Testing MPLS and IP VPNs

Agilent Technologies RouterTester
Whitepaper

Introduction

With the tightening economy in the US and rest of the world, the focus of service providers has shifted to exploring new avenues of revenue generation. As a result, there has been a dramatic increase in interest in both development of new service offerings and in the re-packaging of existing services such as Virtual Private Networking.

VPNs have been used for some time to facilitate mobile worker and telecommuter access to corporate network resources (the well known, used and accepted concept of "remote access".) The earliest remote access implementations provided users with dial-in access through modem banks maintained at the corporate headquarters. It was later recognized that use of a public network, such as the Internet could save telephone costs. This concept of using a public network for private communication is the basis of today's VPN. Today's VPNs encompass various Layer 2 and Layer 3 technologies, in conjunction with enhanced security features, and offer not only remote access, but also site-to-site connectivity.

For the purposes of this paper, the following definitions will be used:

- VPNs based on a Layer 2 (Data Link Layer) technology and managed at that layer are defined as Layer 2 VPNs (MPLS, ATM, Frame Relay) - ref. OSI Layer model
- VPNs based on tunneling above Layer 3 (Transport Layer) are defined as Layer 3 VPNs, (L2TP, IPSec, BGP/MPLS)
- IP-VPNs are a type of Layer 3 VPN, which are managed purely as an IP network (L2TP, IPSec)

Under our generic definition, therefore, VPNs are a "service" that offers secure and private data communications over a public network, through the use of standard tunneling, encryption and authentication techniques.

As various VPN-related technologies and networking protocols began to emerge and be deployed, a new work group (WG), the Provider Provisioned VPN (PPVPN) WG, was created by the Internet Engineering Task Force (IETF). This group's main task is to define the framework and requirements for specification of VPN service offerings. The recent PPVPN requirements document submitted at the PPVPN WG proposed that an SP must be able to offer Quality of Service (QoS) service to a customer for at least two generic service types: managed access VPN services or edge-to-edge QoS services.

These two service types have important distinctions. For example, BGP/MPLS VPNs, a Layer 3 service, are considered to be a managed-access VPN service, since VPN services are fully managed by an SP. However, there are many enterprises who wish to manage their own Layer 3 overlays, and many service providers who wish to provide Layer 2 interfaces to such customers. This leads to an edge-to-edge service.

Edge-to-edge services are not a new concept. Enterprises have long built their own wide area networks by purchasing wide area, point-to-point, data link layer connectivity from SPs, and then building their own Layer 3 infrastructure on top of that. The Pseudo-Wire Edge to Edge (PWE3) WG at the IETF is also defining these services, in part.

While both these approaches have their merits and drawbacks, this paper will focus on methods for verifying implementations of both these types of VPNs, in devices and in networks.

VPN Testing

In order to discuss VPN testing, it is necessary first to briefly review some basic concepts of router testing.

Any exchange of information through frames or packets (referred to collectively in this paper as "packets") requires two packet types - control packets and data packets. Typically, much device intelligence is vested in the handling of the information carried by control packets. This is accomplished through the device's "control software."

Once control packet exchange and control software establish the correct configuration, the device can make proper packet forwarding decisions and data packets can be sent across the network. Thus, device data processing design must include the ability to correctly make a high volume of forwarding decisions, at rates of up to 10 Gigabytes/sec.

Further, VPN testing must focus on both the control and data processing levels (commonly referred to as the Control and Data planes) within switch/routing devices. As will become evident in later discussion, different types of VPNs distribute processing activities differently between these two planes.

Types of Testing

Typically, any testing discussion must cover the various aspects of each type of testing, namely Conformance, Functional, Performance & Scalability, Interoperability and Service testing.

Conformance and Functional testing focuses on the details of the control and data packets themselves. This type of testing is common in the software development and system integration labs of network equipment manufacturers (NEMs).

As software begins to be integrated with hardware, it is possible to conduct performance and scalability testing on both the control and data planes. For example, the PPVPN requirements document estimates that a large service provider will require the ability to support 10,000 VPNs within four years and that the number of site interfaces will range from a few to more than 50,000 site interfaces per VPN. Clearly, any device expected to be part of such a network will need to be tested for its ability to support these levels of activity. Performance and scalability testing are generally performed at system integration and pre-deployment labs of NEMs, as well as Service Provider evaluation labs.

Before a device can be deployed in the network, it is also necessary to evaluate whether it will be used in a single or multi-vendor environment. Multi-vendor networks require interoperability testing to determine the level of interoperability that has been achieved between VPN implementations of various NEMs. Interoperability testing is critical to ensuring that SPs are not limited to use of equipment from a single NEM.

Finally, it is necessary to test the service in the actual network environment. Typically, SPs will initially deploy a particular VPN service in a certain segment of the network, for use between internal sites or by a very small set of customers, for testing purposes. When customers are involved, testing is often based on a Service Level Agreement (SLA) specifying the details of the service and Quality of Service (QoS) levels.

This paper will focus chiefly on "black box" testing (i.e. all observations are made on an external, third party test tool.) The term 'Tester T' will be used to refer to a generic tool for testing VPNs. Tester T must be able to analyze packets and emulate a network of devices, including Customer Edge (CE), Provider Edge (PE) and Provider Core (P) devices.

A CE device is the 'exit point' of the customer or enterprise network. A PE device is the first device on the SP network connected to CE devices and P devices are the core devices in the SP network (Ref. Fig. 1). The device being tested will be referred to as the System under Test (SUT). Since only a PE device is expected to have VPN functionality, the SUT will always be a PE router.

