

Routed Optical Networking

Understanding the value of modern end-to-end IP/MPLS
transport networks

by Emerson Moura, Distinguished Architect

Contents

Introduction	3
Are people skills a real problem for routed optical networking adoption?	4
Beyond more cost-effective DWDM links—efficient wavelength usage	6
Full services convergence—only possible today with IP/MPLS	8
Optimal traffic forwarding to improve customer experience and application performance	9
Easier-to-automate day-to-day operations	10
IP/MPLS silicon is at another level—scale, power, and cost efficiencies	11
How can IP/MPLS meet transport service-level agreements?	11
Conclusion	13
Learn more:	14

Introduction

Routed optical networking embraces mass simplification of the end-to-end network infrastructure to achieve cost savings, operational agility, and improved network efficiency. Simplification is accomplished by de-layering the network switching stack, removing redundancy in both software and hardware components, fully converging transport, L2 and L3 virtual private networks (VPNs) plus internet services at the internet protocol/multiprotocol label switching (IP/MPLS) layer, and integrating technologies at the routing layer whenever possible. The goal is to do more with less—more services, richer service-level agreements (SLAs), and higher service availability with less traffic switching layers, fewer software upgrades, and less network protocols to run. As a result, routed optical networking has been demonstrated to provide up to 46% savings in total cost of ownership (TCO)¹ and independent network operator modeling has demonstrated over 70% TCO savings, with over 80% savings in space and power.

"Thanks to game-changing innovations that span across silicon, optics, and systems, complex layers can finally converge into a simpler and more scalable architecture."

STAFFAN GÖJERYD, CEO, ARELION

In this context, IP and optical integration is a fundamental part of routed optical networking. Leveraging the latest industry innovations in optics, routed optical networking replaces costly transponders, which are otherwise needed to connect the IP/MPLS layer to the dense wavelength-division multiplexing (DWDM) network, with digital coherent optics (DCO) pluggable transceivers. These are available today at speeds of up to [400 Gbps using QSFP-DD transceivers](#) that can be deployed directly in standard 400GE ports in the routers, and potentially even faster in the future (e.g., 800 Gbps).

The use of these pluggable DCO transceivers to replace DWDM transponders is fundamental to routed optical networking and a key source of the dramatic economic savings of the solution, and the fact that 400-Gbps DCO has been standardized and embraced by the optics industry (e.g., through Open ROADM and OpenZR+) adds to the vast list of benefits provided by the technology. However, the value of the solution goes well beyond that. Arguably, the biggest benefit of routed optical networking is the full services convergence and the network efficiencies enabled by packet switching with the intelligence, efficiency, and operational simplicity of a modern, end-to-end routed IP/MPLS network.

This paper discusses in greater detail those additional network efficiency gains enabled by the optimized routing layer and other relevant aspects of routed optical networking that aim to bridge the gap between packet and transport networks. But before we talk about technology, let's address another common topic related to any network transformation.

