A Management Architecture for Layer 1 VPN Services

Neumar Malheiros, Edmundo Madeira,
Institute of Computing
State University of Campinas (UNICAMP)
PO 6176 – 13083-970 – Campinas SP, Brazil
{ncm, edmundo}@ic.unicamp.br

Fábio Verdi, Maurício Magalhães
School of Electrical and Computer Engineering
State University of Campinas (UNICAMP)
PO 6101 – 13083-970 – Campinas SP, Brazil
{verdi, mauricio}@dca.fee.unicamp.br

Abstract

A distributed control plane architecture enhances transport networks with dynamic and flexible connection control. As a result, it allows the provisioning of advanced connectivity services, like Virtual Private Networks (VPNs), on layer 1 switching networks. Such Layer 1 VPN (L1VPN) services enable multiple customer networks to share a single transport network. In this work[1]*, we propose an architecture for L1VPN management. Our approach has been to use Policy-Based Management (PBM) to provide customers with some level of control and management over their L1VPNs. We also present a prototype implemented to validate the proposed architecture and discuss implications of policies for L1VPN configuration management.*

1 Introduction

Traditional transport networks must be enhanced in order to deal with increasing growth in traffic, service network convergence, and the stringent quality of service requirements of new advanced applications. In this context, the Automatic Switched Transport Network (ASTN) architecture, specified by the International Telecommunications Union (ITU), has emerged as a key approach to design the next generation transport networks. ASTN enhances transport networks with a control plane architecture that enables dynamic topology and resource discovery, automated connection provisioning, and efficient recovery mechanisms. One such architecture is the Generalized Multi-Protocol Label Switching (GMPLS) [9], defined within the Internet Engineering Task Force (IETF). GMPLS extends IP-based routing and signaling protocols to build a distributed control plane architecture which supports multiple switching technologies.

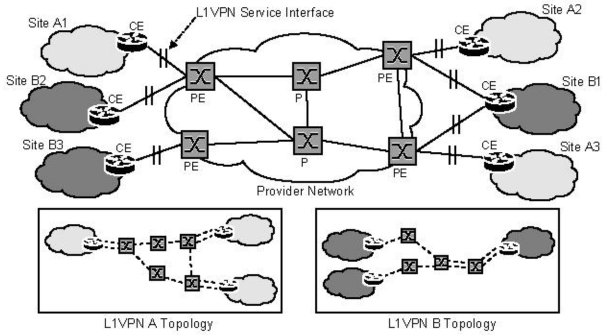

Figure 1. L1VPN service reference model.

The control plane allows transport networks to dynamically provide connections to customer networks. As a result, it is possible for providers to offer more advanced connectivity services such as Virtual Private Networks (VPNs) on layer 1 transport networks like optical and time division multiplexing (TDM) networks. The Layer 1 VPN [16] service enables multiple customer service networks to share a single layer 1 transport network. It allows customers to establish connections between their sites by dynamically allocating resources from the provider network. A primary requirement is that customers should be given some level of management and control over their Layer 1 VPN (L1VPN), which includes modifying the topology.

The Fig. 1 shows the basic elements of the L1VPN reference model [16] and also possible topologies of two L1VPN services (A and B). A Customer Edge (CE) is a device within the customer network that receives L1VPN services from the provider network. A Provider Edge (PE) is a device within the provider to which at least one CE is connected. It provides L1VPN service functionalities to CEs through the L1VPN service interface. A Provider (P) node is a device within the provider network that is not connected to any CE, but only to PE and P devices. Control and data connectivity is restricted to L1VPN membership which is the set of CEs under control of the same customer. L1VPN

[1] The authors would like to thank CAPES, CNPq, FAPESP and Ericsson Brazil for supporting this work.

1-4244-0425-8/06/$20.00 ©2006 IEEE

is a port-based provider provisioned VPN. The Customer Port Identifier (CPI) and the Provider Port Identifier (PPI) are the logical endpoints of the link between CE and PE respectively. A VPN member is identified by a CPI–PPI pair.

With the advances in layer 1 networks and the development of intelligent IP-based control plane architectures, Layer 1 VPN defines an interface by which providers can offer cost effective, flexible and on demand bandwidth services to multiple customers. Recently, standardization organizations have worked on L1VPN services. The ITU has specified generic service requirements and service architectural elements [6], as well as functions and architectures to support L1VPN [7]. Moreover, the IETF has created a specific working group which is aimed at specifying how to provide L1VPN services over GMPLS enabled networks. The first steps concern service requirements and framework [13], and the analysis of applying GMPLS protocols and mechanisms in the support of L1VPN services [14].

We have been investigating how to provide L1VPN services on transport networks enhanced with a distributed control plane. In this paper we are concerned with L1VPN configuration management issues. We propose an architecture for L1VPN service management. The main problem here is how a single provider network supports multiple L1VPN services, while providing customers with independent control and management over their L1VPN. In order to meet such requirement the architecture is built on the Policy-Based Management (PBM) approach [20]. The main focus of interest is to discuss L1VPN policy classes and how the proposed architecture supports PBM instead of defining specific policies for L1VPN management. Furthermore, the design of the architecture makes the assumption that the provider network control plane supports dynamic connection setup and topology information discovery.

First, we present related work. In Section 3, we describe how the IETF Policy Framework has been used in the context of L1VPN management and define major classes of policies for L1VPN service management. Then, we propose an architecture for L1VPN service management in Section 4. We describe the architecture functional model and usage scenarios for different L1VPN service models. Then we discuss the prototype implementation in Section 5 and evaluate the implications of using policies in L1VPN management in Section 6. Finally, we conclude and present future work in Section 7.

