GUEST EDITORIAL

Generalized Multiprotocol Label Switching (GMPLS) is a family of IP-based control plane protocols devised within the Internet Engineering Task Force (IETF), and intended to offer an advanced and uniform control plane for a variety of network technologies from packet switching networks, through layer 2 and time-division multiplexing (TDM) networks, to lambda and fiber switching. Significant drivers for the development of GMPLS were advances in wavelength-division multiplexing (WDM) technology and the promise of a unified control plane for the new optical networks.

A question often asked is why there has not been much traction with GMPLS protocols in the optical market. Is the downturn in capital expenditures the reason for the derailment of GMPLS adoption? Are finances the only concern, or are there technical and organizational issues? There is an argument that signaled end-to-end connections, whether via GMPLS, frame relay, or asynchronous transfer mode (ATM), have never really succeeded in the market. Operators prefer to retain precise manual control over the way traffic is placed within their networks rather than relinquish that control to the software and protocols in the control plane, no matter how "intelligent" that control plane is supposed to be. Furthermore, there has been little incentive for Service Providers to use a dynamic provisioning technique in networks where physical resources are in any case dedicated to the services they support.

However, what customers really would like to have is a pay-per-use service with rapid delivery of the service and without any requirement for the price model to be reviewed by the Service Providers' business development managers. In this service model the Service Provider may have resources that could be unused for a significant proportion of time, and with the increasing popularity of mesh network topologies, it is possible for those resources to be used for other services on a dynamic basis. This turns idle network capacity into revenue-generating resources. A control plane may provide the basic building blocks to facilitate this mode of operation.

GMPLS, because the protocols are extensions of IP/MPLS, and unlike traditional TDM transport protocols, can interact directly with IP and MPLS, allowing the transport network to be optimized for data-centric communication and even eliminating the use of ATM and TDM layers for the transport of IP over dense WDM (DWDM) networks. The promise was that with the new all-optical transparent lambda switches, the design of GMPLS carrier networks would be revolutionized.

The savings from simplifying operations, and the additional revenue from new on-demand optical services such as lambda on demand and optical virtual private networks (VPNs) would be reason enough to cut over from traditional transport networks to this brave new world. However, GMPLS is in direct competition with legacy proprietary protocols. This has created obstacles to deployment in several different ways; for example, the inability to upgrade legacy equipment, different technology requirements, and established operational procedures. One solution that has been proposed to overcome these obstacles is to utilize GMPLS within a network model where protocol interfaces are placed between network elements that use different protocols.

These interfaces (the user–network interface at the edge of the network and the network–network interface within the network) were developed as extensions to the GMPLS protocols. However, while there were several early field trials and interoperability demonstrations, there has been little actual deployment traction for the UNI and NNI. With Service Providers eager to reduce operating expenditures (OPEX), the network management system/operations support system (NMS/OSS) management/provisioning model took center stage again.

But GMPLS is not dead. As deployments of new equipment begin to grow again, and as DWDM gains support to carry more traffic on existing fibers using opto-electronic and transparent optical switching, the cost savings enabled by control plane technologies have encouraged Service Providers to start planning again for all-optical islands using routers running GMPLS to provide the control and data management. It is clear that GMPLS as an established, standardized, and vendor-independent technology does have a place in these future networks.

Research and carrier organizations have been exploring service applicability for GMPLS deployments and building experimental networks, some of which are carrying live traffic. Organizations are looking at protocols to control optical transport networks and leverage new developments in optical technology that optimize data architectures and handle unpredictable data traffic, and they are pursuing how protocols can better integrate IP and optical networks by examining integrated control for data and transport, and producing differentiated service and business models. A service such as bandwidth-on-demand requires resource reservation functions to manage and reserve lambda capacity. The customer must have a simplified order interface for the order, and billing must be automated and simplified for operations.

The objective of this month's Series is to present a collection of six articles that show the current state of thinking with regard to the standardization and deployment of GMPLS.

