Converging optical IP: can GMPLS take control?

By Debanjan Saha
Manager
Interworking and Advanced Applications
Tellium Incorporated
E-mail: dsaha@tellium.com

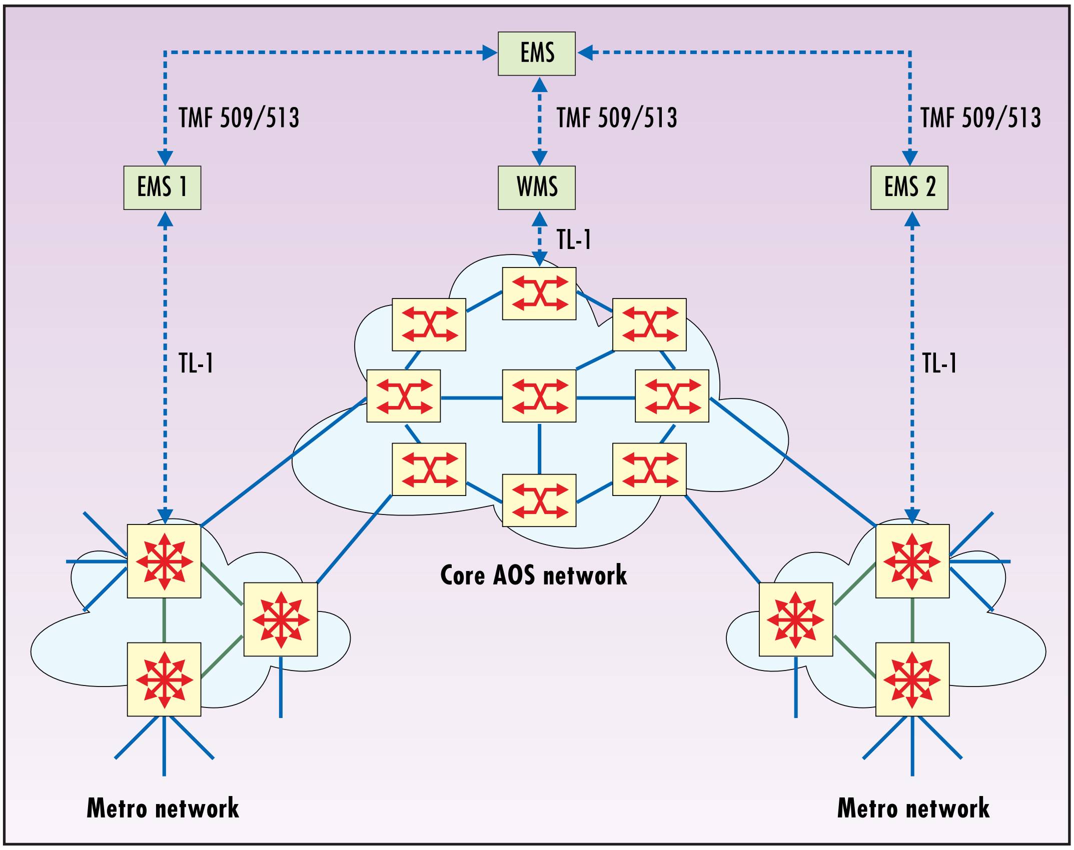

Figure 1: Management-plane approach is used in this mesh provisioning architecture.

Telecommunications equipment designers face a huge task in converging the optical-networking and Internet Protocol (IP) worlds to answer carrier demands for greater efficiency and improved cost-effectiveness. While the task may be huge, so too are the potential rewards, in the form of expanded next-generation services for customers and a more-favorable investment return for operators and providers.

Getting there, however, will require tighter integration of the IP and optical layers, as well as more standardization of protocols in today's telecommunications-switching environment than we have seen to date.

Optical-technology advancements continue to push the limits in terms of capacity and distance for long-haul networks, while ever-declining component prices are rendering optical deployment viable for metro and even access networks. Simultaneously, the demand for bandwidth has surged dramatically, with data having now surpassed legacy voice and private-line aggregates, giving prominence to IP. For that reason, convergence of the IP and optical layers is expected to be the defining theme in the next phase of Internet expansion.

However, to realize their greatest potential benefits, IP services in particular will need to be more intelligent, flexible and scalable. That will allow operators and providers to offer IP services beyond the less-profitable commodity category, where they reside today.

The promise of GMPLS

Efforts to simplify the layers underlie the developments at various standards groups. Generalized multiprotocol label switching (GMPLS), driven by the Internet Engineering Task Force (IETF), is one effort drawing attention.

GMPLS has evolved from the earlier multiprotocol label-switching (MPLS) framework and the optical user network (O-UNI) definitions in the Optical Internetworking Forum (OIF). GMPLS integrates Layer 2 information about network links–such as bandwidth, utilization and latency–into Layer 3 (IP). The tightened integration gives network operators flexibility to divert and route optical-network traffic around link failures, congestion and bottlenecks.

GMPLS is one of the options available for dynamic provisioning in mesh optical networks. The standards body efforts are helping define a common set of protocols for link management, topology discovery, routing, signaling and survivability across Internet Protocol and optical networks.

