CAPITALIZING ON IP/OPTICAL CONTROL INTEGRATION

QUANTIFYING TCO SAVINGS OF MULTILAYER PROTECTION AND RESTORATION LEVERAGING A GMPLS CONTROL PLANE

FINANCIAL WHITE PAPER

IP/optical integration removes operational and technological barriers that inflate overhead cost and impede routing and transport convergence. With IP/optical integration, service providers will be able to scale network capacity more economically, conduct multilayer network operations more efficiently, and effortlessly monetize network assets.

- *GMPLS protection and dynamic restoration* features leverage the flexibility of the OTN and ROADM technology to efficiently protect services and improve network utilization.
- *GMPLS UNI integration* extends these dynamic transport control capabilities to the routing layer, establishing a *unified multilayer control plane* to efficiently coordinate cross-layer operations.

This financial network modeling study from Alcatel-Lucent Bell Labs compares traditional IP/MPLS and optical transport layer protection and restoration approaches with a balanced multilayer resiliency strategy based on new techniques including GMPLS and GMPLS UNI.

TABLE OF CONTENTS

EXECUTIVE SUMMARY

Service providers are always looking for ways to run their networks hotter in order to maximize their returns on network investments. Conventional transport layer protection mechanisms often reserve as much as 50% of provisioned capacity (referred to as 1 + 1 redundancy) to recover from failures. Besides keeping half of the available transport capacity in reserve, backup protection paths are pre-provisioned — unable to take the actual failure location into account for rerouting traffic and unable to recover from multiple failures. IP routing topologies typically have high degrees of connectivity with many alternate paths, but 1 + 1 optical layer protection only supports the use of a single backup path.

While traditional 1 + 1 optical network protection is costly, the alternative approach of only leveraging MPLS-based protection and restoration mechanisms at the routing layer over unprotected and physically disjoint optical transport paths is equally inefficient, even though these inefficiencies are perhaps less apparent at first glance. Nevertheless, in many networks this is the Present Mode of Operation (PMO) for lack of better alternatives.

The good news is that state-of-the-art Optical Transport Network (OTN) and Reconfigurable Optical Add/Drop Multiplexer (ROADM) transport technologies do provide a better alternative by introducing an intelligent control plane that can provide the right amount of failure resiliency and availability for a given service class, with a more cost-effective utilization of networks resources. These protection capabilities leverage the Generalized Multiprotocol Label Switching (GMPLS - RFC 3945) architecture. GMPLS adopts key concepts from the MPLS control plane used in IP routing with functional enhancements to support multilayer transport networks.

GMPLS enables the transport network to dynamically route or reroute traffic around failures or on to optimal paths based on network utilization constraints and/or service level agreements (SLAs). The GMPLS user-network interface (UNI) lets routers dynamically signal transport paths with support of various service protection options. Dynamic restoration capabilities enable efficient sharing of backup resources by moving from a 1 + 1 to a shared alternate path protection model that is also able to recover from multiple failures. As a result a considerable amount of reserved protection resources are freed up, with the remaining being used for revenue-generating traffic with strict SLA criteria. This leaves more capacity for less demanding services and lowers delivery costs for more demanding, mission-critical services.

With the right architecture, GMPLS-based transport layer recovery mechanisms can be applied in combination with recovery mechanisms in the IP/MPLS routing layer to offer and implement differentiated availability SLAs for different classes of service. Differentiated service availability requirements can be subsequently mapped on an appropriate multilayer traffic protection and restoration strategy in order to balance availability, redundancy and resource utilization for the best returns on network investments.

Alcatel-Lucent Bell Labs has quantified the operational cost advantages of levering GMPLS-based transport layer recovery in addition to IP/optical control integration based on GMPLS UNIs for the purpose of implementing a differentiated multilayer protection and restoration strategy, as opposed to traditional approaches solely based on MPLS protection and restoration.

The network total cost of ownership (TCO) study presented in this paper finds that:

- Introducing GMPLS protection and restoration in the photonic switching layer *saves 40% on optical transponders* and *37% on router ports* over a 5-year study period, compared to using MPLS layer protection over an unprotected photonic transport layer.
- Also leveraging GMPLS UNI in an integral MPLS/GMPLS multilayer network protection and restoration strategy *accelerates GMPLS cost savings by 4 to 5 years*. Doing so also yields *incremental 6% cost savings on optical transponders and an improvement on router port savings from 37% to 42%* over 5 years.
- The study also confirms that the efficiency gains achieved from deploying GMPLS and GMPLS UNI integration do not compromise on average service availability requirements. All service availability requirements are fully met.

In support of the network-level TCO analysis, the cost savings and service availability of the individual MPLS- and GMPLS-based protection and restoration mechanisms used in the network TCO analysis are compared in a companion paper.

