# Layer 1 Virtual Private Networks: Driving Forces and Realization by GMPLS

*Tomonori Takeda, NTT; Deborah Brungard, AT&T Labs*
*Dimitri Papadimitriou, Alcatel; Hamid Ould-Brahim, Nortel*

## ABSTRACT

This article describes an emerging service for next-generation networks, layer 1 virtual private networks. L1VPNs allow customers desiring to connect multiple sites to be supported over a single shared layer 1 network. In the article we first describe the transport network's evolution and the shift in expectations of both service providers and customers. We provide an overview of the motivation for L1VPNs and examples of network usage. We follow by reviewing existing GMPLS mechanisms (addressing, discovery, and signaling) for realizing L1VPN functionality and identifying other work areas.

## INTRODUCTION

Traditionally, transport networks provided information transferring links to client (voice) switching networks. With the growth of packet-based services and the accelerating pace of service applications requiring higher bandwidth, transport networks are evolving to support on-demand switched transport services.

Today's networks are based on synchronous optical network/synchronous digital hierarchy (SONET/SDH), controlled and managed typically by element management systems (EMSs) and network management systems (NMSs). Management systems not only have associated development and capital costs for a carrier, they often delay the timely introduction of new capabilities. Furthermore, service provisioning often requires manual support across multiple systems and may require weeks or even months to complete. This traditional type of network operation is at the risk of not only increased cost, but also the loss of business opportunity.

Technologies such as optical crossconnects (OXCs) and SONET/SDH crossconnects equipped with distributed control plane protocols are being evaluated for solving these issues in the next-generation network. Standardized control plane protocols provide operators with automated operations, distributed at the equipment level, to manage transport resources and provide multivendor interoperability. Generalized multiprotocol label switching (GMPLS) is a suite of control and measurement plane protocols that are being defined within the Internet Engineering Task Force (IETF), developed from their MPLS expertise, covering packet, Ethernet, SONET/SDH, and optical [1]. Previous articles (e.g., [2]) have described GMPLS and the value of using a control plane.

Today's large carriers are providing multiple services, such as layer 3 (IP) virtual private network (VPN) service, layer 2 (e.g., Ethernet) VPN service, traditional (point-to-point) private line services, Internet access service, and Internet transit service, typically all over a single, shared transport network. The service networks and the transport network are often owned and operated by different internal organizations of the carrier. Layer 1 VPNs (L1VPNs) are drawing attention in order to flexibly support this type of shared network architecture. L1VPNs enable internal clients to utilize the transport network as if dedicated to themselves, without needing to build and operate their own dedicated network. L1VPNs also enable providers to support novel service offerings, creating new revenue opportunities.

International Telecommunication Union—Telecommunication Standardization Sector (ITU-T) Study Group (SG) 13 has standardized L1VPN service requirements and high-level architecture [3, 4], and new work is being discussed in the IETF mainly on protocol aspects [5, 6]. Indeed, there are various protocol mechanisms in support of L1VPNs already [7, 8]. The work by ITU-T SG13 has been summarized in an existing article [9]. The objective of this article is to introduce L1VPNs in the context of recent work in ITU-T and IETF. The article discusses from a carrier's perspective the driving forces for L1VPNs and reviews the ongoing protocol evaluation work needed to support L1VPNs. This arti-


0163-6804/05/$20.00 © 2005 IEEE

cle first investigates network architecture evolution, and describes the motivation for L1VPNs and network usage scenarios. It then covers how L1VPNs can be realized by GMPLS in terms of addressing, discovery, and signaling aspects.

## Network Evolution and GMPLS

A guiding objective of next-generation networks is the capability to offer innovative services while reducing development, capital costs, and operations across the different services. The traditional network was functionally divided between switching and transport. Switching provided dynamic on-demand connectivity using distributed control protocols. Transport provided fixed bandwidth and fixed point-to-point connectivity. Service provisioning and equipment management capabilities ranged from use of manual provisioning (craft interfaces) to support of semi-automatic functions (e.g., "point-and-click" provisioning) via a carrier's proprietary management interface between a network element and a carrier's operations system. The capital expenditures (CAPEX) and operating expenditures (OPEX) of supporting proprietary interfaces were expensive for both the equipment vendor and the carrier, and often limited both to exclusive relationships. As the transport network expanded geographically and increased in capacity, and more functionalities were added, such as recovery and efficient grooming, network operations became more and more complex. Transport standards took on the task of specifying management interfaces, using the traditional three-tier approach: network element, EMS, and NMS. Standards work is still ongoing, and while providing standard open interfaces, the carrier continues to be very dependent on operations system(s) and the associated CAPEX and OPEX. In order to reduce this dependency, vendors and carriers are now evaluating using distributed (IP-based) control planes to support key operations capabilities such as automated inventory, end-to-end dynamic connection setup, and end-to-end distributed recovery schemes.

