

eBook: The 100G Challenge and Beyond

Introduction: Transitioning to 100G and Beyond: The Big Picture

As the industry moves forward to meet the enormous demand for data with video, mobile, and cloud, the core networks need to transition from 10G to 100G and beyond.

Consider some of these findings from Cisco's 2013 [Visual Networking Index](#) (VNI) report:

- Global network users will generate 3 trillion Internet video minutes per month, which translates to 6 million years of video per month, or 1.2 million video minutes every second, or more than two years' worth of video every second.
- Wi-Fi and mobile-connected devices will generate 68 percent of Internet traffic by 2017.
- More traffic will traverse the Internet in 2017 than the total amount of traffic that traveled the Internet cumulatively from 1984 to 2012.
- The "Internet of Things" has arrived, and applications such as digital health monitors, smart meters, asset and package tracking, chips for pets and livestock, and video surveillance are driving more and more traffic. Globally, machine-to-machine (M2M) connections will grow three-fold, from 2 billion in 2012 to 6 billion by 2017. Annual global M2M IP traffic will grow 20-fold over the same period - from 197 petabytes in 2012 to 3.9 exabytes by 2017.

While this dramatic growth occurs, projections call for the cloud to account for nearly two-thirds of data center traffic and for cloud-based workloads to quadruple over traditional servers. That adds another element to the picture: changing traffic patterns. Under a cloud model, a university, for example, can build its network to handle average traffic volumes but then offload data on heavier trafficked days to a public cloud service when demand dictates, such as when it's time to register for the next semester of classes.

The cloud introduces new stresses to the network and creates new traffic patterns. Data centers traditionally hosted north-south traffic, with data traveling from the data center to the user. Now, however, traffic patterns are shifting, with information flowing from data center to data center in what's considered an east-west route. That increases the number of interconnects and makes the inherent traffic patterns more distributed.

In addition, the difference between average and peak traffic volumes is increasing, and event-specific triggers, such as news events and/or natural disasters, prompt surges in traffic that are impossible to predict. Networks need to react faster than ever before.

In the past, it was enough for service providers to identify major traffic patterns, bolster network infrastructure around those patterns, and periodically adjust as needed. Now, with so many changes to traffic patterns and the unpredictable nature of that traffic - a website suddenly becomes popular or a new mobile device is introduced - service providers need to be able to react within minutes rather than waiting weeks or months. They also need to meet increasingly stringent service-level agreements (SLAs) to retain their top customers.

Clearly, core networks built with 10G interconnects that were considered adequate in the mid- to late 1990s are no longer enough, and while the industry prepares to shift to 100G and beyond, a lot of preparation needs to happen before that's achieved. This newly conceived network needs to be optimized in terms of size, cost, and operational capabilities. In fact, an Evolved Programmable Network (EPN) that addresses this network transition must offer an open, fully virtualized, programmable, and automated networking infrastructure. IP+Optical convergence is one important piece in the puzzle.

IP+Optical Convergence: What Is It, and Why Does It Matter?

For service providers and their customers, 100G is a timely and critical boost in transport services. It's a critical enabler for service provider network and data center interconnects to scale in response to service evolution and increasing bandwidth demands. But the ability to transport data at that speed is just one aspect of the overall scalability story. To receive the greatest benefit from 100G, operators must optimize many aspects of the network, and that optimization always involves reevaluating historic network architectures. Since the introduction of the Internet, IP and optical networks have operated independently, preventing integration and limiting the ability to share network intelligence in real time. But these two layers cannot continue to operate as independent "silos." They need to interact with one another and give operators the ability to react to customer demand within minutes rather than months.