2 Related work

Research work have mainly focused on control plane issues in the provisioning of L1VPN services. Indeed, most work have investigated how to provide L1VPN services on GMPLS enabled transport networks and evaluated the necessary extensions in the control plane protocols in order to support L1VPN. The work presented in [12] proposes the so called Generalized VPNs (GVPNs) services. Such services use the Border Gateway Protocol (BGP) as a VPN membership auto-discovery mechanism and the GMPLS signaling and routing mechanisms to establish VPN connections.

The work presented in [17] compares types of L1VPN architectures namely, centralized, distributed and hybrid architectures. It focus on the management-based service interface as a suitable approach for initial steps in L1VPN service deployment. In this context, it evaluates centralized and hybrid architectures and argues that the last achieves more scalable, resilient and fast operations. Then it proposes a hybrid architecture that combines centralized and distributed functions to support the L1VPN service. In that case, VPN specific functions are centralized and common signaling and routing functions are distributed.

A discussion on existing GMPLS mechanisms for realizing L1VPN functionalities is presented in [15]. That work describes management and control-based L1VPN service models taking into consideration the service interface by which the customer accesses L1VPN functionality. Furthermore, it explains how L1VPN services can be performed by GMPLS in terms of addressing, membership discovery and signaling aspects. Finally, it discusses open issues in L1VPN provisioning. Such ones mainly include management functionalities like per-VPN resource management and VPN configuration management.

Differently from aforementioned work, our contribution focus on configuration management of L1VPN services. Instead of compete with them, the present proposal complement that work in order to deploy L1VPN services in a fully functional way.

3 Policy framework

We have decided on a policy-based approach in order to achieve that customers have some level of control and management over their L1VPN service. Indeed, an L1VPN requirement is that customers must be able to specify policies to control their L1VPN operation. Therefore, Policy-Based Management (PBM) emerges as a suitable approach to meet these requirements.

PBM has been widely used to address the complexity of network and service management. More specifically, some research work have proposed policy-based architectures to manage VPN services on IP networks. PBM provides dynamic network-wide management. In a Policy Framework, an administrator defines policies to be enforced within the network in order to control the behavior of the system as a whole. Policies are a set of rules that define how network resources and services can be used. Each rule consists of conditions and actions. If the conditions are evaluated true, then the actions are executed.

3.1 Framework

We have adapted the IETF policy framework [20] taking into account L1VPN management aspects. The adapted framework is presented in Fig. 2. In this framework, the provider network operator should firstly specify policies to manage the L1VPN services based on administrative and business goals. This is represented by the *Modeling* process. Policies are modeled in accordance with the provider network infrastructure, the L1VPN service framework supported by the provider, and customer requirements. Furthermore, defined policies should be in compliance with a policy information model to assure interoperability. There is a need for a policy management tool that should provide policy modeling, including syntax check, conflict resolution, and so on. The defined policies for L1VPN control and management are represented in the figure as *Service Configuration Policies*.

A Policy Decision Point (PDP) should evaluate the service configuration policies in order to decide which actions must be enforced. Decisions may be requested or triggered by the occurrence of specific events. Such decisions should consider the current network conditions. PDP also may translate configuration policies to configuration information specific to the network nodes based on specific capabilities of each node. Finally, a Policy Enforcement Point (PEP) is responsible for performing configuration changes according to actions sent by the PDP. They are also responsible for the report of device capabilities and network status to the PDP.

The *Decision* process performed by the PDP is logically centralized. The *Enforcement* process is distributed over the network by implementing the PEP functionality on the managed entities. However, *Decision* and *Enforcement* processes may be implemented in a centralized management system.

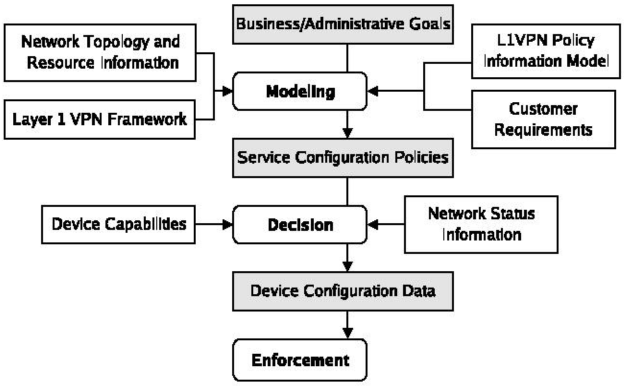
Network Topology and Resource Information
Layer 1 VPN Framework
Business/Administrative Goals
Modeling
L1VPN Policy Information Model
Customer Requirements
Service Configuration Policies
Device Capabilities
Decision
Network Status Information
Device Configuration Data
Enforcement

Figure 2. Policy-based framework for Layer 1 VPN management.

3.2 Policy classes

We have defined three major categories of L1VPN policies, namely, configuration, admission control, and routing policies. *Configuration Policies* are used to define configuration parameters which control L1VPN service operation. Main service operational aspects to which that policies may be applied are described as follows:

- There are two L1VPN resource allocation models. In the dedicated model, provider network resources are exclusively reserved for a specific L1VPN and they can not be allocated by another L1VPN. In the shared model, resources can be allocated by different L1VPNs in a time-sharing manner. Configuration policies can be used to configure resource allocation management and control the specification of allocation models for the several L1VPN customers.
- Configuration policies can also be applied in fault handling [3]. The provider network may support several restoration and protection schemes to recover from link and node failures. Such policies may determine which recovery scheme should be used for each L1VPN or each L1VPN member.
- The provider network may support differentiated classes of L1VPN services. Provider network operators can specify configuration policies to determine the class of service for each L1VPN.
- Configuration policies can also define routing algorithm parameters, since each L1VPN service may use different path computation algorithms, link weights, and other routing attributes in connection routing.

