Lucent Technologies
Bell Labs Innovations

WHITE PAPER

# Next Generation VPNs

*Network-Based services based on Virtual Routing and MPLS deliver truly scaleable, customizable VPNs*

Abbas Bagasrawala

# Network-Based VPNs using Virtual Routers and MPLS
# The Importance to Network Service Providers

*Network Service Providers are faced with challenges in deploying IP-based network services. Among these are flexible service provisioning, scalability, operational management, value-added service enhancements, billing flexibility and cost containment. A single solution that addresses these issues is a Lucent Technologies SpringTide™ IP Service Switch that offers a highly scalable Virtual Router architecture linked to the MPLS protocol suite so that service providers can deploy effective network-based VPNs for their customers. This paper explores how this solution meets the needs and goals of the service provider community.*

# Introduction

Virtual Private Networks (VPNs) are the top network choice for organizations of all types and size. VPNs have steadily achieved a dominant role in enterprise networking as IP-based networks have increased in sophistication and technological dependability.

VPNs enable enterprises to extend their LANs to a variety of off-site locations; they support highly efficient intranets, and they enable the deployment of complex and efficient extranets - all with an attractive cost-performance and flexible connectivity.

Carriers and other Service Providers (SPs), faced with a dramatic change in traditional revenue sources, have moved to support enterprise VPNs with a variety of network enhancements and advanced features. Their ultimate objective is to deploy network-based VPNs (NW- VPN) that enable them to offer a suite of new service offerings, while maximizing their bottom line profitability.

# 1 The Network-based VPN Challenge

The traditional VPN model has a number of handicaps when compared to the SPs' objectives. In a traditional VPN, for instance, the router on the customer edge (CE) ordinarily provides the management of the VPN, including provisioning, support for routing and forwarding tables, traffic control, and security. In some cases, these capabilities are assigned to a router on the provider edge (PE) that is dedicated to the customer's traffic and use, but its functions essentially remain the same. This early approach to SP-supported VPNs, built on a hardware-based platform, provides limited scalability and quite limited opportunities for distinctive service offerings, class of service traffic management or service billing.

These VPNs use an overlay model that consists of an expensive, fully connected mesh of virtual channels (VCs) with the endpoints on the customer premises (or the surrogate router on the provider edge) transported over the controlled carrier network. In this model, packet ingress and egress are the same - the VC mesh is visible to the CE routing and requires labor intensive, skilled management of a virtual backbone. In a more contemporary variation, VCs are replaced with tunnels terminated at the provider edge, but otherwise it has the same limitations.

Network-based VPNs (NW-VPNs) require much more than this simple surrogate functionality. In order for their service offerings to be attractive to their customers and to be operationally feasible for the SPs, the NW-VPNs need to address the following concerns:

- **Scalability** - the platform and architecture used by the SPs must enable them to support thousands of customers, without building massive racks of physical routers. The deployment of racks of physical routers is naturally limited from a cost and space usage perspective. A network-based solution has to provide a different, cost-efficient method for dealing with this.

- **Broad Flexibility in Service Offerings** - Each customer will have many similar needs, but most will have their own unique set of variations on a theme. Some will have a significant need for encryption, for example, while others might have a need to support voice traffic and a converged voice-data architecture. The challenge for providers is to build a platform that enables a relatively rich mix of service offerings that can be easily deployed and billed, so the SPs can meet the unique needs of their customers and increase revenue opportunities.

- **Single Platform** - Thirdly, the architecture should be based on a single, coherent platform that accommodates not only the basic VPN routing and control functions, but also all of its related functions, such as encryption and security. Alternative solutions that build services with a piecemeal assortment of devices are often prone to dysfunctional performance due to interoperability problems, and usually demand more management skill, as well as being more costly.

- **Class of Service and QoS Capabilities** - Historically, VPNs have been used to segregate traffic flows so that performance and privacy can be assured. A key challenge for service providers is to offer bandwidth privacy guarantees tailored to each customer's needs and yet be able to scale to thousands of VPN customers. SPs need to be able to provide application-level performance assurances, as well as support for different VPN types such as a virtual private LAN segment operating at the MAC layer or a routed IP-based VPN.