To address this network transition to EPN, Cisco developed the Cisco[®] Network Convergence System (NCS) family of products. The Cisco NCS family is built to provide simplicity in converging IP and optical networks. Since service providers have many different business models and changing network applications, the NCS family is offered in three primary forms to address various business situations ranging from simple 10G to 100G dense wavelength division multiplexing (DWDM) network migration to full-on network redesign for the Internet of Everything. This is a quick summary of the different NCS solutions:

- Cisco NCS 6000 Series Routers: High-density router, optimized for 100G and beyond, delivering unmatched network agility and petabyte scale at the lowest total cost of ownership (TCO). This intelligent, distributed system provides the network fabric for the EPN. It's built on the industry-leading Cisco nPower X1 network processor, and features the industry's first terabit-per-slot capability.
- Cisco NCS 4000 Series: Multiservice aggregation platform that bridges the IP and optical layers for 100G and higher-rate optical interfaces. It provides multichassis scalability and full multilayer optical transport network (OTN), Carrier Ethernet, and IP/Multiprotocol Label Switching (MPLS) service aggregation in a transport form factor.
- Cisco NCS 2000 Series: Agile reconfigurable optical add/drop multiplexing (ROADM) platform that takes advantage of Cisco nLight™ ROADM technology. It offers touchless optical layer programmability, massive DWDM scale through the use of 100G and super channels, and ultra-long-haul and metro performance options. The NCS 2000 Series is the first platform in the industry to deliver Flex Spectrum capability to the optical layer.

The NCS products use common software, common network management, and the multilayer Cisco nLight Control Plane, which allows service providers to integrate these previously independent networking layers and intelligently share relevant network information, all in the interest of WAN optimization. Using this approach, service providers can reduce network capital expenditures and operational costs while meeting or improving SLAs for mobile, video, and cloud services. While traditional solutions have been focused on layer-by-layer product advancements, Cisco is delivering an opportunity to break down the legacy model of per-layer technology silos and operate the network as one streamlined entity.

Why Converge IP and Optical?

Traditionally, if a large customer put in a request for a new 100G connection between its home office and an office in a distant location, the Layer2/Layer3 network engineer receiving the request would need to fill out paperwork, review the customer's SLA requirements, send the request to the transport group, wait for a reply, and eventually work with the transport group to identify the ideal fiber path, including a disjoint backup path, for their service. That manual process could take weeks or months to accomplish. This extended timeline occurs because the IP and optical layer personnel operate independently, with little shared knowledge, yet they ultimately need to communicate and cooperate to fulfill the customer's service request within their siloed workflows.

Through a series of network enhancements, it's possible to converge the IP and optical layers to create a control plane where the two can dynamically share information, adding intelligence for real-time updates between the two and making each layer aware when the architecture changes. This approach eliminates much of the human intervention, as well as the errors associated with that intervention. It also accelerates service delivery so that service providers can more quickly fulfill their customers' requests and get revenue-generating services to market faster.

As part of its initiative, Cisco developed its Open Network Environment (ONE), a comprehensive solution to help networks become more open, programmable, and application aware. The capabilities of Cisco ONE help meet the needs of various market segments, including evolving concepts such as software-defined networking (SDN). With this solution in place, it is possible to gather network intelligence, orchestrate and evaluate the impact of various changes on the network, and finally, reprogram the network in a way that best suits all the constraints of any number of services.

The Cisco nLight Control Plane is built on the client-server model, in which information flows from the "client," the IP router interface, to the "server," the DWDM network, and conversely. Information is harvested one layer at a time and then shared across layers to automate network deployment and service creation. The Cisco nLight Control Plane supports Packet over DWDM, Packet over Packet Transport, Packet over OTN, and multiple layers of Packet over Any Transport over DWDM.

It's important to recognize that previous generations of DWDM equipment have not supported software-configured "any-to-any" connectivity. Directing any add/drop channel (client) to any multidegree composite signal (trunk) required manual fiber recabling. The recent advent of colorless, contentionless, and omnidirectional capability in ROADM equipment makes software-configured "any-to-any" DWDM connectivity a reality. This added ROADM functionality means that nLight Control Plane technology can be meaningfully extended to the DWDM layer to cover ROADM features for zero-touch optical agility for the first time. Cisco's addition of Flex Spectrum capability is another new "freedom" within the ROADM technology that allows the network operator to match the bandwidth required to the optimal amount of optical spectrum consumed. Automation and flexibility go hand in hand with this new approach.