Admission Control Policies are used in the L1VPN connection admission control which also considers membership information. Beyond common admission control aspects, like resource availability, those policies allow to specify additional constraints on admission control as following:

- Such policies can enhance membership control with additional connectivity restrictions. They allow to define restrictions within an L1VPN by defining which members can establish connections to each other.
- Admission control policies can also be used to limit resources per L1VPN service, as well as to limit the number of connections per L1VPN service or member.
- In the admission of customer connection requests, pre-provisioned connections in the provider network may be used in the establishment of the L1VPN connections. Admission control policies can be used to optimize the selection of the pre-provisioned core connection. In a previous work, we have described an architecture for policy-based grooming responsible for

managing the installation and aggregation of customer traffic flows within core optical connections [19].

Routing Policies aim to control path computation for L1VPN connections. The major applications of these policies are described as follows:

- Routing policies can be applied in support of Constraint-Based Routing (CBR). CBR is mainly used in traffic engineering and fast connection reroute mechanisms. In this case, path computation is subject to resource and administrative constraints like route restrictions [1]. Those policies can be used to specify such constraints on L1VPN connection routing.
- Routing policies can also be used in resource management as a way to support dedicated and shared allocation models when path computation is centralized.
- When there are several suitable routes for a connection, routing policies can be used to optimize route selection, for instance, in terms of resource utilization.

4 Proposed architecture

In this section, we propose a management architecture for L1VPN services. Firstly, we describe the architecture functionalities and how its modules interact with each other. Then we discuss operational aspects of the architecture with respect to different L1VPN service models.

4.1 Functional model

The architecture design is based on the assumption that the transport network is enhanced with a control plane which provides dynamic connection setup and network topology discovery. The proposed architecture is presented in Fig. 3. It defines a user interface represented by the *Access Interface* module. Through this interface the customer or the provider network operator can request L1VPN connections, define policies or receive performance information about the respective L1VPN service. On the other hand, the *Control Plane Interface* module is the interface between the management system and the control plane. The core modules are isolated by the interfaces. This design is aimed at achieving flexibility and making interoperability easier.

The *Access Interface* module is responsible for processing the requests from customers and provider network operators, as well as for authentication and authorization procedures. According to the requests it invokes one of the three modules described as follows. The *Service Monitor* is responsible for providing performance and fault information about L1VPN services for their customers. Some of those information can be obtained from the control plane. The *Service Provider* is the main module. It is responsible for the configuration and accounting management of L1VPN services. It performs admission control on customer connection requests considering policies and membership information. It is also responsible for connection control by requesting the control plane to create, modify or delete layer 1 connections according to the desired L1VPN topology. The *Policy Manager* allows customers and provider network operators to add, edit, remove, and activate policies. It is responsible for managing policies and processing decision requests from the other modules.

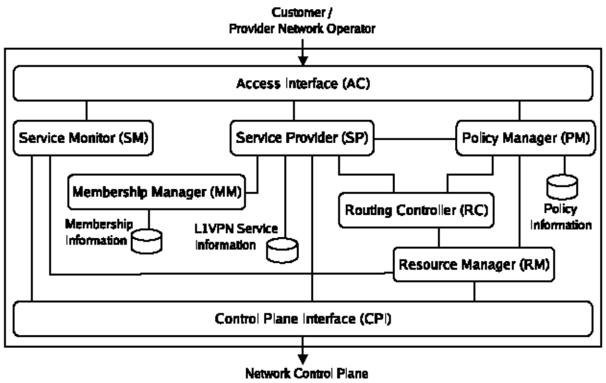

Figure 3. L1VPN management architecture.

The *Membership Manager* is responsible for managing L1VPN membership information. It includes functions for adding or removing L1VPN members and verifying membership. The *Resource Manager* is responsible for managing provider network resources. It should provide support for shared and dedicated resource allocation models. The *Routing Controller* is responsible for computing the route for L1VPN connection requests by making use of resource availability information from the Resource Manager and collecting topology information from the routing mechanism of the control plane.

Some of the functionalities are not VPN specific and may have already been implemented on the provider network. However, such functionalities may need extensions in order to support L1VPN services. Furthermore, L1VPN policies may be specified by provider network operators according to a Service Level Agreement (SLA) and customer requirements. Also, the customers should be allowed to specify high-level policies to configure and control their L1VPN services. As a result, separate L1VPN control and management is provided to customers.

The operation of the *Service Provider*, *Routing Controller*, and *Resource Manager* modules is subjected to the rules defined by the policies from the *Policy Manager*. The Fig. 4 illustrates how the architecture modules interact in the provisioning of L1VPN connections. The *Access Interface* processes the customer request and then invokes the *Ser-*

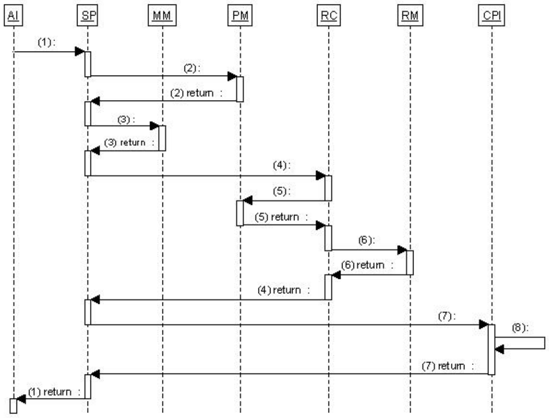

Figure 4. Connection setup.