- **Ease of Management** - Key to a SP's interests is the ease and flexibility of the operations and management of the NW-VPN. It is obvious that supporting a large population of customers represents a significant operational challenge. The right NW-VPN approach has to assign the highest priority to dealing with this by using three different strategies:

- ✣ First, by simplifying the management task - unifying the management functions within a single platform, with a coherent user interface, consistent operational procedures, and integrated reporting.
- ✣ Secondly, by distributing appropriate management functions between the PE and CE routers. The customer need not be concerned with the SP's network, but does need to manage and assign rules-based traffic management tools as well as firewall policies. A thoughtful demarcation between the functions that rightfully belong to the user versus the service provider can really streamline the operation.
- ✣ A third key element is the use of a unique directory-based provisioning model that provides easily managed service definitions. This directory-based provisioning is highly scalable, and it provides a consistent view of services to end users, irrespective of their point of connection to the SP's network.

- **Customized Security Options** - Each customer has varying security requirements. Extranets, for example, require different levels of security for corporate users and non-corporate users who require higher levels of screening. Third party routers may also be part of the VPN, and the VPN may even have segments supported by a different, non-trusted service providers. Easily deployed differing levels of encryption may be necessary for the efficient and secure operation of the VPN. Without some way to quickly apply security measures on an as-needed basis, complex network-based VPNs cannot be supported.
- **Billing Granularity** - Service providers that offer a wide range of flexible applications to a large population of corporate users must have some way to monitor, measure and bill for these differing service options. Without this, the promise of enhanced revenues and long term growth cannot be realized.
- **Cost Effectiveness** - Of course, a healthy bottom line is essential for service providers. This requires a VPN solution that maximizes the use of network and platform resources - one that supports a reasonable migration and interoperability strategy so that the use of existing infrastructure can be continued and maximized while the new architecture is being deployed. Furthermore, the solution must be massively scalable so that it provides 'future-proving', with long-term investment protection.
- **Opaque Transport of Data Between Sites** - VPN traffic should be carried transparently across the SP network and should be unrelated to other traffic carried by the IP backbone. This is needed because the VPN addressing scheme is often unrelated to the IP backbone and private IP addresses might conflict with other IP backbone traffic, and, in some cases, the VPN traffic might use a different protocol than the IP backbone.
- **Migration Path** - Finally, the strategy must have a clear migration path from the customer's existing models to the new model. It must provide the basis for moving from an overlay, tunnel-based model, with its costly and cumbersome fully connected mesh, to a peer-peer VPN model.

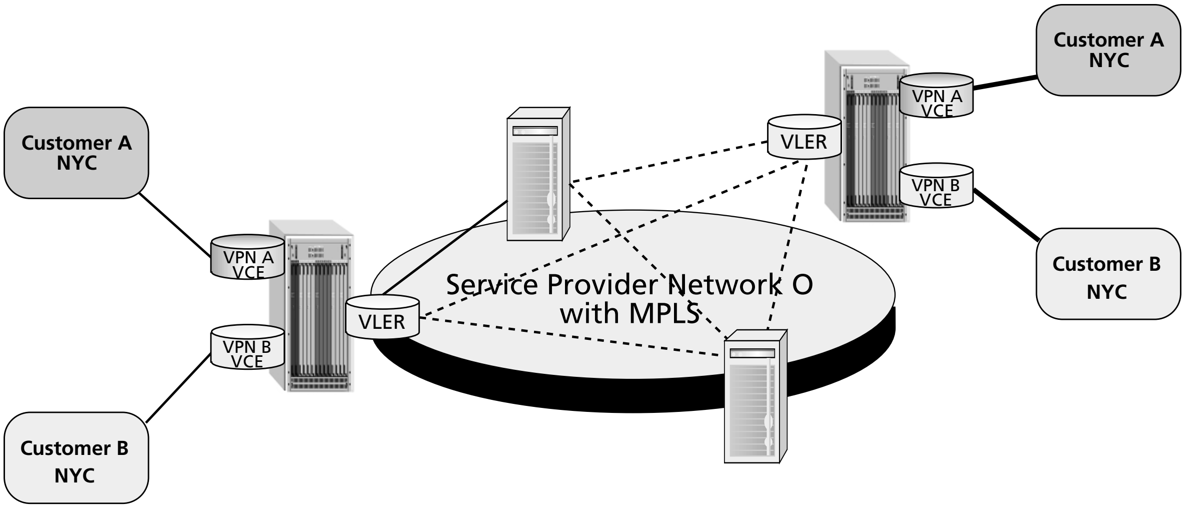

*Figure 1 - Lucent Technologies Virtual Router/MPLS Structure*

## 2 The Lucent Technologies Solution for Providing Network-Based VPNs

The Lucent Technologies answer to providing the capabilities described in the foregoing section is an architecture that is compatible with IETF RFC 2547 and which uses Virtual Routers (VR) and MPLS. A key attribute of this solution is the functional separation of the network layer from the service layer within the SpringTide™ IP Service Switch.