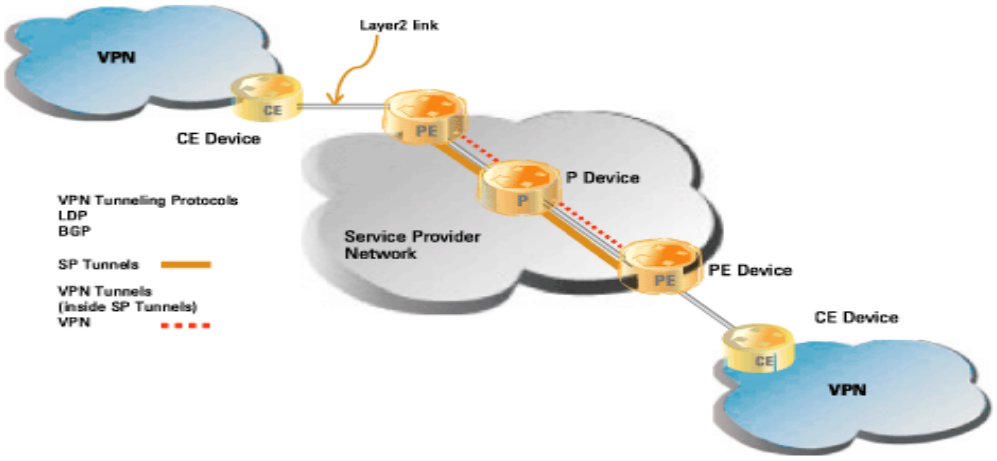

Figure 1: Virtual Private Network Scenario

Layer 2 and Layer 3 VPNs

The remainder of this paper will focus on testing Layer 2 and Layer 3 VPN services managed by SPs over their networks. (The concept of layers is taken from the OSI layer model - Layer 2 is the Data Link Layer, while Layer 3 is the Network Layer.) In this context, the phrase "Layer 3 VPN" will denote a VPN service used to carry Layer 3 traffic end-to-end, while "Layer 2 VPN" will denote a VPN service used to carry Layer 2 traffic over an SP network.

Layer 3 VPN testing will be explored first, as Layer 3 VPNs have been offered for some time by a number of NEMs and SPs. Layer 2 VPNs, specifically in the area of MPLS, are a more recent phenomenon and will be discussed in the final portion of this paper.

Finally, we will not explicitly cover "negative testing" here. Negative testing, or error injection, consists of sending packets into the PE router with incorrect contents or in incorrect sequences to evaluate the robustness of the SUT. This type of testing is a standard part of Conformance and Functionality testing, and can also be used for Stress testing to check the router's ability to handle errors under high control and data load conditions. It is expected that Tester T is capable of conducting these kinds of tests.

Layer 3 IP VPNs: IPSec based VPNs

IPSec is a security architecture [IPSEC] that is mandatory in IPv6 implementations. The inherent nature of the IPSec technology, as it relates to tunneling, encryption etc., however, does not change regardless of whether IPSec is operating in an IPv6 or IPv4 network.

IPSec is a set of protocols designed to provide enhanced security to IP and upper layer protocols (UDP or TCP). IPSec achieves this by describing standard ways to specify traffic, how traffic is to be protected and how to authenticate both the sender and the receiver of the traffic. A key feature of IPSec is that it protects every datagram.

The basic security concepts defined in IPSec consist of the following:

- Authentication of the identities of the communicating parties. This can be done through manual configurations, pre-shared secrets or Digital Certificates issued by a trusted authority and acceptable to both parties.
- Authentication and encryption of the data. The keys for authentication can be dynamically generated through specific protocol exchanges or generated from the pre-shared secret.

Current public key technologies (RSS, DSS etc.) are adequate for authentication but do not have the capability to operate on a per-packet basis.

IPSec VPNs utilize the IPSec tunnel mode for both Remote Access and site-to-site connectivity. The basic network element is the Security Gateway (SG), which can be a dedicated device or implemented in PE/CE routers. The function of the SG is to protect the clients (PCs in a LAN) or a remote user. The procedure and packet exchanges involved in tunnel set-up are described in RFC 2408 [ISAKMP] and RFC 2409 [IKE]. The details of payload encryption and authentication are described in RFC 2406 [ESP] and RFC 2402 [AH].

Setting up an IPSec VPN

Before setting up an IP-VPN, an enterprise must have a security policy detailing the extent and breadth of security within various organizational entities. Based on these requirements and IPSec, VPN authentication, encryption and access control can then be established.

The actual process of setting up a VPN connection involves two broad steps:

1. Establishing an ISAKMP tunnel
2. Establishing an IPSec tunnel for sending data

The purpose of an ISAKMP (Internet Security Association and Key Management Protocol) tunnel is to provide a secure and authenticated tunnel for exchanging control packets. The establishment of this tunnel, known as Phase One exchange, can be done in "Main Mode" through the exchange of six messages. A short cut, known as "Aggressive Mode" which requires exchange of only three of messages, is also used.

Establishing an ISAKMP 'tunnel' requires three steps:

1. The initiating of a Security Gateway (SG1) proposes a ISAKMP Security Association (SA) which is accepted by destination SG (SG2)
2. SG1 exchanges "keying information" with SG2 through "Diffie-Hellman exchange"
3. Both SG1 and SG2 are authenticated using encrypted exchanges and ISAKMP Tunnel set-up is completed

The establishment of an IPSec Tunnel, which can use Encapsulated Security Protocol (ESP) or Authentication Header (AH) or both, is called "Quick Mode" and consists of a single step:

- SG1 proposes IPSec Security Association (SA) and is accepted by SG2. An IPSec tunnel is established

For control plane testing of an IPSec VPN, Tester T must perform the following:

- Emulate a Security Gateway (SG1) behind a tester port, which sets up an ISAKMP tunnel by making SA Proposals to the destination SG - SG2 (SUT)
- For each ISAKMP tunnel set up, the tester sets up one or many IPSec (ESP or AH or both) tunnels
- Periodically, the SG1 sends an ESP or AH encapsulated data packet, to keep the IPSec tunnel alive.

This scenario can be scaled by emulating more than one Security Gateway behind the tester port, thus setting up a very large number of IPSec tunnels to SG2 (refer Fig. 2).

Tester T must have the capability to measure performance related to the number of tunnels/sessions, set up/ deletion rate of tunnels, lifetime of tunnels etc.

