

LIGHTWAVE®

■ TECHNOLOGY

Intelligent optical layer improves broadband disaster recovery

By KEVIN OYE

An unforeseen power failure stretches from the U.S. Northeast to the Midwest, stressing the limits of existing government and enterprise disaster recovery plans. A powerful earthquake off the coast of Taiwan disrupts voice, Internet, and private-line traffic in Asia, leaving millions without communication service for days, and in some cases weeks. A double fiber cut takes down a leading Tier 1 service provider's major service

area, immediately impacting services for its fixed-line and mobile customers. The range of disaster scenarios to be considered in planning for business continuity has grown significantly in a world that has experienced 9/11 and Hurricane Katrina.

At the same time, as globalization of networking accelerates and as applications in enterprise and government networks become increasingly bandwidth-intensive, survivability of the

underlying infrastructure carrying mission-critical broadband traffic has become more important than ever. Financial organizations, for example, are increasingly dependent on global connectivity to support a broad range of mission-critical transaction applications highly sensitive to service interruption. For these companies, the damage is measured not simply in downtime, but in the millions of dollars in potential losses that result from network disruption. For

Optical mesh resiliency and flexibility

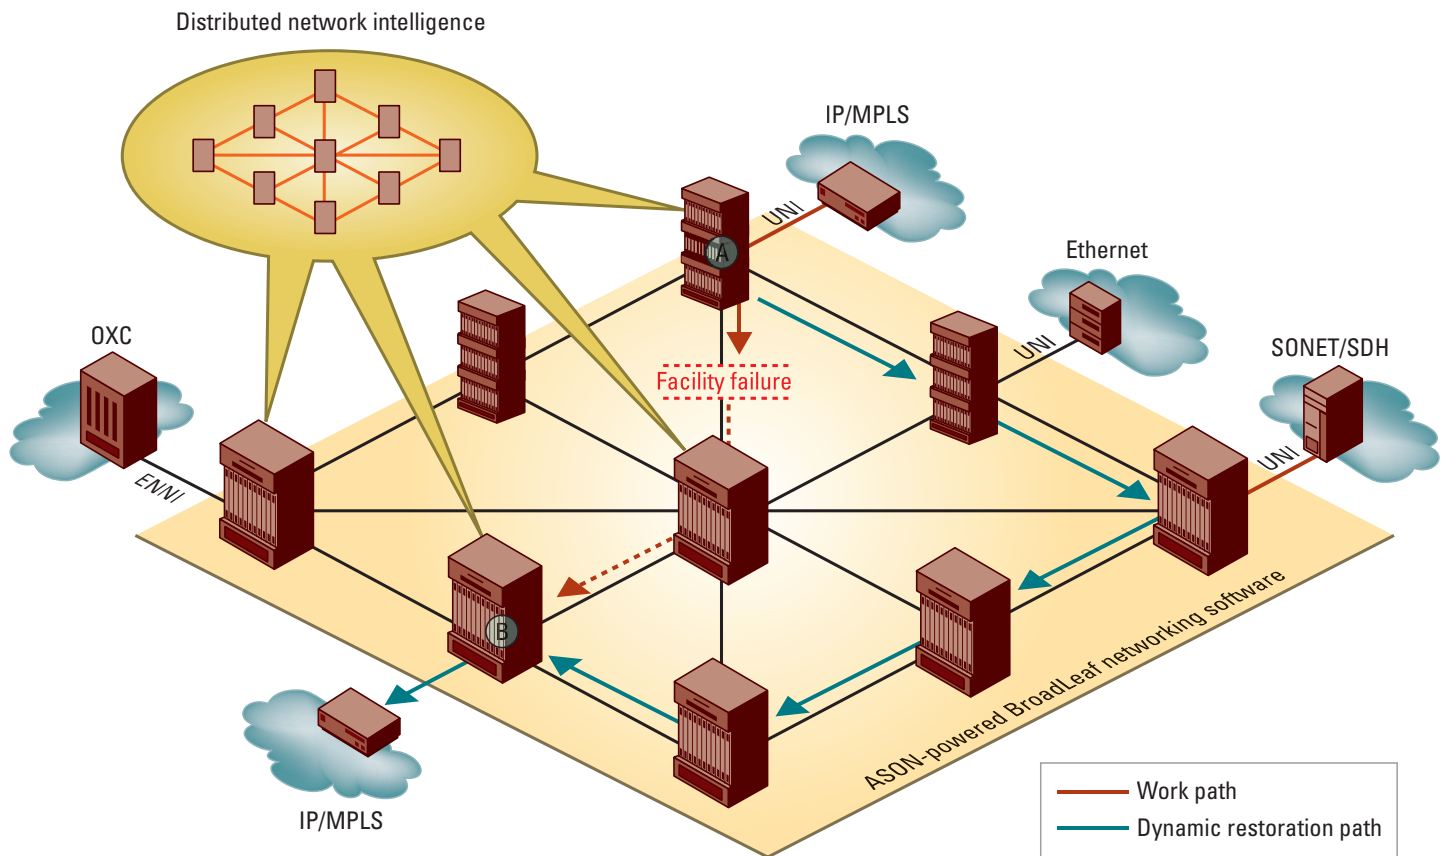


Figure 1. Optical mesh resiliency enables rapid recovery—even from multiple, simultaneous failures—and empowers fully automated, end-to-end broadband networking. Leveraging a common control plane, each node in the intelligent optical mesh network actively participates in provisioning and restoration tasks, sharing real-time information about link status, capacity status, routing tables, and available protection and restoration.

service providers, the same scenario translates into lost revenue and customer churn. The stakes are high and getting higher.

The communications applications that businesses and consumers rely on today are not only more bandwidth-intensive, they are also more dynamic and delay-intolerant, which demands greater levels of flexibility within the infrastructure that allocates network capacity to support end-to-end services. These trends affect disaster recovery planning at multiple layers of the network, with implications that reach beyond traditional topics such as redundant storage of data.

At the optical layer, more advanced service protection and restoration mechanisms are required to ensure true survivability in the event of service disruption from unexpected facility failures, human error, deliberate acts, or natural disaster. The introduction of intelligent dynamic routing and signaling protocols in optical switches and the emergence of flexible, self-healing optical mesh architectures are creating the foundation for a more resilient and scalable broadband infrastructure.

Next-generation optical recovery

A heightened sense of network vulnerability, accelerating rates of traffic growth, and the increasing dependence on high-bandwidth applications have created the need for a much more agile, scalable, and survivable optical network, one that can support greater diversity in recovery routes and new levels of intelligence to allow the optical network to rapidly react and adapt as disaster scenarios unfold.