The SpringTide™ IP Service Switch includes a Virtual Label Edge Router (VLER) that interfaces with the provider core network and which can connect to other RFC2547 devices in the network, and a Virtual CE (VCE) router that connects to the customer premise device, (which could be a simple router or a switch). (Figure 1).

The VCE and VLER are connected internally through the switch fabric and are transparent to the user and administrator. This platform approach enables service providers to deliver positive, profitable, and flexible services such as Firewalls and Encryption with consistent levels of security and quality. The customer VPNs are managed within the service layer of the VR and the Virtual Label Edge Router (VLER) supports the MPLS networking. (See section 3.1 for more information on virtual router functionality.)

The SpringTide™ VR platform can be interconnected to the customer device through any access method including digital subscriber line (DSL), cable, wireless, LAN line, ATM PVC, or dial-up.

## 2.1 VR MPLS VPN Architecture

The fundamental building block in this arrangement is the Virtual Router (VR). Virtual routers provide the secure, segregated environments required for delivering business-quality IP services. Each virtual router has its own routing information base (RIB), policy information base (PIB), management information base (MIB) and a separate MPLS data forwarding engine with its own code address space with memory, which prevents any one virtual router from affecting other virtual routers. Within each virtual router, individualized service definitions for bandwidth, priority, and security are retrieved from policy directories and provisioned on either a per-subscriber or per-traffic flow basis. The PE router platform can contain many unique virtual routers supporting large numbers of VPNs.

Since the virtual router software is not run as a separate task in the underlying operating system, CPU resources and memory utilization are independent of other virtual routers in the same system. The performance of one virtual router does not affect other virtual routers in the system.

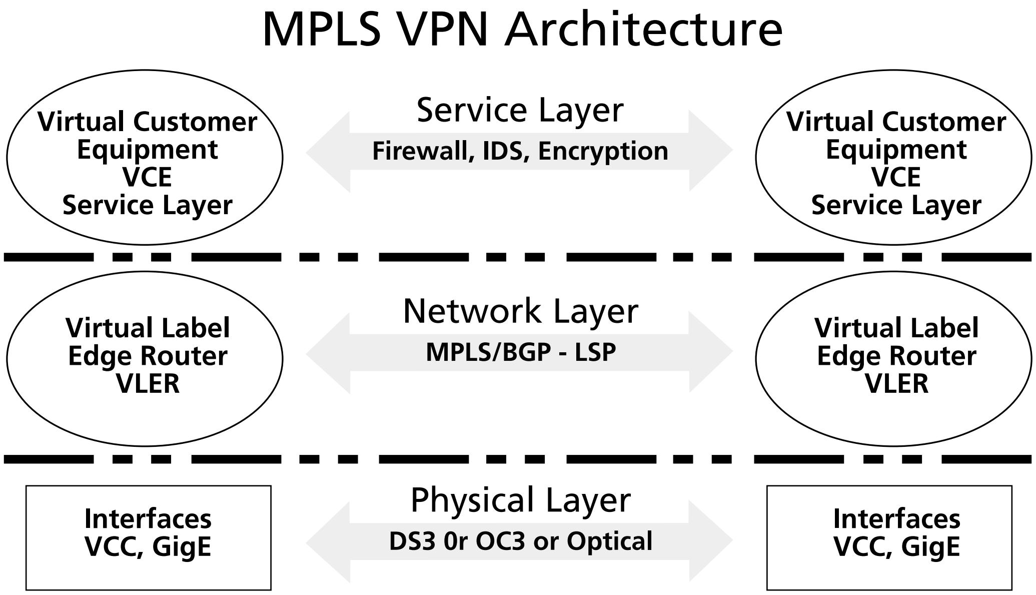
MPLS VPN Architecture
Virtual Customer Equipment VCE Service Layer
Service Layer
Firewall, IDS, Encryption
Virtual Customer Equipment VCE Service Layer
Virtual Label Edge Router VLER
Network Layer
MPLS/BGP - LSP
Virtual Label Edge Router VLER
Interfaces VCC, GigE
Physical Layer
DS3 0r OC3 or Optical
Interfaces VCC, GigE

*Figure 2. - The VR- MPLS Architecture*

## 2.2 Two Logical Layers

The SpringTide architecture consists of two logical layers - The Service Layer and The Network Layer. The CE router supports the Service Layer, while the VRs in the PE router platform acts as virtual label edge routers (VLER), and interface with the SP backbone and other RFC2547 compliant devices. The VLER supports the Network Layer. These two layers are tightly integrated and provide an elegant solution to building VPNs. (Figure 2). It is this clear separation of logical functions, that enables SPs to deploy customized options, application-layer performance and as-needed security solutions.