For testing involving traffic and measurement of QoS parameters, Tester T must have the ability to generate encrypted traffic at line rate and send the traffic through the established IPSec tunnels. This tests whether a particular SG is encrypting or decrypting data correctly. In practice, specialized chips are used for encryption/decryption and testing of encryption algorithms is not necessary.

In the absence of an encrypted traffic generator, a back-to-back test scenario can be employed, where a pair of gateways are connected back-to-back and non-encrypted traffic is used.

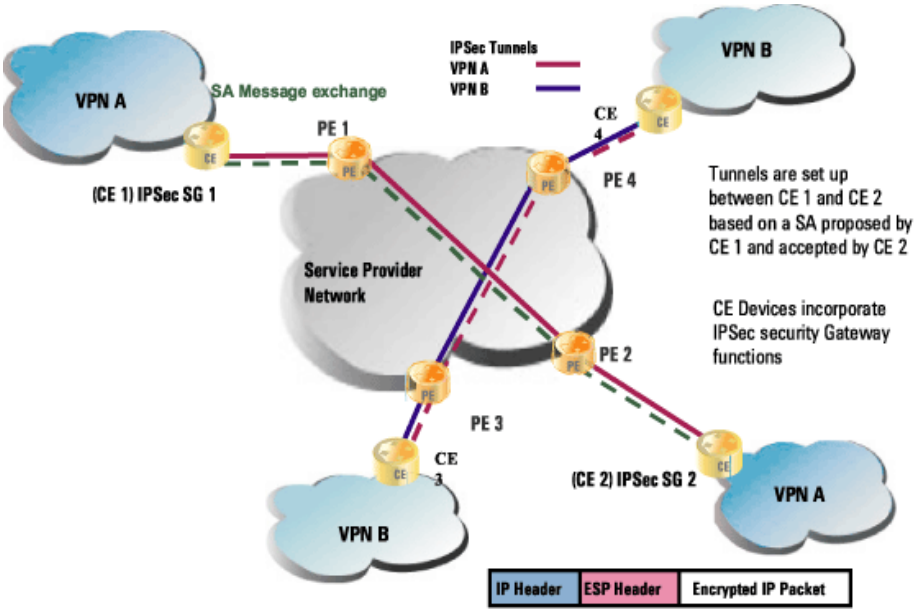

Figure 2: IPSec - IP Security Network Scenario

Back-to-back encryption performance testing

For back-to-back testing, both SG 1 and SG2 should be configured to follow a scenario such as: "If IP traffic is received on Interface I from IP address 100.100.100.X destined for IP address 200.200.200.Y, then send that traffic on an IPSec tunnel to gateway 250.250.250.4, using 3DES to encrypt the traffic.

Subsequent testing involves following steps (Ref. Fig. 3):

1. Emulate a client C1 behind a test port and another client C2 behind a second test port.
2. Connect C1 to the first SUT; connect SG1 and C2 to the second SUT, SG2.
3. Send IP traffic from C1, destined for C2, to SG1.
4. SG1 will examine the source and destination address and, based on that data, will be required to forward the traffic to SG2
5. To accomplish this SG1, will either create a new IPSec tunnel or use an existing IPSec tunnel
6. SG1 will then send encrypted traffic to SG2
7. SG2 will decrypt the traffic and deliver normal IP traffic to C2
8. For the traffic flow, Tester T should conduct the following checks:
   - Standard QoS parameters
   - Data integrity check

The data integrity check will ensure that encryption and decryption occurs correctly. An analyzer may also be used, in between the two gateways, to ensure that the traffic is actually encrypted.

For scalability testing, the tester can emulate many pairs of source and destination clients, forcing the gateways to set up many IPSec tunnels, while sending traffic at required rates on each.

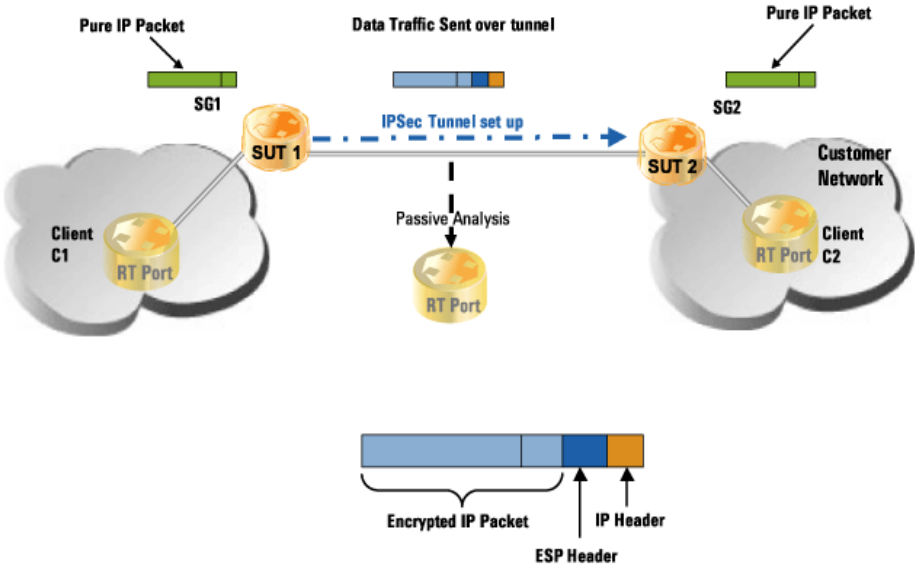

Figure 3. Traffic Encryption Performance Test Scenario

Layer 3 IP VPN: Remote Access VPN using L2TP

The purpose of L2TP (Layer 2 Tunneling Protocol), as defined in RFC 2661 [L2TP], is to "tunnel" PPP [PPP] sessions. It can be used over IP (using UDP) or Frame Relay Permanent Virtual Circuits (PVCs) or ATM PVCs. The context in which L2TP is used over Frame Relay or ATM gives it the name "Layer Two Tunneling Protocol."

