INTERNATIONAL JOURNAL OF NETWORK MANAGEMENT
Int. J. Network Mgmt 2009; **19**: 139–152
Published online 18 September 2008 in Wiley InterScience
(www.interscience.wiley.com) DOI: 10.1002/nem.696

A network management tool for resource-partition based layer 1 virtual private networks

Jing Wu*†, Michel Savoie, Scott Campbell and Hanxi Zhang

Communications Research Centre (CRC) Canada, Ottawa, Ontario, Canada

SUMMARY

A Layer 1 Virtual Private Network (L1-VPN) has two models for service management: the resource-partition based model and the domain-service based model. In this paper, we present a network management tool for resource-partition based L1-VPNs. A Transaction Language One (TL1) proxy is designed to achieve resource partitioning at the network element level. Building on top of a TL1 proxy, we implemented a User-Controlled LightPath (UCLP) system to support physical network brokers to assign and allocate virtually dedicated resources to customers, and to enable customers to directly manage their resources. With such a capability, customers are able to create wide area networks based on their traffic pattern, and to adjust their traffic pattern based on available resources.

1. INTRODUCTION

Our primary design objective for the network management tool discussed here is to provide customers of the Layer 1 Virtual Private Network (L1-VPN) service with detailed information and direct operation of their virtually dedicated resources. Thus customers are able to create an L1-VPN as a wide area network to carry their application or discipline specific traffic. Most importantly, the customers are able to adapt their traffic pattern based on the L1-VPN resource availability. This means that in addition to the conventional traffic engineering approaches that allocate resources to a given traffic pattern, the customers can tailor their traffic pattern by adjusting their application-layer setup. Such a capability is particularly useful to long-lived bandwidth-demanding applications such as e-science applications. The power of the integrated resource allocation and traffic adaptation has not been fully explored in previous studies and existing approaches of traffic engineering. This motivated us to design a new L1-VPN management tool.

With our tool, a physical network broker may provide the L1-VPN service with additional features. A broker collects resources from various sources, e.g., purchasing point-to-point lightpaths from different providers, leasing time slots on certain links, and participating in condominium fibre builds in metropolitan areas. A broker provides additional business values to its customers. For example, it may split a large granular resource into pieces or, as a reverse operation, bundle multiple parallel resources between the same node pair into one entity for larger bandwidth and ease of management. By making cross-connections (i.e., layer 1 switching) at intermediate network elements, it may reconfigure connections, or chain a sequence of compatible resources into one entity. A broker is then able to resell or sublease the derived resources to its customers, which could be end users or another tier of brokers. Thus the L1-VPN service could be provided in a hierarchical fashion.

*Correspondence to: Jing Wu, Communications Research Centre Canada (CRC), 3701 Carling Avenue, Ottawa, Canada K2H 8S2.
†E-mail: jing.wu@crc.ca

Received 23 June 2006
Revised 28 January 2008
Accepted 31 January 2008

The key technologies to support the resource-partition based L1-VPN service include resource isolation, reconfigurable resource partitioning, assignment and allocation of partitioned resources, and secure and accountable transfer of resource management authority. Although separating transmission signals at the physical layer is relatively easy because there is no statistical multiplexing and packet mixing as in Layer 2/3 VPNs, the isolation of control messages and the separation and transfer of access to management functions need special techniques.

Our tool supports the configuration of resource partitioning at the network element (NE) level, dynamic assignment and allocation of resources to customers, and on-demand composition of a functional L1-VPN. A Transaction Language One (TL1) proxy is designed to achieve the NE-level resource partitioning. To increase resource utilization, we develop a technique that allows a broker to separate the assignment and allocation of resources to customers, so that a resource can be assigned to more than one customer and be allocated to a requesting customer on demand. A management software is developed to verify the proposed architecture and techniques.

This paper is organized as follows. In Section 2 we provide an overview of the L1-VPN service and a survey of related work. In Section 3 we present our design of a management tool for a physical network broker to manage its L1-VPN service. In Section 4 we describe the design of a TL1 proxy to achieve resource partitioning at the network element level. We conclude this paper in Section 5.

2. L1-VPN SERVICE AND RELATED WORK

There are two resource allocation models of the L1-VPN service: the resource-partition based model and the domain-service based model [1]. In the former model, a provider partitions resources into disjoint sets. Each L1-VPN virtually has full control over a contracted sub-network. A partitioned resource is exclusively used by its designated L1-VPN; i.e., the partitioned resource is not shared among different L1-VPNs. Thus, it is also called a port-based L1-VPN model [2] or dedicated user-plane model [3–5], since each access port of a provider edge node is explicitly allocated to one single L1-VPN. Using this model, it is possible to offer customers more control over their L1-VPNs. However, a disadvantage of the model is its low resource utilization. In contrast, the domain-service based model allows a provider to dynamically allocate resources and create connections between given access ports of a provider's edge nodes. The resource allocation is transparent to customers. Resources are time shared among different L1-VPNs. This model is also known as the connection-based L1-VPN model [2], since what are visible to a customer are connections, not the component links of the connections. This model is called the shared user-plane model [3–5], which means a resource can be allocated to different L1-VPNs at different times. The disadvantages of this model are contentions and customers' limited information about their allocated resources.

Three service models of the L1-VPN service are defined [6]: the management-based model, the signalling-based model, and the signalling and routing model. In the first model, the service interface is provided as a management interface between the management systems of customers and providers. In the last two models, the service interface is provided as a control plane interface between the edge nodes of customers and providers. In the signalling and routing model, routing information is exchanged across the service interface, revealing partial information about provider network topology and remote site reachability to customers.