First-generation high-speed optical networks relied on self-healing ring architectures and SONET/SDH ring-based protection mechanisms for network re-

Mesh protection options		
Protection option	Definition	SONET/SDH equivalent
Unprotected circuits	No protection allocated; in event of failure, circuit remains inactive until the path is repaired	Unprotected circuit
1+1 path protection	Traffic is dual-cast on dedicated and diversely routed working and protection paths. No single network event can fail a circuit. Each destination node simultaneously receives both paths and selects which path to use based on alarm and failure conditions	UPSR/SNCP
Span restoration	Local, span-level protection scheme that provides rerouting of the affected circuits around the failed span. The protection route may consist of a single parallel span, or a sequence of spans. Protection bandwidth is shared among different failure scenarios	Traditional span protection schemes: linear APS/MSP; BLSR/MS-SPRing
Shared path restoration	End-to-end, path-level protection scheme that provides rerouting of the failed circuits. Protection bandwidth is shared among different failure scenarios	N/A

covery. Originally designed to scale voice networks and widely deployed in metro, regional, and long-haul networks, SONET/SDH ring architectures have served as the underlying optical service infrastructure for most carrier networks since the early to mid-1990s.

SONET/SDH networks provide protection via linear or ring-based protection protocols. In redundant 1+1 protection, information flowing across terminals is bridged at the source and sent along both a “working” and a “protect” facility—on separate fiber connections. At the receiving end, the best signal is chosen. In the event of a failure, a switch is made from the working to the protection facility. This scheme works well—unless, however, a carrier experiences a double fiber cut. Furthermore, ring-based protection schemes limit the fiber utilization to only 50%, since half the fiber capacity is always held in reserve as the backup protect capacity. Ring-based protection, while proven and widely deployed, is emerging as an architectural limitation in core networks as traffic scales and the risk of service disruption due to natural disaster or other causes increases.

Optical network evolution

The evolution from static ring-based architectures to agile, next-generation optical networks has just begun in the past five years, and pragmatic migration remains a key consideration for network operators. The advent of intelligent optical networking devices (e.g., ASON/GM-PLS-compliant optical crossconnects) that concurrently support SONET/SDH protection schemes and a variety of mesh-based protection options offers carriers nondisruptive and cost-effective means to migrate their optical networks while

optimizing their broadband infrastructure for the highest degree of network survivability. In addition, the application of intelligence at the optical layer provides network architects the tools needed to reduce costs through improved network efficiencies, simplified provisioning through the automation of operational tasks, and enhanced service availability through the introduction of optical mesh and diverse protection mechanisms (Fig. 1).