STUDY INPUTS

The network study is based on a multilayer reference network as shown in Figure 1. An MPLS control plane is deployed at the routing layer, while GMPLS provides a multilayer transport control plane to control connectivity in electronic (OTN) and photonic (ROADM) equipment. The GMPLS UNI allows routers (UNI-C or Client side) to communicate with the transport devices (UNI-N or Network Side). GMPLS can be used to create protected optical segments (UNI-N to UNI-N) between transport devices. GMPLS UNI is required to create protected end-to-end generalized label switched paths (gLSPs) between routers (UNI-C to UNI-C) through the optical network.

Figure 1. Multilayer network model with GMPLS

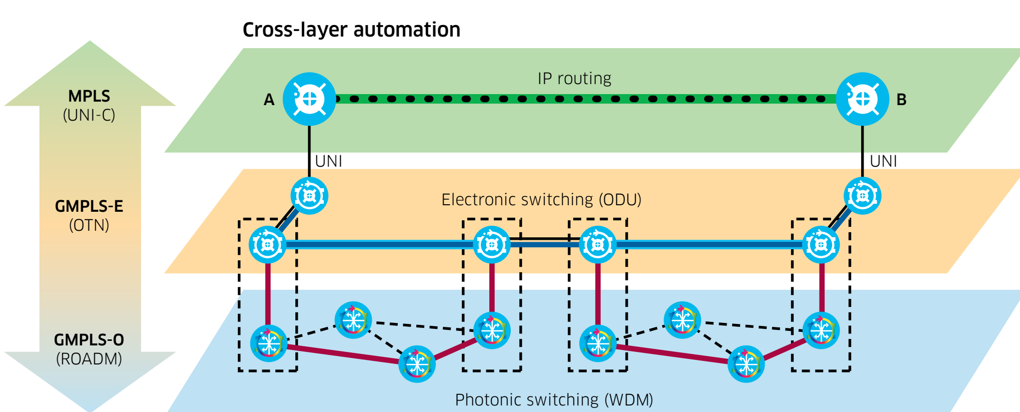

The network TCO analysis deploys a number of basic MPLS, GMPLS and GMPLS + UNI protection and recovery schemes depicted in Table 1 in order to meet service requirements in the various modes of operation.

Table 1. Overview of GMPLS recovery schemes

Optical segment	1+1 Protection (SNCP)	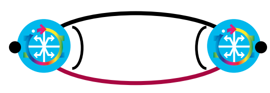	• Classical 1+1 self-healing segment protection • Fast recovery from single failure but full redundancy
	1+1 Protection and Restoration Combined (PRC)	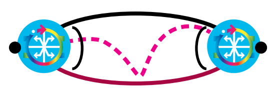	• Improved 1+1 self-healing protection by using dynamic restoration to protect against multiple failures
BOTH	Dynamic Source Based Restoration (SBR) and Guaranteed Restoration (GR)	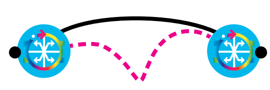	• Slightly slower failure recovery than 1+1 protection but more cost efficient sharing of protection resources
	Unprotected with SRLG constraints	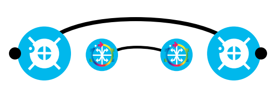	• To establish a backup path that is physically disjoint from a primary path to avoid single points of failure
	gLSP Full Rerouting (Source Based)		• End-to-end protected gLSP (UNI-C to UNI-C) with dynamic restoration to protect point-to-point traffic
End-to-end	gLSP/gLSP group protection (1:1/1:N protection)	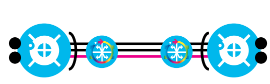	• Cost-efficient 1:N protection of a group of gLSPs
	gLSP tunnel group protection (1+N load-sharing)	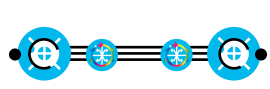	• Cost-efficient protection of gLSPs groups with use of IP interface hashing for transparent IP traffic restoration
PORT	Floating backup ports (protect router interconnect)	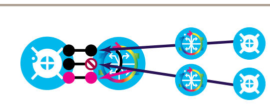	• Cost-efficient 1:N port protection of the router port interconnections to the optical transport devices

Each scheme has different datapath resource requirements that determine the relative cost efficiency and the resulting availability and recovery times that accumulate into network TCO savings. The cost optimization strategy aims to minimize the amount of resources required for protecting revenue-generating traffic, while meeting all availability requirements (including restoration time) of service traffic. The chosen protection schemes for each mode of operation are listed in Table 1 to protect a mix of IP traffic over a photonic network (that is, IP/MPLS over flexible DWDM transport without making use of an OTN grooming/switching layer).

