available at www.sciencedirect.com

ScienceDirect

journal homepage: www.elsevier.com/locate/ijcip

Analysis of security threats to MPLS virtual private networks

*Denise Grayson, Daniel Guernsey, Jonathan Butts, Michael Spainhower, Sujeet Shenoi**

Department of Computer Science, University of Tulsa, 800 S. Tucker Drive, Tulsa, Oklahoma 74104, USA

ARTICLE INFO

Article history:
Received 19 April 2009
Accepted 12 August 2009

Keywords:
Multiprotocol label switching
Virtual private networks
Security analysis

ABSTRACT

Multiprotocol label switching (MPLS) based virtual private networking is one of the fastest-growing network technologies. It provides corporate and government customers with flexible, inexpensive "autobahns" that seamlessly connect multiple, geographically-dispersed sites, enabling voice, video, data and other high-bandwidth applications. The technology is also attractive to service providers because it enables them to flexibly provision resources for a variety of classes of service and applications with excellent quality of service at low cost. This paper analyzes the principal security threats to MPLS virtual private networks (VPNs). Because BGP is crucial to implementing MPLS VPNs, special attention is directed at the protocol and its multiprotocol extensions. This paper describes three classes of exploits on MPLS VPNs: route modification, traffic injection and denial-of-service attacks. It also discusses mitigation strategies that can be implemented by service providers and MPLS VPN customers.

1. Introduction

Multiprotocol label switching (MPLS) leverages circuit-switched and packet-switched network technologies (see, e.g., [1–3]). MPLS networks provide IP packet switching speeds comparable to those in circuit-switched networks by separating IP addresses from the packet forwarding algorithm. A packet entering an MPLS network is assigned a label that indicates a pre-determined path for it to follow. Thus, internal routers have only to examine the label to determine where the packet should go, eliminating the need to implement an expensive longest-match algorithm. Service providers can use routing and traffic engineering protocols to construct MPLS network paths with varying levels of quality of service (QoS) *and* class of service (CoS). The separation of addresses from forwarding provides exceptional flexibility and interoperability – any number of protocols (and applications) can be run over a telecommunications infrastructure while leveraging high-speed MPLS paths. Indeed, within a few years, much of the world's telecommunications and Internet traffic is expected to travel over MPLS networks [4].

Virtual private networking (VPN) is one of the primary applications of MPLS technology [5,14,15]. The Border Gateway Protocol (BGP) and its multiprotocol extensions are used to create Layer 3 VPNs that hide the service provider's infrastructure and implement routing and traffic separation between service provider and customer networks. MPLS-based VPNs are less expensive than leased lines and more flexible than traditional tunneling methods. They are increasingly being used by large corporations and government customers to run high-bandwidth applications over seamless IP-based networks that connect multiple, remote sites.

The reliance on MPLS VPNs by corporations and government agencies means that attacks – ranging from intercepting sensitive data to disrupting data, voice and multimedia services – can significantly impact vital operations. While several works (see, e.g., [6,7]) have discussed BGP/MPLS VPN security, they mainly rely on address separation and internal

* *Corresponding author.*
E-mail address: sujeet@utulsa.edu (S. Shenoi).

1874-5482/$ - see front matter
doi:10.1016/j.ijcip.2009.08.002

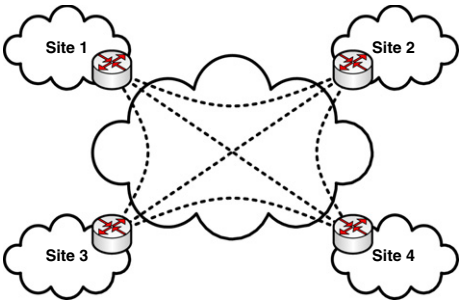

Fig. 1 – Traditional VPN overlay model.

