Layer 1 Virtual Private Network Management by Users

Jing Wu, Michel Savoie, Scott Campbell, and Hanxi Zhang, Communications Research Center Canada
Bill St. Arnaud, CANARIE Inc.

ABSTRACT

The Layer 1 Virtual Private Network (L1VPN) technology supports multiple user networks over a common carrier transport network. Emerging L1VPN services allow: L1VPNs to be built over multiple carrier networks; L1VPNs to lease or trade resources with each other; and users to reconfigure an L1VPN topology, and add or remove bandwidth. The trend is to offer increased flexibility and provide management functions as close to users as possible, while maintaining proper resource access right control. In this article two aspects of the L1VPN service and management architectures are discussed: management of carrier network partitions for L1VPNs, and L1VPN management by users. We present the carrier network partitioning at the network element (NE) and L1VPN levels. As an example, a Transaction Language One (TL1) proxy is developed to achieve carrier network partitioning at the NE level. The TL1 proxy is implemented without any modifications to the existing NE management system. On top of the TL1 proxy, a Web Services (WS)-based L1VPN management tool is implemented. Carriers use the tool to partition resources at the L1VPN level by assigning resources, together with the WS-based management services for the resources, to L1VPNs. L1VPN administrators use the tool to receive resource partitions from multiple carriers and partner L1VPNs. Further resource partitioning or regrouping can be conducted on the received resources, and leasing or trading resources with partner L1VPNs is supported. These services offer a potential business model for a physical network broker. After the L1VPN administrators compose the use scenarios of resources, and make the use scenarios available to the L1VPN end users as WS, the end users reconfigure the L1VPN without intervention from the administrator. The tool accomplishes L1VPN management by users.

INTRODUCTION

The Virtual Private Network (VPN) technology enables the coexistence of multiple user networks over a common infrastructure. In a Layer 1 VPN (L1VPN), a carrier maintains a common transport network, and offers virtually dedicated transmission between a group of users. The L1VPN technology extends layer 2/3 packet-switching VPN concepts to circuit-switching networks [1], for example, Wavelength Division Multiplexing (WDM) networks, Time Division Multiplexing (TDM) networks, and so forth. The benefits of the L1VPN technology are twofold [2]: carriers benefit from reduced operating cost and charges for the L1VPN premium service, and users benefit from rapid network deployment with lower up-front costs and a reduced facility management cost. In this article L1VPN users are classified as super users and regular end users. We call the former L1VPN administrators and the latter L1VPN end users for simplicity.

New services for L1VPN management are emerging. The trend is to offer increased flexibility and provide management functions as close to users as possible, while maintaining proper resource access right control. New services allow L1VPNs to be built over multiple carrier networks, and to lease or trade resources with each other. With these services, L1VPNs are constructed by composing resources from different sources. In the operation of L1VPNs, new services allow users to manage L1VPNs. L1VPN administrators can partition or bond resources, create or delete end-to-end connections, and create complex topologies of interconnected L1VPNs. L1VPN end users can reconfigure an L1VPN topology, and add or remove bandwidth. The L1VPN management by users eliminates time-consuming service orders to carriers, thus greatly increasing the users' capability to manage their own leased resources.

L1VPN management by users is a solution for applications that have continuous large traffic flows between a set of known remote end points, where a mesh of L1VPN across multiple carriers is required. For example, a scientific experiment requires a large amount of continuous data (on the order of 1 Gb/s per source) be transferred from a few satellite dishes to a remote supercomputing centre. Each experiment lasts from a few days to a few weeks. The satel-

0163-6804/06/$20.00 © 2006 IEEE

	Resource-partition-based model	Domain-service-based model
Advantages	• Clear physical separation among different L1VPNs, minimize L1VPN users' concern on interference • L1VPN users are granted visibility and access to the management functions in the carrier networks for the allocated resources • An L1VPN interacts with a user's network based on either a client-server or a peer-to-peer relationship	• Carrier network resources are time-shared among L1VPNs, thus increase the possibility for re-using idle resources • The carriers only need to verify whether a given L1VPN is allowed to connect two carrier's edge node ports at a given time for a given bandwidth requirement. Therefore, it is more scalable.
Disadvantages	• Carrier network resources cannot be directly time-shared among different L1VPNs, thus may result in low resource utilization • Carriers need to keep track of partitions for each network element, that is, which L1VPN can use which resource partition. When there are many L1VPNs and the resource partitioning uses fine granularity, scalability is a challenge.	• L1VPN users do not have visibility into the carrier networks, and completely depend on the carriers to provision connections between carriers' edge node ports • An L1VPN can only interact with a user's network based on a client-server relationship, where the L1VPN functions like a virtual node or link • Contentions on time-sharing resources need to be solved by the carrier administrative policy

Table 1. *A comparison of the two models for the management of carrier network partitions for L1VPNs.*

lite dishes and the supercomputing centre are so geographically distributed that no single carrier has complete physical connections to build such a network. L1VPN management by users provides a solution by composing L1VPNs using resources from multiple carrier networks.