The study compares three different network protection and restoration strategies (see Table 2):

1. Leverage MPLS to protect all IP traffic at the routing layer (PMO)
2. Leverage GMPLS to protect all IP traffic at the photonic switching layer (FMO1)
3. Leverage both MPLS and GMPLS with GMPLS UNI in an integrated multilayer protection strategy (FMO2)

Table 2. Network service protection and restoration strategies for IP over DWDM

Scenario	Expedited forwarding	Assured forwarding	Best effort traffic
PMO: MPLS routing layer protection	1+1 LSPs with fast reroute 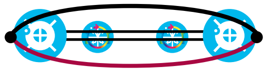	MPLS fast reroute 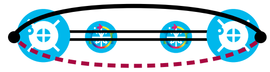	Unprotected, disjoint LSPs 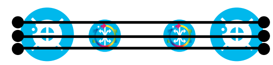
FMO1: GMPLS optical layer protection	1+1 protection and restoration 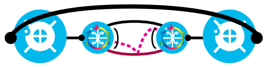	1+1 protection (SNCP) 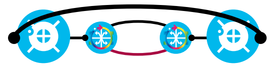	N unprotected gLSPs 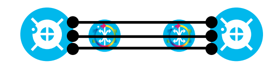
FMO2: MPLS+ GMPLS multi-layer protection	MPLS FRR over GMPLS GR 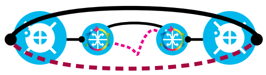	MPLS FRR over GMPLS GR 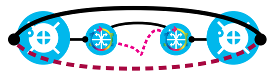	1:N tunnel group protection 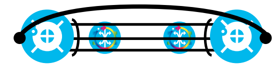

Reference network topology

The network TCO analysis is based on a reference network topology for a backbone network that consists of six core routing nodes and five optical transport nodes (Figure 2). The physical transport network topology is partially meshed in PMO and FMO2 and fully meshed in FMO1. The core routing topology is a logical mesh. Core routers (for example, Alcatel-Lucent 7950 Extensible Routing System (XRS)) are connected to optical transport nodes using either 10GE or 100GE gray optics. The transport nodes (for example, Alcatel-Lucent 1830 Photonic Service Switch (PSS)) support ROADM and optionally OTN functionality. Core routers are either single-homed (node A) or dual-homed (node F).

Figure 2. Reference network topology

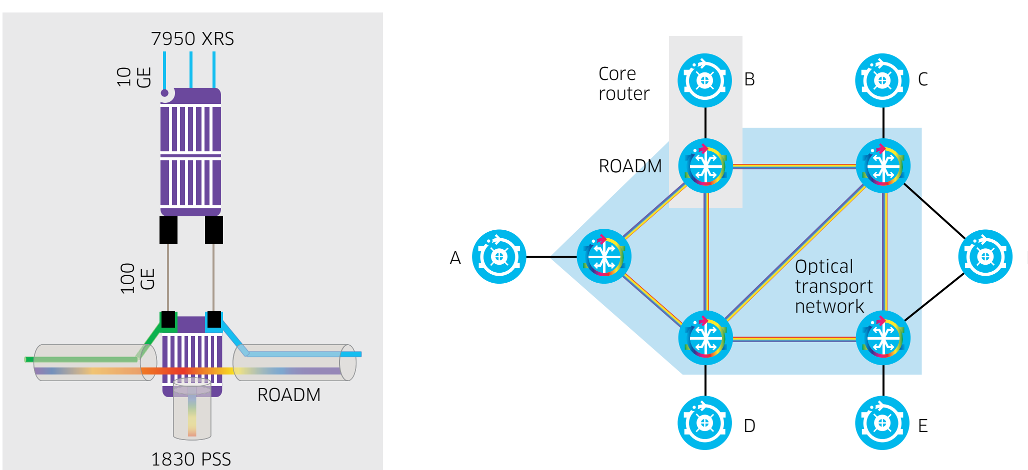

Traffic requirements

All traffic is symmetric and grows evenly at 40% annually over a 5-year study period.

- 10% is Expedited Fowarding (EF) traffic, requiring 50 milliseconds restoration with resiliency against multiple failures.
- 30% is Assured Forwarding (AF) traffic, requiring restoration <500 milliseconds and resiliency against a single failure.
- 60% is loadbalanced Best Effort (BE) traffic, and able to take a single failure without degradation.

The analysis applied a traffic matrix as shown in Table 3 to model cost evolution of the various modes of operation.

