



Testing Edge Services: L2oMPLS VPNs and Service Level Agreements

Introduction

In January 2001, the Multiprotocol Label Switching (MPLS) working group of the Internet Engineering Task Force (IETF) published the first proposed standard for MPLS deployment in the core Internet. Since that first milestone, MPLS has become the basis for a growing number of related technologies and enhancements to existing protocols, such as Label Distribution Protocol (LDP) and Resource Reservation Protocol (RSVP). As development of these MPLS-based technologies continues, so do the many services that capitalize on their traffic engineering abilities. From Virtual Private Networks (VPNs) to Virtual Private LAN Services (VPLS), MPLS-based technologies offer service providers a chance to optimize their networks and differentiate themselves in a tight competitive market.

VPNs are currently a hot topic among network engineers. While much literature exists on the deployment of Layer-3 MPLS VPNs, it typically ignores the traditional workings of existing network deployments. Such deployments often involve ATM or Frame Relay (FR) virtual circuits and a mish-mash of Layer-3 protocols or manually configured databases to perform routing and forwarding functions. A relatively new development incorporating the notions of virtual circuits and edge-to-edge networking into the efficient and effective MPLS domain, the concept of Layer 2 over MPLS VPNs proposes a method of bringing more efficiency into the area of traditional ATM and Frame Relay routing or switching.

In addition to permitting the provisioning of Layer-2 connections through an MPLS network, other benefits of L2oMPLS include advanced traffic engineering and the ability to measure traffic forwarding rates. This last feature is of particular importance: with increased competition in the access market, customers have come to expect the best quality and speed of service for their money. The result of these expectations is the practice of signing Service Level Agreements (SLAs), which aim to provide customers a benchmark of network performance and reliability upon which they can rely for their business critical applications and communications. SLAs typically state specific forwarding performance guarantees and several other maintenance-related parameters (such as maximum network downtime) or special service-specific parameters (such as VPN setup rates, etc.).

Both L2oMPLS VPNs and Service Level Agreements pose specific testing challenges. For L2oMPLS VPNs, these challenges include verification that the appropriate labels have first been exchanged and then are appended to and removed from each packet that crosses the MPLS domain for a given VPN. For Service Level Agreements, the challenge for service providers is to provide realistic guarantees for their traffic forwarding performance. Therefore, thorough testing of individual devices and complete networks is critical.

This white paper introduces the concepts of L2oMPLS VPNs and Service Level Agreements, while providing specific guidelines for testing them.




Figure 1: General MPLS VPN configuration

Layer 2 over MPLS VPNs

L2oMPLS VPNs provide a means of sending specific Layer-2 protocol data units (PDUs) across MPLS-enabled networks by establishing a WAN link from one customer edge (CE) router to another. This technology, then, represents a breakthrough for service providers in terms of cost-efficiency, since legacy networks of ATM and Frame Relay equipment have been traditionally operated separately from MPLS networks in the same domain. Using L2oMPLS means that pseudo-virtual circuits – be they ATM, FR, or Ethernet-based – can be established between sites or devices across newer, faster, automated MPLS networks, as opposed to less flexible pure IP networks.

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 service providers, and then building their own Layer-3 infrastructure on top of it. Surprisingly, however, there are not yet any standards for L2oMPLS VPNs, although some current Internet drafts have gained attention (the so-called “Martini Draft” (<http://www.ietf.org/internet-drafts/draft-martini-l2circuit-trans-mpls-10.txt>) and the “Kompella Draft” (<http://www.ietf.org/internet-drafts/draft-kompella-ppvnp-l2vpn-02.txt>). Point-to-point connections of this type are now being defined by the Pseudo-Wire Edge to Edge (PWE3) working group at the IETF. The new drafts introduced in the PWE3 WG are “Transport of Layer 2 Frames Over MPLS” (<http://www.ietf.org/internet-drafts/draft-ietf-pwe3-control-protocol-00.txt>) and “Encapsulation Methods for Transport of Ethernet Frames Over IP/MPLS Networks” (<http://www.ietf.org/internet-drafts/draft-ietf-pwe3-ethernet-encap-00.txt>).