As the transport network evolved, the switching network was also changing. Voice was no longer the only service. Packet-based data services are being introduced, and a carrier's network architecture is evolving from a transport network providing point-to-point connectivity of voice switches to a network of multiple internal service organizations sharing a common transport network. Similar changes are also occurring at the customer premises. As routers scaled, router interfaces evolved from plesiochronous digital hierarchy (PDH) to SONET/SDH. Data network connectivity with the transport network is evolving from connections via multiple levels of multiplex equipment to converged layer 1 (e.g., SONET/SDH, optical transport network [OTN]) architectures of directly connected data equipment and OXC/wavelength-division multiplexing (WDM). Expectations of customers of the transport network are also changing. Data customers are expecting similar service functionality across layers 1 to 3: flexibility and convenience (e.g., the ability to add/delete capacity with ease). This shift from layer 3 to enhanced layer 1 service functionality is inevitable and will require dramatically different approaches for layer 1 management. Distributed control plane technology provides not only new capabilities for managing the network; it also provides the ability to support innovative layer 1 services. This network evolution is summarized in Fig. 1.

GMPLS [1], as a superset of MPLS, provides an open, multivendor, standard-based approach allowing new end-to-end-based transport architectures with automatic topology discovery and network inventory, dynamic connection setup, and fast and efficient distributed restoration mechanisms. This distribution of control functionality at the network element level is diametrically opposed to today's transport networks; distributed control supports efficient, fast decision mechanisms at the local network element level compared to the traditional three-tier communication chain structures of network element, EMS, and NMS. And similar to MPLS, GMPLS has capabilities, such as routing information exchange based on trust policies, automatic configurations, secure message exchange, and private address support, to allow a carrier to support customized innovative services for its internal data customers and external customers.

This network evolution will redefine transport networks. As with any new technology, it has taken time to develop the standards needed for network deployment. IETF's Common Control and Measurement Plane (CCAMP) working group has focused on providing a comprehensive base set of standards to support vendor interworking. CCAMP is currently extending its work to support multiple applications, for example, recovery schemes, interdomain, automatically switched optical network (ASON), user–network interface (UNI), and L1VPNs. Similar to any new technology introduction, deployment in existing networks will be gradually introduced, with applications varying for each provider.

*GMPLS provides an open, multivendor, standard-based approach, enabling transport networks to provide connections to client networks on demand by distributed automatic control. L1VPNs define a service interface with connection control and management.*

## Overview of L1VPNs

As described above, GMPLS provides an open, multivendor, standard-based approach, enabling transport networks to provide connections to client networks on demand by distributed automatic control. L1VPNs define a service interface with connection control and management. In the following, merits of L1VPNs for customers and providers are summarized, and examples of network usage scenarios are described.

### Motivation

Typically, transport networks and client networks are built as follows:

- Multiple client networks are supported over a single shared transport network (e.g., both a carrier's MPLS network supporting IP VPN service, and their IPv4 and IPv6 networks supporting Internet traffic are supported over a single SONET/SDH network).
- The transport network and each client network are owned and operated by different organizations (either separate companies, e.g., other carriers, or divisions within the same company).

*L1VPNs define a service interface, which enables dynamic Layer 1 connection provisioning across this service interface. A client network becomes a customer network of L1VPN services, while a transport network becomes a provider network of L1VPN services.*

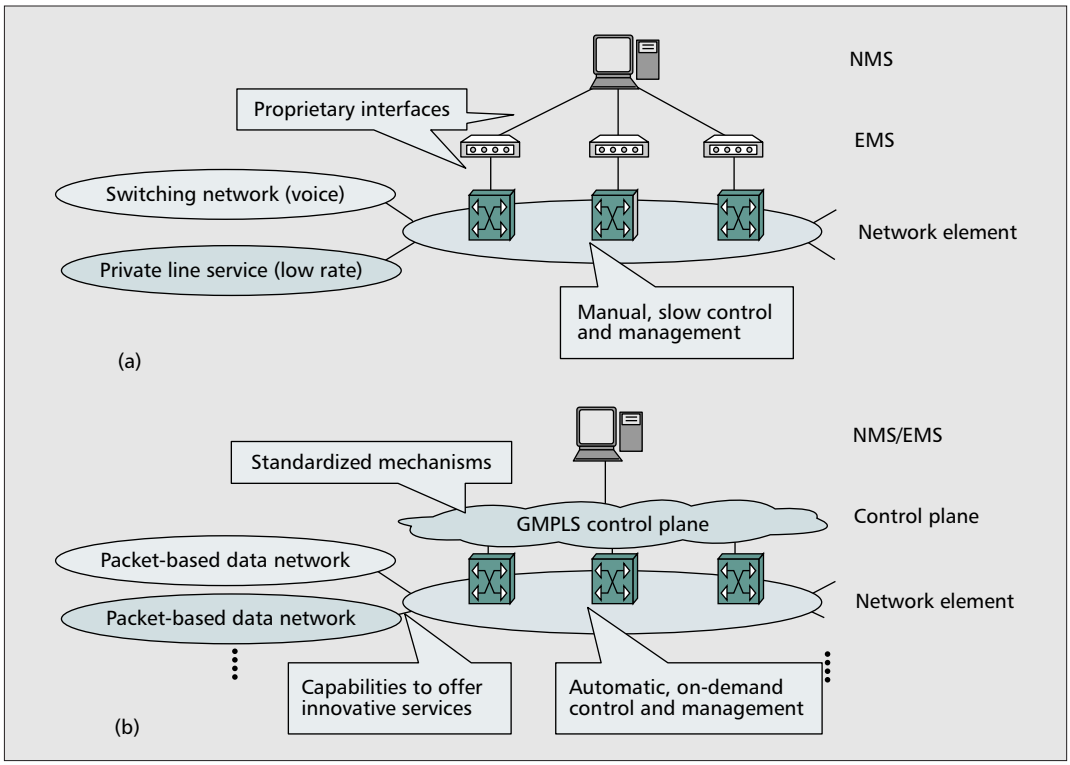


**Figure 1.** *a) Traditional transport network; b) next-generation transport network.*

The first point is obvious in terms of network multiplexing efficiency. The shared transport network provides bandwidth and capacity, and client networks use a portion, increasing their demand as needed.

