five 9s from transport perspective	1 Layer 3 interface 1 DWDM interface DWDM interface is split into 2 paths No protection against Layer 3 interface failure No protection against router failure Not five 9s from transport perspective	2 Layer 3 interfaces 2 DWDM interfaces unprotected Protects against Layer 3 interface failure Protects against router failure

Using optical restoration can significantly reduce overall network costs. However, optical restoration does not mitigate the need for IP or packet-layer coordination in circuit protection. This is because in a pure optical restoration approach, the optical network maintains no awareness or understanding of the underlying affected circuit. What are the constraints on the circuit, and what level of tradeoffs can be made with the restored path? How can the optical restoration mechanism address the circuits after a failure is repaired in a hitless manner? Pure optical-layer restoration mechanisms cannot answer these questions; only the IP layer is aware of the actual payloads, the priority, and the constraints on the traffic. The IP layer decides how to reroute traffic, over which interfaces, and what level of priority each circuit or service requires. Using advanced IP protection mechanisms such as IP Fast Reroute or Proactive Protection, the network can restore traffic in less than 50 milliseconds, and in some cases in an almost hitless manner, whereas traditional optical-layer restoration mechanisms are relatively slow, on the order of minutes. The best approach, therefore, is to combine them – using the IP layer for protection and circuit constraints, while using the optical layer for the restoration. Ultimately, this approach can yield significant network savings.

Cisco nLight technology combines the advantages of the IP layer with those of the optical layer. Using the Cisco nLight Control Plane Protocol and Cisco IPoDWDM allows service providers to implement protection through the IP layer and restoration through the optical layer. As a result, service providers can reuse actual DWDM interfaces and increase the utilization of router interfaces, decreasing overall network costs.

Figure 6 details the Cisco nLight technology approach to link protection and circuit restoration. Diagram A depicts a multinode network encompassing multiple wavelengths, with the illustration focusing on two wavelengths, depicted as red and orange lines. A fiber failure or degradation along the orange path will initiate proactive protection in the IP layer to protect the traffic, as shown in diagram B. The optical layer then restores the capacity back to the network. This process may require a new wavelength to be used over the new path. The DWDM layer then signals to the router interface that there is a new path that requires the routers to tune to a new wavelength (the blue line in diagram C, on the right side of the figure). As diagram C shows, the router changes its source wavelength from orange to blue, and the DWDM layer routes the new wavelength to the destination. Note that from the IP layer perspective, there is no change in the network topology. At time of failure repair, the Cisco nLight Control Plane analyzes the options of reverting back to the original path in a hitless fashion or remaining on the existing path. Without Cisco nLight technology, these decision-making capabilities would not exist, nor is there the ability to revert back to the original path in a hitless manner.

Figure 6. Cisco nLight technology for Protection and Restoration

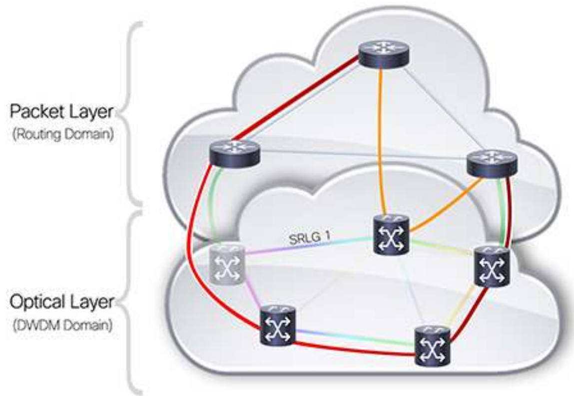

A: This diagram depicts a multi-node network encompassing multiple wavelengths, with the focus on two wavelengths depicted as red and orange lines.

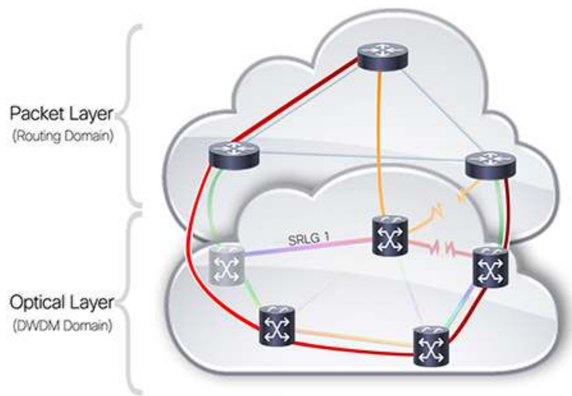

B: A fiber failure or degradation along the orange path will initiate proactive protection in the IP layer to protect the traffic, as shown in the diagram. The optical layer then restores the capacity back to the network. This process may require a new wavelength to be used over the new path. The DWDM layer then signals to the router interface that there is a new path that requires the routers to tune to a new wavelength (the blue line in the next diagram).