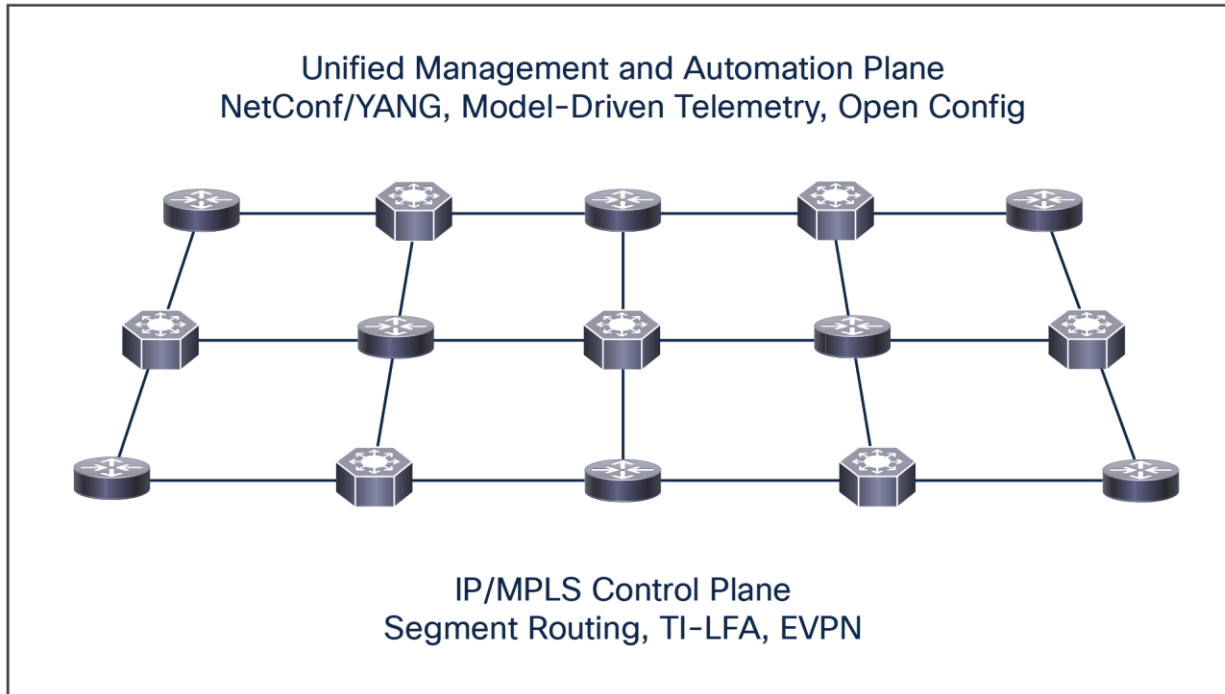


Figure 1. Routed optical networking architecture: A simplified end-to-end, fully converged IP/MPLS network tightly integrated with the DWDM layer built for cost efficiency and business agility.

Are people skills a real problem for routed optical networking adoption?

The technology components of routed optical networking are full of innovations in silicon, optics, network protocols, and software tools. That raises questions about the technical skills required by operators who will run day-to-day network tasks.

Let's remember that this isn't the first time the traditional transport industry has undergone a major transformation. During the mid-2000s, legacy SONET/SDH networks started to become obsolete with the emergence of Ethernet services that demanded packet-based networks to support the full spectrum of services required by the market (e.g., MEF E-Line, E-LAN, and E-Tree services). Despite industry attempts to extend the life of SONET/SDH systems by integrating Ethernet services, eventually packet-based carrier Ethernet solutions became ubiquitous given their superior economic benefits and technical advantages. At the time, the industry also went through a major debate about people skills, as carrier Ethernet was very different from legacy time-division multiplexing (TDM) technologies. Eventually the business drivers pushed the industry to find a solution—transport network operators became fluent in carrier Ethernet technologies at the same time as traditional IP network engineers. Most network operators embraced the new technologies introduced to deliver Ethernet services over IP/MPLS, and today few in the industry would argue there's a lack of technical expertise to operate IP/MPLS-based carrier Ethernet networks in any major equipment vendor or service provider.

With routed optical networking, we're at the start of a similar transition. The new technology components push network engineers from both transport and IP organizations out of their comfort zones. However, there's a major difference this time. With very few exceptions—like private line emulation (PLE) that's truly a new industry development—most components of routed optical networking are optimizations of well-established technologies. DCO pluggable transceivers are compact implementations of DWDM transponders. DWDM line systems are based on the same reconfigurable optical add-drop multiplexer (ROADM) technologies currently used, just with a bigger focus on operational simplicity, automation, and provisioning of "alien" wavelengths (i.e., DWDM signals generated directly by a client of the network). IP/MPLS networks can be simplified by adopting segment routing when possible, while reuse of classical routing paradigms based on IP/MPLS with label distribution protocol (LDP) or resource reservation protocol (RSVP) is also an option, even though it provides less optimal results. Both approaches are well known and widely deployed in the industry. Lastly, at the network operations level, automation can replace legacy EMS/NMS systems, and this can happen over time too, allowing network operators to develop the required skills.

The most fundamental remaining question is: what new skills should you learn or develop for routed optical networking? In other words, if you're a transport engineer, should you learn IP/MPLS? Or if you're an IP/MPLS engineer, should you learn DWDM? In both cases the answer is a "soft" yes. That's good news because learning a new skill in the technology industry should be perceived as a competitive advantage, and the "soft" yes answer means that you need to learn enough to be conversational, but you don't need to become a full expert in every component of routed optical networking to be successful.