The latter is a common approach for large carriers for several reasons. First, service and network operations for different client networks and the transport network are very different. For example, type of equipment, protocols used, and requirements such as recovery time and scalability are different. Second, network isolation is required for scalability and simplified operations. For example, client network 1 does not need to know topology information or failure information of client network 2; it also does not need to know full topology information or detailed failure information of the transport network. Third, evolution of the different service networks can occur independently. Lastly, the clients can share the transport network without building their own transport network, which leads to OPEX/CAPEX reduction.

In order to address this network architecture, L1VPNs define a service interface, which enables dynamic layer 1 connection provisioning across this service interface. A client network becomes a customer network of L1VPN services, while a transport network becomes a provider network of L1VPN services. As shown in Fig. 2, L1VPNs enable operation separation between the transport network and client networks. L1VPNs provide a logically secure separate network to each customer. Similar to closed user group (CUG) type services, customers can request a connection only between devices in the same VPN. In addition, customers are allowed to assign addresses in their own address space, and access topology, resource, and failure information related to their VPN, but not other VPNs.

Merits of L1VPNs for customers and providers can be summarized as follows:

- Customers can utilize a transport network virtually dedicated to them without building their own transport network (CAPEX reduction). At the same time, direct management of this transport network is outsourced to the provider, where operational cost of the transport network can be shared across multiple customers (OPEX reduction). Customers can provision connections on demand to support unpredictable traffic increase between any pair of sites.
- The provider can offer dynamic on-demand service, with the offloading of provisioning procedures to automatic control. They can support internal clients or sell their resources to create new revenue.

Note that in current deployment practice, based on traffic estimation for each pair of sites, equipment is ordered and installed. However, as service bandwidth needs change rapidly, traffic estimation will become more difficult. Deployment of excess spare equipment will be inevitable in order to allow quick response to new service orders. L1VPNs allow more rapid service response and efficient use of the network. A connection can be provisioned between an arbitrary pair of sites (i.e., utilizing mesh end-to-end connectivity) based on actual service needs.

## Network Usage Scenarios

There are various network usage scenarios where L1VPNs can be used. Depending on customer requirements, GMPLS usage within the customer network, and the trust relationship