The issue begins with "The Influence of GMPLS on Network Providers Operational Expenditures: A Quantitative Study" by A. Iselt, S. Pasqualini, A. Kirstädter, R. Chahine, and S. Verbrugge. This article is a valuable contribution to the debate over the value of an intelligent control plane for OPEX reductions. By introducing a model by which operator costs can be quantified, the authors are able to demonstrate the benefits of using GMPLS and go on to claim the potential to halve the OPEX in most operational models.

In the second article, "GMPLS User–Network Interface (UNI) in Support of End-to-End Rerouting" by D. Papadimitriou and D. Verchere, the authors describe a popular deployment architecture that places a protocol interface (the UNI) between the user network and the core optical network. The article examines how the GMPLS UNI developed within the IETF can meet the high level of functional requirements for end-to-end protection and restoration demanded in modern network deployments, and makes a quantitative comparison of the effectiveness and efficiency of the GMPLS UNI with the UNI developed by the Optical Interworking Forum (OIF).

The third and fourth articles deal with an emerging concept of great importance: multilayered networks under the control of a uniform control plane. In "Multilayer Traffic Engineering for GMPLS-Enabled Networks" by B. Berde, M. Vigoureux, L. Andersson, and M. Jäger, the multilayer network is examined as a traffic engineering problem. The authors pose the question of what benefits may be derived from a single GMPLS control plane used to drive distinct switching layers while presenting the network as a coherent whole for the purpose of placing and routing traffic. They answer this by listing the properties of unified traffic engineering and comparing them with previous techniques in separated network layers to reach the conclusion that the deployment of an integrated multilayer network under the control of GMPLS can provide savings in network efficiency and speed of service provisioning.

In "Benefits of GMPLS for Multilayer Recovery," B. Puype, J.-P. Vasseur, A. Groebbens, S. De Maesschalck, and D. Colle explore the impact of a multilayer network on protection and recovery, and show how GMPLS, as a single distributed control plane, can deliver rapid and end-to-end provisioning of services across layer boundaries that can be exploited to deliver higher and more cost-effective service recovery functionality. This fact is supported by example case studies the authors use to demonstrate the potential opportunities and issues still to be resolved.

The fifth article in the series, "Layer 1 Virtual Private Networks: Driving Forces and Realization by GMPLS" by T. Takeda, D. Brungard, D. Papadimitriou, and H. Ould-Brahim, articulates the reasoning behind the development of layer 1 VPNs (L1VPNs). In the past, L1VPNs have been presented as a potential use for a GMPLS control plane, but many people have remained unconvinced by the business case for such network models and services. This article addresses these concerns by detailing the motivations for L1VPNs and describing how they might be used in real network scenarios. The authors conclude by examining the current state of GMPLS: how GMPLS can be used to build L1VPNs, and what further protocol development work is needed.

The final article is "GMPLS Operations and Management: Today's Challenges and Solutions for Tomorrow" by T. Nadeau and H. Rakotoranto. This provides a very welcome examination of the issues and challenges facing the user of GMPLS in deployed networks. In particular, this article looks at management requirements developed by Service Providers who have deployed or are about to deploy GMPLS networks, and examines the motivations that lie behind these requirements. The authors then move on to examine some of the solutions available to meet the provisioning, configuration, performance monitoring, and operations, administration, and maintenance requirements. The article concludes by promoting the role of standardization in the production of a coherent and full-function set of operations and management tools for future GMPLS networks.

In conclusion, we note that far from being abandoned in a theoretical back alley, GMPLS is very much alive and well. Furthermore, GMPLS is experiencing massive interest from vendors and Service Providers where it is seen as the tool that will bring together disparate functions and networks to facilitate the construction of a unified high-function multilayer network operators will use as the foundation of their next-generation networks. Thus, while the emphasis has shifted away from the control of transparent optical networks over the last few years, the very generality of GMPLS and its applicability across a wide range of switching technologies has meant that GMPLS remains at the forefront of innovation within the Internet.