In mesh networks, nodes are interconnected through multiple paths and network elements share information over a common control plane. When multiple network elements (typically geographically diverse) are connected in an irregular topology (in contrast to a structured topology such as a ring or a linear chain), the ability to discover the physical topology of the network and track changes to the topology in real time is very important. Intelligent optical switching enables this kind of self-discovery and provides the distributed intelligence to self-discover the best path between two endpoints given a set of routing constraints, such as bandwidth desired, diversity of path, etc.

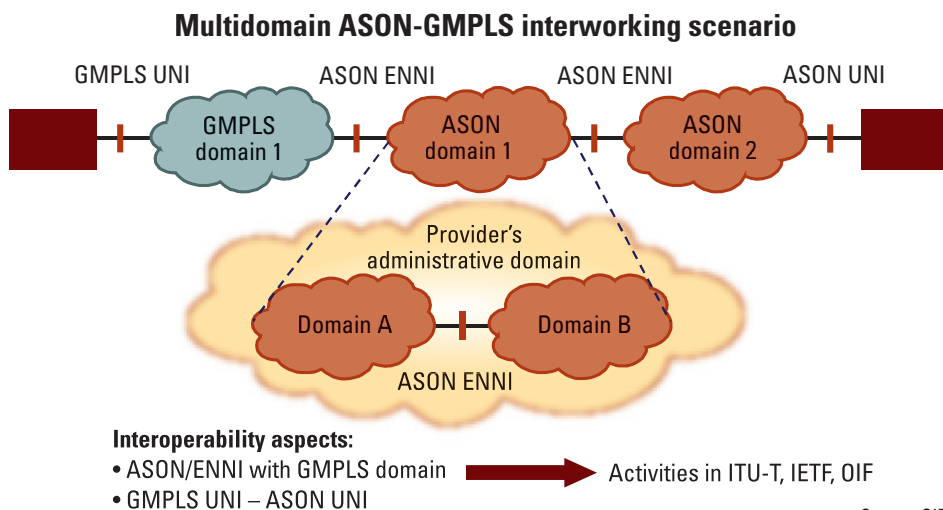


Figure 2. The work of several standards bodies has opened the door for distributed network intelligence as well as more robust optical layer protection in a multivendor environment.

Over the past several years, to enable smooth interoperability between optical networks from multiple vendors, coordinated efforts among the ITU-T, IETF, and OIF standards bodies have resulted in a set of network interface definitions that describe how a third-party user device can request optical bandwidth services (e.g., the user network interface, or UNI) and how two optical networks can talk to each other as peers (e.g., the network to network interface, or NNI).

These industry standards for routing and signaling (ASON/GMPLS) are a key step toward broader deployments of optical networks, as they will encourage the smooth deployment of multivendor infrastructures that are all built to common interface definitions. These protocols, which today are available in a wide range of IP- and optical-layer network equipment, have provided network operators with the tools to enhance multidomain IP/optical interworking (Fig. 2).

The successful application of this technology in multivendor, multicarrier environments was recently demonstrated at an OIF-sponsored “On-Demand Ethernet” interoperability demonstration at the 33rd European Conference and Exhibition on Optical Communication in Berlin. The 2007 OIF Worldwide Interoperability Demonstration showcased end-to-end provisioning of dynamic switched Ethernet services over multiple control-plane-enabled intelligent op-

tical networks through the use of OIF implementation agreements of UNI 2.0 and ENNI. In-service Ethernet bandwidth modification and control plane discovery and failure recovery were new features demonstrated. This event included participation by such global carriers as AT&T, China Telecom, Deutsche Telekom, France Telecom Group, KDDI, Telecom Italia, and Verizon, as well as several equipment vendors.

Unleashing multiple protection options

With the ability to make intelligent decisions and dynamic adjustments, the intelligent optical mesh network automates

traditionally complex, time-consuming operational tasks, and provides the foundation for more efficient capacity utilization. Mesh protection schemes use the distributed network intelligence of optical switches to determine new paths through the network, either on the fly or in advance. The actual protect capacity can be either reserved, shared, or provisioned at the time of fault, all under the control of the network operator, who can offer the various levels of protection as part of a tiered service model.

With mesh-based architectures, many options are available in the event traffic has to be rerouted due to a fiber break or system fault. By definition, mesh topologies can be architected with multiple restoration paths and, unlike traditional rings, are not limited to only one protect path. Network operators can proactively design their networks with whatever level of redundancy they desire to be able to survive multiple network failures. This is in contrast to ring protection protocols, which are not designed to survive multiple failure events.