The Internet Engineering Task Force's (IETF's) activities on L1-VPN focus on providing the L1-VPN service over Generalized Multi-Protocol Label Switching (GMPLS) enabled transport networks. Therefore, the service interface between a provider and its customers is based on the GMPLS control plane. In the first step, IETF defined a framework and requirements for L1-VPNs [6]. In the second step, IETF works towards defining the signalling messages that a customer's control plane exchanges with a provider's control plane [7,8]. At the same time, IETF works on auto-discovery mechanisms, which dynamically discover the set of provider edge nodes that attach to customers of the same L1-VPN [9,10]. In the future, IETF plans to define a service model that in addition to signalling messages, allows customers to exchange routing information with a provider. Features of L1-VPN operation, administration and maintenance, as well as modules of the L1-VPN Management Information Base (MIB), will be defined.

Different mechanisms for resource partitioning have been used in various contexts of networking technologies. By partitioning the resources of an asynchronous transfer mode (ATM) switch and allowing controllers to control different partitioned resources, third parties may lease a virtual network from an ATM network operator and directly control the leased resources [11,12]. In the concept of partitioning an IP router into virtual routers, router resources such as routing tables, bandwidth, buffer space, labels, and CPU control resources are partitioned for instances of virtual routers [13]. The virtual network concept and resource partitioning techniques play a useful role in the management and control of GMPLS-controlled multilayer networks [12]. Active networks as a type of programmable packet-switched network [14] allow data streams to change the policies at the network elements that control the streams. Potentially, programmable networks may offer dynamic resource partitioning. The mechanisms for resource partitioning can be based on the Simple Network Management Protocol (SNMP) and MIB [15], General Switch Management Protocol (GSMP) [16], Common Object Request Broker Architecture (CORBA) [11], and the GMPLS control plane [7,8].

3. A MANAGEMENT TOOL FOR THE L1-VPN SERVICE

We adopt the management-based L1-VPN service model. In this model, the L1-VPN service interface is provided as a management interface between the management systems of customers and providers. Although using a control plane to exchange routing information between customers and providers achieves the same goal, it requires the customers to use a control plane. In long-lived applications, customers prefer a lightweight L1-VPN service interface, because they usually use long-term resource leases instead of short-term resource scheduling and reservations. The scale and dynamics of customers and their L1-VPN resources may not justify the use of a heavyweight control plane such as the GMPLS control plane. We provide customers with a simple L1-VPN service interface, enabling them to use their network resources in a similar way to using computing, storage, and data acquisition devices. Customers use this tool to assemble resources at different NEs into a functional L1-VPN.

3.1 Functional requirements of a physical network broker

A broker that offers the L1-VPN service requires flexible resource assignment and allocation, connection management, and membership management. These three categories of management functions can be further classified as follows:

- Flexible resource assignment and allocation
  - separation of resource assignment and resource allocation, allowing a resource to be assigned to more than one customer;
  - on-demand allocation of resources to a customer;
  - resolution of potential resource contentions;
  - requesting and granting resources between customers, transferring allocated resources from one customer to another, supporting leasing and subleasing mechanisms;
  - transactional secure resource allocations, ensuring that no resource is lost in the case of a failed connection establishment;
  - publishing available resources.
- Connection management
  - searching available resources;
  - operation of L1-VPN resources, such as cross connecting two adjacent compatible resources;
  - notifications to a requesting customer about its failed connection establishment;
  - performance and fault monitoring and notifications.
- Membership management
  - customer authentication;
  - authorizing a customer to access its resources;
  - allowing a member to dynamically join and leave an L1-VPN.

The separation of resource assignment and allocation is a feature of our design. It allows a resource to be assigned to more than one customer and to be allocated to a customer on demand. Because the TL1 proxy requires intensive configurations, resource partitioning at the NE level cannot be practically done in fine granularity and should remain unchanged for a relatively long time. However, a broker may increase its resource utilization by enabling resource oversubscription. A given customer is provided with a temporary dedicated resource on an as-needed basis.

Our tool provides a scalable solution for a broker to manage its customers and their access to the management interface of virtually dedicated resources. It is not scalable to require a customer to manage its interfaces to all its virtually dedicated NEs. From a customer's perspective, our tool provides a single service entry to the management interfaces of all resources. From a broker's perspective, our tool provides access control, including granting, revocation and transfer of a customer's access to a given NE.

3.2 User-controlled lightpath system

The core of our management architecture is based on the brokers that serve L1-VPN customers. Via the coordination of brokers, customers provision Layer 1 (L1) connections across independent provider domains. For simplicity, we call such L1 connections 'lightpaths'. Each broker operates an instance of our management tool, called the User-Controlled LightPath (UCLP) system. The resources that a broker collects from different providers are abstracted as LightPath Objects (LPOs) and presented to customers (see Figure 1). Each LPO is assigned to a list of L1-VPNs. A customer dynamically decides which L1-VPN to join in, and chooses the partner customers to connect to.