GMPLS is a natural extension of MPLS to the optical signaled overlay model and uses an open platform for the dynamic interconnection of multiple client layers, including IP. Its control architecture provides a simple and mature set of protocols. As a result, a full range of network elements (NEs), looking ahead to all-optical switches, will be able to respond to dynamic traffic trends and network reconfiguration, using standardized signaling.

The new intelligence that GMPLS brings to the control plane will also help with network management. With it, many of the advanced intelligent resource-engineering concepts developed for IP networks, such as constraint-based routing and traffic engineering, are finding their way into the optical arena. In addition, the wider universal control plane within GMPLS will allow providers and operators to create a much broader range of dynamic services.

GMPLS can also address two of the key tasks in network operations and growth: provisioning and restoration.

Provisioning methods

There are two paths to mesh provisioning: the centralized management plane, based on Tele Management Forum (TMF) standards, and the distributed control plane, based on GMPLS.

The two approaches contain some important differences in the main provisioning components (neighbor discovery, topology discovery, route computation and light path setup). A key to efficiency in next-generation networks will lie in the optimal use of GMPLS within those provisioning components.

The management-plane approach uses a centralized management architecture that is hierarchical. At the core of the network, element management systems (EMSes) control the cross-connects. Each sub-network level, such as metro, is controlled by a separate EMS. All EMSes report to an overall network management system (NMS) at the top of the hierarchy.

Using the centralized management-plane architecture for provisioning, neighbor discovery is largely a manual process in the EMS. Each connection and port is identified and entered into the system. Once the node- and port-identifier information is available and configured, the topology is compiled from the neighbor discovery data. This process, too, is largely manual in most management systems.

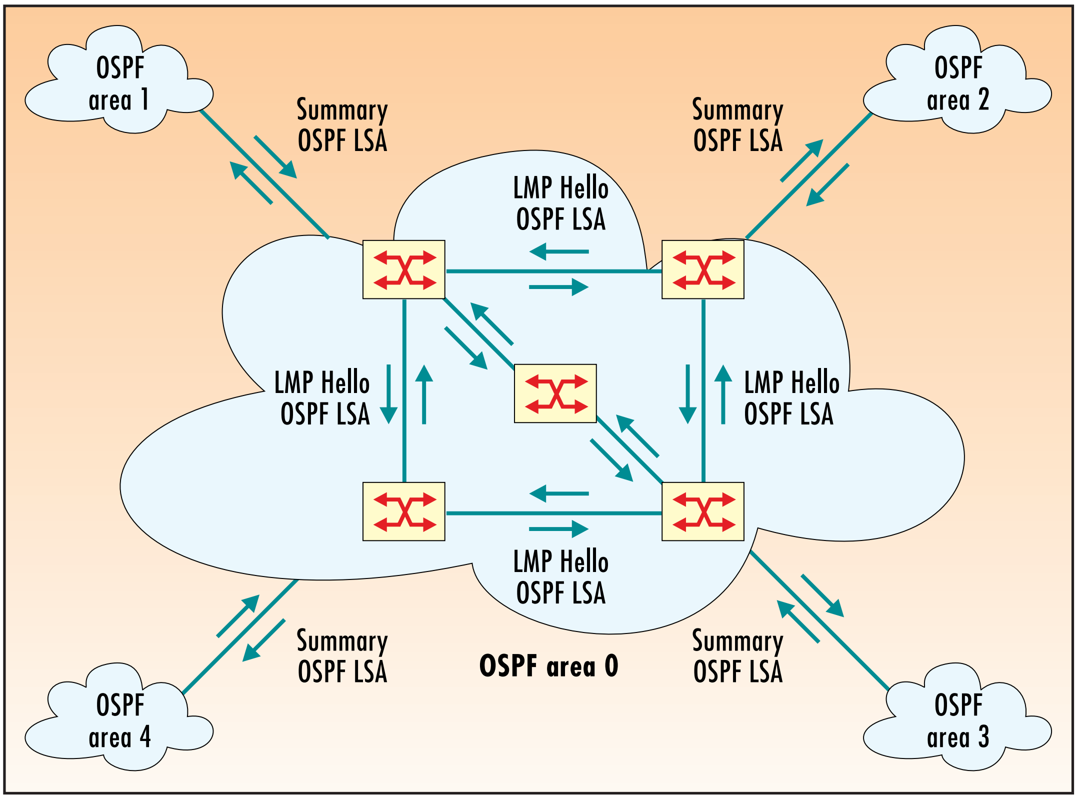

Figure 2: Control-plane approach topology used in a mesh provisioning architecture.

From the topology database, the routes can be computed by the NMS/EMS, which compute the primary and backup data paths. These systems ensure that the paths are physical-resource-independent, so that the two paths do not share the same risk of failure. With the route information determined, the EMS will set up the light path in each network element and configure the network elements using TL-1, SNMP or other messages.

Contrasted with the management plane is the control plane, a distributed architecture for which GMPLS has evolved as a particularly good solution for optical-mesh networks. In addition to offering greater network scalability because of its distributed nature, the control-plane architecture permits automation of labor-intensive, error-prone processes.