The role of the expert—be it an optical expert, IP/MPLS expert, or management and automation expert—will remain crucial. Companies looking at building organizations that will deploy routed optical networking need to ensure those experts are part of a cross-functional team. That multidisciplinary organization will help develop the industry in two fundamental ways:

- Cross-pollination of knowledge about the technology components, driven by team dynamics and collective problem solving
- Breaking technology silos and creating an environment that allows for better alignment between experts, resulting in more efficient network designs and operations that can deliver significantly better economics

The cross-pollination and cross-functional collaboration can drive further simplification of the overall network architecture. Once the network is designed around the key principle that IP services should rely on IP layer mechanisms to provide survivability, fundamental transport technology choices can be revisited. For instance, does the DWDM network still require relatively complex and costly add/drop elements that implement colorless, directionless, and contentionless functions that were crucial for DWDM layer restoration? Or should the DWDM elements be simplified when the IP layer is self-sufficient for service protection and network survivability across all failure conditions?

In summary, once again the industry is going through an inevitable transformation, and as we've seen in the past with carrier Ethernet, those who step up and learn new skills will have the biggest professional opportunities. This doesn't mean giving up on your current expertise, as it's going to remain important and relevant, but this time it's about incorporating new skills. It's like moving from an individual sport to a team sport where you have a specific role as part of a team and the wins are collective. Network professionals will need to collaborate as part of a cross-functional team that includes packet, transport, management, and automation expertise. With regards to technical skills, this is what routed optical networking is all about: it's a team sport.

With that debate out of the way, let's now talk about the benefits of IP/MPLS in routed optical networking.

Beyond more cost-effective DWDM links—efficient wavelength usage

Saving on DWDM transponders, which represents the bulk of a DWDM system cost, is important to reduce overall network costs, and routed optical networking achieves that by integrating the transponder functionality in the routers using DCO pluggable transceivers.¹

However, the benefits enabled by routed optical networking go well beyond DCOs. A modern, SDN-based end-to-end IP/MPLS network deployed as part of routed optical networking also maximizes the use of such expensive and scarce DWDM wavelengths.

First, routed optical networking proposes a router-to-router connectivity model that in many cases significantly reduces the optical distance for the DWDM wavelengths. With shorter optical distances, higher bit rates per wavelength are possible, potentially leading to more 400-Gbps DCO running at full speed in the network.

Besides the higher bit rate enabled by shorter optical distances, higher DWDM wavelength utilization is achieved thanks to three key capabilities of the IP/MPLS network:

1)**Statistical multiplexing:** Traditional "optical bypass" or "hollow-core" network architectures dedicate a DWDM wavelength to directly connect any pair of routers with traffic interest, often resulting in underutilized wavelengths. Conversely, routed optical networking looks for more opportunities to fill wavelengths, and so it's able to use the statistical nature of IP to efficiently multiplex traffic across the routed network. Wavelengths are shared by multiple routers along a path. Thanks to the router-to-router connectivity model, as a result, less wavelengths may be required. Considering that most of the traffic in today's networks is IP, which is bursty by nature, statistical multiplexing is not only desired, it's required to scale the network, keep costs under even tighter control, and reduce overall network costs per bit.

2)**Traffic engineering:** IP networks forward traffic using either shortest paths toward a given destination as calculated by routing protocols (e.g., IS-IS, OSPF) or custom paths. These custom paths are commonly referred to as traffic engineered paths and are used when specific business logic or service requirements demand the traffic takes specific paths across the network. Routed optical networking supports traffic engineering using a classical RSVP-TE based IP/MPLS network. Moreover, advanced traffic engineering can be deployed using segment routing² technology, which is SDN enabled and allows for a centralized path computation engine to distribute traffic across the network to maximize the utilization of available bandwidth while meeting business and application requirements. This is achieved without overwhelming the network or its operations given the simplified protocol stack provided by segment routing and the policy-driven approach applied to service provisioning.

