CISCO SYSTEMS

SECURITY

MPLS VPN Security

A practical guide to hardening MPLS networks

Michael H. Behringer
Monique J. Morrow

ciscopress.com

MPLS VPN Security

Michael H. Behringer

Monique J. Morrow

Cisco Press

800 East 96th Street
Indianapolis, Indiana 46240 USA

outside the scope of MPLS. The task of the MPLS core is to forward packets from the Internet to the VPN and vice versa. This includes potential attacks. It is, however, within the scope of MPLS security to make sure that an attack against a given VPN does not affect other VPNs or the core itself. (This is discussed in Chapter 4.) Also outside the scope of the MPLS architecture is any kind of firewalling required for such cases.

- **Customer network security**—Every attack that originates in a customer VPN and terminates in that same VPN is outside the scope of MPLS security. The MPLS VPN architecture forwards packets between VPN sites; it is not concerned with the nature of these packets, which could also be attack packets. This also includes IP spoofing within a VPN.

NOTE When discussing the security of MPLS VPN networks, take care to maintain a balanced view of the overall risks to a customer. For example, it is in relative terms close to irrelevant to argue about chances of an attacker sniffing a core line, if the customer network has unsecured wireless access points; it is also not important to worry about a service provider misconfiguring a PE, when attackers have uncontrolled physical access to hosts in an enterprise. Security is a question of balance: there is no point in putting extra secure locks on the door of your house if the windows are left open.

Comparison to ATM/FR Security

Many enterprises have been using VPN services based on ATM or Frame Relay (FR) in the past and are considering moving to MPLS VPNs. Unfortunately, the discussion about this topic has often been emotional and unbalanced.

New MPLS users are often concerned about the fact that an MPLS VPN service has a control plane on Layer 3. However, as shown in the previous sections, Layer 3 services can also be correctly secured and are fit to provide VPN services.

ATM/FR might be perceived as more secure because they are mostly not vulnerable to Layer 3 attacks (also ATM/FR switches typically have a Layer 3 control plane such as telnet). However, the security of Layer 2 in those technologies is typically assumed rather than actually proven. As we discuss in various parts of this book, Layer 2 has its own security issues that have to be considered. Many ATM/FR users are asking very hard questions about MPLS VPN security, while never having questioned whether a flood of signaling packets to an ATM switch might not affect that switch. It is good to discuss security of a technology, but it should be discussed in a balanced way.

This section discusses the features of both technologies and compares them.

VPN Separation

A VPN user requires his VPN to be separate from other VPNs and the core. In Layer 2 technologies, this is achieved implicitly by layering: the core exclusively uses Layer 2, so that the Layer 3 information of a VPN is separate. In MPLS VPNs, separation is achieved logically, by maintaining separate contexts on a provider router. Both ways are different, but both achieve the same result: each VPN can use the entire IP address space in their VPN, and it is impossible to send packets into other VPNs on the same core.

Misconfigurations are a problem in both technologies: an ATM circuit can be misconfigured, connecting a VPN router to a router from another company. Because many topologies are hub-and-spoke with default routing, this might lead to serious security breaches.

On the MPLS side, misconfigurations can equally break security. For example, a wrong route target on a PE router can bring a CE into a wrong VPN. You can argue about which misconfigurations are worse or more likely, but the fact is that if the core is misconfigured, VPN separation might be broken—in any VPN technology.

Robustness Against Attacks

VPN users demand a stable service, and most of all a service that cannot be attacked from the outside. For many VPN users, it would not be acceptable if a VPN service could be affected by a DoS attack from the outside. Even worse, an attacker gaining control of a network element could control any VPN. Therefore, any VPN technology must be resistant against attacks.

MPLS VPNs have been heavily scrutinized for the Layer 3 control plane and their frequent accessibility from the Internet. The issue was raised that given enough time, a good hacker would get access to a PE router over the Internet.

As shown previously in this chapter, an MPLS core has few and well-defined interface points to the outside. An MPLS core is not at all comparable to a traditional IP core, where every router was accessible (assuming the MPLS core has no global interfaces to the outside, only VRF interfaces). Rather, only single interfaces can be reached, and those can be very well secured. Therefore, it is very difficult to attack an MPLS network directly. An attack using transit traffic is the only possibility, and it might lead to a DoS condition. However, this can be controlled through appropriate dimensioning of the routers and architectural decisions, as is discussed in Chapter 4, in the section "Internet Access."

ATM or Frame Relay networks are also resistant against attack, assuming correct implementation. However, ATM switches and Frame Relay switches also have Layer 3 control planes (for example, telnet), and can be attacked if not appropriately secured. Frame Relay links depend on correct dimensioning of parameters such as Committed and Extended Information Rate (CIR/EIR). Misconfigurations of any of the above protocols or parameters can also result in a degradation or loss of service.

As long as both types of VPN technologies are configured correctly, they cannot be easily attacked.

