GMPLS Controlled Ethernet: An Emerging Packet-Oriented Transport Technology

Attila Takács, Howard Green, and Benoit Tremblay, Ericsson

ABSTRACT

Most of the world's traffic is now packet-based. While SONET/SDH transport technology continues to evolve, there are now major standardization efforts to develop a native packet-oriented transport technology. Ethernet, long dominant in the enterprise, is one of the fastest developing technologies for the transport layer. GMPLS, based on mature signaling and routing protocols, is gaining traction as a transport control plane, providing fast restoration and supporting automation of provisioning. We will briefly highlight recent Ethernet standardization and then introduce the concepts and standards work allowing GMPLS to control Ethernet.

INTRODUCTION

The dominance of Ethernet in enterprise networking has led to high-volume components and well understood operational practices, and hence to low-cost solutions. This dominance is largely due to its "plug and play" features. On the other hand, the fundamental mechanisms[1] of plug and play have limited the scale of Ethernet network deployments to LANs.

Several solutions have been proposed to improve Ethernet scalability to allow its application in aggregation and metro core networks, and to provide the necessary dependability and operational features required by network operators. As more and more traffic is carried in Ethernet frames, operators are looking for an Ethernet-based packet transport solution that nonetheless retains the major virtues of SONET/SDH. Connection-oriented Ethernet (CO-Ethernet) has been proposed for this application [1]. It also seems a natural step to develop the generalized multiprotocol label switching (GMPLS) control plane, previously deployed in SONET/SDH and optical transport networks (OTNs), to support CO-Ethernet provisioning and reconfiguration.

Today, major standardization organizations — the IEEE, International Telecommunication Union — Telecommunication Standardization Sector (ITU-T), Metro Ethernet Forum (MEF), and Internet Engineering Task Force (IETF) — are active in the development of carrier Ethernet standards supporting metro area network (MAN) and wide area network (WAN) deployment of Ethernet. Currently, the standardization work is focusing on Ethernet features required for deployment in transport networks. In particular, the IEEE is working on provider backbone bridging — traffic engineering (PBB-TE), and the IETF is extending the GMPLS control protocols for Ethernet. In this article we first highlight the basic concept of packet-oriented transport and introduce the current alternatives. Then we turn our attention to CO-Ethernet, briefly describe the developments in the IEEE, and detail the GMPLS extensions under specification in the IETF. Finally, we report proof-of-concept work in this area.

OVERVIEW OF PACKET-ORIENTED TRANSPORT

LAYERED NETWORK ARCHITECTURE

A key concern for next-generation networks (NGNs) is the increased dynamism of the environment, and the unpredictability of new applications and usage patterns. Networks must grow and change more rapidly, but operational costs must fall. Management of this complexity is aided by layers of abstraction above the physical network. Today's networks make use of a "physical" transport layer. NGNs will extend this approach by providing "virtual topologies" allowing separation of concerns and timescales.

At the lowest layers, there are long-lived structures requiring civil works, way-leaves, and negotiations with third parties, which need long-term planning. At the highest layers, there are near-instantaneous creation and destruction of dynamic service instances for many users. In between is the transport layer, which provides a smoothed aggregate view of these demands. The unified architecture of transport networks is described in ITU-T G.800. The layers and their respective properties are illustrated in Fig. 1. The behavior of the IP network is contained within much slower moving "aggregate tunnels"

[1] *In particular, the "learning bridge" (MAC address learning, flooding of unknown frames) and Spanning Tree Protocol (for loop avoidance) do not scale well.*