With the per-layer technology advancements and nLight Control Plane automation, the 100G service that's described in the opening paragraph of this section can now be fully automated across layers. If the service is IP+Optical, as it is in the original example, the Layer 2/Layer 3 network engineer can make a service request into a converged network management system or SDN controller, and then the nLight Control Plane will signal across layers and create the optimum deployment scenario for both the IP and optical layer personnel. Both groups can see and understand the end-to-end workflow and authorize the resources for their respective portions of the network. IP+Optical convergence is built on cross-layer information sharing to optimize network design and network operations.

Elements of IP+Optical Convergence

Cisco recognized years ago that as bandwidth consumption soared, operators would need a cost-effective way to deploy and manage 100G technologies. Many network operators are finding Cisco nLight technology to be a quick and simple solution for resolving their bandwidth needs. Cisco 100G nLight coherent silicon can work with most of the existing fiber in the ground, and has even been deployed in undersea applications.

Key elements in the transition to 100G include Cisco's coherent 100G optics, next-generation ROADMs, complementary metal-oxide semiconductor (CMOS) photonics, and Cisco Quantum™ WAN Orchestration with the nLight Control Plane.

Coherent 100G

Coherent technology is a high-end mathematical algorithm for digital signal processing technology that provides a more sophisticated way of gathering information from pulses of light. Some consider noncoherent systems to be analogous to Morse code, whereas a coherent system receiver works with the local oscillator to tune into the frequency of interest and recover a more clear view of the eye diagram, or signal-to-noise ratio, of the optical signal. Coherent technology has become increasingly necessary as the industry moves from 10G and 40G to 100G and beyond, and as we cope with lower signal-to-noise ratios within these data rate transitions.

As part of the move to coherent technology, the Cisco 100 Gigabit Ethernet Coherent DWDM Interface Modules enable long-haul and metro DWDM transport on Cisco's ROADM platforms, including the NCS 2000 Series, supporting distances of over 4800 kilometers. Also available on the Cisco CRS platforms, the industry's most widely adopted core routing system, and the NCS 4000 Series, Cisco's 100 Gigabit Ethernet Coherent DWDM interface demonstrates the company's continued strong heritage in combined IP+Optical innovations.

DWDM is an optical technology used to increase bandwidth over existing high-capacity fiber-optic backbones. The solution works by combining and transmitting multiple signals simultaneously at different wavelengths on the same fiber. It is protocol independent, so DWDM can carry SONET, SDH, storage protocols, data, and video. Essentially, multiple signals can be transmitted at the same time on one optical fiber, with each signal traveling on a different wavelength and offering the ability to be transmitted over extended distances if needed.

Naturally, service providers want to use as much of their existing infrastructure as possible as they migrate to newer technologies. Cisco's solution set provides the ability to convert 10G networks into 100G networks while taking advantage of more than 95 percent of the existing fiber plant. In fact, Cisco's coherent 100G has been field proven as an alien wavelength, running 100G services seamlessly over the remaining channels of most competitors' legacy 10G DWDM platforms. With only 4 channels remaining in a legacy 40-channel 10G DWDM network, Cisco can more than double the lifespan of that network asset with 100G channels.

Next-Generation ROADMs

Sometimes referred to as optical routers, the reconfigurable optical add-drop multiplexer (ROADM) has matured greatly since it was introduced commercially almost a decade ago.

Cisco manufactures the industry's most widely deployed ROADM and has been in the ROADM business for over a decade. Cisco offers a fully integrated ROADM solution for delivering any wavelength to any location in a metro, regional, or long-haul network. The latest Cisco NCS 2000 Series ROADM cards are part of Cisco's intelligent ROADM architecture, designed to reduce DWDM complexity and speed the deployment of next-generation networking solutions.