Table 3. Traffic evolution from Year 1 and Year 5 in 100G units

	EF	A	B	C	D	E	F	AF	A	B	C	D	E	F	BF	A	B	C	D	E	F
Year 1	A		0.2	0.2	0.2	0.2	0.4	A		0.4	0.4	0.4	0.4	0.8	A		1.2	1.2	1.2	1.2	2.0
	B			0.2	0.2	0.2	0.2	B			0.2	0.2	0.2	0.6	B			0.8	0.8	0.8	1.4
	C				0.2	0.2	0.2	C				0.2	0.2	0.6	C				0.8	0.8	1.4
	D					0.2	0.2	D					0.2	0.6	D					0.8	1.4
	E						0.2	E						0.6	E						1.4
	F							F							F						
	EF	A	B	C	D	E	F	AF	A	B	C	D	E	F	BF	A	B	C	D	E	F
Year 5	A		0.8	0.8	0.8	0.8	1.5	A		1.5	1.5	1.5	1.5	3.1	A		4.6	4.6	4.6	4.6	7.7
	B			0.8	0.8	0.8	0.8	B			0.8	0.8	0.8	2.3	B			3.1	3.1	3.1	5.4
	C				0.8	0.8	0.8	C				0.8	0.8	2.3	C				3.1	3.1	5.4
	D					0.8	0.8	D					0.8	2.3	D					3.1	5.4
	E						0.8	E						2.3	E						5.4
	F							F							F						

Expedited Forwarding traffic

EF traffic is typically "lifeline traffic" such as VoIP, which is very sensitive to delay, jitter, packet loss and outages. Protecting and restoring EF traffic has the highest priority.

The PMO (IP-centric protection) strategy applies 1 + 1 redundant label switched paths (LSPs) with MPLS Fast Reroute (FRR) over unprotected but physically disjoint optical transport links. Service traffic is protected end-to-end against multiple failures with very fast restoration times below 50 milliseconds.

The FMO1 (GMPLS optical layer protection) strategy applies unprotected LSPs over optical segments with GMPLS 1 + 1 protection and restoration combined. As in the PMO, all EF traffic is protected against multiple failures with a restoration time below 50 milliseconds. The difference with PMO is that service protection and restoration is transparent to the IP layer and done at an aggregate level, which means that losing a transport link will not impact the LSPs it carries.

The FMO2 (multilayer protection) strategy applies a combination of MPLS FRR over optical segments with GMPLS Guaranteed Restoration to protect traffic against multiple failures with rapid protection switching within 50 milliseconds. While the restoration time at the optical layer can be up to 10 seconds for photonic layer segments, MPLS FRR will be able to rapidly restore traffic over alternate optical segments while the primary segment is being restored. Optical segments will be able to share spare resources for restoration purposes.

Assured Forwarding traffic

AF traffic is often mission-critical, high-revenue data traffic that requires reliable transport but can compensate for limited packet loss, for instance through retransmission by the Transmission Control Protocol (TCP).

The PMO strategy applies non-redundant LSPs with MPLS FRR over unprotected but physically disjoint optical transport links. Spare resources are provisioned to protect service traffic end-to-end against at least a single failure with very fast restoration times below 50 milliseconds, which would avoid traffic being dropped during the protection switching.

The optical layer protection strategy applies unprotected LSPs over optical segments at the photonic layer with GMPLS 1 + 1 optical segment protection to meet the 300 millisecond restoration requirement and provide resiliency against single failures in the transport layer (optical segments at the OTN layer can use GMPLS Guaranteed Restoration instead). The difference with PMO is that service protection and restoration is transparent to the IP layer and done at an aggregate level, which means that losing a transport link leads to one protection switching action, while in the case of PMO there can potentially be multiple LSPs that need to be restored.

The multilayer protection strategy for AF traffic can be identical to the protection scheme used for EF traffic; it is a combination of MPLS FRR over optical segments with GMPLS Guaranteed Restoration to protect traffic against multiple failures with a rapid protection switching within 50 milliseconds. During the FRR restoration, full bandwidth recovery is guaranteed for EF traffic while limited packet loss may occur for AF traffic. MPLS FRR is necessary to meet restoration requirements for IP transport when using photonic layer restoration. If the transport service is switched through the OTN layer, traffic can be restored within 300 milliseconds, which would allow using unprotected LSPs and would also avoid possible racing conditions between optical and IP layer restoration mechanisms. Using optical segments with GMPLS Guaranteed Restoration for both EF and AF traffic classes eases traffic engineering and leads to a more efficient pooling of capacity for high revenue services.

Best Effort traffic

BE traffic is typically high volume Internet traffic with low revenue per bit and the most relaxed availability requirements. Nevertheless, service providers want to prevent long or frequent outages.

