Chapter 9

SECURITY ANALYSIS OF THE MPLS LABEL DISTRIBUTION PROTOCOL

Daniel Guernsey, Aaron Engel, Jonathan Butts and Sujeet Shenoi

Abstract Since its inception more than a decade ago, multiprotocol label switching (MPLS) has become one of the fastest-growing telecommunications infrastructure technologies. The speed, flexibility, sophisticated traffic management and cost savings offered by MPLS have prompted service providers to converge existing and new technologies onto common MPLS backbones. Indeed, much of the world's data, voice communications, video traffic and military applications traverse an MPLS core at some point.

The rapid adoption of MPLS raises significant concerns – primarily because of the dependence of critical communication services on a technology that has yet to undergo significant security testing. This paper examines security issues associated with the Label Distribution Protocol (LDP), which is the primary route construction protocol in MPLS networks. Our analysis has identified ten attacks that exploit weaknesses in the LDP specification: six attacks that disrupt service and four that divert traffic from intended routes. Details of the attacks are presented along with suggested mitigation strategies and security postures.

Keywords: Multiprotocol label switching, Label Distribution Protocol, security

1. Introduction

Multiprotocol label switching (MPLS) is quickly becoming the *de facto* protocol for transporting traffic in modern telecommunications networks. MPLS networks leverage the performance and availability of circuit-switched networks with the robustness and flexibility of packet-switched networks. Traffic entering an MPLS network is tagged with labels based on customer quality of service (QoS) and class of service (CoS) requirements. This allows traffic to be classified and then routed according to provisioned services (e.g., data type, message source, message destination and bandwidth requirements) instead of destination-only methods employed in traditional IP networks.

T. Moore and S. Shenoi (Eds.): Critical Infrastructure Protection IV, IFIP AICT 342, pp. 127–139, 2010.

In December 2005, the United States Department of Defense (DoD) achieved full operational capability of the Global Information Grid Bandwidth Expansion (GIG-BE) Program. The GIG-BE is designed to deliver global, high-speed classified and unclassified services to meet national security intelligence, surveillance and reconnaissance; and command and control requirements [10]. MPLS was chosen as the network transport backbone primarily due to its efficiency, simplicity and popularity in commercial environments [5, 9]. The DoD's use of MPLS for critical data is by no means unique. Many major telecommunications service providers around the world have invested massively in MPLS technology [2, 4, 16]. In fact, according to one source [13], 84% of enterprises have already transitioned their wide area networks to MPLS.

Despite the massive growth of MPLS networks, very little research has focused on the security aspects of core protocols such as the Label Distribution Protocol (LDP). LDP is the primary mechanism for transforming IP routes into high-speed "autobahns" within the MPLS paradigm. Weaknesses in LDP can be exploited by an attacker to achieve a wide range of strategic effects, including disrupting voice, global data and emergency communications.

This paper examines the security issues related to LDP. In particular, it discusses how LDP can be exploited to isolate network segments, reroute network traffic, disable the routing of network traffic and perform targeted attacks. Ten exploits are discussed: six denial-of-service attacks and four route modification attacks. Denial-of-service attacks target weaknesses in LDP to degrade or deny legitimate traffic delivery. Route modification attacks alter the path of targeted MPLS traffic traversing the network. The paper concludes by outlining mitigation strategies and security postures.

2. Multiprotocol Label Switching Networks

Connection-oriented and connectionless protocols are the two principal paradigms for transporting traffic across large networks [12]. ATM (OSI Layer 2) is an example of a connection-oriented technology that provides low latency and high quality of service (QoS). IP (OSI Layer 3) is a connectionless protocol that supports a multitude of underlying heterogeneous network technologies.

Service providers are eager to leverage the flexibility of IP and the speed of ATM without sacrificing efficiency [8]. In traditional implementations, an overlay model is used to create an ATM virtual circuit between each pair of IP routers. The IP routers are unaware of the ATM infrastructure and the ATM switches are unaware of IP routing. The end result is relatively inefficient: the ATM network must construct a complete mesh of virtual circuits among the IP routers.

