



The five vectors of GMPLS evolution

White paper

Contents

The need for maximum network availability	3
Causes of optical network downtime	3
Limitations of traditional (1+1) protection	3
Distributed GMPLS control plane basics	4
Distributed GMPLS control plane benefits	5
GMPLS Evolution	5
Vector 1: More features	5
Vector 2: More scalable	7
Vector 3: Faster restoration	7
Vector 4: Multi-layer	8
Vector 5: Elastic photonic networking	12
Summary	14

The need for maximum network availability

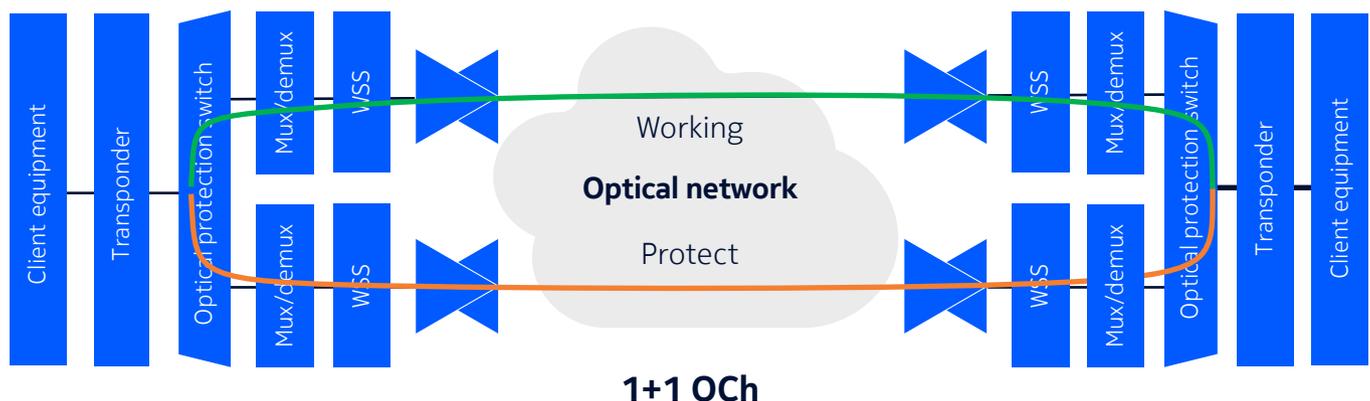
As both consumers and enterprises become ever less tolerant of service downtime, the need for maximum network and service availability has never been higher. For large enterprises the cost of downtime hit \$9,000 per minute in 2024 according to Forbes, with finance and healthcare in particular facing downtime costs of \$5 million per hour, not counting any potential fines and penalties. According to a 2024 report by consultancy.uk, the world’s largest two thousand companies lose \$400 billion per year to downtime. In response telecommunications providers are offering ever more stringent service level agreements (SLAs). For example, wavelength services are being offered with target availability of up to 99.999%, which is equivalent to only five minutes downtime per year. At the same time the 5G Public Private Partnership (5G PPP) is targeting “zero” perceived” downtime.

Causes of optical network downtime

But what are the causes of optical network downtime? The most common potential causes of optical network downtime include fiber cuts, fiber degradation, network equipment failures, site failures and human error. Fiber cuts can occur during construction, during maintenance of electricity, gas or water utilities, or as a result of sabotage, vandalism and theft. According to a DIRT report there were 64,198 telecom/CATV damage incidents in USA in 2023. Fiber can degrade due to splices for installation or repair, fiber additions during repair, dirty or poor connectors, micro and macro bending. Network equipment failures can relate to power supply units, fans, controllers, transponders, amplifiers or wavelength selective switches. The sites (point-of-presence, local exchange, data center, etc.) housing the network equipment can also suffer failures such as loss of or insufficient power and cooling. The final common cause of network downtime is human error such as a provisioning misconfiguration or a “careless elbow” that knocks a fiber connector out of its socket. A key objective of a distributed GMPLS control plane is to enable the network to survive these failures.