ROADM cards typically operate on the fixed-grid International Telecommunications Union (ITU) 50-GHz wavelength plan, with each card integrating automatic per-channel optical power control capabilities for automatic node- and network-based power level management. Per-channel optical path selection is also done automatically through Wavelength Path Provisioning (WPP) at the network level, featuring GUI-based wavelength provisioning and SONET/SDH-like wavelength management. Depending on the desired network management model, the ROADM shelves can also be managed individually or by a Cisco router such as the Cisco ASR 9000 Series. Managing a ROADM shelf with a router is called optical network virtualization (optical nV), and it allows the ROADM shelf to look like a remote line card of the router, from a management perspective. It is important to provide a wide array of management options as operations teams shift from various points in the current network lifecycle toward the evolved programmable network.

The flexibility provided by the ROADM unit simplifies the design of optical networks. With the Cisco Transport Planner optical design tool, network planners can design DWDM networks based on ROADM functions and verify all possible optical paths and DWDM interface types. This method of network design allows network managers to compare different solutions and create “what-if” scenarios.

Early iterations of the ROADM were always “touchless” or software reconfigurable in the middle of the connection path, but not necessarily at the endpoints or add/drop ports. The greatest benefit of the current ROADM iteration is the introduction of complete touchlessness to the optical layer. This software programmable “any-to-any” provisioning is a key ingredient in the IP+Optical value proposition. After all, if you can’t fully reprovise the optical layer based on service churn, how much flexibility and efficiency can you achieve in a dynamic service environment?

Because the Cisco NCS 2000 Series ROADMs feature colorless, contentionless, and omnidirectional provisioning, they remove the constraints of traditional DWDM networks. DWDM networks using nLight technology can signal to reroute any interface or wavelength at any point in the optical span when needed, removing the need for truck rolls, manual intervention, and fiber recabling at the optical add/drop ports.

The initial Cisco nLight Control Plane focused on Layer 3 and Layer 1 networks in which the Layer 3 network used IP over DWDM interfaces for converged efficiencies. The control plane has evolved and is flexible enough to adapt to a number of different deployment options, including IP/MPLS over DWDM and OTN switching over DWDM. ROADM technology combines the wavelengths into a single, composite signal and then transports the composite wavelength across the network. A multidegree ROADM in the NCS 2000 Series can route individual wavelengths in up to 16 or even 32 different directions or degrees. From optical add/drop port to multidegree composite trunk port, it can all be signaled and reconfigured with the nLight Control Plane.

Flex Spectrum: With NCS 2000 Flex Spectrum technology, multiple channels of a specified modulation scheme are effectively overlapped to achieve a higher effective transmission rate. Channels grouped in this fashion are referred to as “super channels.” These super channels are overlapped within the optical spectrum such that they consume less optical spectrum than they would in a traditional fixed-grid design. Whereas fixed-grid designs tune channel boundaries in fixed increments of 50 GHz and 100 GHz, Flex Spectrum boundaries are of arbitrary length and are tuned in variable increments of 12.5 GHz.

Flex Spectrum technology adds another level of freedom in optical networking design. The network operator can now better match signal rate and signal reach with the lowest amount of optical spectrum required, and thus be efficient and well prepared for the 100G-plus era. Optical technologies have a long network lifespan, so the forward-looking optical spectrum design flexibility enabled by Flex Spectrum has gained widespread attention throughout the industry.

The latest generation of Cisco NCS 2000 Flex Spectrum ROADMs also supports colorless, contentionless, and omnidirectional capability.

Colorless: Colorless ROADM add/drop ports are not frequency specific, so any single frequency or channel can add or drop from any physical port. This introduces far more flexibility into the system than older fixed-frequency add/drop solutions.

Contentionless: Contentionless add/drop allows the same frequency to be added or dropped from multiple ports on the same add/drop multiplexer, assuming these same frequency channels will be leaving the ROADM on different degrees or directions. Current ROADMs do not have this capability.