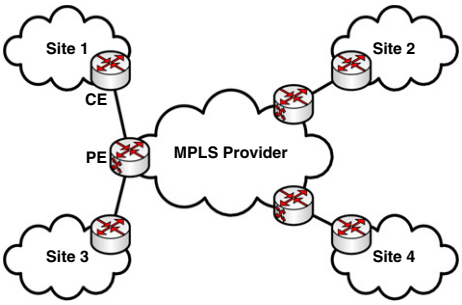

Fig. 2 – MPLS VPN model.

address hiding to protect against external attacks. However, several techniques can be used to gain access to the network core, especially by exploiting trust relationships at service provider boundaries. Once inside the relatively unprotected network core, an attacker can target customer VPNs as well as other MPLS network assets.

This paper analyzes the principal security threats to MPLS VPNs assuming that all nodes and links are susceptible to compromise. Three classes of exploits – route modification, traffic injection and denial-of-service attacks – are discussed. Several representative attacks are also described along with mitigation strategies that can be implemented by MPLS service providers and VPN customers.

2. MPLS-based virtual private networks

Virtual private networking is one of the most popular applications supported by MPLS networks. MPLS VPNs enable corporate and government customers to seamlessly connect multiple offices at distant locations into one network, possibly using the resources of multiple service providers.

Fig. 1 shows a traditional VPN overlay model with four sites. Each site is connected to routers located at other sites via point-to-point links. The links form a virtual backbone typically consisting of leased lines, frame relay circuits or ATM circuits.

Fig. 2 illustrates a VPN with four sites connected via an MPLS backbone. Each office site has a customer edge (CE) router that connects to the service provider at a provider edge (PE) router. Note that Site 1 and Site 3 share a PE router. Using MPLS in the telecommunications core enables traffic engineering protocols to construct high-speed tunnels between VPN sites without the cost and overhead associated with the traditional VPN overlay model [1]. Additionally, the peer model used in MPLS VPNs reduces the amount of routing information maintained compared with the overlay model and eliminates the need for point-to-point connections for every site. Thus, provider networks can be scaled to accommodate large numbers of VPNs. Moreover, VPNs can be scaled to incorporate large numbers of sites.

2.1. *VPN packet forwarding*

Fig. 3 shows a packet originating from Site 1 (VPN A) as it traverses the MPLS VPN provider core to Site 2. Because the packet is destined for Site 2, the CE router (W) for Site 1 forwards the packet to the provider network via the PE router (A). Router A is called a label edge router (LER) because it is located at the perimeter of the provider network. Likewise, Router F is also an LER.

Upon receiving the packet, Router A selects the appropriate VPN routing and forwarding (VRF) table based on the incoming interface [1,5]. This VRF table, which is specific to VPN A, contains labels and routing information for the reachable destinations in VPN A. Router A looks up the destination IP in this table to determine the outbound labels and the next hop. The router may also examine other attributes (e.g., TCP port) and label the packet for a particular QoS. This feature enables providers to guarantee high performance for applications such as video conferencing on a per-VPN or per-customer basis.

At this point, LER A pushes two labels onto the stack. The first (bottom) indicates the VPN to which the packet belongs. The second indicates the path the packet will take and the service guarantees that it will receive. Since the packet is headed to LER F, Router A labels the packet with F and forwards it to the (internal) Router B, called a label switching router (LSR).

Normal MPLS forwarding procedures apply to packet traversal within the network core. Upon receiving the packet, LSR B examines the topmost label, consults its routing tables, swaps the topmost label (F) with the next label along the path to LER F (also F for simplicity) [8], and then forwards the packet to LSR E. LSR E, in turn, examines the topmost label. However, because Router E is the penultimate hop to LER F, Router E pops the path label (F) to expose the VPN label (VPN A) and forwards the packet to LER F.

When LER F receives the packet, it examines the topmost label to determine the VPN with which it is associated (VPN A). Like LER A, LER F has a VRF table for each VPN it connects; the topmost label identifies the relevant VRF. LER F then pops the final label, looks up the destination IP in the VRF and forwards the unlabeled IP packet to Router Y, the CE router for Site 2.

2.2. *VPN routing*

In order to correctly distribute routing information for multiple VPNs, a provider network must both constrain routing information within VPNs and separate the addresses corresponding to different VPNs [1,5]. Routing information