In the arrangement, the nodes use a new link-management protocol (LMP) to automatically detect and collect the port-state information from node neighbors and create the port-state database without manual intervention. This is good because the separation distances within the core, as well as within the subnetworks, create logistical problems for operators. The autodiscovery process takes far less time than manual discovery, allowing rapid service provisioning.

With the proliferation of fiber in central offices (COs) and remote sites, autodiscovery also avoids the errors in manual port data collection. A side benefit of autodiscovery is that it can be used to debug fiber wiring mistakes.

For topology discovery in the GMPLS-based distributed control-plane architecture, the discovery data is broadcast to the other nodes in the network. That is done using an interior gateway protocol (IGP), such as open shortest path first (OSPF) or intermediate-system to intermediate-system (IS-IS), with extensions for GMPLS. As a result, every node has a copy of the entire topology database in its control software.

Route computation is performed at each network element, not from a centralized location. When the request comes to a node to compute a path to any other node, the local node intelligence is able to consult the local copy of the topology database.

The NE then computes the route that the data will take through the node to other nodes, and sets up the light path, using signaling protocols. Two such protocols associated with GMPLS are the resource reservation setup protocol with traffic-engineering extensions (RSVP-TE) and the constraint-based label distribution protocol (CR-LDP).

Comparing architectures

The management-plane architecture and control-plane architecture offer valuable comparisons, which can be used to construct an optimal provisioning control approach that is efficient, fast and cost-effective.

Looking at neighbor discovery, the chief issue with the management-plane approach is that it is manual, long, painstaking and error-prone. With the control-plane architecture, there is no manual intervention, and the process is achieved automatically using the neighbor-discovery protocol LMP between and among network elements. GMPLS and the control-plane architecture have a clear advantage for neighbor discovery. Topology discovery is also automated in the control-plane architecture.

One important factor differentiates route-computation methods for the two architectures. In the management-plane architecture, the NMS computes routing from all the information for the topology, light paths and physical-resources inventory, stored in one central location. That means the light paths computed in this architecture will be more efficient than those computed within a control-plane architecture.

The reason for this has to do with the distributed nature of the control-plane architecture's neighbor, topology and resource data. With distributed databases, not all the light-path-level information required to optimize the route computation is available. As a result, the management-plane architecture will produce a more efficient and optimized route computation that allows operators to better use their available resources around the network.

The two architectures are comparable in the light-path setup process. While the control plane employs signaling protocols among all the network elements, the management plane accomplishes light-path setup with hierarchical configuration down from the umbrella NMS. While the mechanisms are different, the complexities are similar, and the result is the desired light path in either case.

The best of both planes

The distribution of relative strengths among the different parameters suggests an overall solution that can leverage the best of both the control- and management-plane architectures described.

Such an architecture would essentially combine certain control-plane and management-plane components into a hybrid scheme. Neighbor discovery would take place in a control-plane architecture to leverage GMPLS-based self-discovery by the network elements using LMP, with the benefits of automation and speed. That information would then be transmitted to a higher-level EMS in a central repository.

Topology discovery then becomes a centralized management-plane function, instead of data's being broadcast to all the elements in the network. That, in turn, permits routing and light-path setup by the EMS, taking advantage of the greater efficiency and accuracy the management plane offers.

The overall goal is to provide optimal efficiency in deploying switched optical mesh networks, and the GMPLS-based control-plane architecture can

	TMF Management plane	GMPLS/G.ASON Control plane
Scalability	Good	Very good
Robustness	Good	Good
Interoperability	Standards in progress	Standards in progress
Provisioning speed	Good	Very good
Routing efficiency	Very good	Good
Manual intervention	High	Low

Figure 3: How control- and management-plane approaches stack up.

contribute to the provisioning process if a hybrid scheme is used. Either of the centralized management- and distributed-control-plane approaches could be used, depending on the consistency of the network's characteristics. For a consistently high dynamic network, the control-plane architecture may be a good fit. For a network that is consistently less dynamic, the management plane would suffice.

For restoration in next-generation networks, the speed of restoration is consistently most prized, so the higher-speed control plane wins out. This is an area where further standardization will be beneficial. For now, however, the attention is on the provisioning part of the picture, since it is more difficult to solve. That could slow adoption, if the benefits to both provisioning and restoration are not evaluated jointly early on.

Another factor affecting the adoption rate of newer technology has to do with culture. In the past, users were familiar with the management-plane architectures, whereas control-plane techniques have emerged from the IP world. Depending on which technology is moving into which territory, the timeline for adoption could be affected.

Nonetheless, GMPLS provides an important connection between the optical and IP layers to permit open and scalable expansion in next-generation networks. Groups within the OIF and IETF are working together to specify and develop a unified GMPLS standard. The International Telecommunications Union is also defining a control-plane approach, using the proposed G.ASON standard for automatically switched optical networks. The standard bodies need to work closer to achieve more-unified standards.

As the standards evolve and mature, providing increased efficiency and better allocation of resources, operators will be able to make much better use of their investments. They will then be able to provision new kinds of revenue-producing services, such as bandwidth on demand, and offer more-flexible service-level-agreement capabilities. □

[Communication Systems Design]