In this article, following an overview of L1VPN service and management architectures, we present carrier resource partitioning at the network element level. At the L1VPN management level, we implemented a tool for the L1VPN management by users, called the User-Controlled Lightpath Provisioning (UCLP) system. The functions of the UCLP system are presented. This article is then concluded, with a short discussion on open issues.

L1VPN Service and Management Architectures

In this article two aspects of L1VPN service and management architectures are discussed: management of carrier network partitions for L1VPNs, and L1VPN management by users. For carriers, resource partitioning and L1VPN isolation are two key requirements. For L1VPN administrators, composition of L1VPNs using a combination of resources from different carrier and private networks, further resource partitioning, and leasing or trading resources with other L1VPNs are key services. For L1VPN end users, reconfiguration of an L1VPN topology is a key service.

Management of Carrier Network Partitions for L1VPNs

In the management of carrier network partitioning for L1VPNs, two techniques are critical: resource partitioning and L1VPN isolation. Partitioning transport network resources involves allocating the resources to L1VPNs, and granting L1VPN users access to the management modules for their allocated resource portions. If a distributed control plane is used within a carrier network, L1VPN users should be granted access to the control functions for their partitions. Carrier network resource partitioning can be static/permanent (based on a carrier's permanent configuration), dynamic/on-demand (based on L1VPN users' signaling to the carriers), or semi-dynamic/soft-permanent (based on carrier network administrator involved reconfiguration). L1VPN isolation has two aspects: isolation of user traffic/signal, and isolation of management/control messages for different L1VPNs. Ideally, L1VPN isolation should make all activities in one L1VPN invisible to other L1VPNs. The minimal requirement is that activities in one L1VPN should not be interfered by other L1VPNs. L1VPN inherently offers user traffic/signal isolation. However, isolating different L1VPNs' management/control messages requires new mechanisms, and thus is the key to L1VPN solutions.

There are two models for a carrier to manage partitions for L1VPNs: the resource-partition based model, and the domain-service based model [3]. In the resource-partition based model, the carrier partitions resources into disjoint sets. Each L1VPN virtually owns the contracted resources, and has full management over a partitioned subnetwork. A resource partition is exclusively used by a designated L1VPN. The resource-partition based model is also called the port-based L1VPN model [4], since each port on a carrier's edge node is explicitly allocated to one single L1VPN. In the domain-service based model, transport network resources are dynamically allocated to L1VPNs, that is, connections between two L1VPN access ports on carrier's edge nodes are created on-demand. Thus, resources are time-shared among different L1VPNs. However, at any time, a resource can only be used by one L1VPN. The domain-service based model is also called the connection-based L1VPN model [4], since what are visible to L1VPNs are connections, not the component links of the connections. The advantages and disadvantages of the two models are summarized in Table 1.

With the objective of building L1VPN management by users, the resource-partition based model is the natural choice for a carrier to manage partitions for L1VPNs. The reason is that,

Once users construct an L1VPN, from time to time they need to reconfigure the L1VPNs. Such reconfigurations change the L1VPN topology, add or reduce bandwidth allocations. L1VPN users should be able to further partition, lease, or trade resources.

only in this model, L1VPN users are granted visibility and access to the management functions in the carrier networks for the allocated resources. In this article we present a design of the resource-partition-based management that is realized at two levels: the network element (NE) and L1VPN levels. Design choices need to be made about which resource partitioning techniques are used and how different resource partitioning techniques are incorporated into one management system. Our design uses both levels: the NE level resource partitioning is used for safeguarding, while the L1VPN level resource partitioning is used for flexibility.

Resource partitioning at the NE level creates virtual NE management interfaces for an NE. The access to each virtual NE management interface is restricted to only one authorized L1VPN. Thus, each L1VPN manages a separate virtual partition of an NE. The carrier network partitioning at the NE level is configured by a carrier network administrator. The resource partitioning at the NE level remains relatively stable. A resource partition lasts for the period of a maintenance/upgrade cycle of a carrier network, (e.g., a few months or longer). Unfortunately, few commercially available NEs support resource partitioning. Generally, when access to an NE management system is granted to users, the users have complete control on all the NE resources. Virtual resource isolation cannot be met by most existing NE management systems. New mechanisms are required for resource partitioning at the NE level. The mechanisms may modify existing NE management systems, or add a functional layer on top of existing NE management systems. A proxy on top of the existing NE management techniques is presented later in this article.

Compared to manual-configuration-based carrier network partitioning at the NE level, a carrier partitions resources at the L1VPN level based on policy. At the L1VPN level, the resource management authority can be transferred among L1VPNs using leasing mechanisms. L1VPN level partitioning is more dynamic than at the NE level. The carrier usually changes L1VPN level partitions every few weeks. Subsequently, we present an implementation of L1VPN level partitioning based on WS.

L1VPN Management

Developing L1VPNs that coexist over multiple carrier networks is a significantly greater challenge, and generally there are two approaches:

- Using some form of the Network-to-Network Interface (NNI) signaling between carriers
- Direct exchanging management information between users and carriers

The first approach assumes that there is a strong business relationship between the carriers, so that the carriers negotiate and signal across the NNI on behalf of the users. In the second approach, the carriers have a minimal business relationship, but follow the same standards to ensure the signal compatibility and transmission requirements. The business negotiations are conducted between the users and the individual carriers, and no business negotiation is required between the carriers. L1VPN management functions are given to the users. L1VPN management by users requires a configuration and provisioning tool that assigns carrier network resources to L1VPNs. Carriers assign long-term bandwidth to L1VPN users. Then, the users combine these assigned bandwidth allocations into a working network.