The advantage of this architecture is in its flexibility and scalability. Flexibility is provided through VPN management, which enable PE VRs to maintain separate route, forwarding, and MIB databases for each VPN. The database separation of each VR allows the service provider and/or its customers to monitor traffic statistics, and to configure network policies - for example, to build dynamic on-the-fly Extranets with business partners.

## 2.3 MPLS

In this architecture, MultiProtocol Label Switching (MPLS) is used for forwarding packets over the backbone, and BGP (Border Gateway Protocol) is used for distributing routes over the backbone.

MPLS facilitates the high-speed transport of IP traffic across the wide area network. To accomplish this it assigns labels to IP flows, and places them in the IP frames, which can then be transported across either packet or cell-based networks. Traffic is then switched according to those labels rather than being routed using the usual IP address look-up. In a traditional IP network, an IP packet is forwarded through a network on a hop-by-hop basis using interior gateway protocol (IGP) such as RIP or OSPF etc. The forwarding decision is made by looking up the packet-destination layer3 address (IP address) against the routing table to determine the next hop. In contrast, MPLS uses labels to forward IP packets. The PE virtual router at the edge of the MPLS network attaches labels to packets based on a forwarding equivalence class (FEC). Packets are then forwarded through the MPLS network based on their associated FECs by swapping the labels through the core routers.

Using MPLS allows the Service Provider to maximize resource utilization, while assuring guaranteed levels of performance according to assigned Classes of Service (CoS).

| # | Outer Tunnel | Traffic MUX technique | Max CoS supported | Comments |
|---|---|---|---|---|
| 1 | One tunnel to each VLER.<br>Different VPN and COS traffic multiplexed on the same tunnel | Different VPN and COS traffic are multiplexed on the same tunnel | 8 | Highly scalable.Mnimizes the number of tunnels to each VLER. Both VPN and COS are MUX'ed |
| 2 | One tunnel for each VPN;<br>all COS for that VPN carried in the same tunnel. | Individual tunnel for each VPN. COS assigned through EXP bits | 8 | Advantage is VPN traffic is separate. Suitable to create VLL type of VPN. This is similar to tunnel based VPN except MPLS tunnels are used instead of IPSEC or L2TP etc. |
| 3 | One tunnel for each COS;<br>different VPN with same COS multiplexed on the same tunnel | VPN separation by inner label. | 64 | Good if SP wants to provide max. 64 COS. |
| 4 | One tunnel for each VPN and COS. | Individual tunnel for COS and VPN | 64 | Not scalable. Maximum of number of tunnels across the backbone. |

*Table 1. - Class of Service Options*

## 2.4 Class of Service - CoS

The SpringTide™ IP Service Switch uses MPLS to enable service providers to offer bandwidth guarantees that are tailored to each customer's needs, while still scaling to thousands of VPN customers. The service provider has the flexibility to offer a variety of Class of Service options, depending on the range of services its customers demand and the Label Switch Routers (LSPs) that are set up in the network core.

This VR-MPLS architecture supports a Differentiated Service (Diff-Serv) style of CoS provisioning. In the Diff-Serv model, the customer is guaranteed a minimum, pre-negotiated bandwidth. The spare capacity in the network is then shared among the VPN customers according to an algorithm such as Weighted Fair Queuing (WFQ). The table below shows the various combinations of mapping VPN traffic to LSPs.

## 2.5 VR-MPLS VPN Benefits

By separating the Service Layer from the Network Layer, Lucent Technologies' VR-MPLS VPN architecture allows the service provider to easily deploy new, flexible services.

### 2.5.1 Management Flexibility

Management flexibility is assured since each virtual router maintains separate routing, forwarding, service, policy and SNMP MIB databases. This allows the service provider to share VPN management with the customer in any appropriate manner that meets the customer's unique needs. In this environment, for example, the customer can easily add, change or remove policies that allow access to its business partners, without involving the service provider.  With its ability to aggregate dial-in user sessions, the SP can offer remote user VPN services to their VPN customers.

Furthermore, the MPLS-based peer-to-peer VPN allows the SP to easily add new VPN customers at a given site by simply allocating a new VCE to that customer. The MIB databases are maintained per VPN, allowing the SP to monitor traffic uniquely for each VPN.