The PMO typically uses N + 1 unprotected, physically disjoint LSPs in an equal-cost multi-path (ECMP) load sharing model. This scheme protects against capacity degradation when a single LSP or optical link segment fails. FMO1 essentially uses the same model as PMO, N + 1 unprotected gLSPs carrying N + 1 unprotected LSPs.

FMO2 uses an unprotected LSP with IP interface port hashing and 1:N gLSP tunnel group protection. This scheme protects against a single link failure in a manner that is transparent to the IP layer.

NETWORK TCO RESULTS

The first set of results shown in Figure 3 of the network TCO study compares the resource consumption to accommodate the traffic evolution over the 5-year study period. The modeling results demonstrate significant savings on router ports in the scenarios deploying GMPLS optical layer protection and restoration compared to a PMO that only uses MPLS layer protection and restoration.

Figure 3. Comparing resource requirements across modes of operation

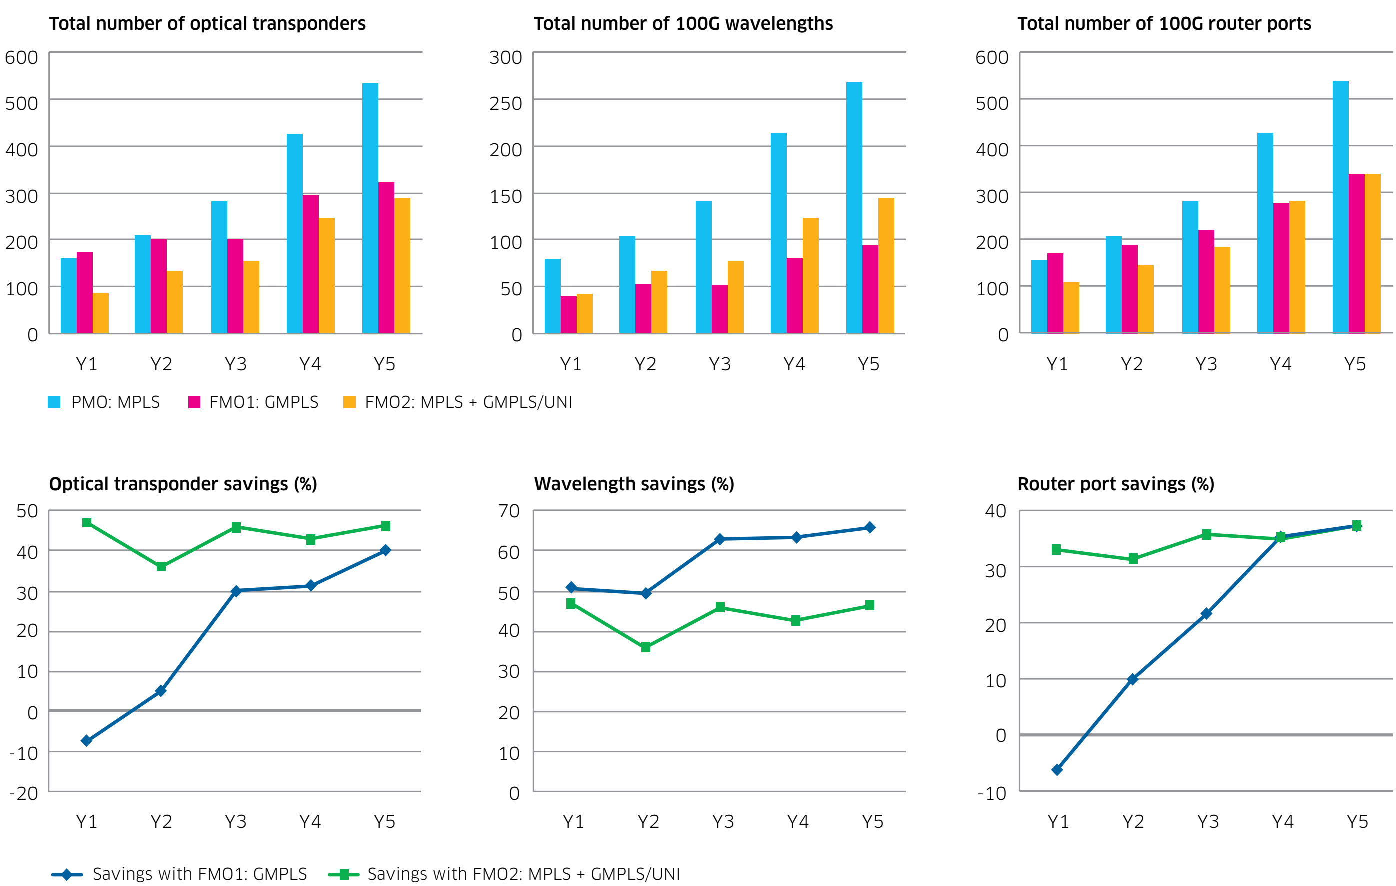

Optical transponders and wavelengths

Optical transport costs are mostly determined by the number of optical transponders required and wavelength consumption. Optical transponder count tracks closely to router port requirements, and are by far the most expensive component in the optical transport path, while wavelength consumption can impact the scaling requirements of intermediate ROADM systems that are switching the wavelengths. The results in Figure 3 (left two diagrams) indicate that FMO2 requires *46% fewer optical transponders* over the 5-year period than the PMO, and even *10% fewer than FMO1*. Again, the FMO2 shows significant cost savings over PMO in the initial years, with 47% savings over PMO and 51% savings over FM01. Both FMO1 and FMO2 consume significantly fewer wavelengths than the PMO (middle two diagrams).

In the initial years, the transport layer build-out is driven by 100GE connectivity requirements for the IP link topology and many wavelengths will be lightly loaded. FMO1 starts out with a larger cost because its connectivity requirements are higher than PMO and FMO2 due to the need for 1 + 1 link redundancy to protect EF and AF traffic, which results in a full mesh. The PMO and FMO2 link topologies on the other hand are only partially meshed because MPLS FRR and GMPLS