In practice, L2TP is almost invariably used in the context of an IP network. As a result, for purposes of this paper, L2TP VPNs have been classified as Layer 3 VPNs.

Setting up a Layer 3 IP VPN

The L2TP Protocol defines two basic protocol elements: LAC (L2TP Access Concentrator) and LNS (L2TP Network Server). The LAC will typically reside on the edge of the Service Provider's network, while the LNS will reside at the edge of the enterprise. The other network element critical to this discussion is the remote user (also referred to as remote "client") who tries to connect to the network residing behind the LNS.

Set up of an L2TP VPN requires a series of protocol packet exchanges involving L2TP, PPP and IPCP protocols. In summary, there are four steps involved:

1. A remote user (client) establishes a PPP session (for example, through a dial up connection) with the LAC
2. The LAC, in turn, sets up an L2TP tunnel with the LNS
3. The LNS establishes a PPP session with the remote client over this L2TP tunnel.
4. Data is sent back and forth over this session.

PPP session establishment involves link establishment through Link Control Protocol (LCP) negotiation, use of the authentication PAP (Password Authentication Protocol), use of the CHAP (Challenge- Handshake Authentication Protocol) to validate the remote client, and invocation of a network control layer protocol (Internet Control Protocol (IPCP) for IP).

This type of tunneling is described as "compulsory tunneling", since once the remote user connects to the LAC, the LAC proceeds to set up the tunnel with the LNS. In "voluntary tunneling", it is the responsibility of the remote user to set up the tunnel with the LNS, requiring the remote user to have LAC functionality.

The L2TP elements, LAC and LNS, can be implemented in routers, although enterprise-class dedicated gateways are still deployed in CPE-based VPNs and can have LAC functionality. The PE router at the edge of the service provider's network typically will provide the LAC functionality, while the LNS will be implemented in a CE router at the edge of the customer's network.

L2TP uses PPP to provide initial encapsulation for the data and then append additional headers for transport through the IP Network. The complete packet has the following format:

To test the LAC, Tester T will use the following steps

1. Emulate one or more LNS elements as L2TP tunnel end points behind another interface.
2. Emulate one or more PPP (PPPoE) Clients behind one interface
3. Initiate PPP sessions from emulated clients with the LAC and cause the LAC to set up L2TP tunnels with the emulated LNS
4. Set up multiple PPP sessions in the same tunnel or different tunnels between the emulated clients and the emulated LNSs through the LAC
5. Once these tunnels and sessions have been set up, Tester T will send IP traffic from the Clients towards the LNS.
6. The SUT is expected to forward the IP traffic to the tester port behind the emulated LNS elements.

This test set-up will simplify measuring performances related to number of tunnels/sessions, tunnel set up/deletion rate, tunnel lifetime. With the tester sending IP traffic, this type of testing can determine if the LAC is forwarding the right traffic to the right destination (LNS). Additionally, QoS measurements such as packet loss, latency, jitter etc. can also be measured.

In a protocol stress test scenario, Tester T can scale the above scenario to set up thousands of L2TP tunnels and PPP sessions through the SUT and then send traffic through the test interface at the rated speed.

To test the LNS, the Tester T use the following steps:

1. Emulate multiple LAC elements and their attached clients behind one of the tester ports.

2. Set up L2TP tunnels from the emulated LACs to the LNS (SUT)

3. Set up PPP sessions between the emulated clients and the SUT

4. Send IP traffic to the LNS

In addition to measurements related to session/tunnel and QoS parameters, this will also test the scalability of the LNS, in terms of its ability to provide service to multiple LACs.