3)**IP layer protection:** Traditional network designs typically deploy protection at the transport layer, even though the IP/MPLS layer is generally designed with traffic protection in mind. When additional protection is added over the already-protected IP/MPLS layer (i.e., at the DWDM layer), it requires the use of more dedicated wavelengths that will be standing by idle in the network until a failure happens. An alternative approach is to use wavelength restoration, which shares spare wavelengths among multiple active ones. However, given the slow convergence of DWDM networks (multiple seconds or even minutes), wavelength restoration is often used in combination with protected wavelengths, meaning that instead of saving wavelengths, restoration requires even more wavelengths that are standing by to cover for failure scenarios. This redundant protection adds up to other inefficiencies of traditional network designs like low average interface and wavelength utilization and results in a "snowball" effect, where relatively low traffic demands, in the order of a few 10s of Gbps, end up consuming multiple 100 Gbps of capacity at the transport layer. Conversely, an optimized IP/MPLS network design that takes both link and node failures into consideration during the network design phase provides enough redundancy without any need for additional transport

resources. Add to that the ability of the IP/MPLS network to switch in less than 50 ms (the transport golden standard) using modern technologies like topology-independent loop-free alternate (TI- LFA), and the IP/MPLS layer can be self-sufficient in terms of traffic protection and network availability. No additional wavelengths may be required for traffic protection or restoration, reducing the number of wavelengths consumed for a given traffic demand.

As a result of maximizing available wavelengths by using these three techniques, routed optical networking may also delay DWDM network upgrades. In some cases, it might even delay the need to deploy additional fiber pairs or L-Band DWDM systems, as the number of wavelengths required for a given traffic demand and related SLAs may be significantly lower.

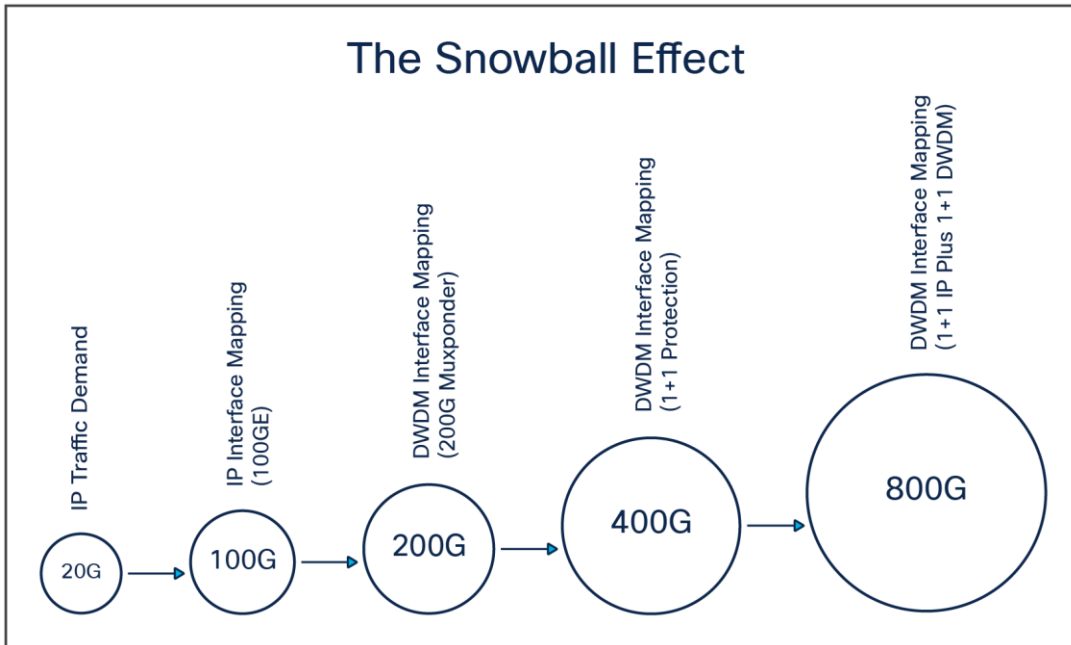


Figure 2.

Remember that routed optical networking with 400-Gbps DCO pluggables works over any 75-GHz-capable DWDM architecture, supporting both fixed optical add/drop multiplexers (FOADMs), mux-demux or ROADMs, with the latter allowing use of tactical bypass of wavelengths as an additional tool for network optimization whenever it brings economical value, without sacrificing other benefits of the simplified architecture.