BIOGRAPHIES

MONIQUE MORROW (mmorrow@cisco.com) is currently CTO consulting engineer at Cisco Systems, Inc. She has over 20 years' experience in IP internetworking that includes design and implementation of complex customer projects and service development for Service Providers. She led the Engineering Project team for one of the first European MPLS-VPN deployments in 1999 for a European Service Provider. She has been a speaker at numerous conferences related to MPLS, and additionally is co-author of the book *Designing IP-Based Services: Solutions for Vendors and Service Providers* (Morgan Kaufmann, 2002). She is currently working on three books, one exploring business aspects of MPLS, another discussing security for MPLSVPN, and the third presenting enterprise drivers and concerns for IP-based service delivery. She is active in both the IETF and ITU-T SG 13 with a focus on OAM. Additionally, she is working on GMPLS, GRID, and NGN topics pertinent to both carriers and NRENs. She has an M.S. in telecommunications management from Golden Gate University, San Francisco, California, and an M.B.A. from City University, Zürich, Switzerland. She is also a Cisco Certified Internetworking Engineer (#1711). She was co-guest editor of a special issue of *IEEE Communications Magazine* on OAM in MPLS-Based Networks (October 2004), and is co-guest editor for a feature topic in the same magazine's July 2005 issue, "GMPLS: The Promise of the Next-Generation Optical Control Plane."

MALLIK TATIPAMULA [SM] (mtatipam@cisco.com) received a B.Tech in electronics and communications engineering from Regional Engineering College, Warangal, India, and an M. Tech in communication systems and high frequency technologies from Indian Institute of Technology, Chennai (Madras). He is currently a senior product manager for next-generation networks in the Routing Technologies group at Cisco Systems. His expertise includes VoIP, mobile wireless, optical networking, IPv6, and GMPLS technologies. He works closely with Service Providers and national research networks around the world in deploying these advanced technologies in their next-generation networks. He has been with Cisco since 1998. Previously he worked at Motorola as a principal engineer, responsible for design of next-generation wireless networks. From 1993 to 1997 he was with Nortel Networks, Ottawa, Canada, as Senior Member of Scientific Staff, working on Nortel's optical and wireless products. He has over 15 years of experience in the telecom and networking industry. He has served on technical program and review committees of several leading IEEE and SPIE international conferences. He has appeared in *Who's Who in the World*, *Who's Who in Engineering*, and *Who's Who in America*. He is an editor of *Multimedia Communication Networks: Technologies and Services* (Artech House), and has authored several papers in IEEE and IEICE (Japan) journals as well as international conference proceedings. He has offered a number of tutorials at various international conferences and lectures at leading universities such as Stanford University, the University of California Berkeley, and the University of Tokyo.

ADRIAN FARREL (adrian@olddog.co.uk) is co-chair of the IETF's Common Control and Management Plane (CCAMP) Working Group that is responsible for the development of the GMPLS family of protocols. He also chairs the Path Computation Element (PCE) Working Group that is applying remote path computation techniques to MPLS and GMPLS networks. Building on his 20 years' experience designing and developing portable communications software, he runs a successful consultancy company, Old Dog Consulting, providing advice on implementation, deployment, and standardization of IP-based solutions, especially in the arena of MPLS and GMPLS. Before this he was MPLS architect and development manager at software house Data Connection Ltd., and director of protocol development for Movaz Networks Inc. As well as his activity within the IETF, where he has co-authored and contributed to numerous Internet drafts and RFCs on MPLS, GMPLS, and related technologies, he is active in many industry fora. A regular attendee at ITU-T meetings that discuss the optical control plane, he was also a founding board member of the MPLS Forum, is co-chair of the technical committee for MPLS 2005, and has served on many technical committees and review panels including the 2004 IEEE International Workshop on IP Operations and Management (IPOM 2004), and the IP and Optical Network Conference (iPOP 2005). As well as frequently speaking at conferences, giving tutorials on MPLS and GMPLS, and authoring several white papers on GMPLS, he is the author of *The Internet and Its Protocols: A Comparative Approach* (Morgan-Kaufmann 2004), which explains many of the IP-based protocols including those that make up MPLS and GMPLS. He is currently working on the text of a new book, *GMPLS: Architecture and Applications*, which is due out later in 2005.