Mesh topologies allow network operators to allocate only as much bandwidth (working and protection) as is required for current demand conditions. Intelligent switching algorithms control the flow of traffic, enabling providers with the flexibility to assign a range of priorities to various levels of service. As carriers transition to mesh-enabled architectures, they can realize significant

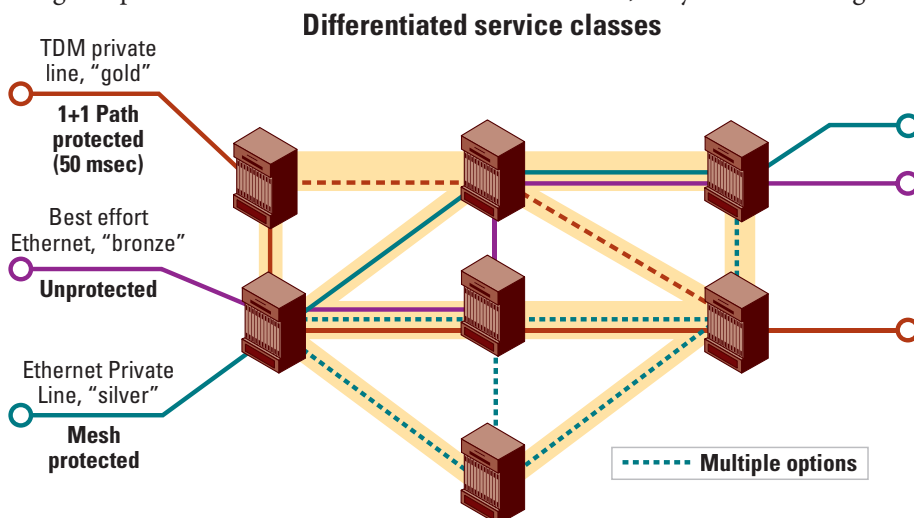


Figure 3. Customized protection and restoration schemes allow network operators to meet diverse application and service-level agreement requirements for customers’ circuit (e.g., TDM) or packet (e.g., IP, Ethernet) traffic.

savings by optimizing the optical infrastructure required—leveraging the richer protection capabilities of a mesh network, in contrast to a traditional SONET/SDH-based ring network, while also reducing operational complexity through the automation of network provisioning. Furthermore, since the allocation of protection capacity is totally under the control of the network operator, fiber utilization greater than 50% can be achieved.

Flexible service protection

With intelligent optical mesh architectures, network operators can choose from a variety of protection schemes, some of which can mimic traditional ring-based schemes, while also offering additional, flexible options. Advanced

intelligent optical switch systems can also simultaneously support traditional ring protection schemes alongside the mesh restoration options highlighted in the table.

By leveraging the diverse protection and restoration capabilities of the intelligent optical mesh, network operators can also configure and deliver differentiated class-of-service offerings (Fig. 3) to enhance their TDM, IP, or Ethernet service portfolios and more fully optimize the performance of their broadband infrastructure.

Optical layer survivability becomes ever more imperative as the high-bandwidth, mission-critical applications driving enterprise and government network growth stress the limitations of legacy protection and restoration schemes. Op-

erationally proven in Tier 1 networks around the world, intelligent optical mesh architectures enhance network resiliency and service availability with advanced technologies including ASON/GMPLS optical signaling and routing, intelligent optical routing algorithms, and multilayer (IP, Ethernet, and optical) interoperability. By deploying these optical networking technologies, network operators can be better prepared to maintain continuity of service—even in the face of disaster scenarios that have already starkly demonstrated the vulnerabilities of legacy technologies and protection schemes. ●

Kevin Oye is vice president of systems and technology at Sycamore Networks Inc. (www.sycamorenet.com).

Sycamore Networks, Inc. • 220 Mill Road • Chelmsford, MA 01824-4144, USA
Phone: 978-250-2900 • Fax: 978-256-3434
www.sycamorenet.com

Sycamore Networks, Inc. (NASDAQ: SCMR) is a leading provider of intelligent bandwidth management solutions for fixed line and mobile network operators worldwide. From multiservice access networks to the optical core, Sycamore products enable network operators to lower overall network costs, increase operational efficiencies, and rapidly deploy new revenue-generating services.

Sycamore assumes no responsibility for the accuracy of the information presented, which is subject to change without notice. Sycamore and Sycamore Networks are trademarks or registered trademarks of Sycamore Networks, Inc. in the United States and/or other countries.

Copyright © 2007 Sycamore Networks, Inc. All Rights Reserved.