Once users construct an L1VPN, from time to time they need to reconfigure the L1VPNs. Such reconfigurations change the L1VPN topology, add or reduce bandwidth allocations. L1VPN users should be able to further partition, lease, or trade resources. Temporary spare resources may be leased to other L1VPNs or traded with other L1VPNs for other required resources. When an L1VPN only partially uses a resource, the L1VPN administrator may further partition such a resource into smaller granularity pieces, and then lease or trade unused portions with other L1VPNs. For example, when a band of wavelengths is allocated to an L1VPN, the wavelength band can be partitioned into individual wavelengths, and then leased to other L1VPNs. In a Synchronous Optical NETwork (SONET) or a Synchronous Digital Hierarchy (SDH) network, an L1VPN can partition the bandwidth on one link into smaller granularity time slots, and trade them with other L1VPNs for time slots on other links.

Interactions Between Carriers' Resource Partition Management and Users' L1VPN Management

L1VPNs exchange management/control messages with carriers by connecting their management systems, or letting user's edge nodes signal to the carrier's edge nodes. Unlike a layer 2/3 VPN, the users' edge nodes may not have control plane connectivity with the carriers' edge nodes. In this situation, an L1VPN administrator (or the L1VPN management system) needs to connect to the carriers' management systems. By exchanging messages between the L1VPN and carrier network management systems, the L1VPN can be provisioned or reconfigured. In this article the L1VPN management is explained using this management architecture. This architecture is suitable to semi-dynamic L1VPNs, where connections remain unchanged for hours or longer time periods. Although the bandwidth usage is not optimized because of relatively large bandwidth granularity and slow provisioning process (in the order of up to minutes), its simplicity is attractive to some applications.

With the progress of the IETF Generalized Multi-Protocol Label Switching (GMPLS) and the ITU-T Automatically Switched Optical Network (ASON) standards, a distributed control plane may be used to control users' edge nodes and their allocated resources in carrier networks. L1VPN management using a distributed control plane enables advanced bandwidth on-demand or scheduled bandwidth applications, but it requires complex inter-domain GMPLS technology. L1VPN management using a distributed control plane is beyond the scope of this article and further discussions on the subject can be found in [5, 6].

Carrier Network Partitioning at the Network Element Level

Carrier network partitioning at the NE level requires four key functions: NE resource isolation, reconfigurable NE resource partition, management information protection, and message logging. These functions can be implemented on top of various NE management techniques, such as the Transaction Language One (TL1), the Simple Network Management Protocol (SNMP), and so forth. For example, a TL1 proxy is developed without any modifications to the existing NE management system [7]. First of all, with the TL1 proxy, different L1VPN users are able to manage their own resource partitions on an NE without interfering with each other. Second, the reconfiguration of NE resource partitions may be done online (i.e., without rebooting the TL1 proxy). Third, the TL1 proxy protects sensitive network management information such as the IP address, TCP port number, login identifier, and password for an NE management interface. An L1VPN administrator or management system logs into the TL1 proxy. Then, the TL1 proxy delegates the login to the NE management systems. Finally, the message logging can be used for debugging and administrative purposes. The carrier may use message logging to resolve disputes on resource access. Resource utilization may be monitored and audited via the TL1 proxy.

The TL1 proxy maps the TL1 sessions from L1VPN users to the TL1 sessions towards NEs. The TL1 proxy is a software application running on top of the TCP/IP protocols. The transport can be optionally encrypted using the Secure Socket Layer (SSL). The TL1 proxy has two types of management interfaces: the North Bound Interface (NBI) to a L1VPN management system, and the South Bound Interface (SBI) to an existing NE management system. Each L1VPN's virtual NE has a unique combination of an IP address and a TCP port number at the NBI. An L1VPN user identifies different virtual NEs by using different NBIs. Based on L1VPN identifiers, NE identifiers and NE partitions, the TL1 proxy verifies an L1VPN user's access right on an NE partition, forwards TL1 commands from the L1VPN user to the NE, and relays alarms from the NE to the L1VPN user (Fig. 1). To ensure that the TL1 proxy properly verifies every TL1 command, the TL1 command-forwarding function in all NEs is disabled, that is, every TL1 command from the TL1 proxy is directly destined to the final NEs. Web Services (WS) are created for each virtual NE, as illustrated in Fig. 1. The NE-WS are part of the L1VPN management system, which is explained in the next section.