Limitations of traditional (1+1) protection

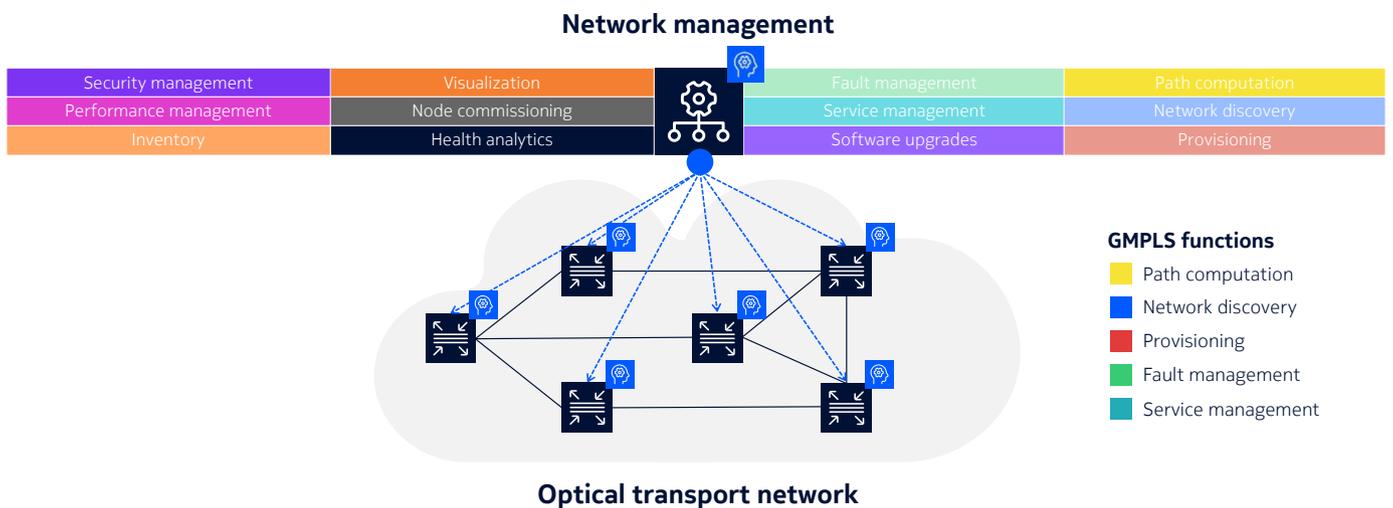
Figure 1 - 1+1 protection example (1+1 OCh)



One option to maximize network availability is traditional 1+1 protection. Types of 1+1 protection include 1+1 OTS, 1+1 OMS, 1+1 OCh (shown in Figure 1), 1+1 SNCP, client protection and Y-cable. These protection mechanisms are relatively simple and provide guaranteed protection against a single failure. Furthermore, 1+1 protection can be very fast with sub 50 millisecond failover. However, these mechanisms can be expensive as protection resources cannot be shared or used for other services when not in use. Also they typically only protect against a single failure. Multi-failure protection is possible but very costly. A distributed GMPLS control plane can provide the benefits of 1+1 protection but without these limitations.

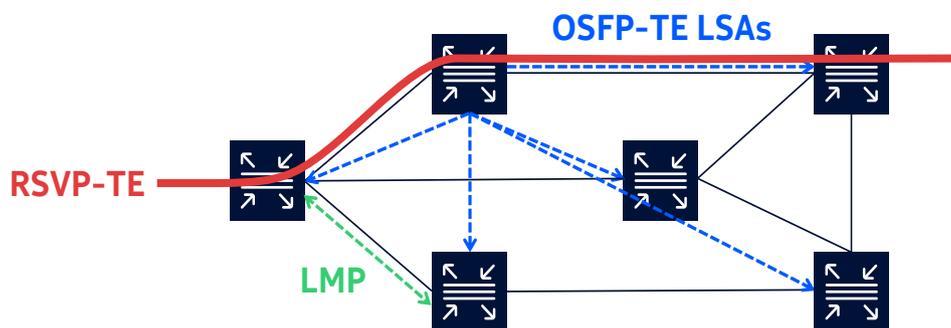
Distributed GMPLS control plane basics

Figure 2 - GMPLS distributes key functions from management to network elements



In a traditional non-GMPLS network, the intelligence is centralized in the network management system (NMS) which provides key functions including fault management, path computation, service management, network discovery, provisioning, software upgrades, security management, visualization, performance monitoring and management, node commissioning, inventory, and health analytics. As shown in Figure 2, with a distributed GMPLS control plane intelligence, including path computation, network discovery, provisioning, fault management and service management, is also distributed to the network elements.

Figure 3 - Key GMPLS protocols



In terms of how a distributed GMPLS control plane works, it leveraged and extended several protocols from the IP world, as listed in Figure 3. OSPF-TE (Open Shortest Path First – Traffic Engineering) with optical layer extensions is used to distribute optical network information through the network. Each node tells all other nodes about its neighbors and links with link state advertisements (LSAs). RSVP-TE (Resource Reservation Protocol - Traffic Engineering) with optical layer extensions provides signalling and is used to set up end-to-end paths, requesting and reserving resources. One additional protocol used by Nokia's GMPLS, that does not come from the IP world, is LMP (Link Management Protocol) which runs between neighboring network elements and is used to manage their shared links.

Distributed GMPLS control plane benefits

GMPLS offers a number of important network operator benefits. These include:

- **Higher availability:** The ability to survive multiple simultaneous failures/cuts/degradations, including with fast recovery.
- **Reduced costs:** Shared protection/restoration resources, efficient routing, multi-layer protection and restoration.
- **Faster provisioning:** Real-time path computation based on an accurate real-time network view with sophisticated modelling of optical penalties.
- **Simplified operations:** Automated network and resource discovery, automated route optimization, and real-time network resource inventory synchronization.
- **Differentiated services:** Provides the ability to offer enhanced services and differentiated SLAs. Minimizes SLA penalties with high service availability and optimized latency.