vice Provider module to establish the connection (1). This module retrieves configuration information from the *Policy Manager* (2). Such information is used to configure connection setup. The *Service provider* also requests policy decisions for connection admission control. In addition, it communicates with the *Membership Manager* in order to verify whether the endpoints specified in the connection request are members of the same VPN (3). Then, the *Service provider* requests the *Routing Controller* to compute a route for the requested connection (4). The *Routing Controller* requests the *Policy Manager* for routing policy decisions which will control path computation (5). This computation requires resource availability information which is obtained from the the *Resource Manager* (6). Since the *Service Provider* has received a route for the connection, it requests connection setup to the network control plane through the *Control Plane Interface* (7). The connection setup by the control plane signaling is represented by step (8). Finally, the customer is reported on the connection establishment through the *Access Interface*.

4.2 Usage scenarios

There are two L1VPN service models based on the service interface. In the management-based service model, L1VPN service is provided on a management interface by which the customer management system communicates with the provider management system in order to control and manage the respective L1VPN. In this case, there is no exchange of control plane messages between customer and provider. In the control-based model, L1VPN service provisioning is achieved through control plane communication between CE and PE. In this case there are two approaches. First, the service interface is based only on control plane signaling between CE and PE devices. Second, routing mechanisms can be supported on the service interface in addition to signaling mechanisms. In the last case, there is exchange of routing information between CE and PE. Thus CE may obtain provider network topology and remote site reachability information.

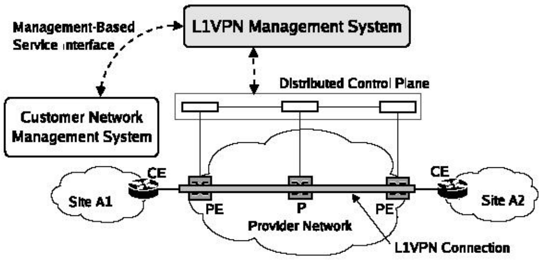
Management-Based Service Interface
L1VPN Management System
Distributed Control Plane
Customer Network Management System
CE
CE
Site A1
Site A2
PE
P
PE
Provider Network
L1VPN Connection

Figure 5. Management-based L1VPN service.

Here we describe usage scenarios in terms of the two L1VPN service models take into account the proposed architecture functionalities. In both scenarios we consider a GMPLS enabled provider network. In this context, dynamic connection setup can be performed by GMPLS RSVP-TE (Resource Reservation Protocol-Traffic Engineering) signaling mechanism [2] and automatic topology discovery can be achieved through GMPLS OSPF-TE (Open Shortest Path First-Traffic Engineering) routing mechanism [8].

The Fig. 5 shows L1VPN service deployment scenario considering the management-based service model. In this case, the service provisioning is based on a hybrid architecture which combines distributed and centralized functions. The control plane functions are distributed since they are performed by the GMPLS architecture. Distributed connection setup signaling and topology discovery contribute for scalability and fast connection recovery mechanisms. The management functionalities are centralized and may be implemented in an L1VPN management system or integrated into the provider Network Management System (NMS).

In order to establish an L1VPN connection, the customer sends a connection request to the L1VPN management system. After processing the request, the management system invokes the GMPLS control plane to set up a connection across the provider network on behalf of the customer. Then the CE nodes can establish routing adjacencies over such connection, as in an overlay scenario. GMPLS signaling mechanisms are used to establish the connection between PE devices. Such connection initiated by a management system is named a soft permanent connection (SPC). After a route is computed for the connection request, the L1VPN management system requests the SPC to the ingress PE.

In the control-based service model scenario, several of the proposed architecture functionalities are also distributed beyond the control plane ones. This usage scenario improves scalability and robustness in service provisioning. However it is a long term solution since control plane proto-

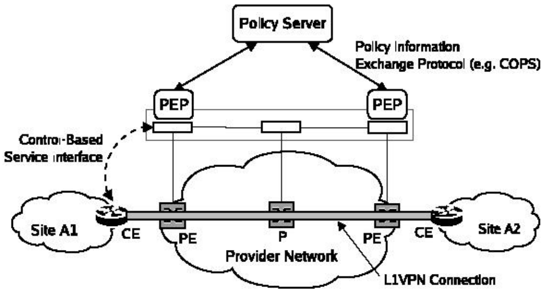
Policy Server
Policy Information Exchange Protocol (e.g. COPS)
PEP
PEP
Control-Based Service Interface
Site A1
CE
PE
P
PE
CE
Site A2
Provider Network
L1VPN Connection

Figure 6. Control-based L1VPN service.

col extensions and ever new solutions are needed. Namely, membership information management, admission and connection control, route computation and resource management functions, before implemented in the centralized system, are now performed on the PE device. Some functions like policy management remain centralized. Also, a centralized management system may be used to provide customer with service monitoring functions as reporting L1VPN service performance and fault information.

Furthermore, in this L1VPN deployment scenario, PE devices need to implement PEP functionality in order to support policy-based management, as illustrated in Fig. 6. PE devices communicate with a Policy Server to request policy decisions. This server supports policy modeling and performs PDP functionality. It is logically centralized and can be implemented in a centralized management system. In this context, L1VPN control and management are achieved through policies since operations on PE devices are subject to policy rules specified in the policy server. In order to support communication between PEPs on PE devices and a policy server there is a need for a policy information exchange protocol like COPS (Common Open Policy Service) [5]. This is a query and response protocol which allows on demand policy decision requests. Furthermore, such protocol could be used to provision PE devices with L1VPN service configuration information [4].