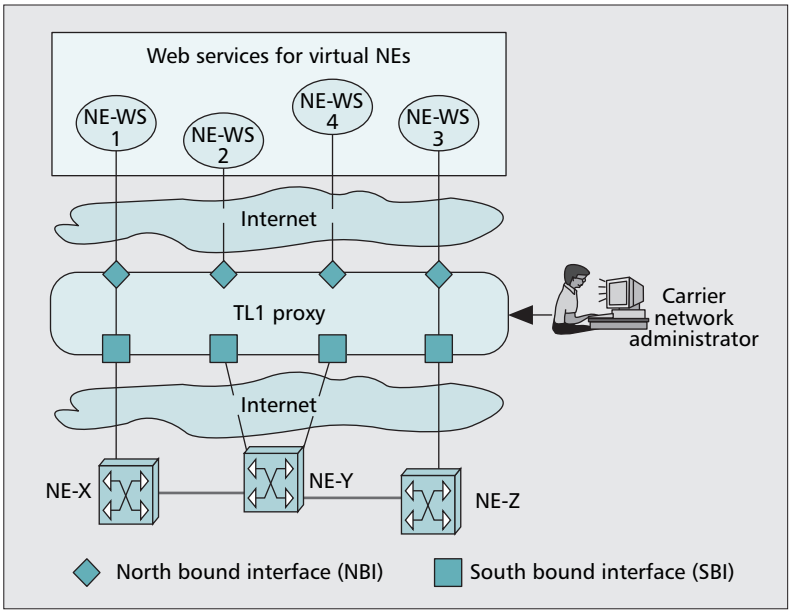

Figure 1. *The TL1 proxy communicates to L1VPN management systems via its north bound interfaces (NBIs), and communicates to existing NE management systems via its south bound interfaces (SBIs).*

WS Based L1VPN Management

In the WS-based L1VPN management, network resources are treated and managed by WS. The WS architecture defines the description, discovery, and interoperability of distributed, heterogeneous applications as services. The building blocks defined in the WS architecture include the eXtensible Markup Language (XML, a flexible and easy-to-extend data format), the Web Services Description Language (WSDL, an interface description of a service in an XML format), the Simple Object Access Protocol (SOAP, a means for communication between WS and client applications), and the Universal Description, Discovery and Integration (UDDI, a standard for the registration and publication of WS and their characteristics so that they can be found by potential clients) [8]. We will explain how WS are used to design an L1VPN management system from an operational point of view.

The L1VPN management functions are implemented as WS. To manage NEs, interfaces (e.g., NE ports, network-enabled instruments), and Light Paths (LPs), a carrier network administrator creates NE-WS, I-WS, and LP-WS, respectively. Then, the carrier network administrator leases or advertises resources in resource lists, which contains links to LP-WS and I-WS. An L1VPN administrator imports resource lists to compose an L1VPN, and creates new L1VPN-WS to manage the composed L1VPN. The L1VPN administrator can create super LPs by concatenating a chain of LPs, partition LPs into smaller bandwidth resources, and create end-to-end connections by attaching two users' edge nodes to an LP. When an L1VPN end user activates L1VPN-WS, the manipulations of the resources are executed, and an operational network is created. The roles of carrier network administrator, L1VPN administrator, and L1VPN end user are summarized in Table 2. The mark ✓ denotes that a player is able to use a service, while the mark ✗ denotes that a player is unable to use a service.

Carrier-Administered Resource-Partitioning at the L1VPN Level

A carrier network administrator uses a physical network editor to create a logical view of an L1VPN (Fig. 2). The logical view of an L1VPN is a collection of LP-WS, which an L1VPN adminis-

The L1VPN resource lists are portable from the carrier network management system to the L1VPN management systems. In this way, the management services for the partitioned resources can be transferred from one administrative domain to another.

	Carrier network administrator	L1VPN administrator	L1VPN end user
Create a physical network (NE-WS and link topology)	✓	✗	✗
View statistics of owned switches	✓	✗	✗
Create or delete LP-WS and I-WS	✓	✗	✗
Lease or advertise resource lists (LP-WS and I-WS)	✓	✓	✗
Import resource lists (LP-WS and I-WS)	✓	✓	✗
Create or dismantle super LPs	✓	✓	✗
Partition or bond LPs	✓	✓	✗
Create or delete end-to-end connections	✓	✓	✗
Create or delete L1VPNs	✓	✓	✗
Modify L1VPN topology	✓	✓	✓
Deploy or undeploy L1VPN-WS	✓	✓	✓
Activate or deactivate L1VPNs	✓	✓	✓
Query owned resources	✓	✓	✓
View statistics of owned LPs	✓	✓	✓
Receive alarms	✓	✓	✓

Table 2. *Roles of carrier network administrator, L1VPN administrator, and L1VPN end user.*

trator uses to compose L1VPNs. First, the physical network editor populates the virtual NE-WS list by either allowing the administrator to manually type in the Universal Resource Identifiers (URIs) of the virtual NE-WS, or relying on autodiscovery mechanisms to do so. If a virtual NE partition is not yet managed by WS, the physical network editor will create the WS. Second, the physical network editor assists the carrier network administrator to create a logical view of an L1VPN. An LP is an abstraction of a transport link. An LP can represent a fiber, a group of wavelengths, a single wavelength, a group of TDM time slots, and so on. The logical view of an L1VPN only consists of LP-WS, which link to the WS managing the two end points of an LP, that is, the I-WS, and link to the corresponding NE-WS. I-WS and LP-WS are dynamically created in the physical network editor. An LP factory service supports the creation of LP-WS. Upon being called and fed with the proper attributes, the LP factory service creates and deploys LP-WS. When creating LP-WS, the physical connectivity is verified. The NE-WS, I-WS, and LP-WS are deployed on the carrier management servers.