As shown in Figure 1, L2oMPLS VPNs are similar to their Layer-3 counterparts in several ways. First, a provider edge (PE) router provides customer access to the network backbone. Also, tunnel security and traffic forwarding across the MPLS backbone are assured by using LDP or RSVP-TE and Label Switched Paths

(LSPs). A distinct difference between Layer-2 and Layer-3 VPNs, however, is that the CE in a Layer-2 VPN simply acts as an access point for the LAN behind it. In this case, the service provider supplies only a Layer-2 interface to its customer, and the customer is responsible for creating and managing the Layer-3 overlay. Thus, for L2oMPLS VPNs, the routing protocol functionality of only the PEs need be tested – however, Layer-2 interactions between the CE and PE should also be verified.

The following test scenarios are based on the Internet Drafts proposed by the IETF PWE3 WG.

L2oMPLS VPN Setup

The setup of an L2oMPLS VPN is similar to that of a Circuit Emulation Service. The ultimate connection between CE routers is considered a virtual circuit (VC) and, typically, a “VC label” is used to identify each one. When testing a device's L2oMPLS VPN capability, the device under test (DUT) is expected to act as a PE router in a provider network. After completing all the necessary setup phases and protocol exchanges, the device's ability to both append labels to (or push them on) and strip (pop) them from pure Layer-2 PDUs is verified. As shown in Figure 2, this test configuration uses two test ports: one to simulate a CE router, and another to simulate both a mesh of provider core (P) routers and a provider edge (PE) router with a CE router and LAN behind it. The simulated next-hop P router behind the second test port should establish an LSP with the device under test. The simulated PE router should then establish an LDP session using the extended discovery mode to directly connect to the device under test. The LDP protocol exchange (running in downstream unsolicited mode) then distributes a few important pieces of information: the VC label, the virtual circuit FEC – a new information element that carries additional information about the VC – and, optionally, the Control Word – an encapsulation that contains necessary information about the enclosed PDU specific to the corresponding Layer-2 protocol. The PE router pushes the VC label onto the Layer-2 PDU (the header of which may be

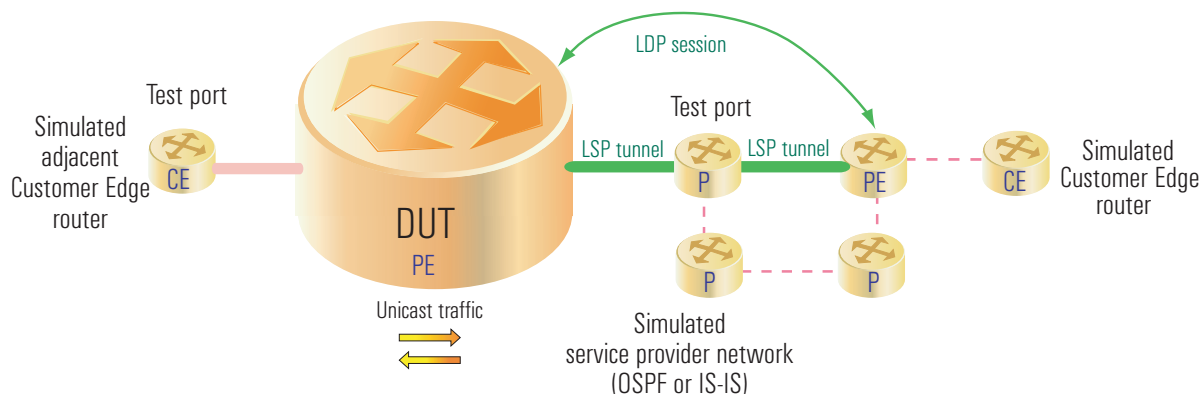


Figure 2: L2oMPLS VPN setup test configuration




Figure 3: Encapsulation of Layer 2 PDUs

replaced by the optional Control Word) before sending them across the MPLS network to the other PE router. For more information about Ethernet data packet encapsulation, see "Encapsulation Methods for Transport of Ethernet Frames Over IP/MPLS Networks" (<http://www.ietf.org/internet-drafts/draft-ietf-pwe3-ethernet-encap-00.txt>).

The encapsulation of a Layer-2 packet is shown in Figure 3. It is important to note both that the content of the Control Word will vary depending on the PDU type (ATM AAL5 PDU, ATM cell or FR packet) and that Ethernet packets use the Ethernet Header instead of the Control Word.