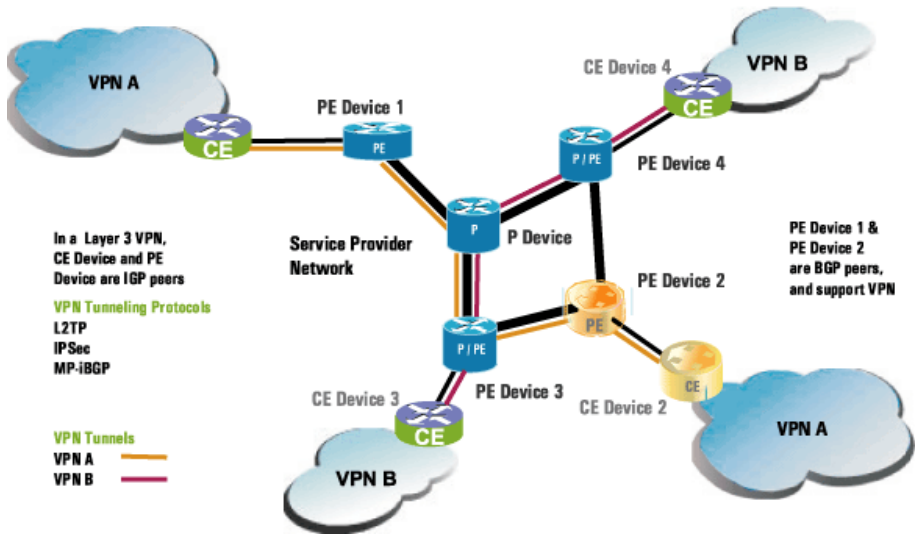

Figure 4: Layer 3 VPN Network Scenario

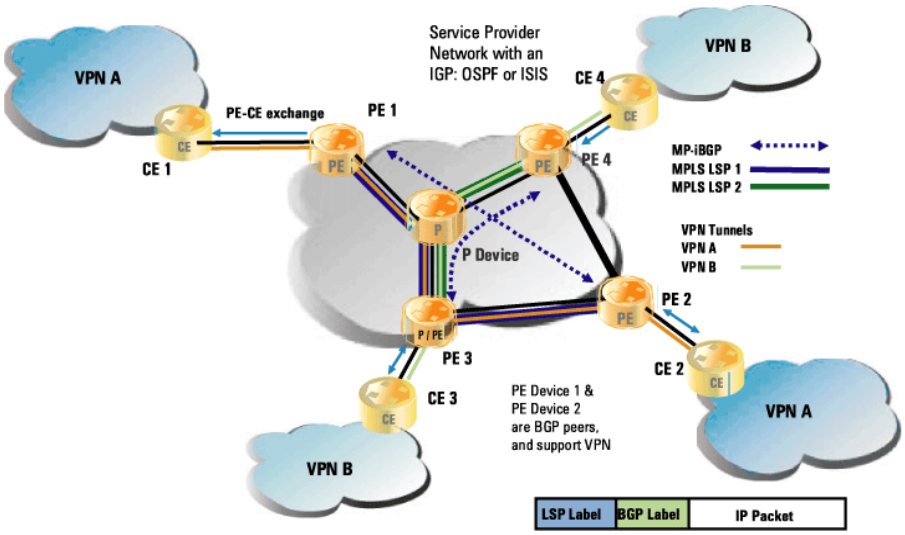

Figure 5: BGP/MPLS VPN Network Scenario

Layer 3 MPLS VPN: BGP/MPLS VPN

The Layer 3 focus of this paper is BGP/MPLS VPNs, as defined in RFC2547bis [BGP-VPN]. In this VPN type, MPLS is used for packet forwarding over the backbone, while BGP is used for route distribution over the backbone.

This section will first briefly discuss the BGP/MPLS VPN concept and then review testing issues for the most common VPN case, where a VPN is set up between two geographically separated sites, connected to two different PE routers (ref. Fig. 5).

PE-CE Connectivity

As with most VPNs, the PE router must manage most of the information processing required for BGP/MPLS VPNs. An enterprise interested in obtaining BGP/MPLS VPN service from an SP is referred to as a VPN site. The router in the enterprise network, connected to the SP network, is called the Customer Edge (CE) router. Thus, a VPN site connects to other sites on the VPN via one or more CE routers (over the SP network). A PE router is attached to a site (S) as the endpoint of an interface or "sub-interface" (e.g., PVC, VLAN, GRE tunnel, etc.) whose other endpoint is a CE device.

This phase of testing checks for interface or sub-interface connectivity.

Tester T must provide the following capabilities for each interface type -

- Ability to send and receive data traffic (on each sub-interface if necessary)
- Ability to run routing protocols, such as BGP4, OSPF, IS-IS, RIP (on each sub-interface if necessary). Static routing is also allowed.

Additionally, Tester T must support all necessary Layer 2 configuration, addressing, alarms, etc. and offer the ability to report on the health of the Layer 2 link (thus confirming the ability to send and receive data packets).

Maintaining VPN information on PE

The PE router maintains independent VPN routing and forwarding (VRF) tables for each VPN site it connects. The VRF table associated with a particular site "S" is populated only with routes that lead to other sites which have at least one VPN in common with "S."

A common problem is the overlapping of address space in different customer networks, resulting from enterprises using the private address space recommended by RFC1918. For example, two independent sites connected to the same PE but belonging to separate VPNs might have an overlapping address space. As a result, traffic destined for one site could be sent to the other.

To avoid this, a new type of address was created, called the VPN-IPv4 address. This is a 12-byte quantity, beginning with an 8-byte "Route Distinguisher (RD)" and ending with a 4-byte IPv4 address. Since the RD can be made unique by its structure, the VPN-IPv4 address for every site can be made globally unique.

Because VRF tables are maintained inside the PE router, an external test tool will not typically have access to them; the implementation of VRF tables, however, can be tested through the test for Data Forwarding for VPNs described below.

VPN Site Reachability

One of the extended attributes of BGP, the "Route Target," is used by BGP to distribute information on sites participating in a particular VPN. When a VPN-IPv4 route is created by a PE router, it is associated with one or more Route Target attributes. Any route associated with a Route Target ("T") must be distributed to every PE router that has a VRF associated with "T". In this way, the route is installed in a particular VRF, which is both that particular route's Route Target and is also one of the "import" targets of that VRF. As a result, controls exist, at both ends, for two sites participating in a VPN.

If two sites of a VPN attach to PEs which are within the same Autonomous System, the PEs can distribute VPN-IPv4 routes to each other by means of an IBGP connection, either directly or through a route reflector. When a PE router distributes a VPN-IPv4 route via BGP, it uses its own address as the BGP "next hop". It also assigns and distributes an MPLS label associated with the VPN-IPv4 route, as specified in RFC 3107 [BGP-LBL]. We will refer to this as the "BGP label".

All of these operations ensure that the VPN reachability information of the participating CE devices is available, in an unambiguous, unique fashion, to the PE devices. In addition to these steps, it is also

necessary to ensure that the PE devices (which are the BGP next hops) are able to send traffic to each other. This is done using an MPLS LSP between them. The [BGP-VPN] document states that, to ensure interoperability among systems which implement this VPN architecture, all systems must support LDP [MPLS-LDP].

Testing of these capabilities is done through basic data forwarding and is discussed in the following section.

Data Forwarding for VPNs

The CE device in a particular VPN will send an unlabeled IP packet to the PE. When a PE receives a data packet from a CE device, it chooses a particular VRF for look-up of the packet's destination address. This choice is based on the packet's incoming sub-interface, since that identifies the site to a VRF association.

The incoming packet will have a BGP next hop and there will be a BGP label for it. This label is pushed onto the packet's label stack, and becomes the bottom label. This BGP next hop, or the PE router, will also have an MPLS label associated with it, representing the LSP set up between the source and destination PE routers. This label is also pushed on to the packet, creating a two-deep label stack.

MPLS will then carry the packet across the backbone. Once the packet reaches the egress PE, it pops off the MPLS label. The egress PE router's treatment of the packet will depend on the BGP Label. In many cases, the PE will be able to determine, from this label, the sub-interface over which the packet should be transmitted (to a CE device), as well as the proper data link layer header for that interface.