A carrier network administrator partitions a network by creating resource lists for different L1VPNs. The resource list editor assists the carrier network administrator to compile L1VPN resource lists, which are in the format of XML files, and composed of links to LP-WS. The L1VPN resource lists represent resource partitions for L1VPNs and contents links to the management services for the allocated resources. The L1VPN resource lists are portable from the carrier network management system to the L1VPN management systems. In this way, the management services for the partitioned resources can be transferred from one administrative domain to another.

WS are extensively used in the L1VPN management by users. As a result, the management system has a uniform way of receiving information and sending management messages, and the managed objects and components of the management system are hidden behind WS interfaces. The WS-based L1VPN management is in line with the vision of the Global Grid Forum that has defined a framework for managing resources as services [9]. The NE-WS, I-WS, and LP-WS are all loosely coupled in the sense that only when executing WS, are the supporting WS located and called. All the NE-WS, I-WS, and LP-WS for one carrier network operate within the carrier administrative domain. LP-WS are called by L1VPN-WS, which are under L1VPN administrations. The service calls may involve communication across network domain boundaries or in a public network. The service calls should be secured to protect the management system. Proper access control should be applied to verify the identity of carriers or users, and authenticate the use of management services.

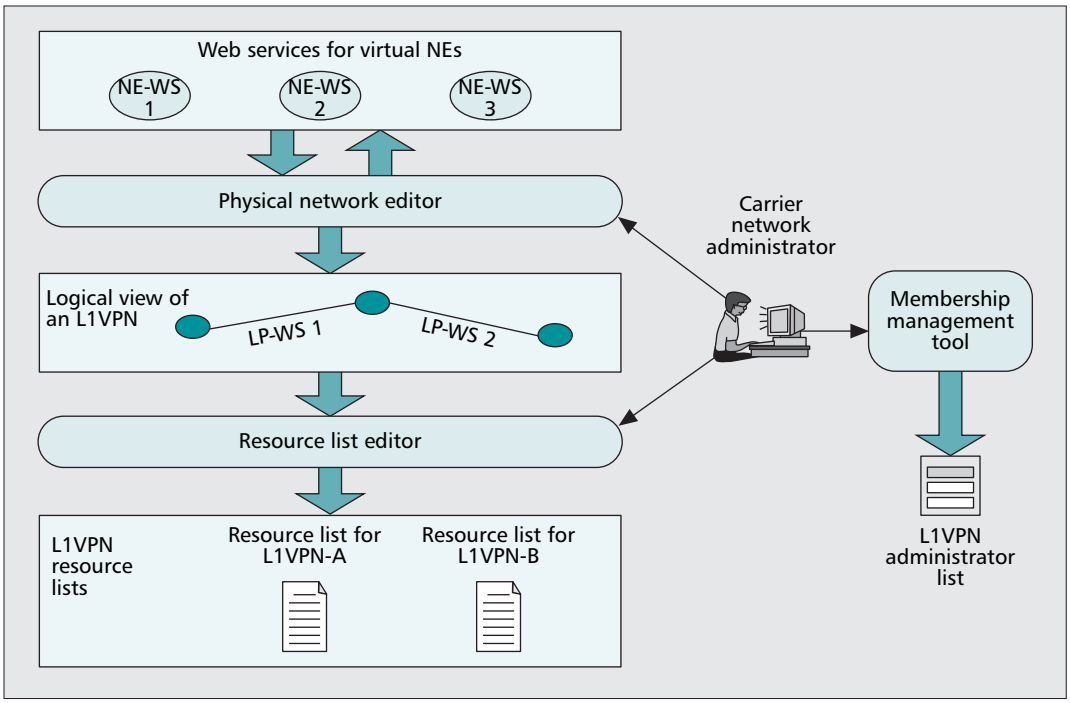

Figure 2. *A carrier network administrator builds a management view of the network by listing the WS for the virtual NEs, then further partitions the network into L1VPNs by creating resource lists for different L1VPNs.*

The most important service that the L1VPN administrator offers to L1VPN end users is L1VPN configuration. L1VPN-WS are workflows of service calls to component LP-WS, and through the LP-WS to I-WS and ultimately to NE-WS. By calling L1VPN-WS, an L1VPN end user completes the L1VPN configuration.