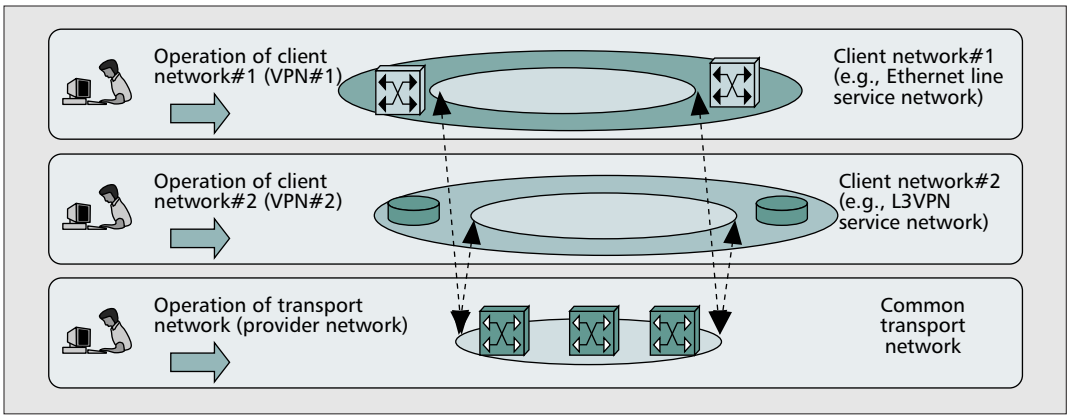


**Figure 2.** *Concept of L1VPNs.*

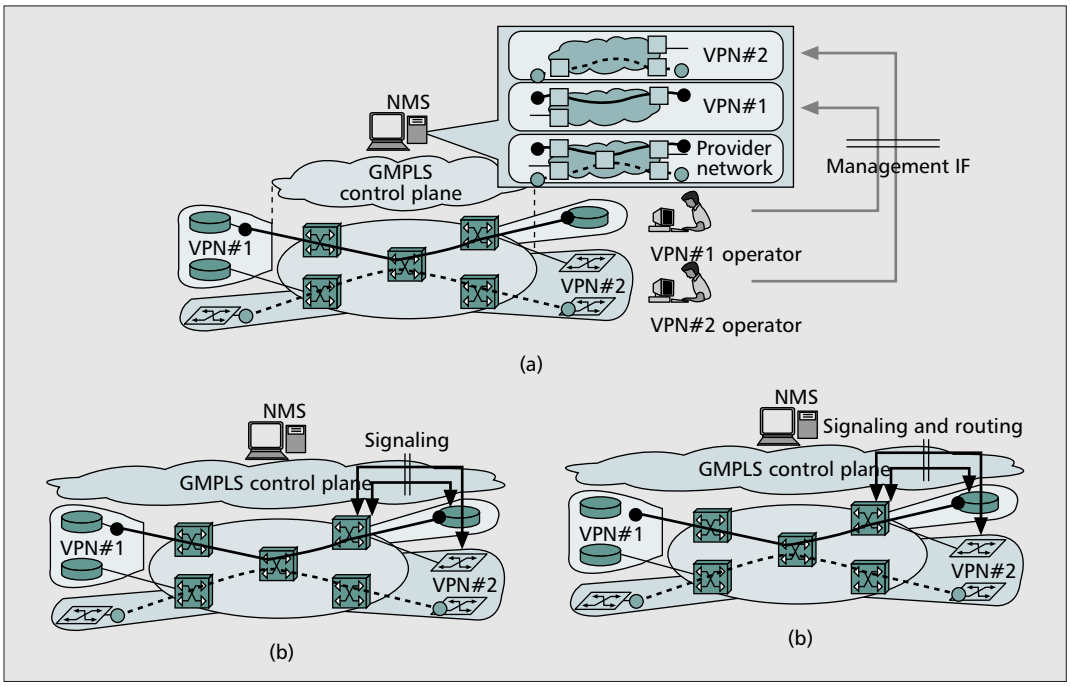


**Figure 3.** *a) Management-based service model; b) signaling-based service model; c) signaling and routing service model.*

between customer and provider, services may be accessed through different methods, and different functionalities may be given to customers. There are three major service models, as shown in Fig. 3.

As an initial step, even when GMPLS is not implemented in customer networks, customers may wish to utilize L1VPN services (e.g., they may wish not to change operational procedures drastically from current practice). In such a case, by enhancing the NMS in the provider network, L1VPN service can be provided. Customers access restricted resources of the provider network via a management interface (e.g., Web interface), and control and manage their portion of the network. This model is called a management-based service model.

By implementing GMPLS signaling functions in customer edges (CEs), customers can request connection setup/deletion/modification to the provider by signaling. In GMPLS a connection is called a label switched path (LSP). By using GMPLS between the CE and the provider edge (PE), interoperability can be secured. GMPLS also provides mechanisms for rapid notification upon LSP failure, even if the data plane technology does not support an equivalent mechanism. This model is called a signaling-based service model.

The first and second models are an overlay approach, where CEs sets up routing adjacencies over connections. This leads to a so-called *N*-square routing problem. If not only CEs, but also customer sites are operated by GMPLS, this problem can be solved by forming GMPLS routing adjacencies between the CE and PE (or more precisely the L1VPN private context instantiated on the PE). The CE can obtain remote site routing information from the PE. Note that in this scenario, provider routing infrastructure is completely hidden from the client routing plane. In addition, abstracted topology information of the provider network may be provided to the CE. If GMPLS is used within cus-

*As an initial step, even when GMPLS is not implemented in customer networks, customers may wish to utilize L1VPN services, for example, they may wish not to change operational procedures drastically from current practice. In such a case, by enhancing the NMS in the provider network, L1VPN service can be provided.*

A customer can seamlessly operate its VPN using end-to-end GMPLS. This model is called a signaling and routing service model. Forming routing adjacencies between the CE and the PE may risk confidentiality, thus use of this model will depend on the trust relationship between the customer and the provider.

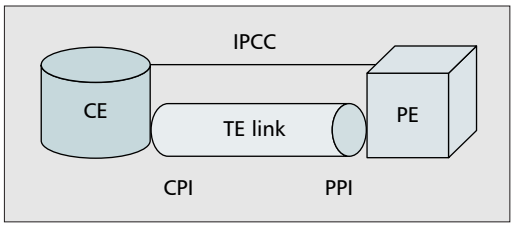


**Figure 4.** *Addressing scheme in L1VPNs.*

tomer sites, all customer internal nodes, not just the CE, can benefit from the L1VPN topology information. This allows greater and more effective traffic engineering (TE) on the customer network. A customer can seamlessly operate its VPN using end-to-end GMPLS. This model is called a signaling and routing service model. Note that forming routing adjacencies between the CE and PE may risk confidentiality (especially when TE-related information is exchanged between the CE and PE); thus, use of this model depends on the trust relationship between the customer and the provider.

In the first model, connections are soft-permanent. Since the CE and PE do not exchange any control plane protocol message, VPN functions can be implemented just in the management plane, not in the control plane.

In the second and third models, connections are switched. Since the CE and PE exchange control plane protocol messages, VPN functions need to be implemented in the control plane.

## GMPLS Mechanisms for L1VPNs

As mentioned in previous sections, there are already defined mechanisms for implementing VPN functions using IP/MPLS and GMPLS protocols. In order to set up a VPN connection, VPN reachability (addresses/discovery) information needs to be propagated, followed by connection setup procedures (signaling). VPN reachability information propagation can be accomplished by using a VPN auto-discovery mechanism based on L3VPN mechanisms, and connection setup can be accomplished using GMPLS Resource Reservation Protocol — TE (RSVP-TE) signaling.

### Addressing

A CE is identified by one or more TE links that connect the CE to the PE. A TE link is a control plane representation of (an aggregated set of) physical resources. Each TE link connecting the CE to the PE is associated with a unique identifier within a given VPN, the customer port identifier (CPI), and a unique identifier within the provider network, the provider port identifier (PPI). The CPI and PPI can be numbered or unnumbered. By using CPI-PPI associations, the provider network can map a CPI to a PPI, which corresponds to a unique address within the provider network.