A customer's management of LPOs is the essential function of the UCLP system. An LPO represents two end-points and the L1 transmission media in between. An LPO is associated with a set of attributes

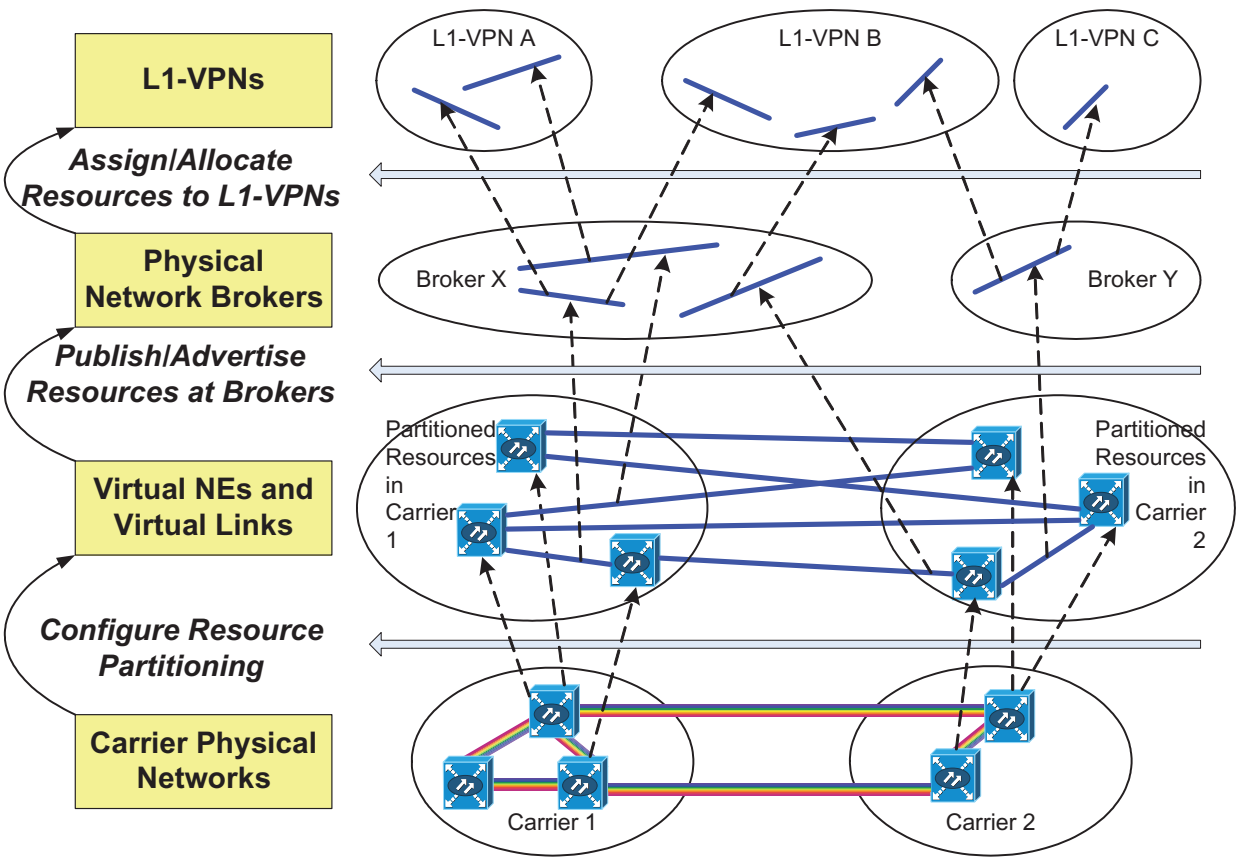

Figure 1. Our management architecture of the L1-VPN service

DOI: 10.1002/nem

and the Java methods that enable it to be connected to other LPOs, or to be connected to a customer at the end-point NEs. Each broker has a separate LPO storage to publish the LPOs that the broker collects. An LPO can be stored in any broker's LPO storage. Therefore, the LPO storage of all collaborating brokers needs to be searched for a given LPO. The distributed deployment of two instances of the UCLP system is shown in Figure 2, one for each broker. The two UCLP systems are linked. Two customers communicate to the access point of their corresponding UCLP system via the Internet.

Using the Jini technology, we designed and prototyped our UCLP system. The Jini technology offers rich functions to build service-oriented systems [17]. A Jini system extends the Java application environment from a single virtual machine to a network. The basic building blocks in a Jini system are Jini services, which may dynamically collaborate with each other to achieve complicated tasks. The organization of Jini services, their awareness of each other's existence and the communications between them are well supported by the Jini functions. The Jini technology is suitable for our management tool. First of all, a service-oriented architecture for a management system is a clear trend [18,19], and the Jini technology supports a service-oriented architecture very well. Second, because instances of the management tool for different brokers are located in the corresponding domains and therefore remote to each other, the network-related properties that Jini are particularly designed for help to simplify the overall system.

There are five key components in our UCLP system [20]: a Jini lookup service, an instance of JavaSpaces [21] for LPO storage, a Jini service access point, an LPO service, and an instance of switch communication service for each virtual NE. The LPO service provides the following functions: create/remove an end-to-end lightpath, add/delete an LPO, concatenate/partition an LPO, find/inquire an instance of JavaSpaces, and receive notifications from an NE.