User Managed L1VPN Reconfiguration

An L1VPN administrator composes working network configurations by utilizing the resources that are either partitioned to this L1VPN by the carrier network administrator or leased/traded from partner L1VPNs. The L1VPN administrator provisions the initial interconnections of resources. L1VPN composition, resource trading, and leasing between L1VPNs can be conducted every few days or even possibly every few hours. Although in the L1VPN management that we present an L1VPN administrator's manual operations are involved in L1VPN composition, resource trading, and leasing, further development can be done to automate such operations, so that L1VPNs can be reconfigured more frequently. The L1VPN administrator uses the resource list editor to receive resources from three types of sources (Fig. 3):

- Carried over from the L1VPN's home carrier network
- Imported from external carrier networks
- Imported from partner L1VPNs

Each L1VPN has a home carrier, which initially authorizes an L1VPN administrator in the L1VPN administrator list. Usually, the home carrier provides an initial resource list to the L1VPN. The L1VPN expands its resources by leasing from other carrier's networks, or leasing/trading with partner L1VPNs. The acquired resources can be repartitioned into smaller granularity resources, where the derived smaller granularity resources are added into the resource list. For example, when a band of wavelengths is repartitioned into individual wavelengths, the LP-WS for the band of wavelengths is removed from the resource list and new LP-WS for wavelengths are added to the resource list. For the reverse operation, smaller granularity resources can be grouped into one single resource. For example, several time slots on one TDM link are grouped into one entity. New developments in the next-generation SONET/SDH supports Virtual Concatenation (VCAT) and the Link Capacity Admustment Scheme (LCAS), which make grouping time slots and "hitless" increasing or decreasing of the link capacity possible. By concatenating a chain of LPs, a super LP can be created by the L1VPN administrator.

The resource partitioning and bonding, as well as the leasing and trading between L1VPNs, are the foundation of the physical network broker business. An L1VPN administrator offers additional business values by negotiating and collecting network resources for L1VPN end users, although an L1VPN administrator is not the real owner or operator of network resources. The L1VPN administrator maintains L1VPN servers, which store resource lists, and run L1VPN-WS.

The most important service that the L1VPN administrator offers to L1VPN end users is L1VPN configuration. L1VPN-WS are workflows of service calls to component LP-WS, and through the LP-WS to I-WS and ultimately to NE-WS. By calling L1VPN-WS, an L1VPN end user completes the L1VPN configuration. That means the cross-connections in NEs are made, and end-to-end connections are created and ready to transport user traffic/signals.

The L1VPN composer is a graphical tool to assist L1VPN administrators to create L1VPN-WS. In our UCLP system, the Business Process Execution Language (BPEL) is used to compose LP-WS into L1VPN-WS. BPEL is an XML with operational semantics (e.g., loops, conditions and fault handlers, etc.), comparable to pro-

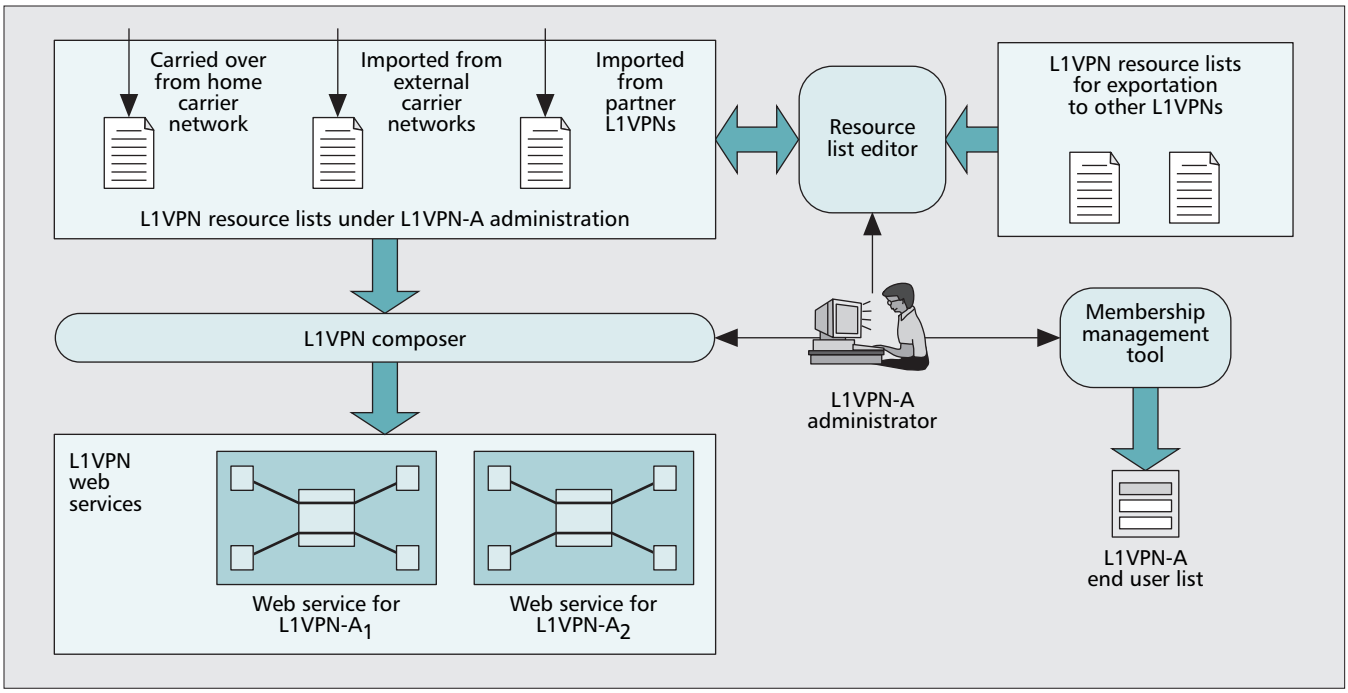

Figure 3. *An L1VPN administrator uses the resource list editor to receive resources. By calling L1VPN-WS, an L1VPN end user completes the L1VPN configuration.*

gramming languages such as Java and C++ [10]. BPEL turns discrete units of business functions (e.g., activating individual resources) into a business process (e.g., creating end-to-end connections, monitoring performance, billing, and auditing).