In other cases, the PE may only be able to determine that the packet's destination address needs to be looked up in a particular VRF before being forwarded to a CE device. Information in the MPLS header itself, and/or information associated with the label, may also be used to provide QoS data on the interface to the CE. When the packet finally gets to a CE device, it will again be an ordinary unlabeled IP packet.

Tester T must perform the following steps to check VRF table implementations and VPN connectivity as part of a basic VPN Setup functional test.

1. Emulate multiple sites (and hence multiple CE routers) in Tester T on one set of interfaces and connect to the SUT
2. Emulate a network of devices and sites on another set of interfaces and connect to the SUT

c. Advertise VPN routes from one set of sites to another (populate the SUT's VRF tables), using iBGP

3. Set up LSPs between the PEs (which are iBGP peers)
4. Send traffic over these VPNs from an emulated PE to a particular set of addresses from a simulated VRF. An MPLS label stack will be automatically inserted, based on the entry in the simulated VRF and the LSP label
5. Send traffic from an emulated CE and receive packets at an emulated PE containing a two-level label stack. The label stack can be validated in real time, using per-stream statistics.
6. Capture and decode incoming data packets including the label stack
7. Capture and decode control packets

Both the VRF table and the BGP label for a site should be distinct. Tester T will analyze the traffic streams to and from different sites. For traffic coming to the emulated P routers, Tester T will check to determine that the correct label stacks have been used. For traffic coming to the CE devices, Tester T will check whether traffic is being sent to the correct destination. (Ref. Fig. 3 and 4)

At this point, Tester T will have all the information and capabilities required to set up a BGP/MPLSVPN. By repeating some of these steps, Tester T can set up a large number of VPN connections to the SUT, thus stressing the control plane. (Ref. Fig. 5) This can be achieved in two ways:

1. Connecting a large number of emulated CE routers as independent VPN sites, thus forcing large number of VRF tables on the SUT.
2. Making one site a member of a large number of VPNs, thus forcing a large VRF table on the SUT.

Data stressing can also be performed. Once a particular VPN is set up, data traffic can be sent over this VPN at up to the line rate allowed by the physical connection.

A series of QoS measurements must also be enabled for data traffic. These measurements include:

- Packet loss
- Packet delay
- Packet jitter

These measurements may need to be made on a per- VPN or per- interface basis.

VPN services will typically be covered under a Service Level Agreement. Service parameters specified may include total allowed bandwidth, availability, traffic QoS parameters etc. QoS details may include total network delay, packet loss, total jitter (if the VPN carries voice traffic), etc. Tester T should have the capability to provide the required measurements at the required granularity level for line rate traffic and high control loads. More detailed discussion on service testing is included later in this paper.

Layer 2 MPLS VPN

As our discussion moves to Layer 2 MPLS VPN testing, it is important to note the distinct difference between Layer 2 VPNs and Layer 3 VPNs, namely that the CE is not involved in Layer 2 VPN creation at all.

In the Layer 2 case, the SP provides only a Layer 2 interface to its customer, and the customer is responsible for creating and managing the Layer 3 overlay. Thus, for Layer 2 MPLS VPNs, only the PE needs to be tested, but Layer 2 interactions should be examined in greater depth.

There are no standards yet for Layer 2 MPLS VPNs, but enterprises have long built their own wide-area networks by purchasing wide-area, point-to-point, data link layer connectivity from service providers, and then building their own Layer 3 infrastructure on top of it. Some of the current Internet drafts, which have gained attention, are [Martini-TRANS] and [Kompella]. There have already been a number of public announcements of NEM implementations of the [Martini-TRANS] draft.

The following section describes a test methodology for the [Martini-TRANS], followed by a brief discussion of one for [Kompella].

PE-CE Connectivity

In this case, the PE-CE connection need not include routing protocol support. The Layer 2 link must be up and data flow verified. A permanent virtual circuit (PVC) may be used. Testing may simply be used to ensure that data can flow between the PE and CE devices.

Setting up a Layer 2 VPN

In the Layer 2 case, setting up a VPN is similar to establishing a Circuit Emulation Service. The CE to CE connection is considered a Virtual Circuit (VC). Typically, a "VC label" is defined to identify this VC.

The [Martini-TRANS] draft deals with point-to-point transport over MPLS. To transport Layer 2 packets (Protocol Data Units or PDUs) from an ingress PE router("R1") to an egress PE routers ("R2"), the following control plane steps are required (Ref. Fig. 4):

1. An MPLS LSP is set up between R1 and R2
2. R1 and R2 establish another LDP session, using the extended discovery mode
3. The VC label is distributed by LDP, running in downstream unsolicited mode
4. As part of the LDP protocol exchange, a new information element is passed to carry additional information about the VC. This is called the Virtual Circuit FEC (Forwarding Equivalence Class) element.
5. At this point, the PE's are ready to start transporting dat

This control plane exchange leads to two-level encapsulation of the data packets - by the VC label and the MPLS LSP label. The [Martini-ENCAP] draft defines data packet encapsulation. In addition, the draft requires a third level of encapsulation, called "emulated VC encapsulation". This level contains the information about the enclosed Layer 2 PDU which is necessary to properly emulate the corresponding Layer 2 protocol. This type of encapsulation is called the 'Control Word'

Thus, when a data packet is sent, it is encapsulated in the following sequence -

For Ethernet packets, the Ethernet Header itself is used as the third level of encapsulation instead of the Control Word.

Without going into the complete details of the process used to convert a normal Layer 2 packet into the above format, it is important to note that the content of the Control Word will vary depending on the PDU type (ATM AAL5 PDU, ATM cell or FR packet).

To test this type of Layer 2 VPN, Tester T will first emulate a network very similar to the one emulated for Layer 3 VPN testing - a collection of CE, P and PE devices. Based on Fig. 6, for example, one of the test ports needs to emulate a CE ("C1") and the other test port will emulate an MPLS network of P and PE devices, with another emulated CE connected to the PE ("R2"). The SUT is a PE device ("R1").