In addition to its PPI, each TE link terminating on a PE also has an identifier that is unique within the VPNs, the VPN-PPI. The PE IP address used for VPN-PPI need not be the same as that used for the PPI.

Between any CE-PE pair, at least one channel allowing IP connectivity between the CE and PE is supported, referred to as an *IP control channel* (IPCC). Figure 4 summarizes the addressing scheme used in L1VPNs.

### L1VPN Auto-Discovery Mechanism

With recent developments in provider-based VPN architectures such as MPLS-based layer 2 and layer 3 VPNs, new techniques have been defined that considerably reduce the operational complexity in managing and configuring VPN services. A key component of these techniques currently being standardized in IETF is called the VPN auto-discovery mechanism. The main objective of a VPN auto-discovery mechanism is to allow VPN members to dynamically discover appropriate information to be used for setting up intersite connectivity. Indeed, previous approaches (e.g., CUG configuration models) usually required intensive configurations when a new site was added to a VPN. All the devices in the network having ports that are members of that VPN needed to be configured with a list of VPN members with which it needed to connect, the connectivity matrix, and client-provider address information. A VPN auto-discovery mechanism allows that the configuration of addition/changes of a new site are limited to only the devices that are attached to that site. The auto-discovery mechanism distributes the information (i.e., addition of a new site) to all PEs that need to be aware of the new member. The distribution requires no operator intervention.

For the purpose of using an auto-discovery mechanism in the context of L1VPNs, PEs that have VPN configured will be given locally:

- The VPN membership of each (logical) port attached to a CE (identification of address space of CPI)
- The port information or, to be more precise, the port addressing identifiers expressed within both the VPN addressing space (CPI) and the provider addressing space (PPI)
- Optionally, the connectivity topology attribute of each port (i.e., whether this port is a spoke, hub, full mesh, or part of arbitrary connectivity topology)
- Optional information related to port compatibility, bandwidth information, and so on

The VPN auto-discovery mechanism for L1VPNs can be implemented using centralized server-based techniques or distributed control-plane-based techniques. Centralized server-based techniques require that a server(s) is configured with a list of VPN members and their corresponding list of provider-client addresses (e.g., CPI-PPI associations); each CE and PE will have to access this server in order to request new connections. Distributed control plane techniques piggyback VPN discovery information onto the client and/or provider-based control plane. For example, the distributed control-plane-based technique using the Multiprotocol Border Gateway Protocol (MP-BGP) auto-discovery mechanism, initially used for L3 MPLS-based VPNs and then extended for MPLS and

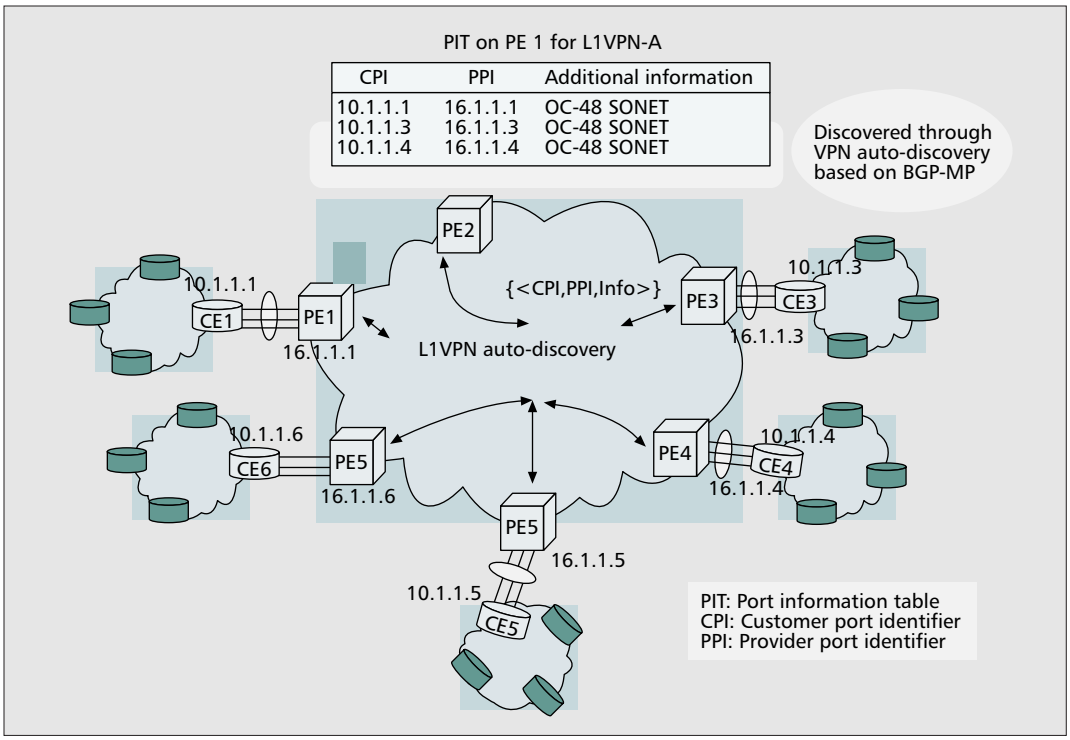


**Figure 5.** *Example of VPN auto-discovery based on BGP-MP [8].*

*The VPN auto-discovery mechanism for L1VPNs can be implemented using centralized server-based techniques or distributed control-plane based techniques. Centralized server-based techniques require that a server(s) is configured with a list of VPN members and their corresponding list of provider-client addresses.*

IP-based L2VPNs, can be used for L1VPNs (defined in [8]). One characteristic of using MP-BGP for auto-discovery function is that it inherits the mechanisms already existing within the BGP standard. MP-BGP provides for selective distribution by not distributing the VPN information to devices that have no VPN in common or no VPN configured on them. Another advantage is fast reaction to new changes in a VPN as additions or changes are propagated via the efficient information distribution of BGP.