5 Prototype implementation

We have implemented and tested an L1VPN management prototype system in order to validate the proposed architecture. The implementation considers the management-based service model and L1VPN configuration policies. The main modules of the architecture were developed using the Java Remote Method Invocation (Java RMI) technology. The *Service Provider* module includes two submodules which are responsible for connection control and admission control. Route computation for requested connections are performed by the *Routing Controller*. It implements the Dijkstra's Shortest Path algorithm. The Fig. 7 depicts the structure of the prototype. Currently, the *Service Monitor* is being implemented and it is not showed in the figure. Also the policies are specified in a static way (in compilation time) within the *Service Provider* which is then responsible for the policy management and enforcement.

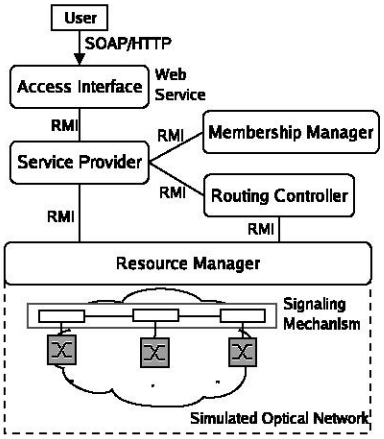

Figure 7. Implemented prototype.

In order to provide a high level of flexibility to access the management system, the *Access Interface* module was implemented as a Web Service. The main objective of the Web Services technology is to provide interoperability and automated communication between distributed and heterogeneous applications using XML standards and Internet protocols. The web interface is illustrated in Fig. 8 which shows the definition of a configuration policy. In this case, the configuration policy specifies a resource allocation model and a path computation scheme for the L1VPN service whose identifier is "100". If the dedicated model is chosen, then it is possible to define how many wavelengths must be reserved in each link. Two path computation schemes are possible: find the shortest path in terms of number of hops or the path with most available bandwidth in terms of number of available wavelengths.

We have also developed a simulation environment where L1VPN services are provided over an optical transport network. In this environment, L1VPN service customers concurrently send connection requests to the L1VPN management system. Then the system attempts to establish the requested connections over the optical network. Therefore, the L1VPN connections are optical connections through a single provider network. An optical connection is established by allocating an available wavelength in each link from the source to the destination node. The connection establishment is done by a simplified distributed signaling mechanism. Such mechanism is performed by control plane agents which are implemented in each node. It includes messages to request and confirm wavelength allocation. After the route is calculated, the *Service Provider* triggers the

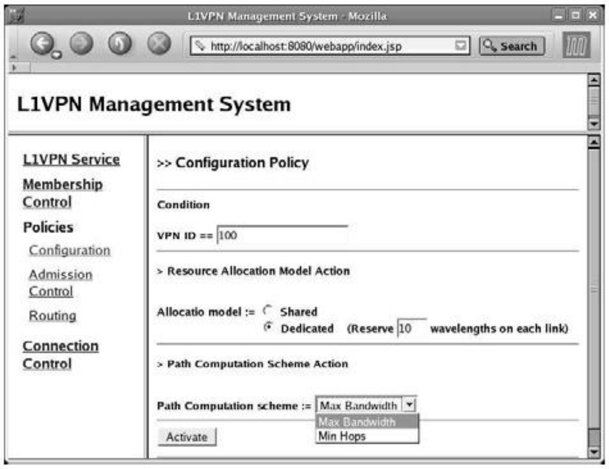
L1VPN Management System - Mozilla
http://localhost:8080/webapp/index.jsp
Search

L1VPN Management System

L1VPN Service
Membership Control
Policies
Configuration
Admission Control
Routing
Connection Control

>> Configuration Policy

Condition

VPN ID == 100

> Resource Allocation Model Action

Allocatio model := ○ Shared
● Dedicated (Reserve 10 wavelengths on each link)

> Path Computation Scheme Action

Path Computation scheme := Max Bandwidth ▾
Max Bandwidth
Min Hops

Activate

Figure 8. Web-based management interface.

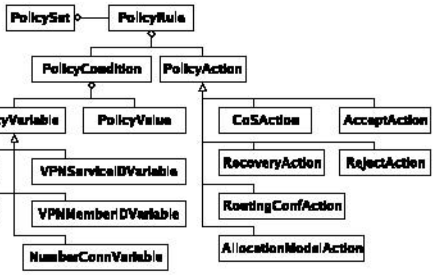

Figure 9. L1VPN policy model.

```
1 <?xml version='1.0'?>
2 <!DOCTYPE Policy SYSTEM ``l1vpnPolicy.dtd''>
3 <Policy id=``001''>
4   <PolicySet>
5     <PolicyRule type=``configuration''>
6       <PolicyCondition>
7         <VPNServiceIDVariable/>
8         <PolicyValue>vpnA</PolicyValue>
9       </PolicyCondition>
10      <PolicyAction>
11        <ResourceAllocationAction allocationModel
              =``dedicated''/>
12        <CoSAction model=``basic'' class=``gold
              ''/>
13        <RecoveryAction recoveryScheme=``
              protection:1+1''/>
14      </PolicyAction>
15    </PolicyRule>
16    <PolicyRule type=``admissionControl''>
17      <PolicyCondition>
18        <VPNServiceIDVariable/>
19        <PolicyValue>vpnA</PolicyValue>
20        <IngressMemberIDVariable/>
21        <PolicyValue>CPI1−PPI1</PolicyValue>
22        <EgressMemberIDVariable/>
23        <PolicyValue>CPI2−PPI2</PolicyValue>
24      </PolicyCondition>
25      <PolicyAction>
26        <RejectAction/>
27      </PolicyAction>
28    </PolicyRule>
29  </PolicySet>
30 </Policy>
```

Figure 10. XML policy example.

connection setup by sending a request to the control plane agent in the ingress node. Then the control agents communicate with each other to set up the optical connections.