Once these elements are emulated, Tester T will perform the following actions:

1. Set up two LSPs between the PE devices R1 and R2 across the MPLS network, one from R1 to R2 and another from R2 to R1 (this is because MPLS LSPs are unidirectional)

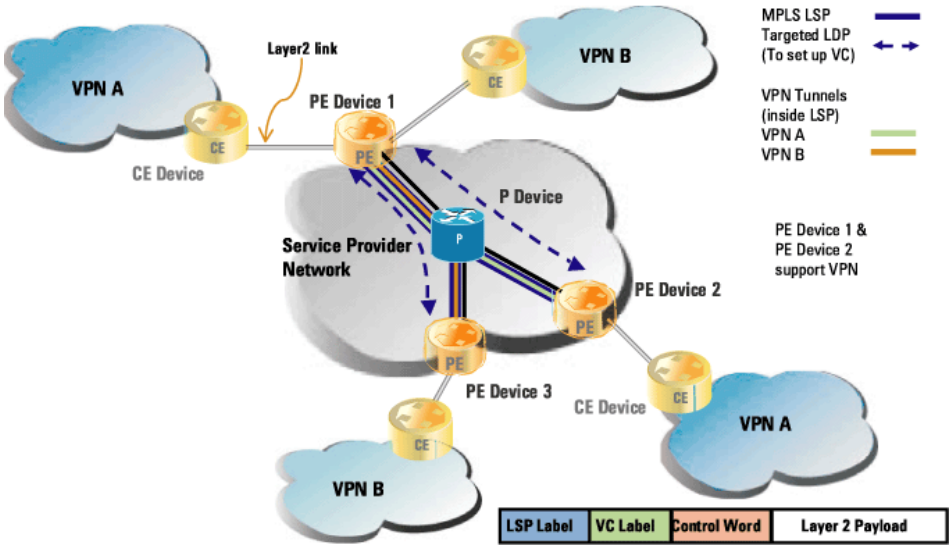

Figure 6: Layer 2 over MPLS ("Martini") Network Scenario

2. Establish an LDP session between R1 and R2 (using extended discovery), and exchange VC information. At this stage, the VC label and the MPLS labels are available to R1 and R2 and the endpoints are ready to exchange data traffic.
3. Send Layer 2 traffic (Ethernet packets, in this example) from emulated CE1 to R1
4. R1 should forward the Ethernet packets to R2 with two labels attached on top (the Ethernet header remains as is)
5. Tester T should verify the two-deep label stack for correctness
6. Tester T will send traffic with the two-deep label stack from R2 to R1, destined for C1
7. The SUT (R1) should strip these two labels and send this traffic to C1
8. Tester T will then check traffic arriving on the port for correctness

The above scenario can be scaled without the need to significantly increase the number of CE ports, since a very large number of Virtual Circuits can be emulated through a single CE-PE connection.

Assume, for example, that Tester T includes Gigabit Ethernet ports, which can emulate a CE with 2000 VLAN sessions. In that case, using just a few ports, a large number of VCs can be set up, per CE port, through the SUT ("R1"), as illustrated in Fig. 7. Tester T can also be configured so that some of the VCs

(say, 1000) from R1 must be set up to a PE (R2), while others (rest 1000) must be set up to R3.

This will then require four LSPs (two per each PE-PE connection, one in each direction). The important stress scenario here is that the SUT, R1, will be required to be part of 2000x2 LDP extended discovery sessions for setting up the VCs. If more than one emulated CE is connected to R1, this increases the number of sessions by that multiple. Finally, verification data traffic can be sent over each of these VCs.

The [Kompella] draft offers another method of setting up a Layer 2 VPN. It defines usage of the BGP protocol to communicate VPN reachability and VC label information, instead of the LDP protocol, as defined in the [Martini-TRANS]. It also contains a number of recommendations on implementation of such VPNs. The steps defined in the [Kompella] draft will not be discussed in detail here, as the draft is still evolving.

Verifying Service Level Agreements

Thus far, all of the discussion in this paper has centered on testing of individual devices, with the remainder of the system being emulated by Tester T. It is important to note that some of the elements in the emulated network discussed may be replaced by actual devices, allowing the user to test combinations or networks of various manufacturers' equipment.

The key application for network testing is verifying Service Level Agreements (SLAs) between SPs and their customers. Various metrics can be measured to determine whether SLA parameters are met.

Service Level Agreements may be defined per access network connection, per VPN, per VPN site, and/or per VPN route. The SP may define both levels and measurement intervals for any or all of the following parameters.

- QoS and traffic parameters
  - PE to PE round trip delay
  - Jitter (variation of delay)
  - Packet Loss ratio
- Availability for the site, VPN, or access connection
- Access Link Load
- Service activation interval (e.g., time to turn up a new site)
- Time to repair interval
- Total traffic offered to the site, route or VPN
- Measurement of non-conforming traffic for the site, route or VPN
- Router utilization

Using the mechanisms defined previously in this paper, Tester T can be used to set up VPNs, send traffic at up to line rate on the VPNs, and then provide measurements at specified intervals for each of these SLA parameters.

Service Verification

Tester T must have the following capabilities for SLA measurement:

- Highly granular traffic measurement, to the degree required by network QoS levels and traffic content.
- Distributed traffic measurement, across geographically distant VPN end points.
- Synchronized measurement capable of correlating traffic sent and received from the test end-points, and creating realistic reports of the state of the VPN and the network.
- Per VPN site, or per VPN route, measurement capabilities

It is important to note that black box testing does not permit measurement of SUT utilization. This must be done by other means, and this brief preliminary discussion of SLA verification is specific to MPLS VPNs. While this subject requires much more attention, it is also important to include some discussion of Network or Service Restoration testing in this paper.

Service Restoration Verification

When an SP offers a commercial VPN service to an enterprise, it will come with certain availability and QoS guarantees. These require implementation of restoration capabilities in the device and the network, since the link carrying VPN traffic may go down for a variety of reasons.