Available resources are published in brokers' instances of JavaSpaces for customers to browse and use. JavaSpaces technology provides a distributed data storage for Java objects. Each broker has at least one

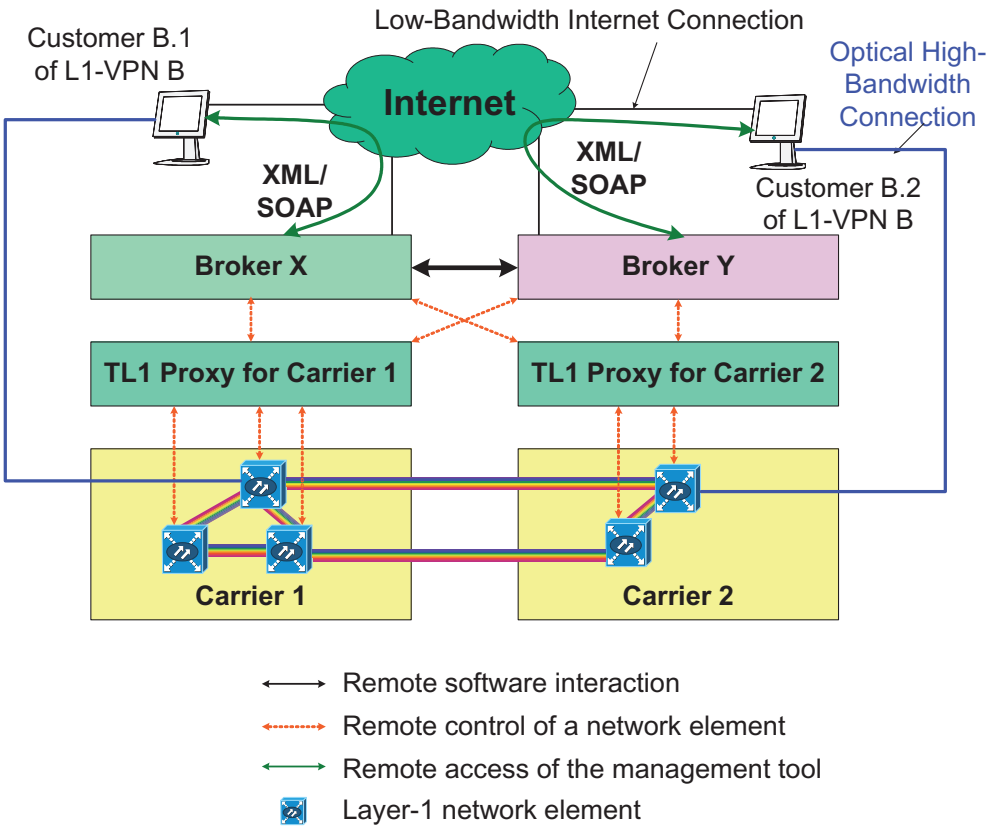

Figure 2. Distributed deployment of two instances of the UCLP system for two brokers

instance of JavaSpaces to store LPOs. Searching and browsing of published LPOs are controlled by the attributes of the LPOs. Such attributes include the assigned L1-VPN, available time, physical parameters of the resource, etc. Operations on an instance of JavaSpaces are transactionally secure. That means either all the service calls in a transaction are committed, or none at all. Transactions are supported for operations on a single JaveSpaces instance, as well as multiple JaveSpaces instances that are spread out in different brokers. In this way, contention is resolved for concurrent access of the same resource by different customers.

The management services provided by our system are classified into two groups: those only available to brokers, and those available to L1-VPN customers. The latter include in particular the 'ConnectionRequest', by which the establishment of a customer-to-customer connection may be requested. One of the functions reserved for brokers is the addition/deletion of new resources. How our UCLP implementation meets the L1-VPN service requirements is summarized in Table 1.

Service requirement	Service description	Mandatory/ optional required by ITU [4]	UCLP implementation description
Basic L1 service features	Connectivity, capacity, transparency, availability, performance	Mandatory	Entire resource-partition based L1-VPN management tool
Dynamic control of L1 connections	Soft-permanent, or switched	Mandatory	Customer-initiated connection establishment and reconfiguration. Connections can be added or removed without bringing an entire L1-VPN as a whole out of service.
Notification of connection rejection	When a service request cannot be completed, the network notifies the requesting customer	Mandatory	When searching for resources, a customer receives a notification if the requested resources are unavailable; when setting up a connection, the requesting customer receives a notification if a cross-connection operation fails
Subscription of multiple L1-VPNs at the service interface	Enabling a customer to simultaneously have connections within different L1-VPNs	Optional	A customer may use different ports/channels to simultaneously connect to different L1-VPNs. It is the customer's responsibility to isolate the traffic of the different L1-VPNs that it simultaneously connects to. A customer may use the same port/channel to connect to different L1-VPNs at different times
Parallel connection with public network	A customer is connected to public networks, as well as L1-VPNs	Optional	A customer may have multiple ports/ channels, connecting either to an L1-VPN or a public network
Authentication	Validation of a customer's identity prior to granting it access to service	Mandatory	A customer's identifier and password are verified, when it logs in to a broker's UCLP system
Authorization	Restricting a customer's control to its authorized L1-VPN	Mandatory	A customer can only access the resources that belong to its L1-VPN(s). This is implemented as a resource searching restriction
Accounting	Recording the quantitative info on usage	Optional	The info on a broker's resource usage can be logged at the TL1 proxy. A broker can also keep a record of an L1-VPN's resource usage
Connectivity restrictions	Restricting the source and destination of a connection to the members of the same L1-VPN	Mandatory	Each resource is assigned to one or more L1-VPNs. The resources that belong to an L1-VPN can only be used by the customers that are authorized to access the L1-VPN. This ensures that when using the resources of an L1-VPN to connect customers, the customers belong to the same L1-VPN