MPLS offers an alternative solution that enables connection-oriented nodes to peer directly with connectionless technologies by transforming ATM switches into IP routers. ATM switches participate directly in IP routing protocols (e.g., RIP and OSPF) to construct label switched paths (LSPs). LSPs are implemented in ATM switches as virtual circuits, enabling existing ATM technology to support the MPLS forwarding mechanism. Conversely, MPLS enables

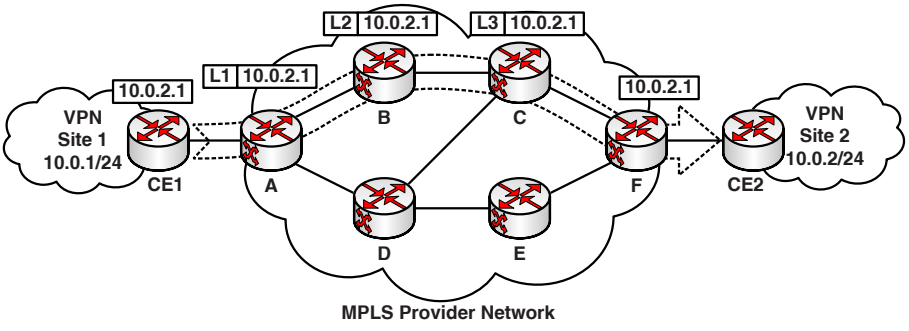

Figure 1. MPLS packet forwarding.

connectionless technologies to behave in a connection-oriented manner by augmenting IP addresses and routing protocols with relatively short, fixed-length labels.

Each label is a 32-bit (fixed length) tag, which is inserted in the Layer 2 header (e.g., for ATM VCI and Frame Relay DLCI) or in a separate "shim" between Layers 2 and 3 [14]. A label works much like an IP address; it dictates the path the router uses to forward the packet. Unlike an IP address, however, an MPLS label only has local significance. When a router receives a labeled packet, the label informs the router (and that router only) about the operations to be performed on the packet. Typically, a router pops the label on an incoming packet and pushes a new label for the router at the next hop in the MPLS network; the network address in Layer 3 is unchanged.

Figure 1 illustrates a typical MPLS architecture that interconnects two customer VPN sites. Routers A through F in the MPLS network are called label switched routers (LSRs). Customer edge routers, CE1 and CE2, sit at the edge of the customer network and provide connectivity to the MPLS core.

Consider the LSP from VPN Site 1 to VPN Site 2 (Routers A, B, C and F). Router A is designated as the "ingress node" and Router F is designated as the "egress node." The ingress and egress nodes are often called label edge routers (LERs) because they are at the edge of the MPLS network [14].

When an IP packet reaches the ingress of the MPLS network, LER A consults a forwarding information base (FIB) and assigns the packet to a forwarding equivalence class (FEC). The FEC maps to a designated label that supports QoS and CoS requirements based on IP parameters in the packet (e.g., source IP address, destination IP address, application).

In this example, LER A pushes Label L1 onto the packet and forwards it to LSR B. LSR B reads the label, consults its local label information base (LIB) to identify the next hop, pops the previous label and pushes a new label (L2), and forwards the packet to LSR C. LSR C behaves similarly, forwarding the packet to LER F. LER F then pops Label L3, examines the destination IP address and forwards the packet to VPN Site 2.

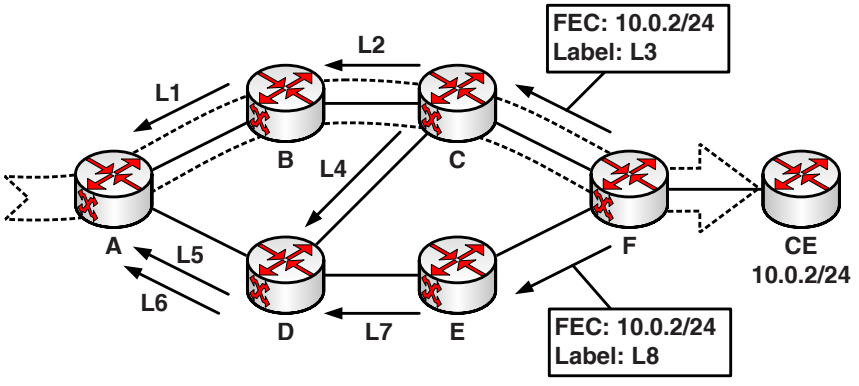

Figure 2. LDP routing information flow.

3. MPLS Routing Information

MPLS defines a forwarding mechanism designed to emulate IP routes using labels and paths. IP networks rely on routing protocols such as RIP and OSPF to populate the IP forwarding table [12]. Similarly, MPLS networks engage label distribution protocols to populate the FIB and LIB and establish end-to-end LSPs.