 0163-6804/08/$25.00 © 2008 IEEE

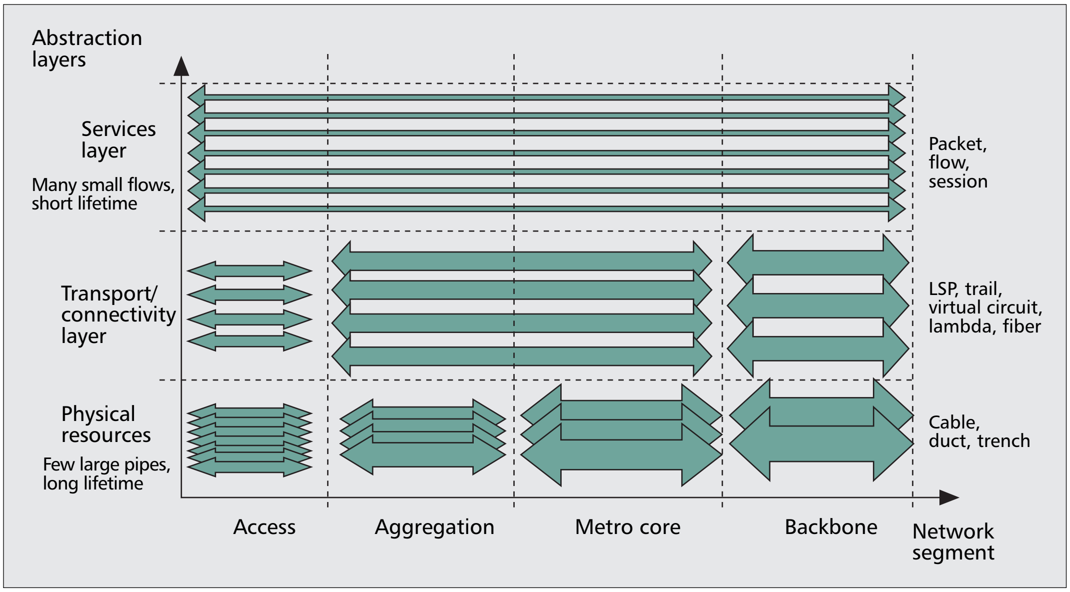

Figure 1. *Network abstraction layers.*

in the transport layer. The existence of a layer of aggregate virtual pipes allows the deployment of bulk recovery and restoration techniques, which would otherwise involve the destruction and recreation of much fine-grained flow state. This contributes to fast and predictable recovery.

A Control Plane for the Transport Layer

Today, transport layers are controlled by vendor-specific management systems. Such networks have achieved a high degree of stability and dependability. Nonetheless, there are disadvantages:

- There are significant constraints on the mixing of vendor equipment, because of the consequent requirement to interwork their management systems.
- Management systems cannot respond in real time. Hence management-based restoration of connectivity after failures involves significant downtime.
- Manual database input is prone to error. Thus, maximum use of auto-discovery is desirable.

There has long been discussion of the need for a standardized "transport control plane" to address these issues. ITU-T specified a general control plane architecture in G.8080. GMPLS is evolving to fulfill the requirements. GMPLS builds on the routing and signaling extensions of MPLS and specifies a strictly out-of-band control plane to support many data plane technologies, from wavelengths in WDM networks to asynchronous transfer mode (ATM) virtual circuits. It is currently deployed in large SDH networks/OTNs. The current justification of these deployments is automatic mesh network restoration, either as an alternative to or supplementing configured protection.

In the longer term, the job of a transport control plane is to support and automate the provisioning process to improve the responsiveness and dependability of the network, and to reduce its operating costs [2].

Current Packet-Oriented Transport Technologies

There are many packet transport alternatives in the marketplace, driven by technology and competition. Network operators are making different choices for reasons of installed base and service mix.

The SONET/SDH hierarchy of synchronous transport systems, and the associated performance monitoring, fast protection, and management infrastructure, is the backbone of current networks. It continues to be the benchmark in terms of protection switching and timing stability, and currently provides the fastest available network bearers at 40 Gb/s. SDH is now extensively used to carry Ethernet traffic, using the virtual concatenation and generic framing procedure (ITU-T G.7041). The OTN hierarchy and physical layer overhead, defined by ITU-T G.709, extend the key benefits of SDH to high-bandwidth optical bearers, and will be the base for many packet transport solutions.

MPLS started life to simplify forwarding of IP packets by adding a label, and from there it has developed in several directions. There is now a set of standard encapsulations, called pseudowires (PWs), to carry non-IP protocols. MPLS traffic engineering (MPLS-TE) is coming into widespread use (in particular for virtual private network [VPN] services). The ITU-T has initiated work on a specific data plane profile for transport applications. The ITU-T and IETF are now jointly working on the specification of the MPLS transport profile (MPLS TP), based on the PW and MPLS specifications.

At the same time as these developments, Ethernet has also been evolving. Recent standards making has focused on resolving "traditional" bridged Ethernet limitations for MAN and WAN deployment.

The job of a transport control plane is to support and automate the provisioning process to improve the resonsiveness and dependability of the network, and to reduce its operating costs.

Ethernet Developments in IEEE

The Ethernet bridging standard, 802.1Q, is being amended in the IEEE to equip Ethernet for large-scale network deployment. The new addi-

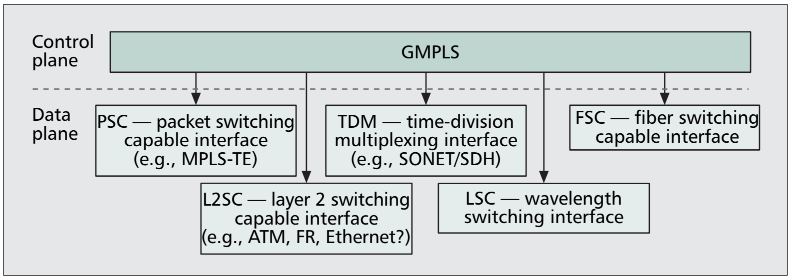

■ **Figure 2.** *The common GMPLS control plane.*

tions are provider bridging (PB), provider backbone bridging (PBB) [3], provider backbone bridging — traffic engineering (PBB-TE) [4], and connectivity fault management (CFM) [5].