### 2.5.2 Directory Driven Model

Implementing new services for VPN users is accomplished easily through a directory-based set of service definitions. The directory-based approach is highly scalable and easy to deploy. The value of a directory-based method of deploying services cannot be over emphasized. Because provisioning options are already predetermined according to certain user classes, provisioning does not have to be uniquely developed for each new user. This assures consistency and interoperability among all of the users. Furthermore, services can be deployed quickly from a central location, and are easily scalable since the service permutations are manageable, even for large numbers of user.

### 2.5.3 Accounting Support

The SpringTide™ IP Service Switch supports the provisioning of subscriber and user services through a central repository of policy definitions. This allows the centrally located system to determines what services a user has, who the service provider is, and what billing rates apply.

### 2.5.4 Security

The SpringTide™ IP Service Switch has built-in hardware encryption engines that enable encryption services based on IPsec standards to be deployed as needed. With this inclusive platform, there is no need to add an external device for traffic encryption. The switch's policy driven stateful firewall provides granular firewall policies, which have "follow me" characteristics - so where ever the network is accessed, the SpringTide switch will gather and implement the customer's specific security policies. The SpringTide IP Service Switch also supports intrusion detection and denial of service protection, assuring the customer of complete security confidence.

# 3 Virtual Routers and MPLS

## 3.1 Virtual Routers

The SpringTide™ IP Service Switch is a virtual router platform that logically subdivides a physical router into multiple software-based routers, each of which has its own dedicated I/O ports, buffer memory, routing table and network management software. This virtual routing architecture uses modular software with multiple implementations operating simultaneously in a true multiprocessing environment.

While the routing functions are implemented in software, the packet-forwarding functions are implemented in hardware, in order to deliver wire-speed performance when connected to high-speed broadband facilities. This hardware capability is flexibly assigned to each virtual router on an as needed basis.

The virtual router software controls the wire-speed hardware, the physical I/O ports or label switched paths used to transmit and receive packets.

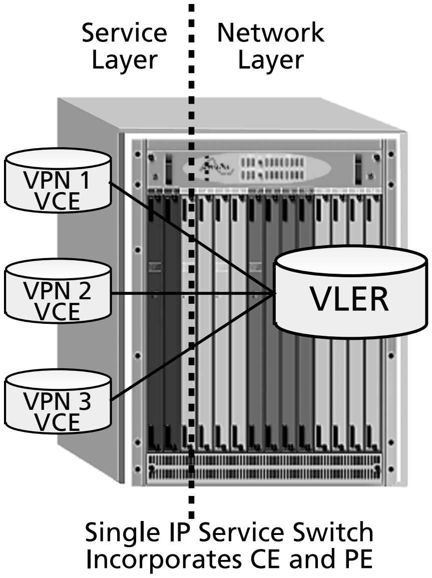

*Figure 4 - The SpringTide™ Virtual Router Platform*

Virtual routers can be configured so that each one has a managed resource capacity. This assures the service provider, for example, that each subscriber's packet-buffer memory allocation and forwarding tables do not adversely interfere with the operation of other virtual routers.

### 3.1.1 Separate Protocols

Each virtual router can execute separate instances of routing protocols such as RIP, Open Shortest Path First (OSPF) and Border Gateway Protocol (BGP) as well as network management software such as SNMP. With this, the Service Provider can let customers independently monitor their own network elements. Customers can also independently configure their own IP address domain, manage a user-based security model, and exercise the VR's other management functions.  Because each VR's packet forwarding path is unique, and allows independent traffic policing, users can independently oversee and manage the performance of each virtual router.

### 3.1.2 Client by Client Customization

Virtual routers give Service Providers the ability to deploy new Internet services on a client-by-client basis while enabling comprehensive management and administration, network security and fine-tuned network performance.

Virtual routing capabilities provide the ability to dynamically provision exact end-user bandwidth needs to match specific applications, while offering maximum end-user control and management of that bandwidth.  This delivers to Service Providers the capability of developing suites of competitively priced, highly customized IP services.

A statement in preliminary RFC2547 documentation say, in part: "If every router in an SP's backbone had to maintain routing information for all the VPNs supported by the SP, this model would have severe scalability problems; the number of sites that could be supported would be limited by the amount of routing information that could be held in a single router. It is important to require therefore that the routing information about a particular VPN be present ONLY in those PE routers that attach to that VPN. In particular, the P routers should not need to have ANY per-VPN routing information whatsoever."

The SpringTide™ solution does this very well, enabling the SP to deploy highly efficient, flexible services.