Explicit link selection	Customer specifying an explicit link or series of links (e.g. route)	Optional	When a customer searches resources to build a connection, it can specify preference. It gives the endpoints to be connected. Then, available routes are displayed for it to choose from
Distribution of membership info	Network distributing list of current members	Optional	UCLP does not distribute the membership info. It assumes such info is distributed to customers by other means, e.g. application-layer mechanisms
Distribution of member availability info	Network distributing ability or willingness of current members to participate in the L1-VPN	Optional	UCLP does not distinguish the registered and available members of an L1-VPN. In UCLP, there is no notion of whether a customer is online or not
Transfer of resource info	Network providing network topology view, performance, utilization, resource status	Optional	A customer builds its view of network topology by searching published resources. A customer also has an option to search only available resources. Through the interface to the NE management system (via the TL1 proxy), a customer may poll the performance and utilization statistics of a resource
Transfer of connectivity info	Network providing list of current active connections within the L1-VPN	Optional	A customer may view active connections in its L1-VPN
Transparent transfer of control info between customer entities	For example, topologies of isolated sites are shared when connected over a provider	Optional	UCLP only provides connections to customer edge nodes. The topologies of customer internal networks are behind customer edge nodes. Therefore, UCLP does not manage them
Network participation in customer domain routing	For example, a provider may use topology and status info of isolated sites to optimize routing	Optional	In UCLP, a provider does not participate in the routing of customer internal networks. Customers optimize routing, since they know their internal topology (and traffic pattern), as well as available external resources from a provider
Per L1-VPN policy	Ability to enforce policies (i.e., link selection policy) to each L1-VPN	Mandatory	Each L1-VPN independently sets up its policies on link selection, connection routing, etc.
Selection of L1 class of service (e.g. availability level)	Survivability mechanisms are offered corresponding to the class of service that is requested by customers	Optional	In UCLP, all resources are of the same class of service. A customer may build backup routes for its own survivability requirement
Customer network management	The customer's ability to view and control the service across an interface to the management system	Optional	A broker assigns resources to L1-VPNs. Customers can view their assigned resources. This is one of the key design objectives of the UCLP
Per customer-edge policy and its management	The customer's ability to modify the policy for each customer edge	Optional	UCLP does not manage the policy for customer edges
Transfer of performance info	The network providing the performance info of a L1 connection provided by the service provider (may include control plane status)	Optional	A provider generally does not monitor the end-to-end performance of L1 connections. A customer may poll limited performance info of the resources via the interface to the NE management system
Transfer of fault info	The customer can receive fault info for the resources in the user and control planes	Optional	A customer receives the alarms for the resources that it is using

Table 1. UCLP implementation descriptions

4. A RESOURCE-PARTITION BASED NETWORK ELEMENT MANAGEMENT SYSTEM

A TL1 proxy is developed to achieve resource-partition based NE management. Currently, commercially available NE management systems do not support resource partitioning. In general, a human operator or a software tool can access an NE management interface. For simplicity, we do not distinguish between them. The meaning is clear in the context. When a broker is granted access to an NE management interface, the broker has complete control on all the resources on the NE. That means every resource on the NE can be operated by the broker. To remedy this all-or-none resource management problem, our partner developed a TL1 proxy [22]. TL1 is a network management protocol [23], widely used by transport network elements as the primary management protocol. Our solution achieves resource partitioning without any modifications to existing NE management software. The TL1 proxy wraps around existing NE management software. Although the different brokers that have been allocated to certain resources on an NE access the same NE management interface through the TL1 proxy, the management of each piece of the NE resources is restricted to only one authorized broker; i.e., each broker manages a virtual partition of the NE.

4.1 Operation and features of a TL1 proxy

The TL1 proxy has four key features: NE resource isolation, reconfigurable NE resource partitioning, management information protection, and message logging. First, with the TL1 proxy, different brokers are able to manage their own resource partitions on an NE, and not interfering with each other. Reconfiguration of NE resource partitioning may be done online, i.e., without rebooting NE management modules. The NE resource partitioning is configured by a provider administrator. Then, connection control and resource management are directly accomplished by brokers or any entity that a broker authorizes. Third, the TL1 proxy protects sensitive management information such as IP address, TCP port number, login identifier and password for an NE management interface. A broker logs into the TL1 proxy. Then, the TL1 proxy delegates the broker's login to the NE management interface. Finally, the message logging can be used for debugging and administrative purposes. The provider 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 brokers onto the TL1 sessions towards NEs. The TL1 proxy runs 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 interface: northbound interfaces (NBIs) to a broker, and southbound interfaces (SBIs) to NE management interfaces. Each virtual NE has a unique combination of an IP address and a TCP port number at the NBI. A broker identifies different virtual NEs by using different NBI IP addresses and TCP port numbers of the TL1 proxy. Based on a broker's identifiers, NE identifiers and NE partitions, the TL1 proxy is able to verify the broker's access rights on an NE partition, and thus forward TL1 commands from a broker to an NE and relay alarms from an NE to a broker (Figure 3). To ensure that the TL1 proxy properly verifies every TL1 command, the TL1 command forwarding function in all NEs is disabled, i.e., every TL1 command from the TL1 proxy is directly destined to the final NEs.