Full services convergence—only possible today with IP/MPLS

IP/MPLS networks have proven their multiservice capabilities, and for decades, network operators across the globe deployed networks to deliver L3, L2, and L1 services based on this technology. While L3 and L2 services are well covered by a variety of implementation options, with EVPN being the latest and most integrated Ethernet service control plane technology for IP/MPLS networks today, L1 has been limited to lower speed services with PDH and SONET/SDH services emulated over the packet network using CEOPS and SAToP technologies. Key L1 services had not been delivered by IP/MPLS, including transparent services demanding higher speeds (gigabits/sec and beyond), such as Ethernet private lines with deterministic paths and bandwidth guarantees, and higher levels of transparency for control protocols, leased lines for common storage protocols, and optical transport networking (OTN) leased lines. This gap in the technology has been used as the main reason why communication service providers (CSPs) still deploy electrical OTN switching and DWDM networks with wavelength and subwavelength services.




Figure 3.

Routed Optical Networking innovations: Digital Coherent Optics (DCO) and Private Line Emulation

To address that gap and achieve full services convergence, routed optical networking introduces private line emulation (PLE) services over the IP/MPLS network. PLE is being defined as a standards-based solution within the IETF3 and provides bit-level transparency for transport-like L1 services with a variety of client protocols, including Ethernet, OTN, and fiber channel at gigabits/second speeds.

While PLE provides the client encapsulation to transport those bit-stream services using standard pseudowire techniques, the control plane for the emulated services, which controls the behavior of the service in the network including path definition and protection, is also based on modern industry innovations. Circuit-style segment routing (CS-SR) is an SDN-based solution that allows PLE to mimic the behavior of traditional transport technologies over an IP/MPLS network. This means that PLE services with CS-SR are co-routed (upstream and downstream paths are congruent across the network) and protected using common transport schemes (1+1, 1+R, and 1+1+R). Since the bit stream service is transparent, the bandwidth is consumed all the

time (there's always some bits on the wire even when information is not carried). This combination of features allows PLE to have a predictable behavior across the IP/MPLS network to meet traditional transport operational and SLA requirements.

Finally, it goes without saying that no technology available today except IP/MPLS can provide an optimal transport solution to any L1, L2, or L3 service in terms of scale, efficiency, and flexibility.

Optimal traffic forwarding to improve customer experience and application performance

One of the key benefits of using an IP network is the ability to route traffic based on different routing algorithms or custom, traffic-engineered paths. As the industry evolves and applications, data collection/processing, and content are pushed toward the edge of the network, closer to the end user, this unique capability of IP networks becomes crucial and advantageous.

Forwarding traffic based on routing versus circuit switching provides a set of network efficiency and business benefits to CSPs, including but not limited to:

- **Traffic off-loading and improved network efficiency:** The IP network can route traffic to a given destination using optimal paths, avoiding hair-pinning or longer-than-necessary routes in the network. For applications that support a distributed architecture (e.g., content delivery networks, where an application or content is available in multiple locations), the IP network can forward the traffic to the optimum location as well as load balance traffic using equal-cost multipathing (ECMP). This not only improves application/traffic performance, but also reduces network backhaul bandwidth requirements.
- **Improved application performance:** While traffic off-loading helps optimize network utilization, some applications will have their performance influenced by specific network behaviors or traffic forwarding logic beyond the lowest cost IP path. A common example is an application that is sensitive to latency, such as virtual reality or online gaming. Latency in modern networks is more a result of fiber distances than traffic switching, independently of the switching technology. To minimize traffic latency, a proper measurement of the end-to-end latency, including that imposed by the fiber, is required. IP/MPLS networks based on segment routing can run those measurements in real-time and in-band, while implementing routing algorithms that take them into account using flex-algo technology, which allows for multiple, concurrent routing algorithms running in the same router, creating different traffic forwarding logic based on different views of the same physical network topology. For instance, a routing algorithm can be defined to provide lowest network latency, another for lowest IP cost, and so on. Another common example is for sensitive traffic that must go through secure physical links. Once again this is supported by segment routing with flex-algo. The network is also capable of automatically steering traffic to the proper routes based on the predefined segment routing traffic policy by classifying the packets and marking (or coloring) them so the router can steer it onto the proper network path.