## 3.2 MPLS

Multiprotocol Label Switching (MPLS) was originally intended to improve the packet forwarding efficiency of routers. It has also evolved as the best, currently available way to deploy some key applications, including VPNs. Traffic engineering with MPLS allows network operators to dictate the path that traffic takes through their network in order to improve performance.

MPLS has become a key technology and an effective means of deploying IP networks across WAN backbones. When the Internet Engineering Task Force (IETF) developed MPLS in 1998, it was expected that it would be one of the most important network developments of the 1990's.

The essence of MPLS is the generation of a short fixed-length 'label' that acts as a shorthand representation of an IP packet's header, and the use of that label to make forwarding decisions about the packet. As part of MPLS, packets have a field in their 'header' that contains the address to which the packet is to be routed. Then networks process this information at every router in a packet's path through the network (hop by hop routing).

In MPLS, the IP packets are encapsulated with these labels by the first MPLS device they encounter (in this case the PE Virtual router), as they enter the network. The MPLS PE router analyses the contents of the IP header and selects an appropriate label with which to encapsulate the packet. Part of the great power of MPLS comes from the fact that, in contrast to conventional IP routing, this analysis can be based on more than just the destination address carried in the IP header. At all the subsequent nodes within the network the MPLS label, and not the IP header, is used to make the forwarding decision for the packet.

The packet handling nodes, or routers, are called Label Switched Routers (LSRs) and are usually the Provider routers (P routers) in the core network.

This is different from conventional IP routers that contain 'routing tables' that are reference using the IP header from a packet to decide how to forward that packet. These tables are built by IP routing protocols (e.g., RIP or OSPF), which carry around IP information in the form of IP addresses. In practice, these forwarding (IP header lookup) and control planes (generation of the routing tables) are tightly coupled. Since MPLS forwarding is based on labels it is possible to cleanly separate the (label-based) forwarding plane from the routing protocol control plane. By separating the two, each can be modified independently. With such a separation, there is no need to change the forwarding machinery.

There are two broad categories of Label routers. At the edge of the network are high performance packet classifiers that can apply (and remove) the requisite labels: These are the MPLS Label Edge Routers (LER) and exist as virtual routers in the SpringTide™ platform. Core LSRs are the P routers and need to be capable of processing the labeled packets at extremely high bandwidths.