The TL1 proxy operations can be illustrated in an example, where two brokers (e.g., Brokers X and Y in Figure 3) manage their partitions on a common SONET switch (e.g., NE X in Figure 3). The SONET resources on the switch are partitioned for the two brokers. Partitions of Brokers X and Y are specified in a configuration file by listing all their accessible slots/ports/channels. All allowed TL1 commands are configured for each broker as well. The following information is also maintained in the TL1 proxy configuration file: (i) the identifiers and passwords that the brokers use to establish TL1 sessions to NBIs; (ii) the IP address and TCP port number of the management interface of the switch; and (iii) the login identifier and password that the TL1 proxy uses to set up TL1 sessions from the SBIs to the switch. When a broker opens a TL1 session to an NBI at the TL1 proxy, the TL1 proxy automatically opens a corresponding TL1 session from a corresponding SBI to the switch. The TL1 proxy's operation is transparent

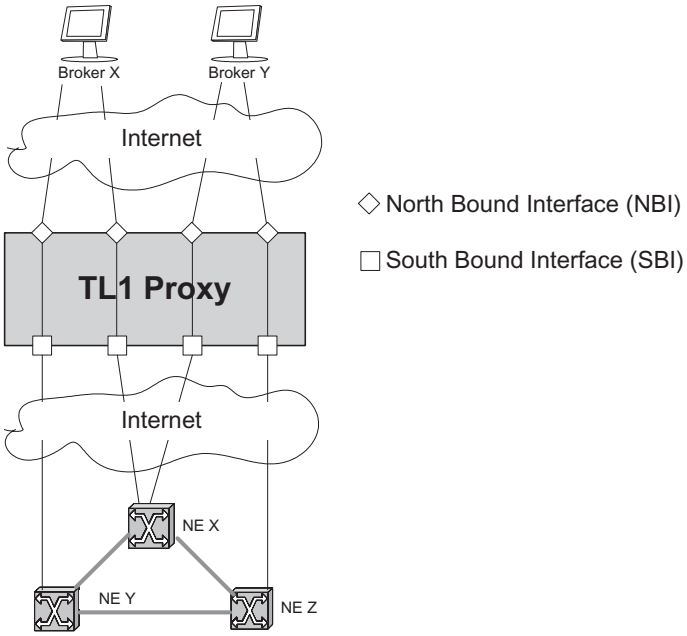

Figure 3. The TL1 proxy communicates to a broker via its northbound interfaces (NBIs), and communicates to NE management interfaces via its southbound interfaces (SBIs)

to the brokers. The brokers do not know the existence of the TL1 proxy. The NBI is used by the brokers as if the NBI was the management interface of the switch. The TL1 proxy's interception of the TL1 commands that the brokers send to the switch is unknown to the brokers and the switch. If the TL1 proxy fails to open a TL1 session to the switch, or an opened TL1 session to the switch is closed, the corresponding TL1 session from the broker to the NBI is closed accordingly. If the TL1 commands that the brokers send to the switch pass the TL1 proxy's permission verification, the TL1 commands are forwarded to the switch. The TL1 proxy delivers the returned text messages from the switch back to the brokers, and meanwhile the switch takes actions in response to the TL1 commands. If the TL1 commands that the brokers send to the switch violate the resource partition policies, an error message of 'resource access denied' is replied to the brokers, and the TL1 commands are discarded at the TL1 proxy. For the brokers, the TL1 proxy emulates the responses of the switch. For the switch, the TL1 proxy sends the TL1 commands satisfying the resource partition policies on behalf of the brokers. When the switch sends notifications to either broker or both, the TL1 proxy needs to make intelligent decisions to whom to relay the notification. When a notification is specific to a particular resource, the TL1 proxy relays the notification only to the resource owner. When a notification is platform-wide, the TL1 proxy replicates the notification to both brokers.

The TL1 proxy is scalable, although it could be a single point of failure bottleneck. A provider runs one or more TL1 proxies. Each TL1 proxy is responsible for one or more NEs. The ratio of TL1 proxies to NEs can be based on load-balancing and response time requirements.

4.2 Configurations of a TL1 proxy

In a TL1 proxy configuration file, the following information is manually configured: the virtual NEs, the valid NE resource partitions, the registered brokers and the permitted TL1 commands for each broker [24]. The TL1 configuration file is in the Extensible Markup Language (XML) format.

Each virtual NE is described as a separate entry in the TL1 proxy configuration file, where NBIs and SBIs of the TL1 proxy for each virtual NE are stored. The configuration of virtual NEs in the TL1 proxy

configuration file is shown in Table 2. In the example, any XML attribute whose value begins with a '@' character is to be encrypted, when the TL1 proxy configuration file is compiled using a utility software. The encrypted values begin with a '$' character.

An NE resource partition is described as a range of bandwidths, ports and channels. For example, a SONET resource partition is shown in Table 3.

All the allowed TL1 commands are specified in the <command> element in the TL1 proxy configuration file (shown in Table 4). A TL1 command can be restricted to a group of brokers. When an L1-VPN sends a TL1 command to the TL1 proxy, the Access Identifier (AID) elements of the TL1 command will be validated against the corresponding broker's element listed in the TL1 proxy configuration file.

Each <user> element defines a broker (shown in Table 5). The password that the broker uses to access the TL1 proxy and the passwords that the TL1 proxy sends to the NEs for the broker are stored in the <user> elements. The NE resource partitions that are allocated to the broker are specified as <lightpath> elements.

So far, the TL1 proxy has been tested on a variety of NEs, including the Nortel Optical Multiservice Edge (OME 6500), Nortel Optical Cross Connect (HDXc), and Cisco ONS 15454 Multiservice Provisioning Platform.

```
<neProxyList>
   <neProxy>
        tid                    = "ons-mon01"
             <!– TL1 Target Identifier (TID) for the NE –>
        neAddress      = "@192.168.1.1"
             <!–IP address of the NE management interface –>
        nePort           = "@1234"
             <!– TCP port of the NE management interface –>
        proxyAddress = "@192.168.1.2"
             <!– IP address of the TL1 proxy's NBI for the NE –>
        proxyPort       = "@4321"
             <!– TCP port of the TL1 proxy's NBI for this NE –>
        proxyType      = "ssl"
             <!– NBI transport type (plain|ssl|vt100) –>
   </neProxy>
       <!– define other virtual NEs here –>
</neProxyList>
```