The Label Distribution Protocol (LDP) is the primary MPLS protocol for exchanging label mapping information [7]. LDP relies on underlying IP routing information to construct a set of LSPs using best-effort routes [6]. LSPs, in turn, can be optimized by employing traffic engineering protocols. MPLS traffic engineering protocols (e.g., RSVP-TE and MP-BGP) use topology information, constraints, specialized algorithms and signaling protocols to create LSPs to match customer QoS and CoS requirements [3, 15]. Traffic engineering protocols rely on LSPs constructed by LDP to discover the underlying routing structure. As such, exploiting a weakness in LDP can be leveraged to affect LSPs generated through traffic engineering.

4. Label Distribution Protocol (LDP)

LDP is designed to distribute information about available routes within an MPLS network. The edge routers begin the process by distributing label information about their adjacent external networks. FECs are created for each network based on IP addresses or prefixes [1].

Consider the example in Figure 2. LER F defines an FEC F1 for 10.0.2/24 and binds it to Labels L3 and L8. Next, it distributes the mappings (L3, F1) and (L8, F1) to its upstream peers (LSR C and LSR E, respectively) to update their LIBs. Upon receiving the mapping, LSR C binds a label to FEC F1 for each of its upstream interfaces and distributes these labels to LSR B and LSR D. Similarly, LSR E distributes the mapping (L7, F1) to LSR D. The process terminates when the information reaches an ingress router (LER A).

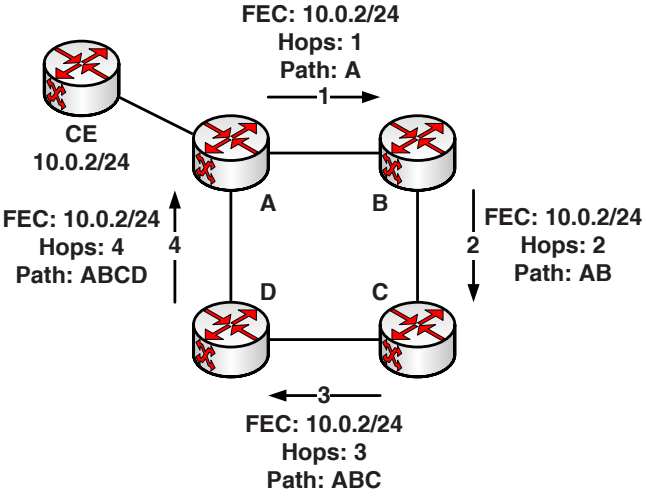

Figure 3. Loop detection using path vectors and hop counts.

The mapping distribution provides LER A with three distinct paths for 10.0.2/24. For example, if LER A receives an IP packet addressed to 10.0.2.1, it consults its FIB for an FEC with the longest matching prefix. Because three paths exist for 10.0.2/24, LER A selects the least cost path determined by its IP table. To meet customer requirements, FECs can be generated through traffic engineering for distinct destinations or applications to ensure that specific bandwidth, latency and other services are adequately provided.

4.1 Label Merging

It is common for two LSPs to converge prior to reaching a common egress [1]. To save memory and label space, LSRs may merge the LSPs at the point of convergence. When a merge-capable LSR receives a label request for an existing FEC and label mapping, it does not forward the request. Rather, it distributes the existing mapping to its upstream neighbors, effectively merging the two requested LSPs.

4.2 Loop Detection

The recursive nature of label request and label mapping messages creates the potential for message loops [1]. LDP uses hop counts and path vectors to detect loops. When a mapping request is forwarded, the LSR increments the message hop count and appends its own ID to the path vector. If the hop count exceeds a configured limit or an LSR discovers its ID in the path vector, the LSR sends a notification to the sender that a loop has been detected. In Figure 3, LSR A detects ID A in the path vector, implying that a loop exists. LSR A

stops forwarding the message and sends a notification to LSR D to prevent the construction of an LSP that contains a loop.

4.3 LDP Messages

Four message classes in LDP are used to facilitate session management and label distribution [1]: (i) Discovery messages that establish network adjacencies; (ii) Session messages that initialize and maintain LDP connections; (iii) Advertise messages that establish and remove LSPs; and (iv) Notification messages that specify advisories and errors.