-
- **Flexible application placement:** There was a time when a given application was statically placed in a well-known location in the network. That's not the case anymore. As applications are modernized, become cloud centric, and adopt more distributed architectures, their network placement becomes more dynamic and defined by the business needs—scale, user experience, business partnerships, and others. IP networks on one side allow the optimal traffic forwarding to improve application experience. Conversely, IP networks also allow applications to be deployed anywhere in the network at any time. That flexibility means a business can evolve its application environments without having to redesign the network to optimally connect them to users and other businesses.
 - **Flexible support for evolved peering strategies:** Internet peering infrastructure has evolved significantly over the years. While in the past it was more centralized and limited to fewer locations, the modern internet has a highly distributed peering architecture to meet the business needs of CSPs and public cloud providers, including bandwidth optimization, improved user experience, and reduced transit costs. As peering continues to evolve and get closer to the end user, end-to-end IP networks again provide a flexible infrastructure that's future proof to deal with any changes.

Easier-to-automate day-to-day operations

Web and cloud operators have successfully automated their massive-scale data centers and transport networks, but automation requires discipline in keeping the network simple and consistent in using standard APIs that are built for modern machine-to-machine communications. This means a departure from command-line interfaces and legacy protocols like SYSLOG and SNMP, to name just a couple, and full adoption of model-driven, open APIs for automated provisioning, configuration, and telemetry data gathering.

Routed optical networking provides the necessary simplification to achieve full network automation. There are only two technology layers—the IP/MPLS layer, where most of the services and advanced features are delivered, and a streamlined and simplified open DWDM layer to maximize fiber utilization, address optical impairments for long-distance applications, and support eventual demands for super-high-speed services (100 to 400 Gbps and beyond). Another simplification available with routed optical networking is that the IP/MPLS logical network topology can match the fiber topology with the router-to-router connectivity model. This congruency of topologies removes the extra work required to maintain path diversity and avoid shared risk groups across IP links that are deployed to protect each other but may end up in the same fiber, fiber ducts, or DWDM line due to errors resulting from excessive network complexity. The simplified network topology enabled by routed optical networking also simplifies network design and planning.

In addition to the simplified technology stack and network topology, routed optical networking deployed over greenfield infrastructure (it can also be deployed over brownfield and classic IP/MPLS networks) also provides open and standard APIs for network automation, which can be used to integrate to a variety of vendor-provided or open-source software tools. For instance, NETCONF/YANG interfaces and OpenConfig models are available for model-driven automation and network programmability at both the IP/MPLS and DWDM layers, and modern protocol stacks including gNMI are also available for efficient machine-to-machine communication including network telemetry. As a result, network automation becomes feasible and sustainable in the long run.

IP/MPLS silicon is at another level—scale, power, and cost efficiencies

Routing and switching silicon technology have evolved rapidly and consistently over the years. [Cisco Silicon One™](#) P100 for example provides 19.2 Tbps of full duplex capacity in a single network processor unit (NPU) and consumes a fraction of the power per Gbps of previous generations of NPUs. Just to put it in a broader perspective, while many routers using technologies like Cisco Silicon One can provide more than 12-Tbps capacity per slot, electrical OTN switches currently deployed typically provide a fraction of that capacity in a full chassis (e.g., 3.84 Tbps per chassis for a high-end OTN switch), for instance 120 Gbps per slot, which is orders of magnitude less than a high-end router.

Novel NPU technologies have also contributed to the simplification of router designs, bringing additional gains in cost efficiency of IP/MPLS networks.

As routing and switching silicon continues to evolve at a pace much faster than any other network switching technology, providing much more capacity at lower footprints and power consumption per traffic unit, any company looking at building scalable, sustainable, and future-proof networks optimized for power, density, and cost efficiency should focus their investments on IP/MPLS-based networks.

How can IP/MPLS meet transport service-level agreements?

There are fundamental technical capabilities provided by traditional transport networks that must be met to support SLAs required by private lines, commonly tied to strict services with high availability metrics.

While SLA is a broader term that involves the network as well as operational processes, technical support response times, and other contractual agreements, when we look at the required network capabilities to meet the strict high availability goals, IP/MPLS networks have never been so well equipped to fulfill them.