GMPLS Evolution

GMPLS has been around for over 20 years. Key IETF standards related to GMPLS include RFC3471 for signalling which was published in 2003, RFC 3945 for architecture published in 2004, and RFC 4426 for recovery which was published in 2006. ITU-T also published Automatic Switched Optical Network (ASON) standards including G.8080/Y.1304 architecture in 2001. At the same time, Nokia has a long and proven history of GMPLS development and deployment, starting with GMPLS for SDH in 2003. To date, over one hundred Nokia GMPLS networks are in operation around the globe. Furthermore, Nokia's GMPLS solution continues to evolve along the following five vectors:

Vector 1: More features

One vector along which distributed control planes such as Nokia's GMPLS are evolving is with more features. Since introducing GMPLS control plane for SONET/SDH in 2003, followed by WDM in 2008 and OTN in 2010, Nokia has continued to enhance its GMPLS with support for both new hardware and new features based on feedback from its large and global GMPLS installed base. For example, Table 1, lists the key GMPLS features added between 2016 and 2025.

Table 1 – Key new features added to the Nokia GMPLS control plane (2016-2025)

2016	2018	2019	2020
<ul style="list-style-type: none"> • Flexible grid • IETF UNI 	<ul style="list-style-type: none"> • L-band support (i.e. C+L) • Explicit restoration paths 	<ul style="list-style-type: none"> • Most disjoint mode • Per LSP configurable revertive/non-revertive 	<ul style="list-style-type: none"> • Subsea configurations
2022	2023	2024	2025
<ul style="list-style-type: none"> • Optical intrusion alarms as restoration triggers • Restoration based on end-to-end path alarms • Alien wavelength restoration 	<ul style="list-style-type: none"> • Configurable maximum number of restoration attempts per service • Mixed software versions in one GMPLS network 	<ul style="list-style-type: none"> • SRLG diversity awareness • Dynamic regeneration • Most disjoint option 	<ul style="list-style-type: none"> • Multi-homed diversity • Automatic route optimization • Latency-aware routing

Today, the Nokia GMPLS implementation provides a broad array of features as summarized in Table 2. These features include a variety of protection options including unprotected, 1+1 protection, on-the-fly source based restoration, guaranteed restoration where an inactive backup path is pre-calculated by source node and resources are reserved, and protection and restoration combined which is similar to 1+1 SNCP but with restoration on both main and spare paths in the event of any failures. Initial active paths can be defined explicitly in the management system while restoration paths are calculated by the source network element. Path computation can consider latency and operator-assigned resource colors which can be used to divide a GMPLS domain into areas (e.g., metro A, backbone, metro Z). Nominal and backup paths can be full disjoint or partially/most disjoint, as can two LSPs that require diversity.

Table 2 - Key features of Nokia’s GMPLS control plane (2025)

<ul style="list-style-type: none"> • Unprotected • 1+1 protection • Guaranteed restoration (GR) • Source based restoration (SBR) • Protection and restoration combined (PRC) • Explicit path (nominal + 1 optional restoration path) • Latency-aware routing • Most/partial disjoint • LSP diversity • Single/dual homed topologies • Resource coloring • Reversion (automatic, manual) • Multi-layer networking (LO+L1) 	<ul style="list-style-type: none"> • Wait to revert timer (WRT) • Priority • Preemption • Crankback • Shared risk link groups (SRLG) • Flexible grid support • Alien wavelength support • Brownfield migration • Soft-reroute based on alarms • C+L support • Graceful ODUflex resizing • Open/external SNCP • Elastic photonic networking <ul style="list-style-type: none"> – Elastic restoration – Elastic regeneration – Elastic channel optimization 	<ul style="list-style-type: none"> • Optical intrusion alarms as restoration trigger • GCC protection • Joint GMPLS/NMS control of network resources: <ul style="list-style-type: none"> – Split NE – Shared ports • Maintenance operations: <ul style="list-style-type: none"> – Move traffic – Lock a port – Link maintenance – LSP freeze – Manual restoration • Automatic route optimization
---	--	--

A pair of LSPs can be dual-homed at one end and single-homed at the other end. Restoration enhancements include reversion, a wait to revert timer, priority settings on each service for both the initial setup and restoration, in cases of contention. Preemption lets lower priority services to be removed in order to enable higher priority services that would otherwise go down. Where two LSPs compete for the same resource the losing LSP can “crankback” to its source and try an alternative path.

Shared risk link groups (SRLG) avoid inadvertently sharing the same physical infrastructure (e.g., two fibers that share a duct). Brownfield non-GMPLS Nokia 1830 PSS networks can be migrated to GMPLS. Maintenance features include the ability to move traffic (e.g., reroute services to manually defrag the network), lock a port (no new connections but existing connections continue), link maintenance (e.g., firmware upgrades) with the source node not notified of alarms, LSP freeze which prevents rerouting in order to test a faulty LSP, and manual restoration. Automatic route optimization periodically searches for routes with fewer or no diversity violations and fewer or no optical-electrical-optical (OEO) regenerations.

Vector 2: More scalable

A second vector for GMPLS evolution is scalability. Early GMPLS networks were limited to tens of nodes. In 2006, Nokia increased its GMPLS scalability to 100+ nodes. In 2021 this was increased to 600 nodes for layer 0 and 400 nodes for layer 1. As of today, Nokia GMPLS supports up to 3,000 nodes at layer 0 and 1,000+ at layer 1. Another important GMPLS scalability metric is the number of neighbors, with Nokia GMPLS supporting up to 140 neighbors. In addition, GMPLS scalability is enabled by scalability in the underlying hardware. Scalability in terms of the spectrum (9.6THz with C+L), the number of ROADMs degrees per node (16 with CDC, 28 with colorless-only), WSS scalability (twin 1x32), CDC add drop scalability with high port count multicast switch and MxN WSS options, and OTN switching with the 1830 PSS-X supporting up to 48 Tb/s.