L1VPN end users independently reconfigure their L1VPNs without intervention of L1VPN administrators or carriers. The L1VPN administrator composes different scenarios for using an L1VPN. Such scenarios are all separate BPEL format L1VPN-WS, which are deployed on the L1VPN servers. For example, in Fig. 3, the L1VPN administrator describes the two operational modes of the middle switch in two different scenarios: L1VPN-A1 and L1VPN-A2 WS. The final action of the switchover is accomplished only when the L1VPN end users call an L1VPN-WS. Then, the operational mode of the switch is locked until the L1VPN-WS is released. After that, the L1VPN end users are able to call the other L1VPN-WS to reconfigure the L1VPN. Although the L1VPN administrator prepared different L1VPN-WS, the preparation is an offline process, and no real-time intervention of the L1VPN administrator is required. The L1VPN end users enjoy great flexibility in reconfiguring their L1VPN. An L1VPN user's IP routers are operating in a single domain environment, although the leased carrier network resources are provided by different carriers.

Conclusions and Future Work

The L1VPN technology offers flexible and cost-effective network operations. It enables new business models such as physical network resource brokers. The management complexity of L1VPNs is moderate compared to layer 2/3 packet-switching VPNs, especially in the resource-partition-based L1VPNs. As examples, we have illustrated how carrier network partitioning can be achieved at the network element and the L1VPN levels. Certainly, other design options exist to achieve the same goal. The use of WS in building an L1VPN configuration and provisioning tool has demonstrated many advantages, for example, the flexibility of transferring the management services and the modularity of the software architecture. The innovative use of workflow composition to prepare L1VPN use scenarios and handover of the management to the L1VPN end users make the UCLP system a truly user-managed system.

Future research needs to address resource discovery for new carriers or resources, optimization of resource allocation, enhanced access control and system security, scalability, and performance analyses. The functions of UDDI and other registration/publication services should be further explored. The naming, addressing, and property descriptions of resources should be standardized. Resource searching in a peer-to-peer manner among partner L1VPNs demands more research. Optimization of resource allocation and usage in noncooperative carrier networks is a new open issue. The wide use of network communication for sensitive management information requires in-depth research on the access control and system security.

Our UCLP system was demonstrated at the CANARIE workshop (March 14–15, 2006, Ottawa, Canada). The L1VPN technology has potentials to add new values to the connection-oriented transport network. The R&D in the UCLP system and L1VPNs have many common properties with the proposed Global Environment for Networking Innovations (GENI) architecture.

ACKNOWLEDGMENTS

This research is partially funded by CANARIE's directed research program on UCLP. We thank Prof. Gregor von Bochmann and his team at the University of Ottawa for their contributions in the discussions and system design. We thank Mathieu Lemay (Inocybe, Montréal, Canada), and Sergi Figuerola, Eduard Grasa, Joaquim Recio and Albert López (i2CAT, Spain) for their participation in the discussions and implementations. We thank Hervé Guy from CANARIE for his discussions.

REFERENCES

[1] 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.
[2] S. French and D. Pendarakis, "Optical Virtual Private Networks: Applications, Functionality and Implementation," *Photonic Network Commun.*, vol. 7, no. 3, May 2004, pp. 227–38.
[3] Y. Xue and L. Dunbar, "Viable Virtual Private Optical Network (VPON) Service Models for IP over Optical," *Proc. NFOEC 2001*, vol. 1, Baltimore, MD, July 8–12, 2001, pp. 212–20.
[4] Z. Zhang *et al*, "An Overview of Virtual Private Network (VPN): IP VPN and Optical VPN," *Photonic Network Commun.*, vol. 7, no. 3, May 2004, pp. 213–25.
[5] *IEEE Commun. Mag.* feature topic, Optical Control Plane for Grid Networks, vol. 44, no. 3, Mar. 2006, pp. 62–131.
[6] T. Takeda, Ed., "Framework and Requirements for Layer 1 Virtual Private Networks," Internet draft, draft-ietf-l1vpn-framework-02.txt, Mar. 2006, work in progress.
[7] TL1 LightPath Proxy; http://www.canarie.ca/canet4/uclp/tl1lightpathproxy.html
[8] B. St. Arnaud *et al*, "Web Services Architecture for User Control and Management of Optical Internet Networks," *Proc. IEEE*, vol. 92, no. 9, Sept. 2004, pp. 1490–1500.
[9] Global Grid Forum, Open Grid Services Architecture; http://www.ggf.org
[10] Business Process Execution Language for Web Services v. 1.1; www.ibm.com/developerworks/library/specification/ws-bpel