Discovery Class Messages

- **Hello:** Hello messages are exchanged among LSRs during the discovery process using UDP. There are two types of messages: (i) Link Hello messages and (ii) Extended Hello messages. Link Hello messages are sent between directly-linked LSRs by addressing the messages to the subnet broadcast address. Extended Hello messages are exchanged between non-directly-linked LSRs by addressing the messages directly to peers.

Session Class Messages

- **Initialization:** Once an adjacency is discovered, the LSR peers establish a TCP connection. Initialization messages are then used to exchange session parameters (e.g., retention mode or label distribution mode) between the LSRs.

- **KeepAlive:** KeepAlive messages facilitate the detection of network errors. LSRs periodically transmit these messages to indicate that a link is still working. An error condition is assumed to have occurred when an LSR does not receive a message from a peer within an allotted timeout period; this results in the termination of the established session and the removal of associated labels.

Advertise Class Messages

- **Address:** Address messages provide neighboring LSRs with mapping information about LSR IDs to interface IP addresses. This information is used to identify the label mappings that correspond to the least cost path.

- **Address Withdraw:** Address Withdraw messages notify neighboring LSRs of disabled interfaces or broken links. Receipt of this message causes an LSR to remove the withdrawn address from its LIB mapping.

- **Label Mapping:** Label Mapping messages are used to distribute FEC-to-label bindings from a downstream LSR to an upstream peer. This message is the primary mechanism for constructing LSPs.

- **Label Withdraw:** Label Withdraw messages are used to notify peers that a particular FEC-to-label mapping is no longer valid (e.g., an egress interface goes offline or the network topology changes). When an LSR receives this message, it removes the label from its LIB and sends subsequent Label Withdraw messages to upstream peers.

- **Label Release:** Label Release messages notify downstream peers that an LSR has removed a particular label mapping. An LSR may remove bindings, for example, when an IP table changes or a Label Withdraw message is received.

Notification Class Messages

- **Notification:** Notification messages convey errors and advisories among peer LSRs. If the message indicates a fatal error, the sending and receiving LSRs terminate the LDP session and remove all associated label bindings.

5. LDP Vulnerabilities

In general, attacks may exploit weaknesses in: (i) the LDP specification; (ii) service provider implementations; and (iii) underlying infrastructure. Attacks on the LDP specification leverage inherent weaknesses in the design of the protocol. Any network that conforms with the protocol standard is susceptible to this class of attacks.

Attacks on service provider implementations exploit configuration errors or code flaws. LDP includes several undefined and reserved fields that can be exploited in attacks [1]. LDP also uses a nested structure of Type-Length-Value fields, which offers numerous opportunities for buffer overflow attacks. Our analysis does not focus on implementation vulnerabilities; nevertheless, we note that all implementations should undergo extensive fuzz testing.

Attacks on the underlying infrastructure exploit vulnerabilities in information technology and network assets or weak security policies. For example, LDP relies on IP to provide session communication and routing information. An attack on the underlying IP protocols may be used to reroute a target LSP or hijack a session. Because these attacks do not explicitly exploit LDP messages, they are not considered in this paper.

Our analysis focuses primarily on how an attacker can use LDP messages to exploit MPLS networks. Given only link access, we discuss several vulnerabilities in the LDP specification that could enable an attacker to deny service to various network assets or to reroute traffic.

5.1 Denial-of-Service Attacks

Denial-of-service (DoS) attacks target network resources or capabilities in order to degrade performance or prevent a provider from delivering services to its customers. Our analysis has uncovered six DoS attacks.

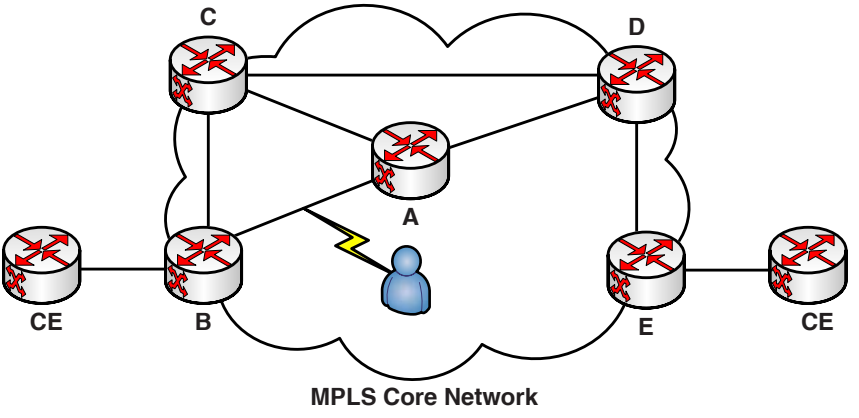

Figure 4. A portion of an MPLS network under attack.