PB and PBB enhance Ethernet scalability by introducing network hierarchy. With PB, a new virtual LAN (VLAN) tag, the service VLAN (S-VLAN), is introduced. It allows providers to use a separate VLAN space while transparently maintaining the customer VLAN (C-VLAN) information. PBB allows a full separation of the customer and provider address spaces by encapsulating customer frames, adding a "backbone" medium access control (MAC) header. Both the MAC addresses and the whole VLAN space are thus controlled by the provider. In addition to the backbone MAC header, a new tag, the service instance tag (I-TAG), is added when customer frames are encapsulated. The I-TAG includes a 24-bit service instance identifier (I-SID) field. The I-SID unambiguously identifies customer services. In PBB we can distinguish edge bridges (backbone edge bridge — BEB), which process customer frames and add the backbone MAC header and I-TAG; and core bridges (backbone core bridge — BCB), which forward frames using the backbone MAC header.

PBB-TE decouples the Ethernet data and control planes by explicitly supporting external control/management mechanisms to configure static filtering entries in bridges and create explicitly routed Ethernet connections. PBB-TE defines Ethernet connections, denoted Ethernet switched paths (ESPs). An ESP is identified by its destination (DA) and source (SA) addresses and VLAN identifier (VID); in short, a 3-tuple of <ESP-MAC DA, ESP-MAC SA, ESP-VID>. In addition, PBB-TE defines 1:1 protection switching of bidirectional Ethernet connections.

CFM specifies operations, administration, and maintenance (OAM) mechanisms for Ethernet networks. It defines continuity check (CC) to allow periodic liveness monitoring; loopback (LB) for on-demand failure verification; and link-trace (LT) for failure localization. CFM is also the basis of the Ethernet OAM extensions defined in the ITU-T Y.1731 specification.

GMPLS Controlled Ethernet

The GMPLS Control Plane

In MPLS the control plane and data plane are tightly coupled. The structure, processing, and location of the label in data packets are specific to the MPLS technology. Each IP packet is tagged with a 32-bit shim header used by MPLS routers to make forwarding decisions. There are two groups of MPLS signaling protocols: those concerned with label switched path (LSP) signaling, and those delivering routing information. Two protocols have been defined to signal LSPs: Label Distribution Protocol (LDP) and the Traffic Engineering extension of the Resource Reservation Protocol (RSVP-TE). Both the Open Shortest Path First (OSPF) and Intermediate System to Intermediate System (IS-IS) routing protocols have been extended to advertise traffic engineering (TE) information.

GMPLS supports various data plane types: packet switching (PSC), layer 2 switching (L2SC), time-division multiplexing (TDM), wavelength switching (LSC), and fiber switching (FSC) (Fig. 2). It exploits the capability of the RSVP-TE and OSPF-TE or ISIS-TE protocols specified for MPLS. In order to generalize the label switching concept, the data and control planes need to be separated. The control plane is based on IP protocols over an IP control channel, and carries the information necessary to control LSPs in the underlying data plane. Transport technologies may not support the addition of explicit labels in the data plane. Nonetheless, they have characteristics similar to the label in MPLS forwarding. For instance, in optical cross-connects (OXCs) the logical identifiers of distinct fibers or lambdas within a fiber can be viewed as out-of-band labels. SONET/SDH has a multiplexing structure that is built on time slots: here the level in the multiplexing tree can be interpreted as the label. In GMPLS these generalized labels are signaled in the control plane and used to configure the forwarding of data plane switching elements.

Since the control and data planes are separated, GMPLS adds a new protocol: the Link Management Protocol (LMP). LMP is responsible for neighbor discovery, maintaining control channel connectivity, verifying data link connectivity, correlating link property information, suppressing downstream alarms, and localizing link failures.

GMPLS Extensions for Ethernet

In the IETF the Common Control and Measurement Plane (CCAMP) working group is responsible for the definition of GMPLS. The CCAMP working group is extending the GMPLS control plane for PBB-TE Ethernet networks: the work is known as GMPLS controlled Ethernet label switching (GELS). It will enable the application of MPLS-TE and GMPLS provisioning and recovery features in Ethernet networks.

Ethernet is in origin a connectionless packet-switched technology: Spanning tree is used to create loop-free topology, and MAC learning with broadcast of unknown frames maintains dynamic forwarding entries. PBB-TE decouples the data and control planes, allowing other control mechanisms and forwarding paradigms in Ethernet networks. For packet-oriented transport applications, GMPLS is a natural choice. It creates explicitly routed Ethernet connections in the network, which thus operates in a connection-oriented packet-switched mode.