Vector 3: Faster restoration

A third vector is restoration speed, with GMPLS vendors chasing the <50 milliseconds gold standard of 1+1 protection. Nokia's protection and restoration combined delivers fast <50ms failovers with restoration used to reestablish 1+1 protection. So, the real question is how quickly can schemes that rely purely on restoration recover from a failure? For a given network size and complexity, multiple factors impact restoration times as listed in Table 3.

Table 3 - Key factors impacting restoration times

	Layer 0	Layer 1
Failure detection time	✓	✓
Signalling latency	✓	✓
Control plane routing protocol convergence	✓	✓
Path computation time	✓	✓
Transponder re-tuning	✓	
WSS switching times	✓	
Amplifier and wavelength power balancing	✓	
OTN switching reconfiguration		✓

Nokia has worked to improve all of these factors driving down restoration times. For example, guaranteed restoration where the restoration path is pre-computed and pre-signalled with resources reserved can speed up path computation and signalling. Layer 1 restoration is now under 300 milliseconds for a single ODUk LSP and under 1 second for eighty ODUk LSPs. However, layer 0 restoration times are dominated by amplifier and wavelength power balancing which are a function of the optical line system's link control software. Re-tuning the transponder to a different frequency and a new path with different chromatic dispersion values can also be a factor. These factors result in longer restoration times for layer 0 compared to layer 1. However, Nokia has significantly reduced layer 0 restoration times by optimizing the link control algorithms of the 1830 PSS/PSI-L. This has taken layer 0 restoration time from several minutes to less

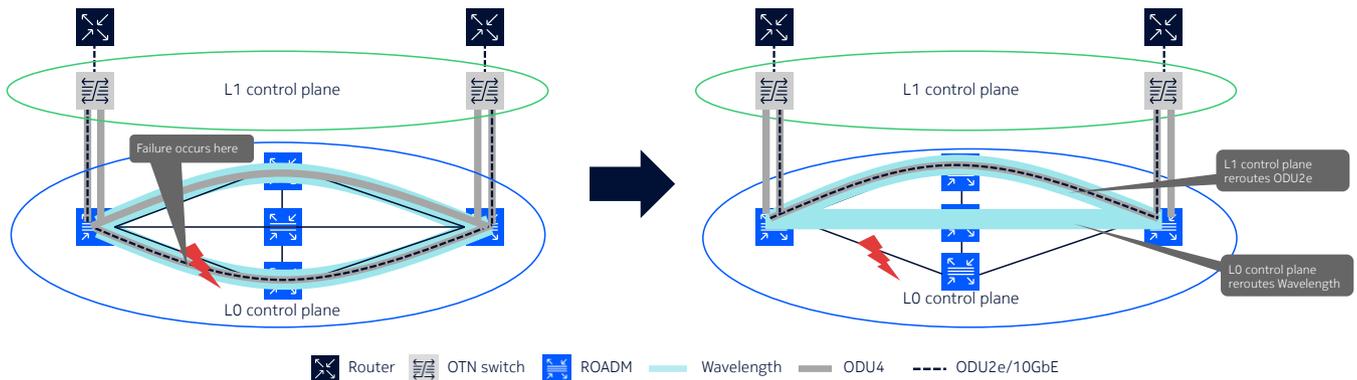
than 10 seconds for up to 10 channels (metro and long haul). 96 C-band channels can be restored in less than 20 seconds for a metro network and less than 1 minute for a long-haul network. Future GMPLS innovation will look to close the gap between layer 0 and Layer 1 restoration times, getting us ever closer to <50 milliseconds.

Vector 4: Multi-layer

A fourth vector of GMPLS evolution, and a key point of differentiation between GMPLS vendors is multi-layer, that is layer 0 (DWDM) and layer 1 (OTN). Without multi-layer GMPLS, network operators would need to run two independent control planes: one for layer 0 and one for layer 1. This results in two independent databases, independent decisions and uncoordinated action resulting in higher costs and/or lower availability. A multi-layer GMPLS control plane leverages a single multi-layer database with traffic engineering link attributes and priority awareness between layers. This enables coordinated protection and restoration, and coordinated maintenance resulting in better cost vs. availability trade-offs, as illustrated by the following four examples.