Omnidirectional: A wavelength can be routed from any ROADM port in any direction. Earlier solutions linked physical add/drop ports to single composite signal degrees or directions. If you wanted to send an add/drop wavelength in the east direction when it was built to go west, recabling was required.

CMOS Photonics

The height of the dot-com boom saw major advances in both photonic and electrical transmission technologies. But while the electronics industry advanced, predictably reducing size and cost in accordance with Moore's Law, the same could not be said for optics. Even as optical interfaces reach 100G, the physical size of the modules and the excessive heat and power they draw are limiting the ability to scale network elements, data centers, and networks. Use of exotic materials such as gallium arsenide, lithium niobate, and indium phosphide in other equipment vendors' 100G client optics also contributes to added cost, heat, and complexity in manufacturing.

Application-specific integrated circuits (ASICs) used in data center routers, switches, and servers are becoming smaller, faster, and more power efficient, benefiting from the CMOS fabrication processes. In fact, almost every computing and electronic product benefits from this drive to produce smaller, faster, and cheaper CMOS chips.

In optical solutions, however, most of the photonic elements, such as modulators and switches, are independently manufactured and assembled from exotic compounds and complex packaging, with each part optimized for a specific function. As a result, optical solutions remain expensive, time-consuming to create, and inefficient in terms of power consumption.

Cisco is changing the optical components space with CMOS photonics, integrating silicon-based photonic elements within the patent-pending CMOS manufacturing process and allowing silicon manufacturers to print optical componentry directly onto a small CMOS die in the same way they fabricate ASICs. Cisco is leading the way in making CMOS-photonics pluggable transceivers a viable optical solution for optical networks and data centers. Finally, Moore's Law can be applied to optics, and service providers can be assured that optical solutions will meet the IP traffic demands of the future.

Cisco CPAK

In March 2013, Cisco announced the Cisco CPAK™. Initially available on the 100G coherent transponder for the ROADM products, it was added a few months later to the Cisco CRS-X router and the Cisco Nexus® 7700 for the data center. The NCS family now also uses CPAK across the board. The Cisco CPAK represents a reduction in space and power requirements of more than 70 percent compared with alternative 100G transceiver form factors, such as CFP, while remaining fully interoperable at SR-10 and LR-4.

Cisco is building on CPAK to provide a full product portfolio, allowing service providers to use advances in CMOS photonics in a variety of pluggable client and network-side applications, from transport to routing to switching. It's also building these solutions with reduced cost, size, and power requirements.

Ten CPAKs are equivalent to 1 terabit and use less power than a 60-watt light bulb. While a typical CFP houses four modules in a single-row front panel measuring about 400 mm wide, the CPAK configuration allows for 10 modules in the same amount of space, producing up to 1000 G. CPAK technology is a key building block and differentiator for many of Cisco's next-generation products.

The award-winning Cisco CPAK solutions provide the foundation to meet the capacity demands of the next generation of IP+Optical networking.

Cisco Quantum WAN Orchestration

The Cisco Quantum product suite integrates network analytics, policy control, self-optimizing network, and WAN orchestration capabilities.

Cisco Quantum WAN Orchestration is one element in the Cisco Quantum solution portfolio, which itself is a subset of the larger Cisco ONE solution. Certain elements of Cisco Quantum are applicable to any service provider's network to orchestrate voice, video, data, and mobile services.

Cisco Quantum WAN Orchestration uses the nLight Control Plane, which shares network intelligence between the IP and optical layers and takes full advantage of the capabilities of the recently announced nPower network processors, thus giving Quantum the unique ability to make highly scalable and informed path computation decisions with the world's highest degree of network intelligence.