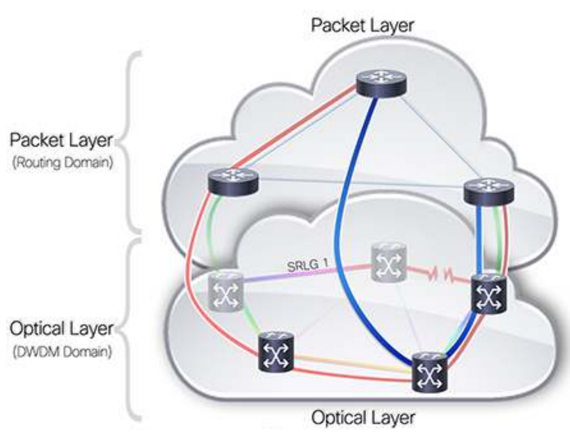

C: As this diagram shows, the router changes its source wavelength from orange to blue, and the DWDM layer routes the new wavelength to the destination. Note that from the IP layer perspective, there is no change in the network topology. At the time of failure repair, the Cisco nLight Control Plane analyzes the options of reverting back to the original path in a hitless fashion or remaining on the existing path. Without Cisco nLight technology, these decision-making capabilities would not exist, nor is there the ability to revert back to the original path in a hitless manner.

Figure 7. Cisco nLight technology for Protection and Restoration

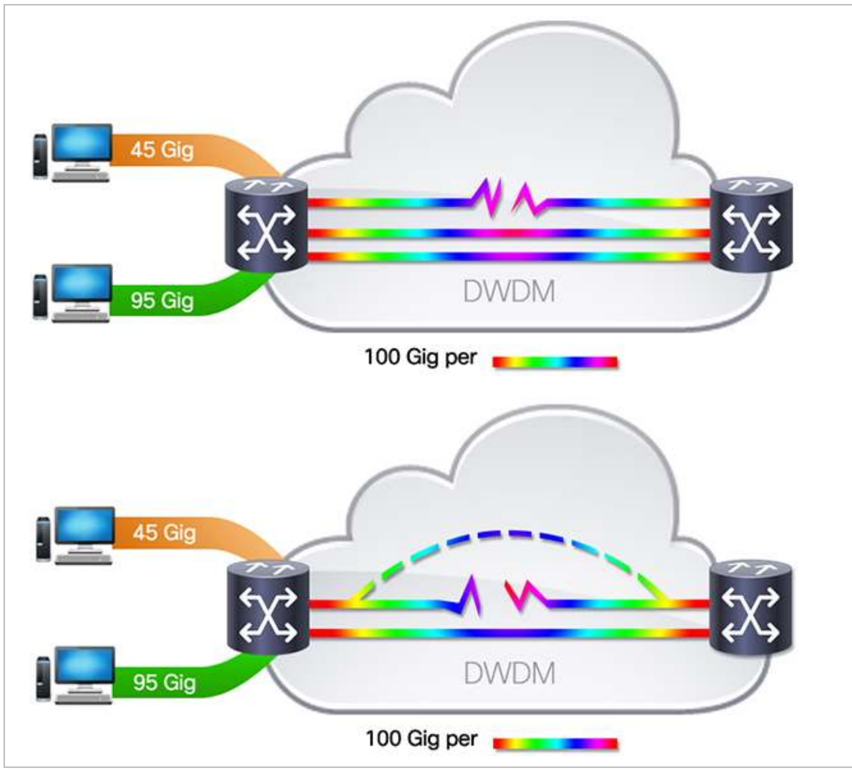

Layer 3 Protection

- Six 100G Interfaces
- 300G Capacity
- 140G Traffic
- 47% Normal Utilization
- 70% Failure Utilization

Cisco nLight technology Restoration

- Four 100G interfaces
- 200G capacity
- 140G traffic
- 70% interface utilization

The first diagram in Figure 7 illustrates the Layer 3 protection case, in which bandwidth is provisioned in reserve in case of a failure, requiring additional router and DWDM resources. The second diagram in Figure 7 shows how Cisco nLight technology restoration increases utilization, reduces the overall number of interfaces, and provides full bandwidth support for a fiber failure.

Cisco nLight technology circuit restoration increases utilization on an interface while providing the level of resiliency necessary for mission-critical traffic. As a result, it allows service providers to realize significant savings in overall network costs, even as they support the most demanding SLAs.

Conclusion

Integrating the intelligence of network layers on a multi-layer control plane can provide network-wide improvements. Cisco nLight technology provides a multi-layer control plane that addresses many of the inefficiencies plaguing conventional networks by sharing information between the Layer 3 network elements and the Layer 1 DWDM elements. With Cisco nLight technology, service providers can automate complex processes and take advantage of advanced features that reduce overall network costs by reducing the number of interfaces required on routers and the number of transponders required. Cisco nLight technology allows service providers to realize significant CapEx and OpEx savings, while increasing the longevity of their DWDM infrastructure.

For More Information

To find out more about the Cisco nLight technology, visit the Cisco Optical Networking Solutions page.

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.**

Printed in USA C11-718852-00 12/12