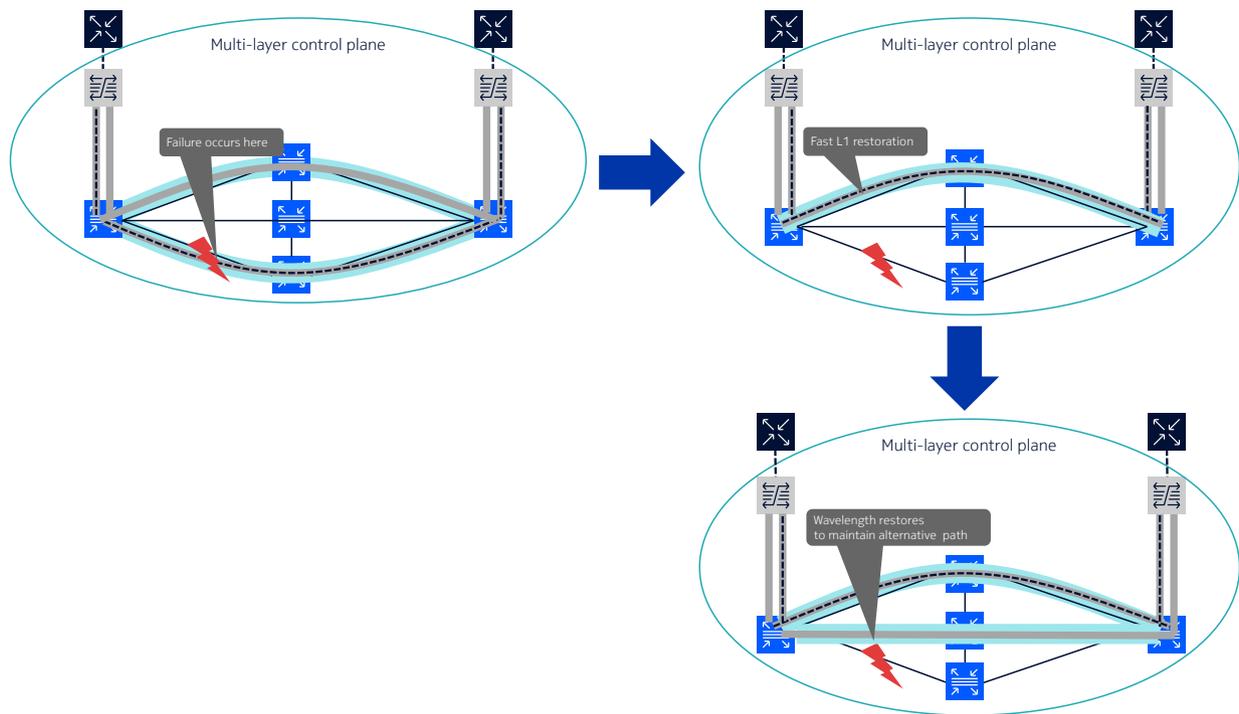
Example 1: Coordinated restoration

Figure 4 – Uncoordinated restoration with separate L0 and L1 control planes



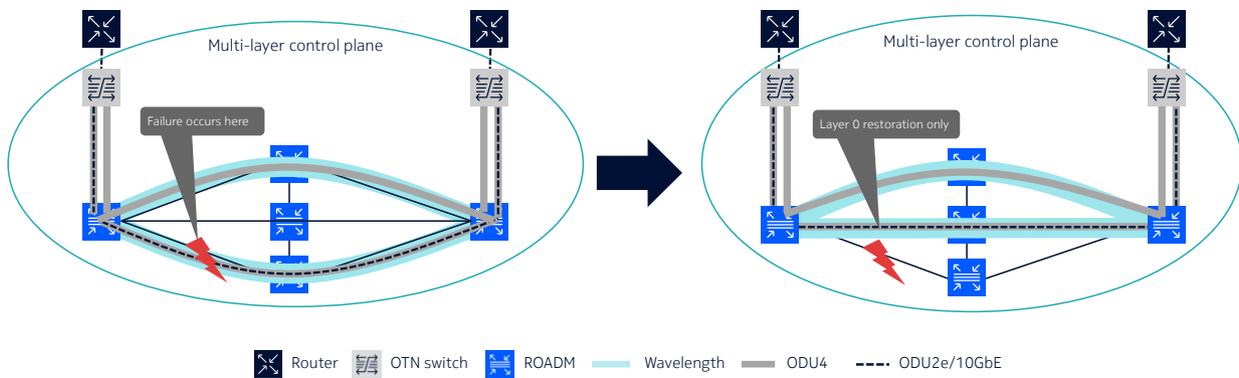
For example, in the event of a failure with two independent control planes, both layers would initiate restoration independently, as shown in Figure 4, thus wasting valuable resources. With a multi-layer control plane if sufficient resources already existed at layer 1, faster layer 1 restoration could occur, as illustrated in Figure 5.

Figure 5 - Multi-layer restoration with L1 prioritized



Alternatively, if layer 1 restoration was not possible or layer 0 restoration was preferred, restoration could instead occur at layer 0, as shown in Figure 6.