In addition, we have investigated the use of XML technologies for policy representation mainly due to the flexibility, interoperability and availability of syntax check tools. To illustrate XML policies we consider the simplified policy model presented in Fig. 9. This model is based on the Policy Core Information Model (PCIM) defined within IETF [11, 10]. It was specified according to the policy classes defined in Section 3 and covers configuration policies and admission control policies related to connection restriction and control of number of connections. In this context, variables are associated with values to define conditions and conditions are associated with actions to define policy rules. The specified policy actions can be used to reject or accept connections and to define VPN configuration parameters as those discussed in the description of configuration policies.

An XML policy example is illustrated in Fig. 10 considering the policy model. Such policy specifies a configuration and an admission control rule. Both rules should be applied to "L1VPN service A" as expressed in the respective conditions (lines 7-8, and lines 18-19). The first rule defines actions to configure the resource allocation model as dedicated, the class of service, and the recovery scheme (lines 11 to 13). The second one is an example of connectivity restriction. It specifies that connections from member CPI1-PPI1 to member CPI2-PPI2, as expressed in the condition (lines 20 to 23), must be rejected (line 26).

6 Evaluation

The simplification and automation of the service management process are the main advantages of the policy-based approach [20]. Simplification is achieved through two key factors. First, all configuration is defined in a centralized way instead of configuring each device itself. In this way, when a new CE device is connected to a PE, the PE can be provisioned with respective L1VPN configuration through the policy protocol. Second, high-level abstractions simplify policy definition. Administrators can specify service-level policies taking into consideration L1VPN service aspects rather then technology specific details. Automation is achieved since administrators do not need to configure the service themselves. They only state system-wide policies which should guide the entities involved in the provisioning of L1VPN services.

We have performed simulations in order to evaluate the effects and implications of configuration policies. From the perspective of the L1VPN customer we measure the con-

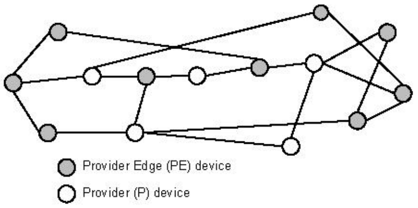

Figure 11. Simulated topology.

nection blocking rate. A connection is blocked if there are not enough available wavelengths. From the perspective of the service provider we measure the network resource utilization rate, where resource means wavelengths.

In all simulated scenarios, the connection request arrival rate is based on a Poisson distribution where the average number of connections per second is 100 or a fraction of 100 when explicitly mentioned (e.g. "Arrival Rate = 0.4" means the average is 40). The connection holding time and the interval between two connection requests are based on an exponential distribution. The topology of the simulated provider optical network is presented in Fig. 11 which shows the PE and P devices. For each L1VPN customer there is one CE connected to each PE. The source and destination nodes of the connections are randomly selected according to a uniform distribution. The simulations are repeated 100 times and the presented results are the average.

A customer should be able to specify policies to manage their L1VPN services. However providers must supervise that task since a customer policy influences the overall provider network performance and may affect another customer service. Therefore L1VPN service provisioning is managed by a combination of provider and customer policies. Furthermore, providers can use policies to define differentiated classes of L1VPN services. The following scenarios illustrate these aspects.

In a first scenario we consider four L1VPN services (L1VPN 0-3). Each customer requests a total of 2500 connections and each optical link has 32 wavelengths. Configurations policies define high priority and low priority services: (1) *High priority service*: Resource allocation model is dedicated and 10 wavelengths on each link are reserved for the L1VPN service. Route computation involves only dedicated resources and must select the path with the maximum available resource. This is achieved by considering the weight of a link is $\frac{1}{w}$, where w is the number of available wavelengths. (2) *Low priority service*: Resource allocation model is shared. Path computation involves shared provider network resources and must find the shortest path in terms of the number of hops.

The Fig. 12 shows the blocking rate of the L1VPN 0 and the average blocking rate of the other L1VPNs when all L1VPNs are defined as low priority services. The Fig. 13 shows new rates when L1VPN 0 is assigned as high priority and the other L1VPNs remain as low priority services. The results demonstrate that the service provider may define configuration policies to differentiate L1VPN services. Moreover, the changes on blocking rate show how policies for a service can affect other ones.

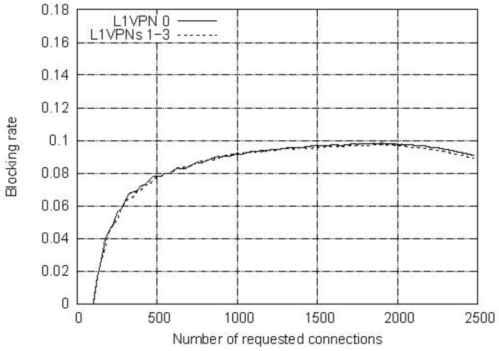

Figure 12. Blocking rate.

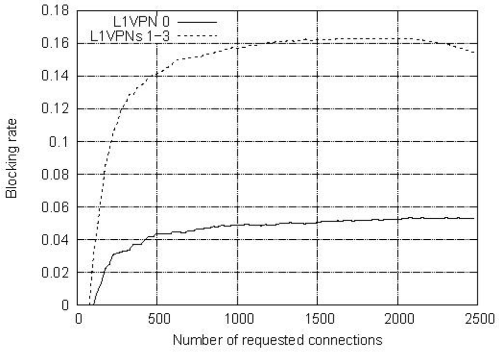

Figure 13. Differentiated blocking rate.