L2oMPLS VPN Scalability

Once an edge device's ability to create and maintain L2oMPLS VPNs has been established, its ability to handle large numbers of simultaneous VPNs must be verified. When testing a device's L2oMPLS VPN scalability capabilities, the test configuration is similar to that of the previous test. However, as shown in Figure 4, the tester should simulate more than one CE adjacent to the DUT. If the device under test does not support multiple sub-interfaces, one test port will have to be used for each CE connection to the DUT. Assume, for example, that one test port

uses a Gigabit Ethernet interface, which simulates one CE router with 4095 VLAN sessions. If that CE is configured so that 2047 of those VCs are configured to reach CE2 (simulated by a second test port across the provider network) and 2048 others are configured to reach CE3 (simulated by a third test port across the provider network), four LSPs (two per PE-PE connection – one in each direction), will need to be established. That means that the DUT will be required to perform 4095 LDP extended discovery sessions. Traffic forwarding is then verified on each VPN session. Finally, the number of CE routers connected to the DUT should be increased incrementally and traffic forwarding and labeling should be verified until a maximum number of VPN sessions is reached. The test should report the number of VPNs successfully created and the number of lost and mislabeled packets for each VPN.

Verifying Service Level Agreements

Thus far this white paper has centered on testing individual devices, with the remainder of the system being emulated by the tester. It is important to note, however, that any element of the emulated network may be replaced by an actual device, allowing




Figure 4: L2oMPLS VPN scalability test configuration

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

Figure 5: VPN-specific Service Level Agreement parameters

the user to test combinations of equipment. When testing Service Level Agreements (SLAs), it is particularly important to incorporate more than one network device in the system under test (SUT) in order to emulate the exact level of service that customers will experience.

Service Verification

The inclusion of SLAs in contracts between business customers and service providers is now standard practice. Before setting service levels for new technologies, however, service providers must test them on individual devices for functionality and scalability, and then on the network for overall performance.

Common means of verifying service levels include measuring basic forwarding performance – as described in RFC 2544 – and service-specific testing which incorporates these basic test techniques. Service verification of L2oMPLS VPNs should define both acceptable values and measurement intervals for any or all of the parameters listed in Figure 5. These measurements can be made per access network connection, per VPN, per VPN site, and/or per VPN route.

The demands on test equipment when measuring system performance for SLAs call for a blend of several qualities. First, the tester must be able to measure traffic sent and received at geographically distant VPN end-points on a per-VPN site or per-VPN basis to the degree of granularity required by network QoS levels and traffic content. From those measurements, it should be able to generate realistic reports of the state of the VPN and the network. It is important to note, however, that this type of testing does not measure SUT utilization, which must be done by other means.

There are many ways to use test equipment for service verification. If, for example, a service provider were asked to provide a connection between two Ethernet networks across an MPLS network, they could use a test port to simulate a virtual circuit with the device under test and another to act as the CE adjacent to the SUT. As shown in Figure 6, the first test port would set up the Layer-2 virtual circuit and send traffic at the required rates, while the second takes QoS measurements. The test results will show whether the particular DUT meets the QoS requirements.