Network restoration can be implemented in multiple ways. Typically, it will be handled using the underlying protocol infrastructure (MPLS in this case.) While various techniques and algorithms may be used, two general methods are device-based and network-based restoration. Device-based techniques offer fast reroute methods. Network-based techniques would provide one or more backup LSPs. Regardless of the method used, it must be assumed that VPN traffic has some alternative path through the network.

Clearly, a key testing issue would be to determine whether the device or network could restore the VPN service within the parametric limits described in the SLA. The following steps describe a method for conducting this type of testing:

1. Tester T emulates a network of devices, where a VPN is to be set up from CE1 adjacent to the SUT (PE router - R1) through an SP network to CE2 behind R2 (ref. Fig. 8)
2. The VPN tunnel is nested inside an LSP ("LSP1") set up on the path R1-P1-P3-PE2
3. Assume the VPN label is L1, and LSP label is L2 when the traffic leaves the SUT
4. VPN traffic flows from CE1 to CE2
5. Label stack on the traffic is - Data+L1+L2
6. Another path is available for an LSP from SUT to PE2: R1-P2-PE2 (An LSP may not be set up beforehand, if the SUT supports fast reroute.)
7. The link between SUT and P1 is brought down by some method (laser-off may be the fastest way). This results in tear down of LSP1.
8. R1 sets up a new LSP, ("LSP2"), through P2 for PE2 using the label "L3" and starts transmitting CE1-CE2 VPN traffic over that LSP

Tester T should perform the following checks:

- Examination of traffic coming to P2 to ensure that it looks like Data+L1+L3
- Measurement of the performance characteristics of the flow and any changes to it, e.g. measurement of the time difference (delay) between the receipt of the last packet over LSP1 and the first packet over LSP, measurement of the jitter of the traffic
- Comparison of the traffic sent out of CE1 and received on CE2 to check for packet loss

The above scenario could also be executed while sending a high volume of control or data traffic to the SUT through other ports. This will emulate a realistic network traffic load on the SUT, and test the restoration capacity of the SUT under such conditions.

Conclusion

It's clear from the above discussion that MPLS VPN testing is a complicated process, chiefly because of the large number of components involved. However, by extending existing tools, the testing capabilities required for NEM and SP MPLS VPN development and deployment testing can be made available.

To summarize, the Tester T must offer the following capabilities:

Functionality

- Ability to emulate all the network elements of the SP network (P and PE routers) and the Enterprise edge (CE)
- Ability to emulate all the required protocols used in the SP network and for the CE-PE connection
- Ability to fully use all the protocol activity e.g. to set up MPLS LSPs and VPN tunnels
- Ability to send traffic up to line-rate over VPN tunnels

Usability

- Ability to perform highly granular measurement of the packet exchanges in the control and data plane in real-time
- Ability to decode all the packet exchanges in the control and data plane in real-time
- Ability to interpret the flow of control and data packets and create usable information
- Ability to represent all the information gathered in an intuitive graphical interface and reports for easy interpretation by a user
- Ability to provide automated ways to conduct all the required functions

Glossary

AH	Authentication header
BGP	Border Gateway Protocol
iBGP	BGP for internal peers in a Service Provider cloud
CHAP	Challenge- Handshake Authentication Protocol
ESP	Encapsulated Security Payload
IKE	Internet Key Exchange
ISAKMP	Internet Security Association and Key Management Protocol
IPSec	IP Security
MP-iBGP	iBGP with Multi-Protocol extensions
GRE	Generic Route Encapsulation (a legacy tunneling technique)
IPCP	Internet Control Protocol for IP.
L2TP	Layer Two Tunneling Protocol
LAC	L2TP Access Concentrator
LCP	Link Control Protocol
LNS	L2TP Network Server
LSP	Label Switched Path
MPLS	Multi-Protocol Label Switching
PAP	Password Authentication Protocol
PVC	Permanent Virtual Circuit (commonly used term in ATM and Frame Relay world)
PPP	Point-to-point Protocol
QoS	Quality of Service
TE	Traffic Engineering
VLAN	Virtual Local Area Networks
VPN	Virtual Private Networks - a way to set up private networks over shared infrastructure

References

1. [REQMT]Service requirements for Provider Provisioned Virtual Private Networks (draft-ietf-ppvpn-requirements-xx, August 2001, work in progress)
2. [BGP-VPN]BGP/MPLS VPNs (draft-ietf-ppvpn-rfc2547bis-xx.txt, July 2001, work in progress)
3. [BGP-LBL]Carrying Label Information in BGP4 (RFC 3107, May 2001)
4. [MPLS-LDP]LDP Specification (RFC 3036, January 2001)
5. [Martini-TRANS] Transport of Layer 2 Frames Over MPLS (draft-martini-l2circuit-trans-mpls-xx, July 2001)
6. [Martini-ENCAP]Encapsulation Methods for Transport of Layer 2 Frames Over IP and MPLS Networks (draft-martini-l2circuit-encap-mpls-xx, July 2001)
7. Kompella]MPLS-based Layer 2 VPNs (draft-kompella-ppvpn-l2vpn-xx, June 2001)
8. MPLS and VPN Architectures (Jim Guichard & Ivan Pepelnjak, Cisco Press, Dec 2000)
9. [IPSEC]Security Architecture for the Internet Protocol (RFC 2401, November 1998)
10. [ISAKMP]The Internet Security Association and Key Management Protocol (ISAKMP) (RFC 2408, November 1998))
11. [ESP] IP Encapsulating Security Payload (ESP) (RFC 2406, November 1998)
12. [AH] IP Authentication Header (RFC 2402, November 1998)
13. [IKE] The Internet Key Exchange (IKE) (RFC 2409, November 1998)
13. [PPP] The Point-to-Point Protocol (RFC 1661, July 1994)
14. [L2TP] Layer Two Tunneling Protocol (L2TP) (RFC 2661, August 1998)
15. IPSec (N Doraswamy & Dan Harkins, Prentice Hall, 1999)

November 1, 2001
5988-4888EN

Agilent Technologies