As an added value feature, the provider network may allow a CE to dynamically discover all (or a subset of) remote CE members of the same VPN. This is referred to as CE-based discovery. A CE-based discovery can be used in conjunction with a provider-based discovery scheme. Once a PE discovers the set of remote CE ports, it will pass that information to the attached CE. The CEs will use this information to select the CEs to which they want to initiate connectivity. Protocols that can convey to the CE the set of remote CE information can be BGP, Link Management Protocol (LMP), and so on.

### Signaling Mechanism

Connections between a pair of CEs are point-to-point, and can be soft-permanent (established between local PPI and remote PPI) or switched (established between local CPI and remote CPI). For a customer-driven soft-permanent connection, the provider network is responsible for the switched provider network connection. A L1VPN can comprise a set of soft permanent connections and/or switched connections. In the following, details for switched connections are provided.

Once a CE obtains the information about the remote CPIs belonging to the same VPN, the CE initiates connection requests using GMPLS RSVP-TE signaling for one or more LSPs to the desired CPI. In GMPLS RSVP-TE, Path and Resv messages are used to set up a connection, the former from the ingress to the egress to request an LSP and the latter from the egress to the ingress in reply.

A fundamental capability for L1VPNs is to support use of private addresses for CPIs. This raises the problem of how to uniquely identify an LSP. Specifically, GMPLS RSVP-TE identifies an LSP by combining the ingress and egress addresses contained in SESSION and SENDER_TEMPLATE/FILTER_SPEC objects. If private addresses are assigned to the CPIs, it is not possible for the provider network to uniquely identify an LSP. To support private address assignment, either of two approaches may be used:

- Shuffling (Fig. 6a): Information carried in RSVP messages identifying an LSP (i.e., SESSION and SENDER_TEMPLATE/FILTER_SPEC objects) is translated by the ingress and egress PE [8].
- Nesting (Fig. 6b): When a Path message arrives at the ingress PE, the ingress PE checks whether there is appropriate PE-to-PE connectivity. If there is not, it initiates a PE-to-PE LSP. This LSP is called a forwarding adjacency (FA) LSP. On top of this FA-LSP, a CE-to-CE LSP is set up [7–10].

A particular case of LSP nesting is LSP stitching, where the properties of an LSP segment are such that exactly one end-to-end LSP can be stitched with the LSP segment (i.e., the PE-to-PE LSP and CE-to-CE LSP correspond exactly one to one). This implies that no label exchange occurs between the head-end and tail-end of the LSP stitching segment compared to the manda-

*A particular case of LSP nesting is LSP stitching, where the properties of a LSP segment are such that exactly one end-to-end LSP can be stitched with the LSP segment, that is, the PE-to-PE LSP and the CE-to-CE LSP correspond exactly one to one.*

tory label exchange that happens in the LSP nesting case.

Detailed procedures for shuffling are described as follows:

In the Path message originated by the ingress CE, the ingress CPI value is inserted into the SENDER_TEMPLATE object and the egress CPI value is inserted into the SESSION object. When the ingress PE receives the request, the PE identifies the appropriate PIT (e.g., from IPCC receiving the request), and then uses the information in that PIT to find out the appropriate PPI associated with the CPI of the egress CE. The PPI should be sufficient for the PE to establish an LSP. Once the mapping is retrieved, the ingress PE replaces the ingress/egress CPI values with the PPI values. As a result, the SESSION and SENDER_TEMPLATE objects included in the GMPLS RSVP-TE Path message carry PPIs, not CPIs. At the egress PE, the reverse mapping operation is performed.

When the Path message reaches the egress CE and gets processed, the latter initiates toward the ingress the exchange of Resv message. Resv messages are processed similar to Path messages. Once the Resv message reaches the ingress CE, the switched connection is established. The provider network may remove/filter information (e.g., recording route information contained in Record Route Object [RRO]) to allow hiding of the internal topology of the provider network for security reasons [7].