For many years, this type of decision making was performed by separate departments within a service provider's engineering teams in a process that required consultations between personnel over the course of weeks or months. With Cisco's solutions, network services can be reconfigured and provisioned within minutes. Using analytics to simulate IP and optical performance and recommend optimal circuit provisioning paths based on path computations (such as average and peak traffic patterns, SLAs, link failure analysis, network architecture, and more), Cisco offers the ability to make informed path computation decisions quickly. And in a world of constant and unpredictable network change requests, this multilayer view of the network is clearly long overdue.

Automating the exchange of information between the IP and optical layers is leading to a paradigm shift in operating efficiencies for operators and represents radical simplification of existing operations to accelerate service delivery.

nLight Technology = IP → Optical, and Optical → IP Information Sharing
Quantum WAN Orchestration = Dynamic Multilayer Path Computation

Benefits of IP+Optical Convergence

Cisco's nLight technology, zero-touch ROADMs, network virtualization, nPower Network Processor Units (NPUs), coherent 100G technology, CPAK, and other operational innovations such as Cisco ONE represent significant contributions for IP+Optical convergence. Over the past few years, Cisco has been inventing and rolling out these building blocks to meet customer requirements for a more agile and automated EPN.

The IP+Optical value proposition for service providers' customers falls into two categories:

- **Improving the network services they offer today:** A more agile network that provides quicker access to highly customized services, with the granularity to turn them on and off and change as needed, makes service providers more competitive and responsive to their customers while accelerating their time to revenue.
- **The near future potential for new, innovative applications:** The ability to create new, innovative network-aware applications can improve service providers' business operations and/or create new revenue-generating services. Information sharing across layers can have a major impact on overall network efficiency through periods of change and growth.

For service providers, the advantages of IP+Optical convergence include reduced capital expenditures (CapEx) and operational expenses (OpEx), improved capacity planning, more efficient traffic engineering, quicker access to revenue, and more environmentally friendly solutions.

Reduced CapEx and OpEx

Moving to an IP+Optical converged model leads to reductions in CapEx and OpEx. These reductions are achieved by providing a higher level of network and system integration and greater service scalability while delivering more features and capacity within a smaller footprint. With less equipment and higher levels of cross-layer system integration comes lower power consumption and corresponding reductions in building cooling, and maintenance costs.

Capacity Planning

Capacity planning also benefits from nLight technology. Research firm ACG Research examined the ramifications of nLight technology and determined that nLight significantly reduces TCO by eliminating even small overestimates of required bandwidth capacity. For example, in a design using unprotected wavelengths, nLight produces a 24 percent reduction in TCO by eliminating a 10 percent overestimation of required capacity.

The deployment of nLight on a network with protected wavelengths produces more than a 50 percent reduction in TCO by eliminating a 10 percent overestimation of required bandwidth capacity. ACG says these savings are achieved because nLight enables a just-in-time capacity provisioning process that reduces the risk of being caught with inadequate circuit capacity.

Traffic Engineering

The Cisco nLight Control Plane automates many of the manual processes that are part of a traditional architecture. When turning up a new circuit or service, for instance, the packet-layer request can include parameters such as matching path, disjoint path, specific latency, and lowest optical cost circuit.

As a result, the Cisco nLight process requires just two steps:

1. The client IP router signals a new circuit request, along with the requirements for the new circuit, to the server DWDM layer.
2. The DWDM layer signals a valid wavelength to the packet layer or signals a path error message if the request cannot be filled without optical layer network growth, and updates its internal network database.

With nLight, all network information is updated automatically between layers and maintained by the network. If a high-priority multilayer service fails within the network, Cisco nLight technology is capable of restoring the circuit without manual intervention.

Overall, nLight enables a higher degree of accuracy and flexibility in planning for and reacting to unpredictable traffic patterns, at both average and peak times. Having real-time visibility into all the bits of unused capacity across multiple services and multiple layers of the network produces savings due to improved traffic engineering. This improvement has the potential to be the greatest advancement of all.