Guaranteed Restoration can dynamically create detours around link failures. As traffic grows and wavelengths fill up, the incremental network build-out is primarily driven by capacity growth and FMO1 catches up in cost over the PMO due to its greater efficiency. Adding OTN layer grooming may provide additional cost savings to FMO1 and FMO2 in the initial years of build-out because it allows capacity provisioning with more granular transport pipes.

Router port savings

Over the 5-year period, the FMO1 and FMO2 are virtually tied in the number of router ports required, both consuming 37% fewer 100GE ports than the PMO. FMO1 and FMO2 can deploy transport layer shortcuts and build direct adjacencies between routers, which reduces the number of router hops in the data path and consequently the number of router ports and optical transponders required. However the FMO2, using multilayer protection and restoration, is far more cost efficient in the initial build-out years of the network, with 33% savings of router ports in Year 1 compared to PMO and 37% savings over FM01. The cost savings are exclusively obtained from the way that EF and AF traffic is being carried in the various modes of operation, because BE traffic is unprotected in each mode of operation (but has 1 + N redundancy).

Figure 4 gives the breakdown in router port requirements per topology node. In the case of deploying the Alcatel-Lucent 7950 XRS core router, all node requirements in Year 1 can be met in a single chassis (offering 80 ports of 100GE). However, in year 5 the PMO would require a dual chassis 7950 XRS-40 configuration to accommodate port requirements for nodes B, D and E, while FMO1 and FMO2 can still grow capacity needs for all nodes in a single XRS-20 chassis.

Figure 4. Router port requirements per node

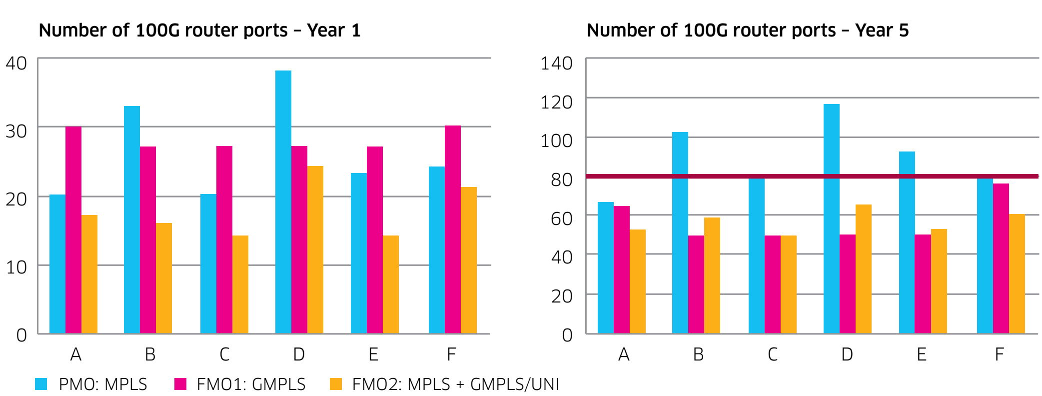

The results in Figure 4 take into account the need for additional redundancy to protect against failures of router interfaces and their connection to optical transport devices. FMO2 can benefit from a further cost optimization by deploying floating backup port protection.

With floating backup port protection applied to FMO1 and FMO2 (which requires GMPLS UNI integration), the router port savings of FMO2 increase by an additional 7% over the PMO (Figure 5, right side) over the 5-year study period.