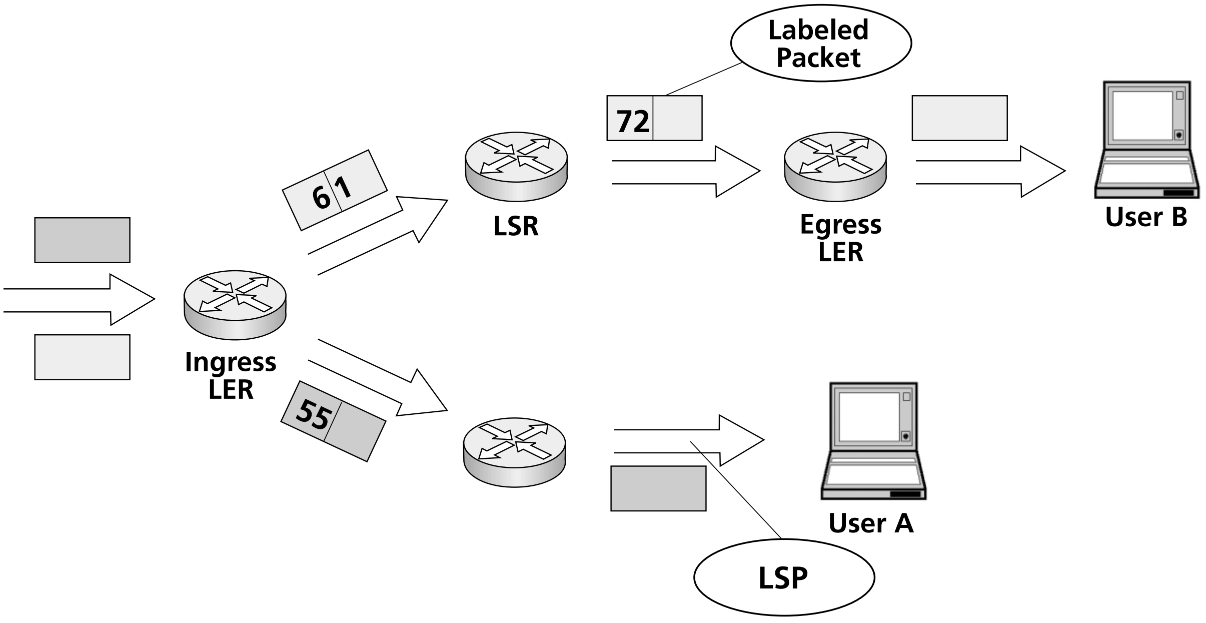

*Figure 5 - an MPLS Route Example*

The above figure 5 illustrates two MPLS traffic flows, coming into an ingress router. The ingress router policy assigns label 61 to one traffic flow and label 55 to the other. In the top path example, the LSRs replace the label 61 with the new label 72 and forwards the traffic to an egress router that removes the label and forwards the traffic to its destination.

## 3.3 Security for VPNs:

A key VPN requirement is data security. Depending on the SP's trust model, the SpringTide solution offers different security levels.

### 3.3.1 MPLS Label Security

If all VPN sites are within single SP network then MPLS provides enough security to separate data traffic from other VPNs in the same system. Because the data traffic is forwarded according to the label, and not the IP header, the data separation is achieved by giving different labels to different VPNs. At the ingress VCE, the data is associated with a VPN and a unique label is assigned to that traffic. The VPN label ensures that data is delivered only to the target destination.

### 3.3.2 Third-party LER

With a third-party Label Edge Router the trust model is different and the customer VPN site not under the SP's network control could potentially create a security threat. For example, if it was a rogue router, it might falsely announce label routes to a VPN.

To prevent this, the SpringTide solution uses MD5 cryptographic algorithm to validate BGP peers before accepting label and route distribution from the third-party LER.

### 3.3.3 VPN Sites Across Different SP

In some cases, VPN sites are distributed among more than one SP backbone. One option is to establish a reasonable Trust Model between the SPs in which MPLS LSP is still used to forward data across the SP network.