## Other Work Areas

In order to realize L1VPN services in a fully functional manner, there are several other work areas related to the control plane protocols,

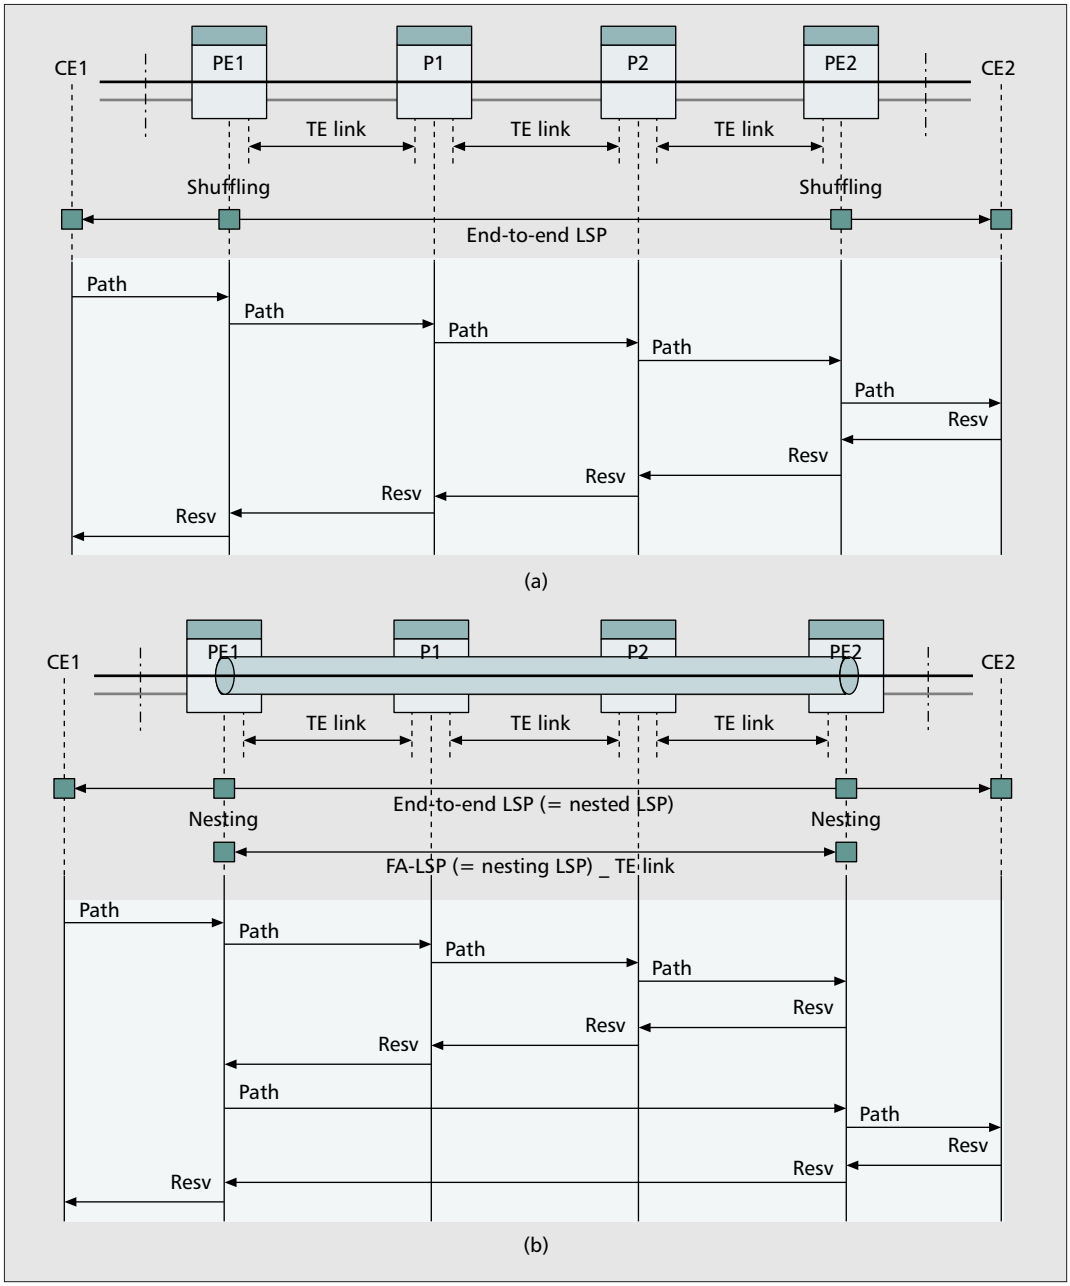


**Figure 6.** *a) Shuffling approach; b) nesting approach.*

management functionalities, and operational procedures.

From the control plane standpoint, most of the protocol components are already available as described above, but there are some remaining issues. One is to clarify under which conditions signaling approaches (shuffling or nesting) should be used. Another fundamental issue is that BGP is not widely deployed in optical equipment. For optical networks that do not want to use BGP for the purpose of L1VPN auto-discovery and still want to use a distributed mechanism, a potential solution is to piggyback L1VPN information onto a link state protocol such as Open Shortest Path First (OSPF) or Intermediate System to Intermediate System (ISIS).

These are also several issues around management functionalities, such as:

- Per-VPN resource management
- Security management
- Accounting
- VPN configuration management

Lastly, in operational aspects GMPLS and automatic control protocols will change network operations, for example, from semi-automatic (management) provisioning to distributed automatic provisioning, and from network planning based on point-to-point topology and pre-ordered connection-specific provisioning to network planning based on mesh topology and on-demand dynamic provisioning. Network operators will need new operational procedures, training, and support tools to achieve the benefits of this new technology.

## Conclusion

This article describes how L1VPNs can bring merits to customers and providers in the next-generation network. It also describes how L1VPNs can be realized by GMPLS, taking into consideration the standardization activities in IETF.