PBB-TE defines point-to-point Ethernet Switched Paths (ESPs) as provisioned traffic engineered unidirectional connections. To form a bidirectional PBB-TE connection, two co-routed point-to-point ESPs are combined. The combined ESPs must have the same ESP-MAC addresses but may have different ESP-VIDs. A PBB-TE connection established using the GMPLS control plane is called a bidirectional Ethernet-LSP. Figure 3 shows the relation of the 3-tuple identifying an ESP and the Ethernet label used in GMPLS for signaling an Ethernet-LSP.

To recall, GMPLS consists of the protocols LMP, RSVP-TE, and OSPF/ISIS-TE. To establish an Ethernet-LSP, only RSVP-TE requires extension. OSPF/ISIS-TE may be used unmodified, while LMP may not be needed at all. As with other technologies applying GMPLS, Ethernet control plane entities will use IP addresses and (usually) unnumbered interface identification. Hence, there is no need to carry MAC address and VLAN information in the routing protocols. However, optional advertisements in routing to optimize for Ethernet specifics may be proposed at a later time. The link fault management and discovery features of LMP are provided by native IEEE protocols, which have much richer functionality than the generic LMP set. Hence, even if LMP is deployed for Ethernet, it should rely on IEEE CFM and Link Layer Discovery Protocol (LLDP). LMP provides mechanisms to create and manage unnumbered interfaces; this is beneficial for large-scale deployment and can be used unmodified for PBB-TE.

RSVP-TE Extensions

The GELS framework [8] describes the general principles of GMPLS for Ethernet, while the technology-specific protocol extensions for PBB-TE are specified in [11].

Ethernet is in the scope of the L2SC switching type; however, most of the existing work around L2SC has been for ATM connections. In order to avoid incompatibilities, the IETF is considering the specification of a new switching type for CO-Ethernet. This discussion also needs to consider the Ethernet hierarchy introduced by the IEEE PB and PBB specifications. Although current standardization of PBB-TE only focuses on single-layer/single-domain networks, the introduction of a set of layer 2 switching types similar to PSC types used for packet switching may enable future flexibility. This issue has no significant impact on the other RSVP-TE mechanisms used for Ethernet-LSP signaling.

The vital part of the RSVP-TE extensions is the Ethernet label. For PBB-TE the label is defined by the tuple <ESP-MAC DA, ESP-VID>, as depicted in Fig. 3. It is processed according to GMPLS procedures, with the limitation that the same value must be assigned at each hop. That is, no label swapping is performed, and the Ethernet label has domain-wide significance. This is similar to the global data link connection identifiers (DLCIs) used in frame relay. Also, the same limitation applies in transparent optical networks, where wavelength continuity must be ensured throughout the network.

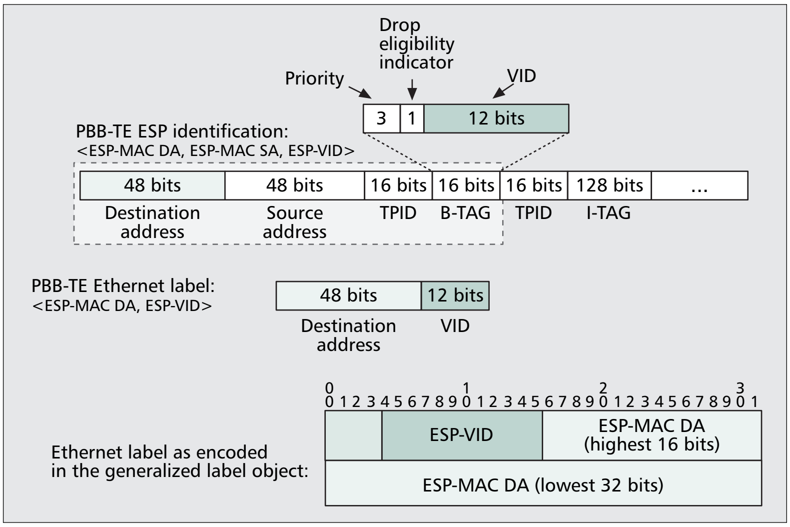

Figure 3. *The PBB-TE Ethernet label.*

Label allocation for PBB-TE must ensure the domain-wide uniqueness of Ethernet labels. Interestingly, this requires only a simple and efficient mechanism. Each BEB must maintain a local label allocation table for its local ESP-MAC address. This table includes all the ESP-VIDs already assigned for Ethernet-LSPs terminating at the BEB. If a new Ethernet-LSP is requested BEBs can locally allocate unused ESP-VIDs for the connection. This method requires about 4000 bit flags per local ESP-MAC address at each BEB.