A similar scenario illustrates how policies can be used in reaction to specific network conditions. Fig. 14 shows the improvement on L1VPN 0 blocking rate when a policy is activated in order to assign L1VPN 0 as high priority. In this case, before such policy is activated, wavelengths in each link are dedicated for the L1VPN 0. The policy is activated after the L1VPN 0 had requested half of the total number of connections. More elaborated conditions are possible, for instance, such policies can be activated in the occurrence of specific events or when a performance degradation threshold is reached. This way, L1VPN management automation is improved.

In previous scenarios, policy configuration defined dedicated model to high priority L1VPN services in order to improve service blocking rate. However, dedicated resources may degrade the overall performance. This is illustrated in Fig. 15. Here, it is compared the average blocking rate of all L1VPN services when they use the same allocation model

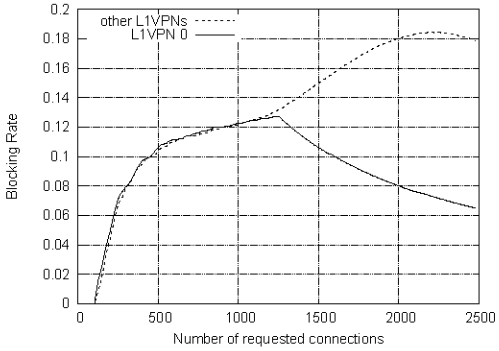

Figure 14. Policy activation effect.

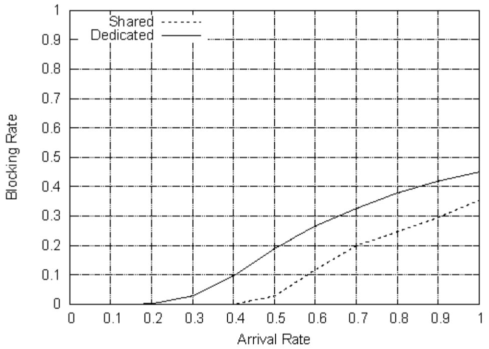

Figure 15. Blocking rate: shared x dedicated.

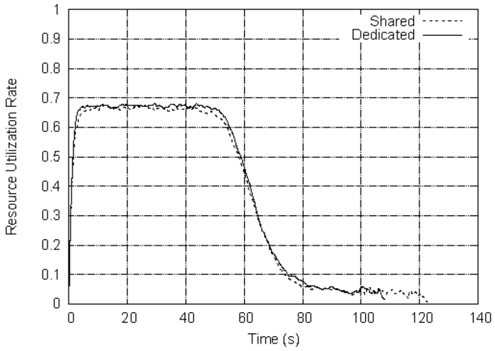

Figure 16. Arrival rate = 0.4.

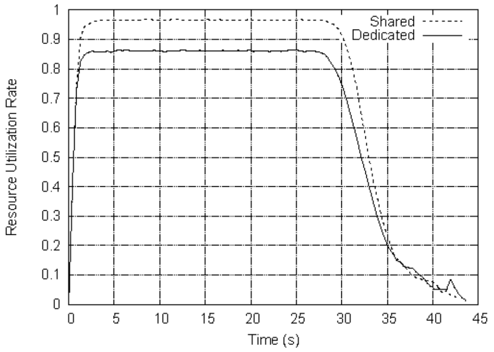

Figure 17. Arrival rate = 1.0.

at different connection request arrival rates. The blocking rate is measured with dedicated allocation model for all L1VPNs and then the blocking rate is measured again but with shared allocation model for all L1VPNs. The graph shows the average blocking rate of all L1VPNs in both cases. When all L1VPNs use the dedicated model the network resources (wavelengths) are equally distributed and reserved to each one. The results show the shared resource allocation model outperforms the dedicated model.

Moreover, the customer configuration policies can impact on the overall provider network performance. We evaluated the effect on the provider network resource utilization at different connection request arrival rates. Here, the resource utilization rate is measured in terms of the number of allocated wavelengths. For each arrival rate, the utilization rate is first measured with dedicated allocation model for all L1VPNs and then with shared allocation model for all L1VPNs. As shown in Fig. 16 and 17, the results demonstrate that the dedicated model is less efficient than the shared model. With low arrival rate there is no significant difference, as shown in Fig. 16. However, when the network load increases, the shared model outperforms the dedicated model, as shown in Fig. 17.

The results show how policies may be used to offer differentiated classes for L1VPN services and to control service behavior and performance. Providers must analyze and elaborate on which configuration policies customers should be allowed to specify, since the same policy may benefit a customer while degrading the overall provider network performance. Moreover, a policy that improves the performance of an L1VPN service may degrade the performance of another service, depending on the network conditions and the configuration of other services. For instance, in the case of the resource allocation configuration policy, the results showed advantages and disadvantages at different perspectives. In the scenario presented in Fig. 13 and Fig. 14, the dedicated model was successfully used to improve the performance of the L1VPN 0 by decreasing its blocking rate. Nevertheless, the average blocking rate of the other services was increased. Also, we saw that the dedicated resource allocation model proved to be less efficient with respect to the blocking rate when all services were configured to use the same model, as shown in Fig. 15. However, in dedicated model the path computation algorithm can optimize resource allocation since resources are not shared and detailed availability information is possible. The work presented in [18] proposes path computation algorithms and describe a multilayer scenario in which dedicated model outperforms shared model in terms of blocking rate. The authors argue that the efficiency of optimized path computation algorithms in dedicated model can overcome the disadvantage of not sharing provider network resources.