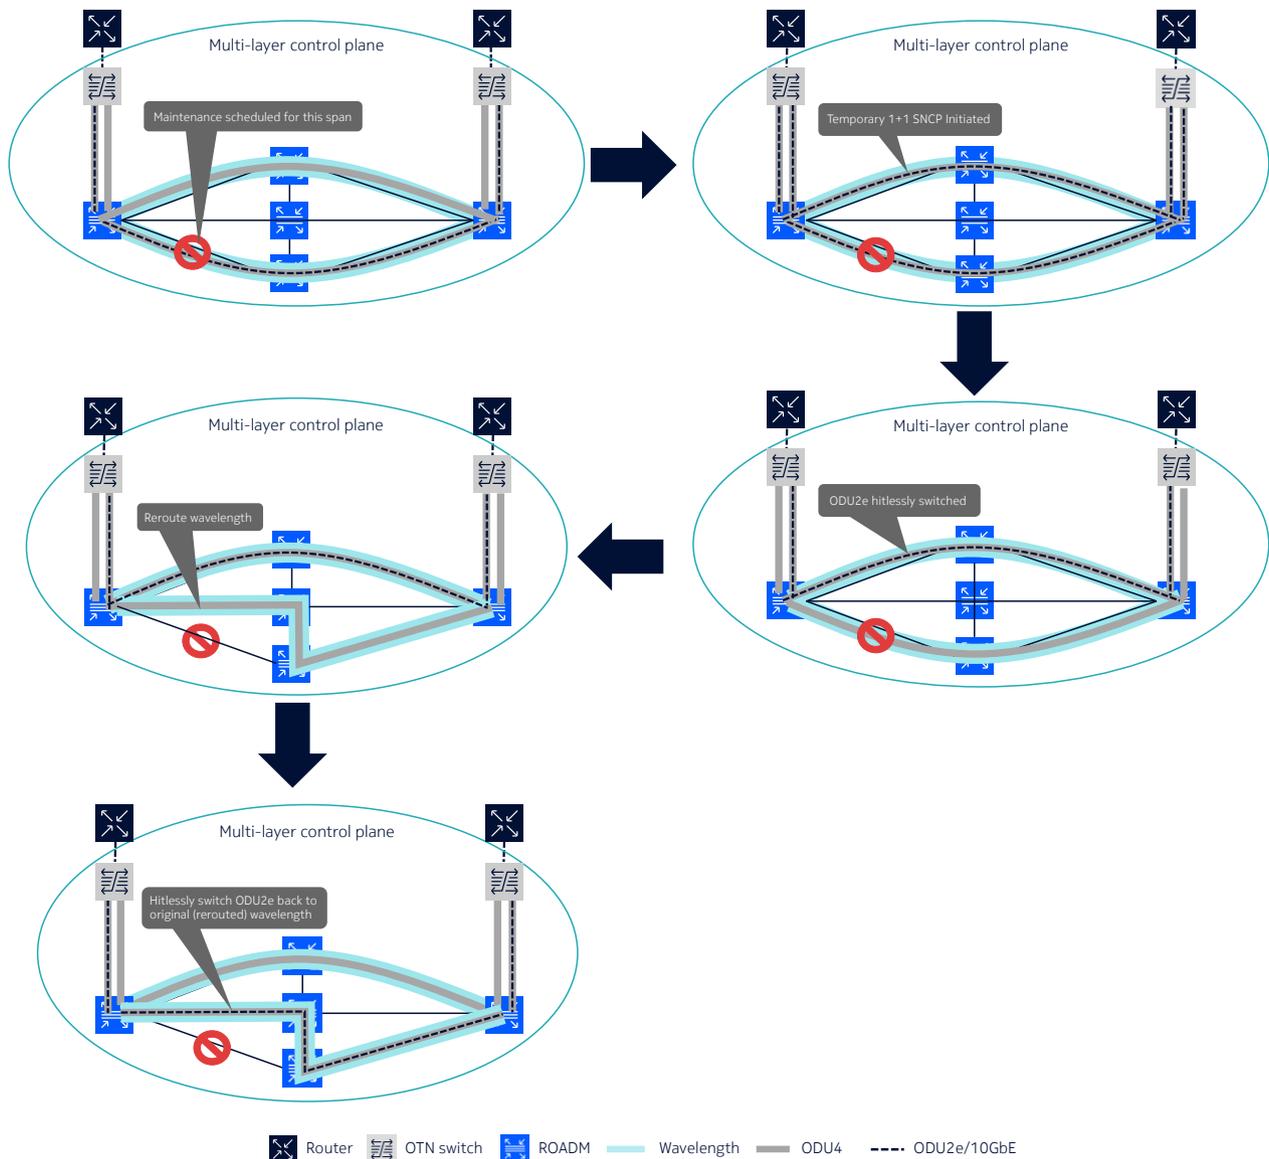
Figure 6 - Multi-layer restoration with L0 prioritized



Router
 OTN switch
 ROADM
 Wavelength
 ODU4
 ODU2e/10GbE

Example 2: Co-ordinated maintenance

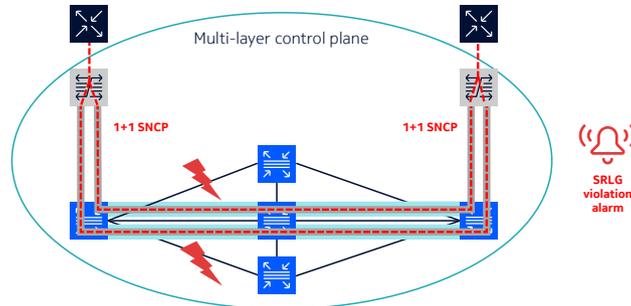
Figure 7 - Coordinated maintenance



Another example of the benefits of a multi-layer control plane is coordinated maintenance. As illustrated in Figure 7, if a link is to be taken out of service for maintenance, a bridge and roll is initiated: a temporary SNCP is created with the protective path over an alternative diverse wavelength, then the layer 1 service is hitlessly switched to this alternative wavelength. Then once the original wavelength has been rerouted around the maintenance link, the service is hitlessly switched back to this original wavelength.