In the case of LSP merging (called shared forwarding in [11]), the same ESP-VID can be reused for multiple LSPs, relaxing the need for unique ESP-VID assignment at BEBs. Accordingly, Ethernet-LSPs that have merged at some point in the network can only be assigned the same Ethernet label at the BEB if they do not subsequently diverge. To ensure this property the BEB needs to consult the explicit route object (ERO) or record route object (RRO) of each LSP when a request for LSP establishment with merging is received. Shared forwarding creates reverse BEB rooted forwarding trees, but at the same time it maintains deterministic resource reservation for each connection. It must be noted that merging as a result of LDP operation in MPLS is different from the controlled and resource aware merging achieved by RSVP-TE signaling [6]. In the latter case deterministic resource allocation for each connection can be guaranteed, while the former breaks the resource guarantees and thus violates the basic properties of connections.

The IETF is also considering how to support the MEF Ethernet service definitions with GMPLS. When creating a TE Ethernet connection, resources need to be set aside for the Ethernet-LSP. The MEF specified a set of traffic descriptors to be used for Ethernet services. Besides explicit bandwidth reservation these parameters can also be used when establishing an Ethernet-LSP supporting a specific service. Reference [7] defines Ethernet-technology-specific Sender_Tspec and Flowspec objects in RSVP-TE signaling to carry MEF-defined traffic

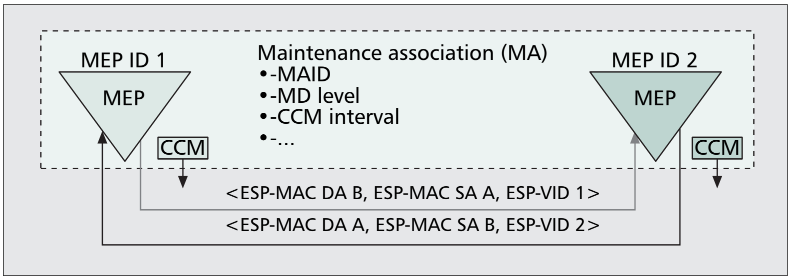

Figure 4. *Illustration of the relevant PBB-TE CFM entities.*

parameters. MEF also defines a user-network interface (UNI) as a generic interface for Ethernet service requests. The UNI 2.0 specification of the Optical Internetworking Forum (OIF) already supports the MEF Ethernet services, while in the IETF two drafts [9, 10] describe the RSVP-TE extensions needed to support the establishment of point-to-point Ethernet services over a GMPLS UNI.

A distinguishing property of transport networks is advanced OAM and protection mechanisms. Transport networks operate with stringent performance and availability requirements. Effective monitoring of service level agreements (SLAs) and rapid location of problems are critical to service quality, and also important to reduce operational cost. In SONET/SDH, since OAM information is carried in dedicated overhead channels, basic OAM mechanisms are inherent to connections and in operation once a connection is set up. We propose in [12] RSVP-TE extensions to equip CO-Ethernet with similar features by configuring the Ethernet OAM entities of Ethernet LSPs during connection setup.

ETHERNET-LSP SIGNALING WITH CONNECTIVITY MONITORING

Connectivity Monitoring of PBB-TE Connections — CFM defines an adjunct connectivity monitoring OAM flow to check the liveness of Ethernet networks. In PBB-TE, CFM is used to detect ESP data plane failures. CFM assigns a dedicated maintenance association (MA) for each PBB-TE connection. An MA is uniquely identified by the maintenance association identifier (MAID) and maintenance domain level (MD Level). Within an MA, maintenance association endpoints (MEPs) are located at the ends of the Ethernet connection. MEPs are responsible for generating and processing connectivity check messages (CCMs), link-trace and loopback messages, and replies. Each MEP has a unique identifier within its own MA: the MEP ID. MEPs assigned to a specific Ethernet connection are configured with the corresponding ESP 3-tuples, as shown in Fig. 4.

To monitor the connectivity of a bidirectional connection, MEPs exchange CCMs at fixed intervals. If a MEP fails to receive three consecutive CCM messages, it declares a connectivity failure and signals the failure to the remote MEP in subsequent CCM messages by setting the remote defect indicator (RDI) bit. If a MEP receives a CCM message with RDI set, it immediately declares failure. The detection of a failure may trigger protection switching mechanisms and be signaled to a management system.

Operators need to balance speed of failure detection against overhead; hence it is beneficial to configure the frequency of CCM messages per Ethernet-LSP. In the case of PBB-TE, data misconnectivity can be detected by examining the ESP-MAC SA and I-SID at connection endpoints. This relaxes the need for frequent OAM flows for LSPs without strict protection requirements.

To simplify network management and reduce the risk (and impact) of misconfiguration, Ethernet-LSP signaling can be used to configure CFM entities at both ends of the LSP. This requires signaling the MA and MEP configuration parameters in RSVP-TE using a new OAM configuration object.

Ethernet-LSP Signaling Overview — We briefly describe below the setup of a point-to-point constrained path bidirectional Ethernet-LSP together with the establishment of Ethernet OAM entities for connection liveness monitoring.