Figure 4.
Cisco Silicon One P100 - 19.2 Tbps router Network Processor Unit

Typical transport networks switch traffic in the event of fiber, node, or link failures under 50 milliseconds, which is the golden standard for any network technology. That switching time is easily met by multiple solutions available for IP/MPLS networks today. For over a decade, classic IP/MPLS networks have supported RSVP-TE-based fast-reroute for link and node failures. More recently, IP/MPLS-based networks started providing sub-50-ms switching even without RSVP-TE, using loop-free alternate (LFA) fast-reroute and topology-independent loop-free alternate (TI-LFA). As the name implies, TI-LFA applies to any network topology, it's very simple to configure, and it's enabled by segment routing, with no need for additional protocols or stateful core tunnels that are operationally demanding to configure and maintain and consume costly router resources.

In addition to the protection mechanisms described above, transport networks often employed restoration mechanisms at the DWDM and/or OTN layers to protect against multiple simultaneous failures. Those restoration mechanisms are costly in the way they're commonly implemented by deploying regenerators throughout the network and mathematically challenging in finding optimum placement for those regenerators. Conversely, a properly engineered IP network using TI-LFA outperforms any such restoration mechanisms with respect to the switching times. Furthermore, well-designed IP networks can provide massive fiber savings where a sizable percentage of restoration fiber routes can be eliminated through proper IP network planning or used to carry paying services instead of sitting idle most of the time.

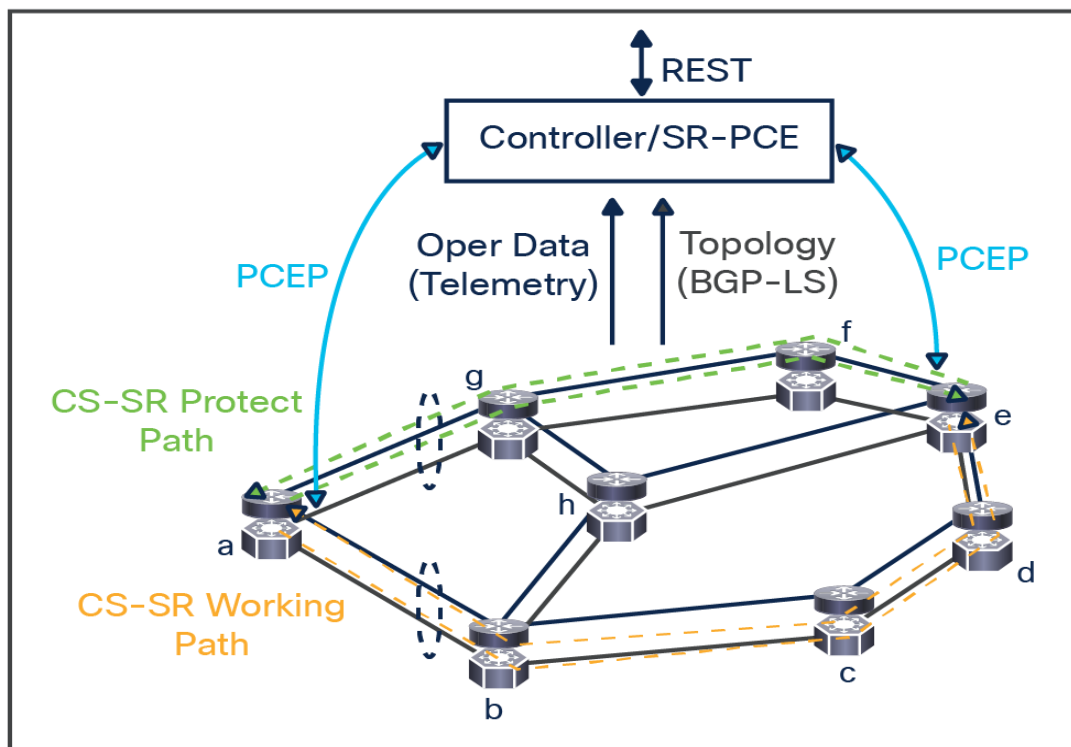


Figure 5. Circuit-Style Segment Routing - Routed Optical Networking innovation that enables transport friendly provisioning of protected services.