Figure 5. Router port savings with floating backup port protection

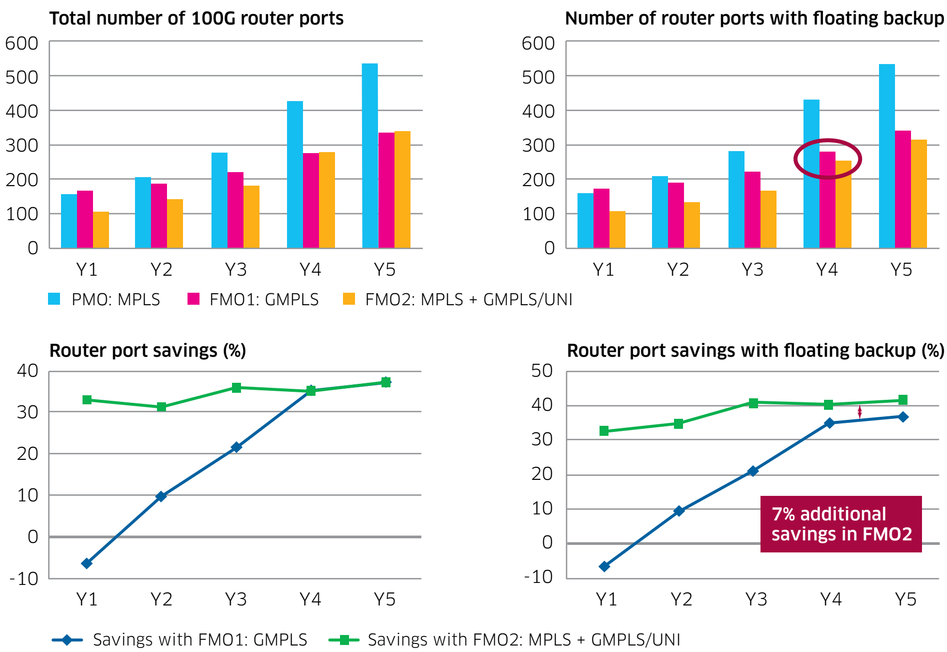

Service availability comparison

The average service availability calculations for each traffic category in each mode of operation verify that the various network protection and restoration schemes do not trade off a lower cost against reduced service availability.

Figure 6. Service availability comparison

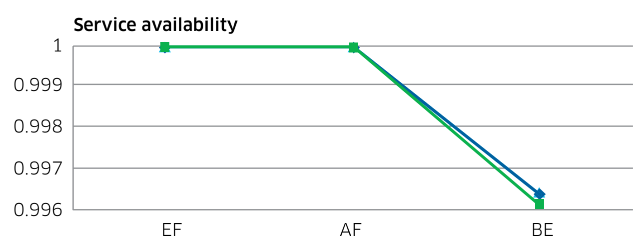

	EF	AF	BE
PMO: MPLS	0.99994	0.99994	0.996432
FMO1: GMPLS	0.99994	0.99994	0.9964297
FMO2: MPLS + GMPLS/UNI	0.99995	0.99995	0.9962201

The assumptions made for the availability calculations are:

- Link availability: 0.998 for all fiber links
- Failure rate: 1 fiber cut/year
- Mean Time to Repair (MTTR): 12 hours

The results are shown in Figure 6 and indicate that the average service availability is similar in each scenario. Each mode of operation performs equally well in terms of service availability, but the cost associated with achieving that level of availability differs.

FMO1 and FMO2 are more cost effective in their use of both routing and transport layer resources than the PMO, with the multilayer protection and restoration FMO2 offering the highest and quickest returns on investment.

Additional TCO savings from OTN layer grooming

The inclusion of an OTN switching layer allows for grooming and cost-effective transport of more granular transport flows (that is, *n* x 10GE). This enables service providers with lower capacity requirements to achieve a better fill rate of their fiber plant. Traffic protected at the OTN layer also benefits from much faster restoration times (in the order of 300 milliseconds) than those achieved by photonic layer restoration. This allows the use of GMPLS Guaranteed Restoration instead of 1 + 1 protection to protect AF traffic in FMO1.

While these savings have not been quantified as part of the study, the expectation is that the inclusion of the OTN layer will yield proportional and similar TCO savings for all modes of operation and all classes of service, with the exception of AF traffic in FMO1. This traffic may see a slightly better network utilization in low capacity backbones due to the option to apply Guaranteed Restoration instead of 1 + 1 protection to protect 10GE and *n* x 10GE links. These links can potentially recover from multiple failures as well. Recovery time will increase from 50 to 300 milliseconds, but this will still meet availability requirements. The study assumes the total amount of AF traffic to be 30% of total network traffic, but only a portion of this traffic (predominantly in the initial years of the build-out) will benefit from more granular traffic grooming at the electronic layer to better fill wavelengths. In later years most traffic can be switched directly at the photonic layer when there is enough volume to fill up 100G wavelengths. Using GMPLS Guaranteed Restoration at the OTN layer instead of 1 + 1 protection switching at the photonic switching layer will reduce but not remove the transport layer redundancy requirements.