First, an ERO specifying the path of the Ethernet-LSP is constructed by a local or remote TE module or path computation element (PCE). For CFM a unique MAID must be allocated for the PBB-TE connection. The MD level and desired CCM interval must be specified by the management system, based on service requirements and operator policies. The initiator BEB sends an RSVP-TE Path message requesting a generalized label for an Ethernet-LSP. It includes the upstream label, formed by its local MAC address (ESP-MAC A) and locally selected VID (ESP-VID 1), the Ethernet-specific Sender_TSpec object, CFM configuration information, and other objects, like the ERO (Fig. 5a).

When the remote node receives the RSVP-TE Path message, it extracts the reachability information for the initiator BEB from the upstream label. Then it allocates a "downstream" Ethernet label formed by its local MAC address (ESP-MAC B) and a VID (ESP-VID 2) selected from the local VLAN allocation table, which it will send in the subsequent RSVP-TE Resv message. By now, all information about the ESPs of the PBB-TE connection is available at the BEB: it is used, together with the CFM configuration information received in the RSVP-TE Path message, to configure the OAM entities (Fig. 5b).

Once the RSVP-TE Resv message successfully arrives at the initiator BEB, it extracts the remote side's reachability information from the Ethernet label. Now this BEB too has obtained all information about the ESPs of the PBB-TE connection and is able to establish related OAM entities (Fig. 5c).

Intermediate nodes process the RSVP-TE messages as usual, according to [13], with the addition that the same Ethernet label must be used unchanged at each hop. The ESP-MAC DA and ESP-VID values contained in the Ethernet label are installed in the bridge filtering tables as static forwarding entries.

PROOF-OF-CONCEPT IMPLEMENTATIONS AND TESTS

OUR EXPERIMENTAL SYSTEM

We have implemented an experimental testbed to validate the feasibility of controlling a PBB-TE network with GMPLS and experiment with the required protocol extensions. Our testbed is depicted in Fig. 6. It contains two PBB-TE edge nodes (BEBs) and two core nodes (BCBs). For demonstration, a streaming server was connected to one of the edge nodes by a set of S-VLANs, and two video clients were connected to the other edge node using port based interfaces.

BCBs are implemented using off-the-shelf PB switches. Since no commercial switches supporting BEB functionality were available when this project was started, we implemented BEBs using network processors. CFM continuity check is implemented on BEBs, allowing the creation of MEPs to monitor the end-to-end status of established ESPs. This is used to trigger protection switching or restoration mechanisms.

The GMPLS stack used in the testbed was derived from Ericsson's existing implementation for the SDH/OTN product line. The actual GMPLS extensions/modifications implemented include the Ethernet label [11], signaling extensions for Ethernet traffic parameters [7], and Ethernet OAM control [12]. The UNI portions of the services are manually configured. On the BEBs, the GMPLS stack runs on a general-purpose CPU included in the bridge. However, for the off-the-shelf BCB an external Linux machine handles GMPLS signaling, and a proprietary protocol is used to control the bridge.

The testbed is used to evaluate and demonstrate different protection scenarios: unprotected, rerouting, and 1:1 protection. Monitoring of the LSP is set up using the OAM control extension. This allows monitoring the path using CCM messages and transmission of events to the GMPLS stack in case of failure.

The 1:1 protection scheme creates a worker and a protection Ethernet-LSP. The ESP-MAC addresses used in the labels are the same for both LSPs, but different ESP-VIDs are assigned. During normal operation traffic is directed to the worker LSP. CFM is used to detect Ethernet-LSP failures and trigger protection switching. The MEP at the other end of the LSP will either detect the failure (missing CCM messages) or receive a CCM message with the RDI bit set. Both cases will trigger automatic protection switching in the data plane. A notification is also sent to the control plane for it to keep track of the actual status of the traffic path. With our testbed we proved the feasibility of Ethernet protection switching within less than 50 ms, which is comparable with the performance of SONET/SDH networks.

The next steps will be to prepare for an interoperability test with other implementations. This will ensure that the extensions to GMPLS are well understood, stable, and provide all the required information to control PBB-TE Ethernet networks.