Monetization

Instead of taking weeks or months to manually fulfill a customer's request in the complex, multilayer workflow described earlier, the service provider can more quickly and accurately provision these services with nLight, thereby accelerating service delivery. The provider can also offer new services at attractive price points enabled by the evolved network technologies such as 100G coherent, CPAK, IP+Optical convergence, nPower NPU's, next-generation ROADMs, and Cisco Quantum WAN Optimization, improving customer retention and identifying new revenue streams.

Sharing information between the optical transport and IP layers opens up entirely new opportunities in services and applications. For example, a service provider may offer a drop-down menu for customers to select the level of service they need for a particular service or application. If the service is the type that can withstand interruptions, the customer can choose a lower and less expensive level of service. In circumstances where a customer cannot afford any interruptions whatsoever, they can choose a higher and more expensive level of service. The inefficient "one size fits all" approach to multilayer service creation now has a viable and intelligent alternative.

If a financial institution needs a 100G service for a massive but short-term backup project, Cisco ONE seamlessly enables that otherwise complex resource allocation task. If a portion of the service provider's network goes down due to a fiber cut or other emergency, the IP+Optical solutions offer service protection at multiple layers, so all is not lost. This can all be accomplished automatically, without a truck roll, using the latest Cisco technology. End users will begin to identify networks that use Cisco IP+Optical solutions and will also monetize the price of flexibility and greater reliability of networks built with those premium solutions.

Green and Greener

Cisco CPAK modules are less than one-third the size of CFP modules and dissipate less than one-third the power. As a result, they can support more than ten 100G CPAK ports in the space allotted for four CFP modules, allowing the optics on a terabit-capacity line card to dissipate less than 55W of transceiver power. Compare that with today's 10G transceivers, which generate about 48W to produce 480 Gbps - less than half the capacity. The use of CMOS technology also means the CPAK is easier to manufacture, with higher yields than the CFP. Even the optical industry's next commercial solution for pluggable 100G client interfaces, the CFP2, is larger than the CPAK. The competition will actually need CFP4 to improve on the CPAK's current, generally available power and space performance. Higher capacity, less power, and greater manufacturability with less waste moves the Cisco IP+Optical solution from green to greener.

Full Solution

Many competitors would like you to think that a feature upgrade here or there will move an existing product or network forward just enough and/or just in time to adapt to future service demands. Consider the architectural capabilities of the NCS product family. Then consider where you think networking will be five or ten years from now.

Building the NCS product family was a major investment for Cisco. Each individual technology and each new platform was scrutinized heavily in the development process to see if we could “use what we had.” Cisco had many resources to draw on, but in many cases the answer was still no. Providing a high-availability multilayer solution that could migrate networks to an EPN required a fresh approach. No matter where you start or where you finish your network migration, consider whether your partner has the technical expertise to deliver on:

- IP/MPLS, including virtualization and scalability
- Carrier Ethernet feature parity across platforms
- OTN functionality to aggregate Layer 2/Layer 3 with legacy time-division multiplexing (TDM)
- DWDM functionality, including proven alien wavelength performance
- Multilayer control plane
- Multilayer SDN and WAN optimization
- Proven approaches to reduce CapEx and OpEx through convergence

Most equipment vendors specialize in only a few of the technologies listed above. The full list is needed to make meaningful change that will scale networks while reducing the cost per bit in both CapEx and OpEx. Going forward, even more importance will be placed on having a full solution and on how easily these multilayer technologies can work together. After all, network convergence can be different for each network operator, and flexibility to understand and achieve your business goals is paramount. Experience and breadth of solution work in Cisco’s favor here.

Conclusion

While the need to converge the IP and optical layers for CapEx and OpEx advantage is well understood, deploying 100G in any network element involves much more than flipping a switch. There are many choices in determining how to proceed with an IP+Optical network convergence project. It’s a multiyear journey, and one in which operators should carefully consider how they want to proceed, as each network operator has its own value proposition and unique customer base and geography. Cisco can deliver an evolved programmable network that will give service providers all the tools they need to reach the customers of today with the rich service mix of tomorrow.




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)