In summary, the simulated scenarios demonstrate that by supporting policy-based management the proposed architecture is a suitable and flexible approach to provide separate L1VPN management for customers. However, a challenge issue is how to estimate the overall effect of policies. It is a difficult task to evaluate how configuration policies for one L1VPN service can affect the behavior of others and the provider network. The providers must develop mechanisms and tools to simulate and evaluate policy effects. Also there can be conflict between policies of customers and provider network administrators. Priority mechanisms can be used to resolve policy conflicts and regulate the level of policy control that is given to customers.

7 Conclusion

We have described a policy-based framework for L1VPN management and proposed major classes for L1VPN policies. Moreover, we have proposed a policy-based L1VPN management architecture. A prototype system has been implemented in order to validate this architecture. The simulations have demonstrated the feasibility of the architecture and the effects and implications of configuration policies considering different points of view. We have demonstrated how Policy-Based Management can be used in the configuration management of L1VPN services. Indeed, the proposed management architecture is a suitable approach to enable multiple customers to manage their L1VPN services. However, it is not easy to service providers the task of supervising and isolating the customer control.

Future work will include to investigate how to provide Layer 1 VPN services over multiple routing domains, evaluate the impact of other L1VPN policies over the provider network, and investigate how L1VPN resource allocation can benefit from policy-based management. Other interesting work is to implement a prototype of the proposed architecture considering the control-based L1VPN service model and evaluate scalability issues.

References

[1] A. Banerjee, J. Drake, J. Lang, B. Turner, K. Kompella, and Y. Rekhter. Generalized Multiprotocol Label Switching: An Overview of Routing and Management Enhancements, 2001.

[2] L. Berger. GMPLS Signaling Resource ReserVation Protocol-Traffic Engineering (RSVP-TE) Extensions. IETF RFC 3473, January 2003.

[3] C. Carvalho, E. Madeira, F. Verdi, and M. Magalhães. Policy-Based Fault Management for Integrating IP over Optical Networks. In *5th IEEE International Workshop on IP Operations and Management, IPOM 2005*, volume 3751 of *Lecture Notes in Computer Science*, pages 88–97. 2005.

[4] K. Chan, D. Durham, S. Gai, S. Herzog, K. McCloghrie, F. Reichmeyer, J. Seligson, A. Smith, and R. Yavatkar. COPS Usage for Policy Provisioning (COPS-PR). IETF RFC 3084, March 2001.

[5] D. Durham, R. Cohen, J. Boyle, S. Herzog, R. Rajan, and A. Sastry. The COPS (Common Open Policy Service) Protocol. IETF RFC 2748, January 2000.

[6] ITU. Layer 1 Virtual Private Network generic requirements and architecture elements. ITU-T Recommendation Y.1312, September 2003.

[7] ITU. Layer 1 Virtual Private Network service and network architectures. ITU-T Recommendation Y.1313, July 2004.

[8] K. Kompella and Y. Rekhter. OSPF Extensions in Support of GMPLS. IETF RFC 4203, October 2005.

[9] E. Mannie. Generalized Multi-Protocol Label Switching Architecture. IETF RFC 3945, October 2004.

[10] B. Moore. Policy Core Information Model (PCIM) Extensions. IETF RFC 3460, January 2003.

[11] B. Moore, E. Ellesson, J. Strassner, and A. Westerinen. Policy Core Information Model – Version 1 Specification. IETF RFC 3060, February 2001.

[12] H. Ould-Brahim. GVPN Services: Generalized VPN Services using BGP and GMPLS Toolkit. IETF Internet-Draft, work in progress, February 2005.

[13] T. Takeda, R. Aubin, M. Carugi, I. Inoue, and H. Ould-Brahim. Framework and Requirements for Layer 1 Virtual Private Networks. IETF Internet-Draft, work in progress, August 2005.

[14] T. Takeda, D. Brungard, A. Farrel, H. Ould-Brahim, and D. Papadimitriou. Applicability analysis of GMPLS protocols to Layer 1 Virtual Private Networks. IETF Internet-Draft, work in progress, September 2005.

[15] T. Takeda, D. Brungard, D. Papadimitriou, and H. Ould-Brahim. Layer 1 Virtual Private Networks: driving forces and realization by GMPLS. *Communications Magazine, IEEE*, 43(7):60–67, 2005.

[16] T. Takeda, I. Inoue, R. Aubin, and M. Carugi. Layer 1 Virtual Private Networks: service concepts, architecture requirements, and related advances in standardization. *Communications Magazine, IEEE*, 42(6):132–138, 2004.

[17] T. Takeda, H. Kojima, and I. Inoue. Layer 1 VPN architecture and its evaluation. In *Communications, and the 5th International Symposium on Multi-Dimensional Mobile Communications. Joint Conference of the 10th Asia-Pacific Conference on*, volume 2, pages 612–616, 2004.

[18] T. Takeda, H. Kojima, N. Matsuura, and I. Inoue. Resource allocation method for optical VPN. In *Optical Fiber Communication Conference, 2004. OFC 2004*, volume 1, 2004.

[19] F. Verdi, C. Carvalho, M. Magalhães, and E. Madeira. Policy-based Grooming in Optical Networks. In *4th Latin American Network Operations and Management Symposium, LANOMS 2005*, pages 125–136, Brazil, August 2005.

[20] D. C. Verma. Simplifying Network Administration Using Policy-Based Management. *Network, IEEE*, 16(2):20–26, 2002.