Table 2. Configuration of virtual NEs in the TL1 proxy configuration file

```
<lightpath name = "lp01">
   <bandwidth value = "[1-3]" />
   <neList>
        <ne tid = "*">
             <resourceList>
                  <resource type = "FAC" slot = "6" port = "1" />
                  <resource type = "STS" slot = "5" port = "1" channel = "[1-12]" />
             </resourceList>
        </ne>
   </neList>
</lightpath>
```

Table 3. Configuration of NE resource partitions in the TL1 proxy configuration file

```
<commandList mode = "allowOnly">
    <command name = "ACT-USER"                    restrictions = "none" />
    <command name = "CANC-USER"                   restrictions = "none" />
    <command name = "RTRV-EQPT"                   restrictions = "none" />
    <command name = "RTRV-CRS*"                restrictions = "none" />
    <command name = "RTRV-NE-IPMAP" restrictions = "none" />
    <command name = "RTRV-MAP-NETWORK"    restrictions = none" />
    <command name = "RTRV-NE-GEN"              restrictions = "none" />
    <command name = "RTRV-HDR"                 restrictions = "none" />
    <!– the following commands are restricted to the L1-VPN "lightpath" –>
    <command name = "ENT-CRS*"                 restrictions = "lightpath" />
    <command name = "ED-CRS*"                  restrictions = "lightpath" />
    <command name = "DLT-CRS*"                 restrictions = "lightpath" />
</commandList>
```

Table 4. Configuration of allowed TL1 commands in the TL1 proxy configuration file

```
<userList>
  <user name = "ProviderAdmin">
      <proxyPassword                value = "$....." />
      <nePassword                   value = "$....."                tid = "*" />
      <!– customized commands for this L1-VPN are listed here –>
      <commandList                  mode = "allowAll" />
  </user>
  <user name = "BROKER-X">
      <proxyPassword          value = "$....." />
      <nePassword             value = "$....."          tid = "ons-lab01" />
      <nePassword             value = "$....."          tid = "ons-lab02" />
      <!– no commandList means the default commandList is used –>
      <lightpathList>
        <lightpath name = "ROOTLPO24STSc">
          <!– attributes for the NE resource partition "ROOTLPO24STSc" are listed here –>
        </lightpath>
        <!– other resource partitions for the Broker X are listed here–>
      </lightpathList>
  </user>
  <!– other brokers are listed here–>
</userList>
```

Table 5. Configuration of brokers' access in the TL1 proxy configuration file

5. CONCLUSIONS

In this paper, we present a management tool for resource-partition based L1-VPN. The management tool is implemented in two layers: the bottom layer is a TL1 proxy, realizing a configuration-based resource partition; the top layer is the UCLP system, supporting a physical network broker to dynamically assign and allocate virtually dedicated resources to customers. In addition, the UCLP system provides connection management and membership management. Our tool provides L1-VPN customers with status information and direct operation on their virtual resources. Via a simple management interface, the

customers are able to use network resources in a similar way to using computing, storage and data acquisition devices. The goal is to enable customers to adapt their traffic pattern based on resource availability, so that they may take advantage of the integrated resource allocation and traffic adaptation.

Our future work will include the security features of the system, and applying a Web Services architecture to the system. The prototype of the tool has been demonstrated at several international events.

ACKNOWLEDGEMENTS

The research is partially funded by CANARIE's directed research program on UCLP. We thank Bill St Arnaud at CANARIE for his leadership and contributions to the UCLP research. We thank Hervé Guy at CANARIE for his discussion on the TL1 proxy. Monfox (www.monfox.com) is acknowledged for its contribution to the design and implementation of the TL1 proxy. 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 and Albert López (i2CAT, Barcelona, Spain) for their participation in the discussions.

REFERENCES

1. Xue Y, Dunbar L. Viable Virtual Private Optical Network (VPON) service models for IP Over Optical. In *Proceedings of the 2001 National Fiber Optic Engineers Conference*, Baltimore, MD, 8–12 July 2001; 212–220.
2. Zhang Z, Zhang YQ, Chu X, Li B. An overview of virtual private network (VPN): IP VPN and optical VPN. *Photonic Network Communications* 2004; **7**(3): 213–225.
3. Takeda T, Inoue I, Aubin R, Carugi M. Layer 1 Virtual Private Network: service concepts, architecture requirements, and related advances in standardization. *IEEE Communications Magazine* 2004; **42**(6): 132–138.
4. Layer 1 Virtual Private Network Generic Requirements and Architectures. *International Telecommunication Union Recommendation Y.1312*, September 2003.
5. Layer 1 Virtual Private Network Service and Network Architectures. *International Telecommunication Union Recommendation Y.1313*, July 2004.
6. Takeda T (ed.). Framework and Requirements for Layer 1 Virtual Private Networks. *IETF RFC 4847*, April 2007.
7. Fedyk D (ed.), Rekhter Y (ed.), Papadimitriou D, Rabbat R, Berger L. Layer 1 VPN Basic Mode. IETF draft draft-ietf-L1-VPN-basic-mode-02.txt (work in progress), July 2007.
8. Takeda T (ed.). Applicability statement for Layer 1 Virtual Private Networks (L1-VPNs) Basic Mode. IETF draft draft-ietf-L1-VPN-applicability-basic-mode-02.txt (work in progress), July 2007.
9. Bryskin I, Berger L. OSPF based L1-VPN Auto-Discovery. IETF draft draft-ietf-L1-VPN-ospf-auto-discovery-02.txt (work in progress), March 2007.
10. Ould-Brahim H, Fedyk D, Rekhter Y. BGP-based Auto-Discovery for L1-VPNs. IETF draft draft-ietf-L1-VPN-bgp-auto-discovery-02.txt (work in progress), June 2007.
11. Rooney S, van der Merwe JE, Crosby SA, Leslie IM. The Tempest: a framework for safe, resource assured, programmable networks. *IEEE Communications Magazine* 1998; **36**(10): 42–53.
12. Leon-Garcia A, Mason LG. Virtual network resource management for next-generation networks. *IEEE Communications Magazine* 2003; **41**(7): 102–109.
13. Bjorkman N, Jiang Y, Lundberg T, Latour-Henner A, Doria A. The movement from monoliths to component-based network elements. *IEEE Communications Magazine* 2001; **39**(1): 86–93.
14. Denazis S, Karnouskos S, Suzuki T, Yoshizawa S. Component-based execution environments of network elements and a protocol for their configuration. *IEEE Transactions on Systems, Man and Cybernetics, Part C: Applications and Reviews* 2004; **34**(1): 82–96.
15. Boutaba R, Ng W, Leon-Garcia A. Web-based customer management of Virtual Private Networks. *Journal of Networks and Systems Management* 2001; **9**(1): 67–87.