Hiding the Core Infrastructure

In Layer 2 networks, the core is usually hidden because the VPN user works on Layer 3. Also, MPLS VPN cores are hidden to the VPN user, although using a different method: most addresses are hidden by architecture; the only visible part is the peering PE address. This address is, however, part of the VPN address space, so that in reality no core information is visible to the outside.

The fact that the PE router is reachable on this single interface is an exception to this rule. However, this usually is not a problem per se, but only in connection with attacks against the PE. As shown previously, this is very difficult if the PE is properly secured.

Impossibility of VPN Spoofing

As shown earlier in this chapter, it is impossible for an outsider to spoof another VPN, or the core, because a VPN user is always treated in his own context. Also, in ATM or Frame Relay, there are no known ways to spoof VPN signaling mechanisms such as the Virtual Path/Circuit Identifier (VPI/VCI) to spoof another VPN.

CE-CE Visibility

There is one area where ATM/Frame Relay point-to-point services do have an advantage over MPLS IP VPNs: Because the former are Layer 2 services, CEs can establish a direct Layer 3 adjacency and "see" the other CE. For example, the *Cisco Discovery Protocol (CDP)* can be used to find out basic properties of the peer router. This includes addressing of the Layer 3 link, so that a CE is able to verify to some extent the identity of the CE on the other end of the point-to-point link.

This is not possible in MPLS IP VPNs, and a given CE has no direct visibility of other CEs in his VPN. The reason for this is the connection model of MPLS IP VPNs: Although ATM and Frame Relay provide mostly point-to-point connections, where such a check is possible, MPLS IP VPNs provide connectivity from a CE to a "cloud." This avoids the overlay issue of having to establish a tunnel between all CEs (the so-called n^2 issue), but it has the disadvantage of losing the direct peering information.

This problem is not only theoretical. There is a real issue when a service provider accidentally or maliciously adds a CE to a wrong VPN by configuring wrong *route targets (RTs)*. The VPN to which this CE has been added has no easy means to find the bogus CE. It can only monitor traffic, control routing, and watch the used IP address space. This must be controlled by the service provider operationally. See Chapter 8, "Secure Operation and Maintenance of an MPLS Core," for more information on this important issue.

Other MPLS services, such as the *Pseudo Wire Emulation (PWE)*, also implement point-to-point services based on Layer 2, where direct CE-CE visibility is possible.

Comparison of VPN Security Technologies

Table 3-2 compares all the aspects of VPN security for the different VPN technologies.

Table 3-2 *Security Comparison Between MPLS and ATM/Frame Relay*

	MPLS	**ATM/Frame Relay**
VPN separation	Yes	Yes
Robustness against attacks	Yes	Yes
Hiding of the core infrastructure	Yes	Yes
Impossibility of VPN spoofing	Yes	Yes
CE-CE visibility	Not in MLPS IP VPNs Yes for MPLS pseudo wire emulation	Yes

Overall, at the time of writing this book, the industry had mostly accepted that both MPLS and ATM/Frame Relay can be operated securely. It was also common understanding that operational issues such as misconfigurations are an issue for any VPN technology.

Summary

In this chapter, we defined common requirements that VPN users have for a VPN service and examined MPLS IP VPNs against these requirements.

The result is that, based on the architecture described in RFC 2547bis, MPLS IP VPNs can be provided securely, meaning that:

- VPNs are separated (addressing and traffic).
- The core cannot be easily attacked.
- VPN spoofing is impossible.
- The core is invisible to the VPN user.

MPLS VPNs provide mostly equivalent security compared to traditional Layer 2 VPNs such as ATM and Frame Relay.

We have also examined Inter-AS and Carrier's Carrier architectures on their architectural security. While CsC networks are quite secure, care must be taken with Inter-AS scenarios when connecting different carriers: not all architectures provide the same level of security between providers.

There are also a number of issues that MPLS VPNs do not address. Among those are the internal security of a VPN, attacks from the Internet into a VPN, and VPN data confidentiality. These issues are independent of MPLS and have to be solved separately.

MPLS VPN networks are only secure when the network implementation is correct and when the network is operated correctly. How to control operations is discussed in Chapter 8, "Secure Operation and Maintenance of an MPLS Core." How to design and implement an MPLS core such that VPN services are secure is the subject of the next chapter.

Footnotes

[1] We will generally refer to IPv4 in this book; IPv6 is supported in the same way as IPv4, as a different address family.

[2] See RFC 2547bis.

[3] See draft-ietf-l3vpn-ipsec-2547 (work in progress).

[4] See draft-townsley-l2tpv3-mpls (work in progress).

[5] See draft-ietf-l3vpn-gre-ip-2547 (work in progress).

[6] For general considerations about route reflectors, consult *Cisco ISP Essentials*, ISBN 1-58705-041-2; *MPLS and VPN Architectures*, ISBN 1-58705-002-1.

In this chapter, you learn about the following:

- How to design an MPLS core for Internet access
- How to provision secure extranet access and firewalling
- How to design a DoS-resistant core
- How to secure Inter-AS and CsC solutions