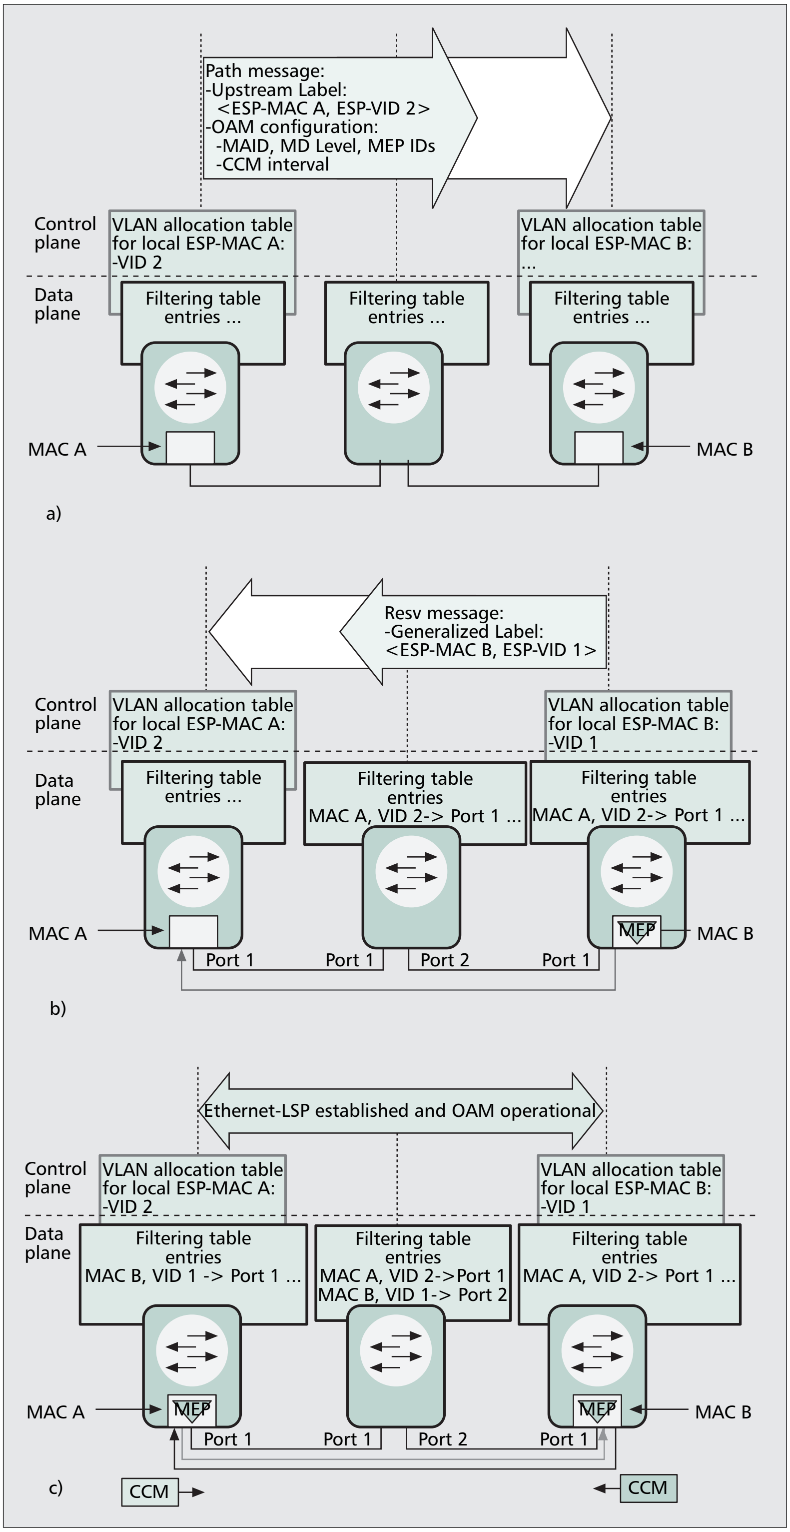

■ **Figure 5.** *Setup of a bidirectional Ethernet-LSP with connectivity monitoring.*

OTHER RELATED RESEARCH PROJECTS

The first adopters of GMPLS controlled Ethernet were research networks. The implementations and evaluations were driven by nationally and internationally funded research projects, including those below.

The Dynamic Resource Allocation in GMPLS Optical Networks (DRAGON) project defines and implements an experimental network to sup-

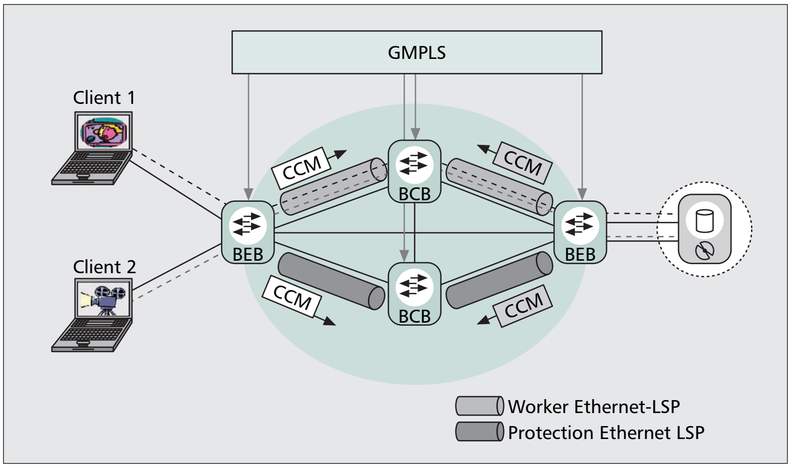

■ **Figure 6.** *Testbed setup.*

port on-demand high-capacity interconnect for grid computing and e-science applications. The DRAGON project utilizes the GMPLS control plane for provisioning of TE paths across heterogeneous network technologies. Although the main focus of DRAGON is optical networks, it also implements GMPLS extensions for VLAN-based Ethernet connections.