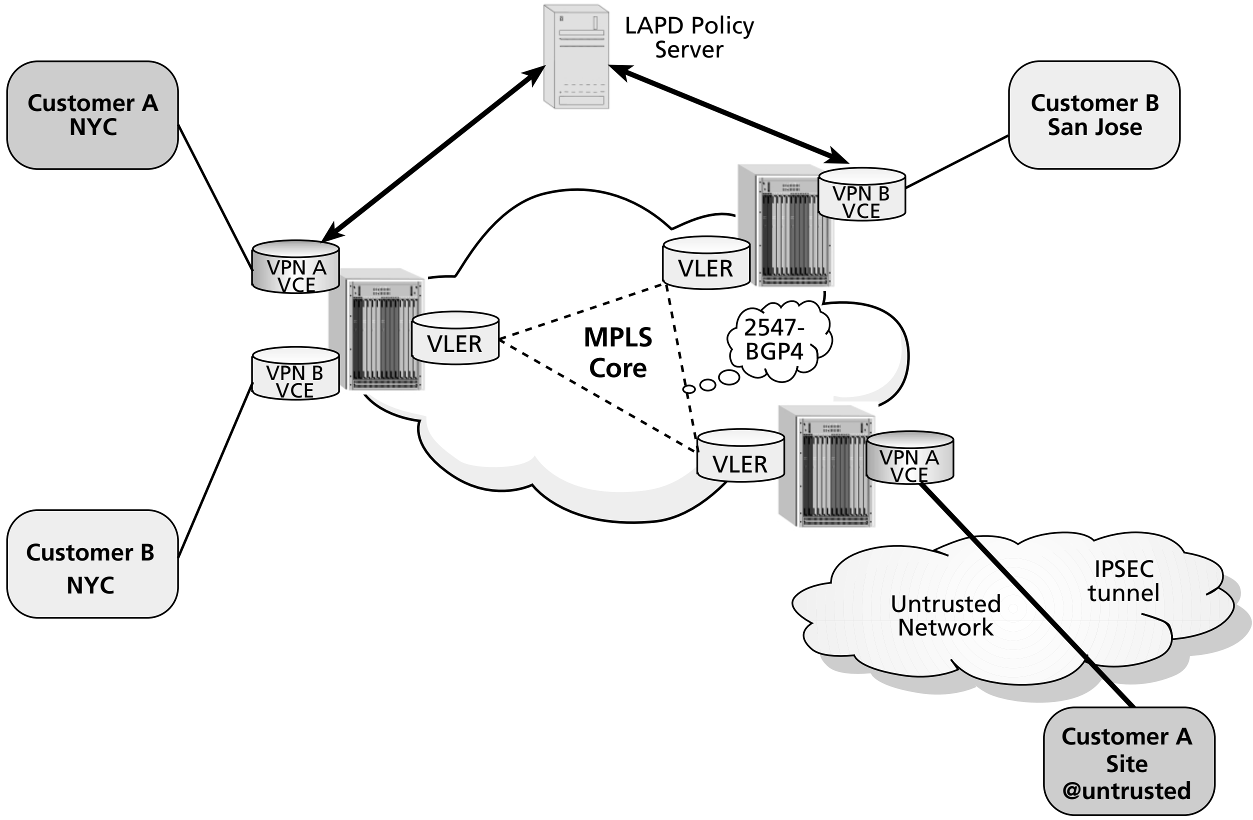

*Figure 6 - VPN Across an Untrusted SP Network*

Often, however, it is not easy or feasible to establish such a trust model between the SP's. The SpringTide solution offers an IPsec-based security level, in which an IPsec tunnel is established from the service provider's VLER directly to the customer VPN site.

### 3.3.4 Secure LSP (IPsec over MPLS)

If there is no trust level and/or the customer requests that certain traffic be strongly encrypted then the SP can provide IPsec over the MPLS-based VPN. In this case, certain customer traffic is encrypted and sent over an MPLS tunnel.  As with other traffic flows, this can be policy driven and individual traffic flow can be given a different security level - (Figure 6).

## 4 Implementation and Management of VPNs

A key value of the SpringTide IP Service Switch for service providers is in the ability to tune the behavior of the core network to better aid their business objectives. SP customers are able to be more productive with value-added networking services that can be deployed by the SPs through this ability to customize, or fine tune, the network. This capability includes:

- **Network-based and managed firewall services** - Enterprise customers as well as smaller organizations that may not have adequate staff to implement security and firewall services, often prefer to outsource this to their service provider.

- **Virtual Private Networks and Routing services** - This solution allows the customer to manage the simple routing or switching equipment on their own premise, while depending on the services of the SPs virtual routers in the core network, which is configured and managed by the SP. This allows all complex routing knowledge to exist in the provider portion of the network.

- **Bandwidth-managed services (the ability to dynamically get more bandwidth on an access link)** - The ability for a user to gain access to bandwidth from the core-networking infrastructure is an important service opportunity for carriers. In this scenario, the customer could be offered the option of gaining access to incremental amounts of bandwidth as a billable service.

- **Application-aware treatment of applications (such as voice**) - For core network providers (providers who own their own facilities and maintain core networks on their own infrastructure) the use of their network assets to transport packet voice is a key service opportunity. By multiplexing voice and data services over the same backbone, the provider is able to provide voice service at little incremental cost.  In the tough voice service market, this can give a carrier key cost advantages over traditional voice service suppliers.

- **Content direction and management** - This enables Service Providers to direct subscribers to certain types of content. A population of users interested in viewing content for their area of interest (financial trading groups for example) may find it valuable to buy such a service from their provider. They would potentially pay for the accelerated viewing of their content rather than the content of the entire world-wide-web.

- **True user roaming** - This allows the user to the same network access privileges and benefits regardless of his or her physical location and is a valuable opportunity for service providers. For example, to allow users to access corporate e-mail and web sites over a wireless medium from anywhere is a great benefit for SPs. To allow firewall and bandwidth protection as described above, from a ubiquitous platform makes this truly unique.

## 5 Summary

Service Providers need a vehicle that enables them to deploy Enterprise VPNs that can be used for extending the LAN or extending the Enterprise with Extranets. This solution has to be scalable, flexible, easily managed, cost-effective, and application sensitive.  It has to assure consistent quality of service and it has to be deployable in a way that evolves and is interoperable with the existing infrastructure.

It must offer a way of implementing sophisticated and billable service offerings to create revenue opportunities for the SPs.

The clear solution is the **Lucent Technologies SpringTide™ IP Service Switch** with two key components:

- A Virtual Router architecture that supports thousands of individual VPNs in a highly flexible, manageable and cost effective manner, and
- A MultiProtocol Label Switching (MPLS) implementation that assures high performance, by application and user, and that assures the highest level of customized security treatment.

With virtual routing and MPLS, the service provider is assured of a vehicle that can be used far into the future to support the most demanding of network opportunities.

This document is for planning purposes only and is not intended to modify or supplement any specifications or warranties relating to these products and services.

To learn more, contact your Lucent Technologies representative, authorized reseller or sales agent. Or, visit our Web site at www.lucent.com.

Specifications subject to change without notice.