Example 3: Traffic engineering link attribute propagation

Figure 8 - Attribute propagation

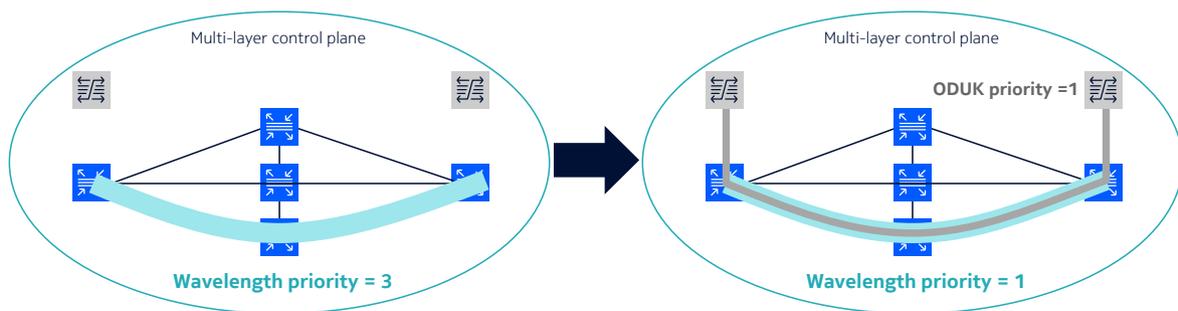


TE link attribute propagation is a key function of Nokia’s multi-layer GMPLS, providing improved visibility of the photonic network in the electrical layer. This enables better restoration and reversion decisions in the electrical layer, especially if a specific diversity requirement must be maintained. As shown in Figure 8, if after two fiber cuts both wavelengths carrying an ODUk 1+1 SNCP protected service share an SRLG, then this can be alarmed to the network operator who would then have the option to prioritize a repair of one or both of these failures or take other actions to ensure 1+1 SNCP without an SRLG violation.

Example 4: Priority Inheritance

If priorities are used in the GMPLS network then it is important to be consistent across layers. For example, preempting a layer 0 wavelength transporting a high priority (e.g., priority 1) layer 1 ODUk client should be avoided. To achieve this the layer 1 client priority must be inherited by the layer 0 server layer. As illustrated in Figure 9, when a higher priority layer 1 service is routed over a lower priority layer 0 wavelength, the layer 0 wavelength priority is upgraded to match the layer 1 service’s high priority.

Figure 9 - Priority inheritance



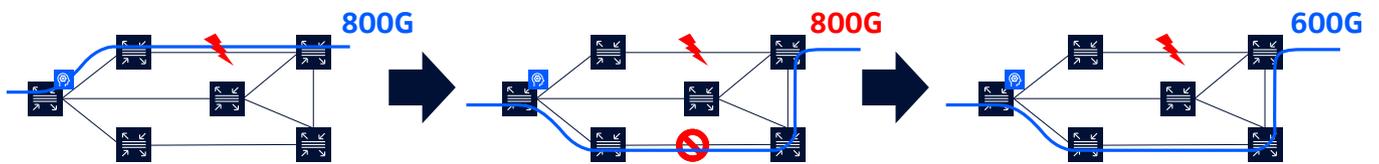
Vector 5: Elastic photonic networking

The latest innovation in distributed GMPLS control planes is Elastic Photonic Networking (EPN). EPN marries the intelligence of a distributed GMPLS control plane with the programmability of advanced coherent technology and flexible grid optical line systems. Optical wavelength data rate, modulation, baud rate, spectral occupancy, forward error correction and other parameters can be changed by the control plane in order to cost-effectively maximize availability, in response to changes in network conditions, as illustrated with the following three examples.

Example 1: Elastic restoration

The first example of EPN is illustrated in Figure 10. After a failure, the only path available cannot support 800Gb/s due to lower OSNR and/or higher nonlinear penalties. The GMPLS control plane calculates that 600Gb/s can be supported and adjusts the probabilistic constellation shaping (PCS) bits per symbol and other parameters then establishes a 600Gb/s wavelength. Another scenario would be where there is not enough spectrum to support 800Gb/s so the control plane reduces the baud rate in order to fit in the more limited available spectrum.