16. Anderson T, Buerkle J. Requirements for the dynamic partitioning of switching elements. *IETF RFC 3532*, May 2003.
17. Scheuring J, Schneider A, Li S(ed.). *Professional Jini*. Wrox Press: Hoboken, NJ, 2000.
18. Boutaba R, Golab W, Iraqi Y, St Arnaud B. Lightpaths on demand: a Web Services based management system. *IEEE Communications Magazine* 2004; **42**(7): 101–107.
19. St Arnaud B, Bjerring A, Cherkaoui O, Boutaba R, Potts M, Hong W. Web Services architecture for user control and management of optical Internet networks. In *Proceedings of the IEEE*, Vol. 92, No. 9, September 2004; 1490–1500.
20. Wu J, Savoie M, Zhang H, Campbell S, Bochmann G, St Arnaud B. Customer-managed end-to-end lightpath provisioning. *International Journal of Network Management* 2005; **15**(5): 349–362.
21. Freeman E, Hupfer S, Arnold K. *JavaSpaces: Principles, Patterns, and Practice*. Addison-Wesley: Reading, MA, 1999.
22. Guy H. TL1 Lightpath Proxy. CAnet 4 Technical Committee Meeting. In *2nd annual Canadian Higher Education and Information Technology Conference (CANHEIT)*, Vancouver, Canada, 1–4 June 2004.
23. Transaction Language -1: Management Protocol for Telecommunication Networks. www.tl1guru.com [Accessed: 12 September 2008]
24. TL1 LightPath Proxy, CANARIE. www.canarie.ca/canet4/uclp/tl1lightpathproxy.html [6 September 2008].

AUTHORS' BIOGRAPHIES

Jing Wu obtained a BSc degree in information science and technology in 1992, and a PhD 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, Canada), an Agency of Industry Canada. In the past, he worked at Beijing University of Posts and Telecommunications (Beijing, China) as a faculty member, Queen's University (Kingston, Canada) as a postdoctoral fellow, and Nortel Networks Corporate (Ottawa, Canada) as a system design engineer. He is also appointed as an Adjunct Professor at the University of Ottawa, School of Information Technology and Engineering. He has contributed to over 70 conference and journal papers. He holds three patents on Internet congestion control, and one patent on control plane failure recovery. His research interests include control and management of optical networks, protocols and algorithms in networking, optical network performance evaluation and optimization. Dr Wu is a member of the technical program committees for ICC 2004–2009 Optical Networking Symposium, GLOBECOM 2007 Optical Networking Symposium, GLOBECOM 2008 Next Generation Networks, Protocols, and Services Symposium, GOSP 2005–2006, BROADNETS 2006–2008, ICCS 2004–2006, ICESS 2005, APOC 2005, 2007 Subcommittee for Network Architectures, Management, and Applications, DRCN 2003–2005, ICCCN 2005, 2007, and HPSR 2004, 2008. He is a Senior Member of IEEE.

Michel Savoie is the 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), an Agency of Industry Canada. 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 to 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 enable 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, while 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 was involved with EUCALYPTUS: A Service-oriented Participatory Design Studio, a project led by Carleton University funded under the CANARIE Intelligent Infrastructure Program (CIIP), which combines SOA and UCLP to provide a community of architects with an on-demand fully collaborative multi-site design capability. He is also involved with PHOSPHORUS: A Lambda User Controlled Infrastructure for European Research integrated project funded by the European Commission under the IST 6th Framework which addresses end-to-end user-empowered service

delivery across heterogeneous worldwide network infrastructures including UCLP systems. Mr Savoie holds a BSc and MSc in electrical engineering from the University of New Brunswick.

Scott Campbell is a network researcher in the Broadband Applications and Optical Networks group at the Communications Research Centre Canada. He has been working there since 2001, when he graduated from Dalhousie University in 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 the development of the User Controlled LightPaths (UCLP) project, which is now in a commercial phase called ARGIA. ARGIA is a resource management and provisioning tool based on service-oriented architectures that allows network administrators and end users to control and manage their own high-speed optical networks.

Hanxi Zhang is a research engineer at the Communications Research Centre Canada. 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 the UCLP system middleware. He obtained his MSc degree from the University of Ottawa, Canada.