As the inclusion of the OTN layer does not favor any of the compared modes of operation in a substantial manner, the conclusions on TCO differentiation drawn from the IP/MPLS over DWDM TCO analysis should apply equally to network scenarios that include an OTN switching layer. It should also be noted that including an OTN grooming layer will carry some additional cost and complexity (for example, OTN switching fabric, muxponders) that will offset some of the cost savings. Adding an OTN layer can save additional cost, but the relative savings are likely to be similar for FMO1 and FMO2. In addition, OTN transport services may be an attractive alternative to MPLS-based services to implement cost-effective end-to-end IP transport services, for instance for large enterprises or data center interconnect applications.

From a GMPLS and GMPLS UNI perspective, there is no difference between photonic and electronic layer control protocols and operations, other than the reaction times involved. GMPLS is a multilayer control protocol that applies to both OTN and DWDM layers, but the photonic layer is slower to react due to the optical mechanics involved.

CONCLUSION

An Agile Optical Network with an intelligent GMPLS control plane saves costs in both the IP routing and transport layer of carrier networks, without compromising on service availability requirements. This Bell Labs study finds cost savings in the order of 40% on router ports and optical transponders for an IP/MPLS over DWDM network with GMPLS protection and restoration in the photonic layer, compared to an approach based on IP/MPLS layer protection and restoration only. These savings are achieved through the implementation of several innovative protection and restoration schemes that have a more dynamic and economical use of network resources. By achieving the same service availability with fewer resource redundancy requirements, more network capacity is effectively applied to carrying revenue-generating traffic. This in turn reduces the cost per bit of service traffic, and enables service providers to run their networks hotter.

The already sizeable cost savings obtained from the introduction of an Agile Optical Network with a GMPLS control plane are further improved when adopting a coordinated multilayer network approach that leverages a GMPLS UNI between routing and transport layers. A multilayer mode of operation enables service providers to maximize the complementary benefits of MPLS and GMPLS protection and restoration in a synergistic multilayer networking strategy. Introducing a GMPLS UNI dramatically accelerates the cost savings achieved by a GMPLS transport layer control plane, and yields additional TCO savings by enabling floating backup port protection.

Closing the control loop between routing and transport reduces the cost and complexity of cross-domain operation support systems and procedures, and streamlines provisioning and assurance processes. A unified multilayer control plane is also required for introducing programmatic network interfaces for SDN control integration and instrumentation.

GMPLS is an operating system feature of the Alcatel-Lucent 1830 Photonic Service Switching platform. At the time of writing both the 1830 PSS operating system and the Alcatel-Lucent Service Router Operating System are being enhanced with GMPLS UNI support. This allows delivering the TCO benefits of IP/optical control plane integration to the installed base of 1830 PSS and Service Router deployments by means of a software upgrade.

ACRONYMS

AF	Assured Forwarding
BE	Best Effort
CAPEX	capital expenditures
DWDM	Dense Wavelength Division Multiplexing
EF	Expedited Forwarding
FMO1	Future Mode of Operation 1 (GMPLS only)
FMO2	Future Mode of Operation 2 (MPLS and GMPLS+UNI)
FRR	Fast Reroute (MPLS)
gLSP	generalized label switched path
GMPLS	Generalized Multiprotocol Label Switching
LSP	label switched path
LSR	Label Switch Router
MPLS	Multiprotocol Label Switching
ODU	optical data unit
ONCP	Optical Network Connection Protection
OPEX	operational expenditures
OTN	Optical Transport Network
PE	provider edge
PMO	Present Mode of Operation
PRC	Protection and Restoration Combined
PSS	Photonic Service Switch
ROADM	Reconfigurable Add/Drop Multiplexer
SBR	source-based restoration
SDN	Software Defined Networking
SLA	service level agreement
SNCP	subnetwork connection protection
SRLG	shared risk link group
TCO	total cost of ownership
UNI	user-network interface
WDM	Wavelength Division Multiplexing
XRS	Extensible Routing System

RESOURCES

- Capitalizing on IP/optical integration. TechZine article
- The urgent need for agility in optical networks. TechZine eBook
- Agile Optical Networking webpage
- Integrated Packet Transport webpage
- IP Core Routing solution webpage
- Cloud-Optimized Metro solution webpage

Alcatel·Lucent