The European project Multi-Partner European Testbeds for Research Networks (MUPBED) implements and maintains a multinational research testbed that integrates and validates large-scale ASON/GMPLS operation and network solutions. One of the various testbeds implemented is based on Ethernet.

The Together IP, GMPLS and Ethernet Reconsidered (TIGER) project sponsored by the CELTIC initiative tries to improve the adaptation capabilities between IP and Ethernet layers utilizing the GMPLS control plane. Besides defining a framework, TIGER also benchmarks and compares the proposed connectionless Ethernet, CO-Ethernet, and IP/MPLS-based solutions. One of the main areas of the project is extending the usage of GMPLS control capabilities to the Ethernet data plane.

Conclusion

The IEEE carrier Ethernet extensions (PB, PBB, PBB-TE, and CFM) address the scalability and dependability requirements of public network applications. GMPLS (already used in transport applications) supports the automation of provisioning and fast reconfiguration after failure, improving the quality and responsiveness of provisioning and reducing operational costs. GMPLS-controlled Ethernet shows promise as a low-cost high-quality packet-oriented transport solution. The basic extensions to GMPLS required for PBB-TE are well understood and in progress in the IETF. Further opportunities exist to enhance the control plane (e.g., by integrating support for Ethernet connection monitoring). Initial proof-of-concept implementations are available, and future interoperability testing will ensure the maturity of the technology.

References

[1] D. Allan *et al.*, "Ethernet as Carrier Transport Infrastructure," *IEEE Commun. Mag.*, Feb. 2006.
[2] S. Pasqualini *et al.*, "Influence of GMPLS on Network Providers' Operational Expenditures: A Quantitative Study," *IEEE Commun. Mag.*, July 2005.
[3] IEEE 802.1Qah, "Provider Backbone Bridging."
[4] IEEE 802.1Qay, "Provider Backbone Bridging Traffic Engineering," work in progress.
[5] IEEE 802.1ag, "Connectivity Fault Management."
[6] S. Yasukawa and A. Farrel, "Supporting Multipoint-to-Point Label Switched Paths in Multiprotocol Label Switching Traffic Engineering," Internet draft, work in progress.
[7] D. Papadimitriou, "Ethernet Traffic Parameters," Internet draft, work in progress.
[8] D. Fedyk, L. Berger, and L. Andersson, "GMPLS Ethernet Label Switching Architecture and Framework," Internet draft, work in progress.
[9] L. Berger, D. Papadimitriou, and D. Fedyk, "Generalized MPLS (GMPLS) Support for Metro Ethernet Forum and G.8011User-Network Interface (UNI)," Internet draft, work in progress.
[10] L. Berger, D. Papadimitriou, and D. Fedyk, "Generalized MPLS (GMPLS) Support For Metro Ethernet Forum and G.8011Ethernet Service Switching," Internet draft, work in progress.
[11] D. Fedyk *et al.*, "GMPLS Control of Ethernet PBB-TE," Internet draft, work in progress.
[12] A. Takacs and B. Gero, "GMPLS RSVP-TE Extensions to Control Ethernet OAM," Internet draft, work in progress.
[13] L. Berger *et al.*, "Generalized Multi-Protocol Label Switching (GMPLS) Signaling Resource ReserVation Protocol-Traffic Engineering (RSVP-TE) Extensions," IETF RFC 3473, Jan. 2003.

Biographies

Attila Takacs (attila.takacs@ericsson.com) joined Ericsson Research in 2001. He has been working with inter- and intradomain traffic engineering, resource management, and cross-layer optimization for media delivery systems. At present, he works with the standardization of GMPLS controlled Ethernet. He has a Master's degree in computer science and a post-graduate degree in banking informatics both from the Budapest University of Technology and Economics.

Howard Green joined Ericsson by way of the Marconi acquisition in 2006, and is currently working in broadband and transport research. He has been in the industry since 1980, working in research, development, and business strategy for many technologies, from public switching to SDH and photonics. He has a Ph.D. in mathematics and an M.B.A., both from Warwick University, United Kingdom.

Benoit Tremblay collaborates with Ericsson for more than 10 years. He has been involved in developing distributed servers for mobile networks. He joined Ericsson Research in 2002, where he focused on establishing experimental systems for the IP edge, access, and broadband networks. He graduated in mathematics with specialization in computer science from the Université du Quebec a Montreal in 1986 and completed a Master's egree in 1988. His current topics of interest are broadband network and packet systems.