BIOGRAPHIES

JING WU [M'97] (jing.wu@crc.ca) has a B.Sc. degree in information science and technology in 1992, and a Ph.D. degree in systems engineering in 1997, both from Xian Jiao Tong University, China. He is now a research scientist at the Communications Research Centre Canada, Ottawa, an Agency of Industry Canada. Formerly he worked at Beijing University of Posts and Telecommunications as a faculty member, Queen's University, Kingston, Canada, as a postdoctoral fellow, and Nortel Networks Corporate, Ottawa, Canada, as a system design engineer. Currently, he is also appointed as an adjunct professor with the University of Ottawa School of Information Technology and Engineering. He has contributed more than 70 conference and journal papers. He holds three patents on Internet congestion control. His research interests mainly include control and management of optical networks, protocols and algorithms in networking, and optical network performance evaluation and optimization. He is a member of the technical program committees for ICC 2004–2007 Optical Networking Symposium, GOSP 2005–2006, BROADNETS 2006, ICCS 2004–2006, ICESS 2005, APOC 2005 Subcommittee for Network Architectures, Management, and Applications, DRCN 2003–2005, ICCCN 2005, and HPSR 2004. He has been a reviewer for numerous conferences and journals.

MICHEL SAVOIE (michel.savoie@crc.ca) is research program manager for the Broadband Applications and Optical Networks group of the Broadband Network Technologies Research Branch at the Communications Research Centre Canada (CRC). He maintains expertise in broadband systems and related technologies such as application-oriented networking (AON), advanced IP, ATM, and WDM-based optical networks in order to provide advice on important national initiatives and demonstrate the application of CRC technologies in a real operational environment. He has managed two User-Controlled LightPath (UCLP) projects funded under CANARIE's directed research program involving teams from the University of Ottawa, the i2CAT Foundation, Inocybe Technologies Inc., and CRC to develop software that enables users to dynamically provision dedicated end-to-end connections over shared network resources, and to provide advanced UCLP services with a graphical resource management tool for creating and managing articulated private networks (APNs). The former is based on Web and grid services, and Jini and JavaSpaces technologies; the latter is based on a service-oriented architecture (SOA) associated with resource lists comprising virtualized networking, computing, software, and instrument resources as Web services and custom workflows using BPEL representing end-to-end services targeting specific user communities. He is also involved with EUCALYPTUS: A Service-Oriented Participatory Design Studio project led by Carleton University and funded under the CANARIE Intelligent Infrastructure Program (CIIP), which combines the SOA and UCLP to provide a community of architects with on-demand fully collaborative multisite design capability; and PHOSPHORUS: A Lambda User-Controlled Infrastructure for European Research integrated project funded by the European Commission under the IST 6th Framework that addresses end-to-end user empowered service delivery across heterogeneous worldwide network infrastructures including UCLP systems. He holds B.Sc. and M.Sc. degrees in electrical engineering from the University of New Brunswick.

SCOTT CAMPBELL (scott.campbell@crc.ca) is a network researcher in the Broadband Applications and Optical Networks group at CRC. He has been working there since 2001, when he graduated from Dalhousie University, Halifax, Nova Scotia, Canada, with a Bachelor of Computer Science degree. At CRC he is involved in the design and development of agent-based network management and control software for all-optical networks. He is currently working on version two of the software associated with the UCLP project, which is a network management tool based on SOA and work flow technologies that allows end users to control and manage their own high-speed optical networks.

HANXI ZHANG (hanxi.zhang@crc.ca) is a research engineer at CRC. His research interests include network and system management, service-oriented software, and distributed applications. He is a key member of the CRC-i2CAT-UofO-Inocybe joint UCLP development team, responsible for UCLP system middleware. He obtained an M.Sc. degree from the University of Ottawa, Canada.

BILL ST. ARNAUD (bill.st.arnaud@canarie.ca) is senior director of advanced networks for CANARIE Inc., Canada's advanced Internet development organization (http://www.canarie.ca). At CANARIE he has been responsible for the coordination and implementation of Canada's next-generation optical Internet initiative, CA*net 4. He has been principal architect of the UCLP (http://www.uclp.ca) concept of applying SOA to network elements to allow users to orchestrate their own Internet network topologies and architectures fully integrated with their specific application needs. Previously, he was president and founder of a network and software engineering firm called TSA ProForma Inc. TSA was a LAN/WAN software company that developed WAN client/server systems for use primarily in the financial and information business fields in the Far East and United States. He is a member of various committees and boards, including the Board of Trustees for ISOC, NomComm committee for ICANN, the UKlight Steering Committee, the GLORIAD policy committee, Neptune Canada Oversight Committee, Globecomm Fellow, and the GLIF policy committee among others. In 2002 he was featured by *Time Magazine Canada* as the engineer who is wiring together advanced Canadian science. In 2005 he also won the World Technology Summit award for Communications. He has authored numerous papers and columns and is a frequent guest speaker at various conferences on the Internet and optical networking. He is a graduate of Carleton University School of Engineering.

Optimization of resource allocation and usage in noncooperative carrier networks is a new open issue. The wide use of network communication for sensitive management information requires in-depth research on the access control and system security.