Another common requirement of transport networks is path predictability. Transport network operators require the ability to co-route traffic, which means map upstream and downstream traffic over a congruent, well-defined path in the network. This applies to the active/protection paths, and it's the way synchronous optical network (SONET) and synchronous digital hierarchy (SDH) networks operated with unidirectional path switched ring (UPSR) and subnetwork connection protection (SNC-P). Electrical OTN and DWDM use similar techniques.

Thanks again to segment routing, IP/MPLS networks provide those transport capabilities by leveraging the new circuit-style segment routing (CS-SR) solution. CS-SR leverages network controllers to configure congruent upstream/downstream active/protection paths with sub-50-ms switching times. When CS-SR is combined with private line emulation, which provides bit-level transparency and predictable bandwidth consumption, the network behavior is fully predictable to meet transport network operator requirements.

Besides meeting required transport switching times, IP/MPLS networks have additional benefits in creating rich SLAs for more advanced services. Segment routing provides additional SLA metrics and advanced capabilities that are implemented through segment routing policies. Those policies can include latency limits, secure paths, and diverse paths, among others. Segment routing makes the definition of such policies, which translate to

SLAs, much simpler, is intent based, and automates traffic steering and policy enforcement by the network, which will honor those SLAs whenever it's possible, even when topology changes or failures happen in the network. Another key aspect of providing high SLAs is the ability to monitor and report the performance of the service across the network using operations, administration, and maintenance (OAM) tools. Once again, IP/MPLS networks have never been so well equipped to provide those capabilities. Besides the rich set of OAM tools available for classic IP/MPLS networks, segment routing adds additional tools for performance management (PM) and data-plane monitoring, and PLE is also adding additional functions to monitor the network performance.

In summary, IP/MPLS networks can be designed to meet strict SLAs required by transport services. In addition, segment routing-based IP/MPLS can provide unique capabilities that support richer and even more-strict SLAs, considering additional network metrics and criteria including latency, path diversity, path security, and possibly others.

Conclusion

IP and optical integration is a fundamental part of routed optical networking given the significant economic benefits provided by the current 400-Gbps digital coherent optics pluggable transceiver technology. However, routed optical networking goes far beyond that, and when deployed with modern IP/MPLS technology, it brings network optimizations and efficiencies to another level. In summary, IP/MPLS provides the following key benefits to routed optical networking:

- **More efficient use of scarce DWDM wavelengths** through statistical multiplexing and traffic aggregation leveraging direct router-to-router designs that better align with fiber topology, plus optional centralized traffic engineering.
- **Full services convergence:** IP/MPLS is the de facto multiservice technology and the only networking technology capable of supporting L1 (emulated TDM circuits), L2 (point-to-point and multipoint Ethernet), and L3 services (internet, IP VPNs). The addition of PLE technology enabled by routed optical networking expands the IP/MPLS capabilities further to support bit-transparent services for high-speed circuits, including OTN, Ethernet, and storage protocols as clients.
- **Optimal traffic forwarding**, improving network efficiencies, application performance, and user experience while providing a flexible, future-proof infrastructure that can deal with any changes in application placement or internet peering strategies.
- **Simpler network architecture**, with higher focus on delayering the network using an IP/MPLS-centric architecture. IP/MPLS technologies are fully capable of meeting the most stringent service requirements including network resiliency and advanced SLAs. It's also truly multiservice. As a result, the network can be designed without overlapping functions like protection at multiple switching layers, which simplifies the requirements to the DWDM network and results in a network with fewer touchpoints that's simpler to operate and automate.

Learn more:

Take a few minutes to learn more about [Cisco® Routed Optical Networking](#).

Citations:

1. ACG Research – The Economic Benefits of IP Transport at 400G

www.cisco.com/c/dam/en/us/solutions/collateral/service-provider/routed-optical-networking/white-paper-sp-acg-400g-ip-transport.pdf

2. Segment routing (SR): www.segment-routing.net

3. Private line emulation (PLE)

<https://datatracker.ietf.org/doc/html/draft-schmutzer-bess-ple>

<https://datatracker.ietf.org/doc/html/draft-schmutzer-bess-ple-vpws-signalling>

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 <https://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: <https://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)