L1VPNs are applicable where a shared transport network supports multiple client networks. L1VPNs enable customers to share the use of a layer 1 network and enable providers to support both internal and external clients with a common service platform. L1VPNs support the provisioning of new services with greater speed and efficiency, creating new revenue opportunities. In order to realize the full set of VPN functionalities, management and operational procedures will also need to be enhanced. With these new capabilities and experience, L1VPNs are expected to be one of the key services in the next-generation transport network.


### Acknowledgments

The authors would like to thank the co-authors of the IETF L1VPN drafts. Special thanks to Adrian Farrel for useful discussions.


## References


[1] E. Mannie, "Generalized Multi-Protocol Label Switching (GMPLS) Architecture," IETF RFC 3945, Oct. 2004.

[2] A. Banerjee *et al.*, "Generalized Multiprotocol Label Switching: An Overview of Signaling Enhancements and Recovery Techniques," *IEEE Commun. Mag.*, vol. 39, no. 7, July 2001, pp.144–51.

[3] ITU-T Recommendation Y.1312, "Layer 1 Virtual Private Network Generic Requirements and Architecture Elements," Sept. 2003.

[4] ITU-T Recommendation Y.1313, "Layer 1 Virtual Private Network Service and Network Architectures," July 2004.

[5] T. Takeda *et al.*, "Framework for Layer 1 Virtual Private Networks," IETF Internet draft, work in progress.

[6] T. Takeda *et al.*, "Applicability Analysis of GMPLS Protocols to Layer 1 Virtual Private Networks," IETF Internet draft, work in progress.

[7] G. Swallow *et al.*, "Generalize Multiprotocol Label Switching (GMPLS) User-Network Interface (UNI): Resource Reservation Protocol-Traffic Engineering (RSVP-TE) Support for the Overlay Model," IETF Internet draft, work in progress.

[8] H. Ould-Brahim and Y. Rekhter, "GVPN Services: Generalized VPN Services Using BGP and GMPLS Toolkit," IETF Internet draft, work in progress.

[9] T. Takeda *et al.*, "Layer 1 Virtual Private Networks: service concepts, architecture requirements, and related advances in standardization," *IEEE Commun. Mag.*, vol. 42, no. 6, June 2004, pp.132–38.

[10] K. Kompella and Y. Rekhter, "LSP Hierarchy with Generalized MPLS TE," IETF Internet draft, work in progress.


## Biographies


TOMONORI TAKEDA (takeda.tomonori@lab.ntt.co.jp) received an M.E. degree from Waseda University, Tokyo, Japan, in 2001. In 2001 he joined NTT, Tokyo, Japan. Currently, he is with NTT Network Service Systems Laboratories, where he is engaged in R&D on the next-generation network architecture, IP optical network architecture, and related protocols. He has been involved in several standardization activities, including ITU-T and IETF.

DEBORAH BRUNGARD [M] (dbrungard@att.com) is a technical consultant at AT&T Labs, Middletown, New Jersey. Since joining AT&T Bell Laboratories in 1984, she has worked for AT&T in the United States and The Netherlands. Her early work included research and development on SDH transmission systems, international network planning, and international private line service development. She is currently working as an optical networking technology architect in support of AT&T's strategic standards development. From 1998 to 1999, she was Vice-Chair of T1X1.5, the ATIS Committee on Optical and Digital Hierarchy Standards. From 2000 to 2003 she was Chair of T1X1.5. She has been an editor for multiple ITU-T Recommendations and IETF drafts, and a member of various IETF GMPLS Design Teams. She received her M.Eng. in electrical engineering in 1984 from Stevens Institute of Technology, Hoboken, New Jersey. She is a member of Tau Beta Pi and Eta Kappa Nu.

DIMITRI PAPADIMITRIOU (dimitri.papadimitriou@alcatel.be) received an M.S. degree in physics from the Universite Libre de Bruxelles and an M.S.degree in computer science from the Universite de Liège. After one year of collaboration on network protocols R&D at the Research Center of Namur, he worked for three years in the design and evaluation of packet (IP/MPLS), Ethernet, and circuit (SDH/optical) transport networks as well as their compliance to IETF/ITU standards. In 2000 he joined Alcatel as an expert research engineer in the Network Technology and Analysis team of the Network Strategy Group. In 2004 he joined the Packet-Transport Interworking (PTI) team of the Research and Innovation (R&I) CTO department. His current areas of interest are focused on distributed IP control planes, IP routing protocols, as well as traffic engineering and QoS mechanisms for connectionless and flow-oriented technologies. He is actively involved in the standardization activities of the IETF Routing Area, where he edits and co-authors numerous working group documents on G/MPLS-based control planes and related protocol aspects. He has also authored numerous technical papers on routing and recovery techniques for sub-IP networks. He is currently involved in the Design and Engineering of the Next Generation Internet Network of Excellence (EuroNGI) and leading the GMPLS control plane architecture as well as the exploitation and dissemination activities of the ITEA TBONES European project.

HAMID OULD-BRAHIM (hbrahim@nortel.com) is a senior protocol architect at Nortel and a renowned industry expert of provider-based VPNs. He contributed to the definition of a wide range of VPN concepts and mechanisms, and has been a key driver of generalized VPNs in IETF.


*In order to realize the full set of VPN functionalities, management and operational procedures will also need to be enhanced. With these new capabilities and experience, L1VPNs are expected to be one of the key services in the next-generation transport network.*