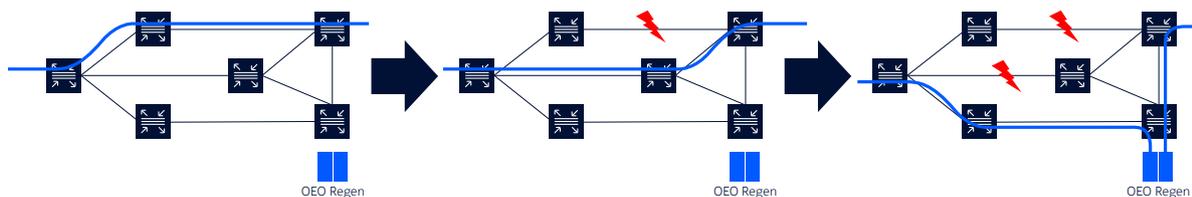
Figure 10 - Elastic restoration



Example 2: Elastic regen pools

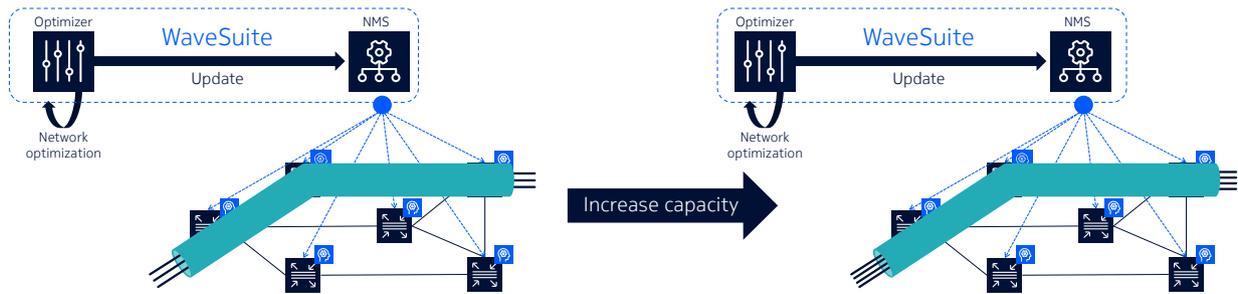
Another example of EPN is elastic regen pools, pools of transponders configured for OEO regeneration, for cases where restored wavelengths cannot be supported without intermediate OEO regeneration. These regen pools can be shared across multiple services. All profiles that are potentially possible on the transponders are usable, with the GMPLS control plane automatically selecting and assigning a matching profile. This is illustrated in Figure 11, where after a second failure, the wavelength leverages a regen pool in order to restore.

Figure 11 - Elastic regen pools



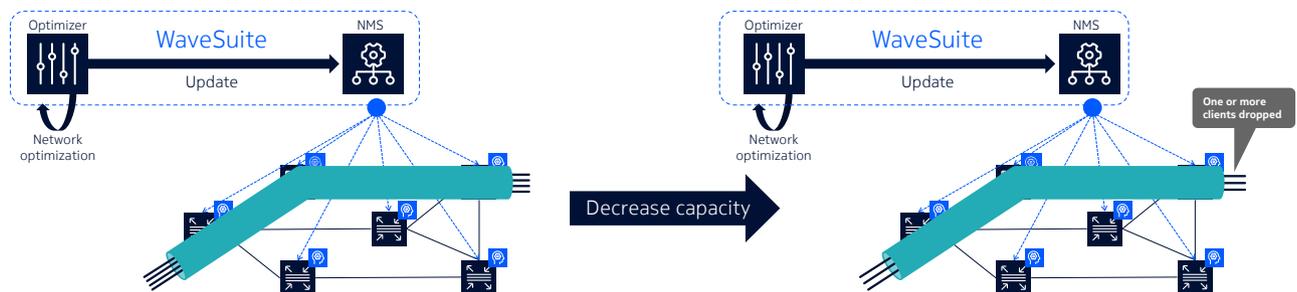
Example 3: Elastic channel optimization

Figure 12 - Elastic channel optimization: increase capacity example



A third example of EPN is elastic channel optimization. Here Optimizer, an application within the Nokia WaveSuite automation suite, analyzes the network searching for optimizations. If it finds an opportunity to increase the capacity of a wavelength it updates the NMS function within WaveSuite, as shown in Figure 12. The NMS then provides a “single touch” instruction to the GMPLS control plane to take all the required actions along the path to implement this capacity increase.

Figure 13 - Elastic channel optimization: decrease capacity example



If WaveSuite optimizer identifies that the capacity needs to be decreased it can also update the NMS which again uses a “single touch” to instruct the GMPLS control plane to make the appropriate changes along the path. As shown in Figure 13, this includes removing one or more clients which the transport capacity can no longer support, with the decision of which clients to drop based on service priority.



Summary

Distributed GMPLS control planes have emerged as a critical technology for enhancing network availability, operational efficiency, and optical-layer intelligence in modern high capacity networks. As downtime costs escalate and service expectations rise, Nokia's GMPLS addresses these challenges by enabling rapid recovery from complex failure scenarios and improving automation across transport layers. With Nokia's extensive history of GMPLS innovation, this technology has evolved and continues to evolve with more features, greater scalability, faster restoration, coordinated multi layer control, and dynamic adaptability through Elastic Photonic Networking. These advancements empower telecommunications providers to deliver differentiated SLAs, reduce complexity, and maintain high service availability.

About Nokia

Nokia is a global leader in connectivity for the AI era. With expertise across fixed, mobile, and transport networks, we're advancing connectivity to secure a brighter world.

Nokia is a registered trademark of Nokia Corporation. Other product and company names mentioned herein may be trademarks or trade names of their respective owners.

© 2026 Nokia

Nokia OYJ
Karakaari 7
02610 Espoo
Finland
Tel. +358 (0) 10 44 88 000

Document code: CID 215313 (February)