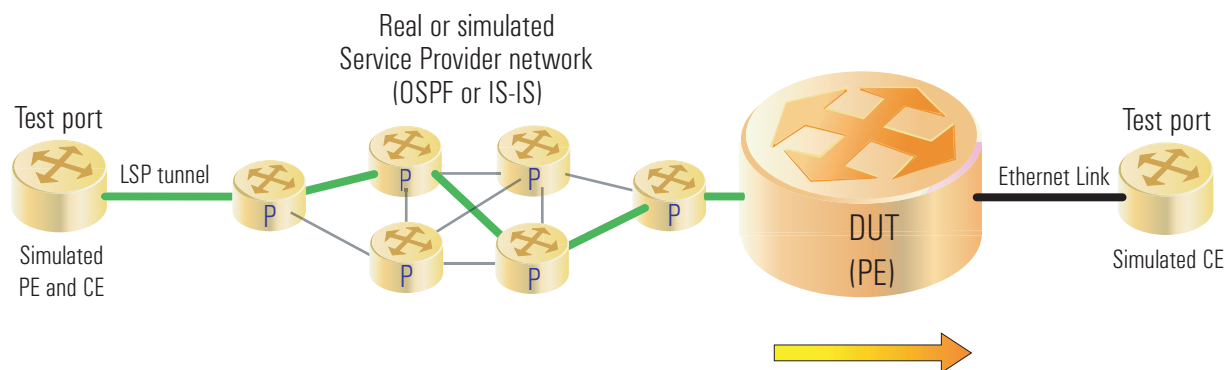


Figure 6: Example of test configuration for Service Level Agreement verification

Service Restoration Verification

There is no denying that networks occasionally go down and links carrying VPN traffic occasionally get preempted. Therefore, when a service provider includes commercial VPN service in an SLA, the guarantees they make inevitably include implementation of restoration capabilities in a specific device and in the whole network.

Network restoration can be implemented by many means. Typically, it is handled using the underlying protocol infrastructure (MPLS in this case). While various techniques and algorithms may be used, device-based and network-based restoration are most often deployed. Device-based techniques offer “fast reroute” methods, while network-based techniques generally provide one or more “backup LSPs”. Regardless of the method used, VPN traffic must have an alternate path through the network.

Clearly, determining whether the device or network can restore a customer’s VPN service within the parametric limits described in the SLA is crucial. After establishing a VPN session (based on the configuration in Figure 2 and the description of L2oMPLS VPNs on page 2), the tester should simulate a network failure in one of the intermediate routers on the LSP between the system under test (SUT) and the simulated PE router. The SUT’s ability to create a new LSP to the other VPN site and to forward traffic is then verified. The tester should analyze the packets received at the destination test port to ensure that the SUT has adapted the appropriate labels to the new LSP. It should also take a series of

measurements, including the interval between the network failure and the time at which traffic flow on the new LSP reaches full offered load, packet loss, and verifications of other packet characteristics. This scenario could also be executed while sending background traffic from other test ports to test the SUT’s restoration capacity under heavy traffic conditions.

Conclusion

Layer 2 over MPLS service has recently been catching the attention of service providers for various reasons. First, L2oMPLS allows network engineers to make use of IP-based networks that have been traditionally kept separate from legacy networks. Next, it provides an opportunity for significant revenue generation when used in a VPN configuration. Finally, these benefits combine to make it possible to include Layer-2 VPNs in Service Level Agreements.

However, any addition of new services adds degrees of complexity to the network, making testing of any device offering such services absolutely necessary. Once individual devices have been tested, the effects of deploying new services on existing networks must also be ascertained.

L2oMPLS represents the first step towards integrating legacy networks into IP-based domains. With thorough device and network testing, this first step will not only benefit service providers, it will also help clear the way for the next step in the development of their network infrastructure.

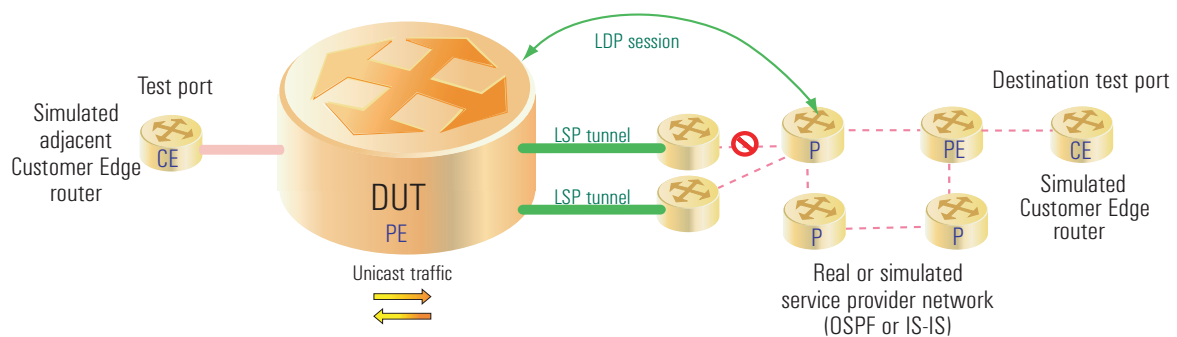


Figure 7: Example of test configuration for service restoration verification