- **Fabricating Notification Messages:** Fabricated Notification messages can be used to target network links. The attack requires read and write access to the target link. In Figure 4, an attacker with access to Link AB fabricates a fatal Notification message from LSR B to LSR A. In response, LSR A and LSR B close the LDP session and remove labels received from the peer. In turn, each router sends Label Withdraw messages to its upstream neighbors to reflect the removed label bindings. Additionally, an attacker with read access to Links AC and AD can intercept TCP sequence numbers and send Notification messages targeting these links via Link AB, which result in the isolation of LSR A.

- **Blocking KeepAlive Messages:** This attack disables a target link. The attacker selectively blocks LDP KeepAlive messages on the target link, which causes the LSRs at either end to terminate the LDP session. The LSRs then remove all the labels associated with the target link as well as the labels from their upstream peers.

- **Fabricating Address Withdraw Messages:** This attack targets three LSRs within an LSP. In Figure 4, an attacker with access to Link AB targets LSPs containing BAC or BAD. To attack BAD, the attacker fabricates an Address Withdraw message from LSR A to LSR B, which withdraws the address associated with the interface for LSR D. LSR B now believes LSR D cannot be reached via LSR A. Subsequently, LSR B tears down any LSPs containing BAD and constructs replacement LSPs.

- **Fabricating Label Withdraw Messages:** This attack targets a specific LSP and requires access to a link along the target path. If the network employs label merging, then the attack also affects all upstream portions of paths merged with the target LSP. Suppose LSR B in Figure 4 binds Label L1 to the LSP EDAB and distributes the binding to LSR A. To tear down EDAB, the attacker fabricates a Label Withdraw message

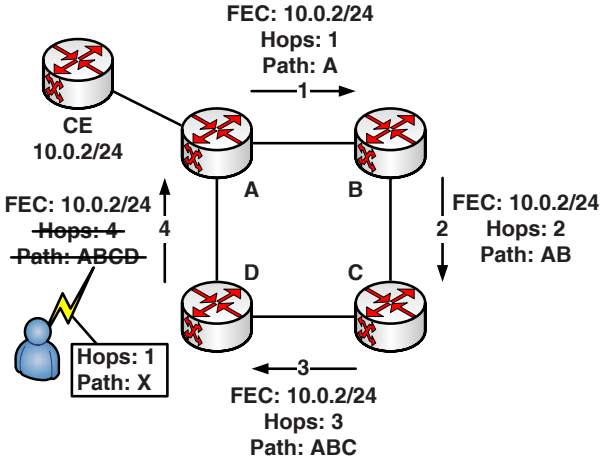

Figure 5. Avoiding loop detection mechanisms.

for L1 from LSR B to LSR A. This causes LSR A to remove L1 from its LIB, delete the label binding for EDAB and send a Label Withdraw message to LSR D. Similarly, LSR D sends a Label Withdraw to LSR E, which completes the destruction of the target path.

- **Exhausting Label Memory:** This attack targets an LSR and requires access to an adjacent link. The attacker floods the target with Label Mapping messages containing random FECs and labels. If the target LSR is configured for the liberal retention mode, it maintains all mappings in its LIB until the memory is exhausted [1]. The target LSR must drop older mappings to replace them with incoming mappings or must refuse all new mappings. In either case, legitimate paths are affected.

- **Creating Loops:** The goal of this attack is to degrade performance within a portion of the network by constructing an LSP loop. The attacker (Figure 5) listens on Link AD for Label Request or Label Mapping messages from LSR D to LSR A. The path vector is modified to reflect one LSR that is not contained within the loop (say LSR X). Any LSR along the path that supports label merging will combine the label request with the existing LSP to create an infinite loop. If no LSR supports label merging, the request completes a full loop, requiring the attacker to perform the modification again. In the absence of label merging, the process continues until the maximum 255 hops exhaust the TTL allowed by MPLS; thus, an infinite loop cannot be created.

5.2 Route Modification Attacks

Route modification attacks change the path of targeted traffic. These attacks enable an attacker to gain access to certain traffic (e.g., maneuver traffic

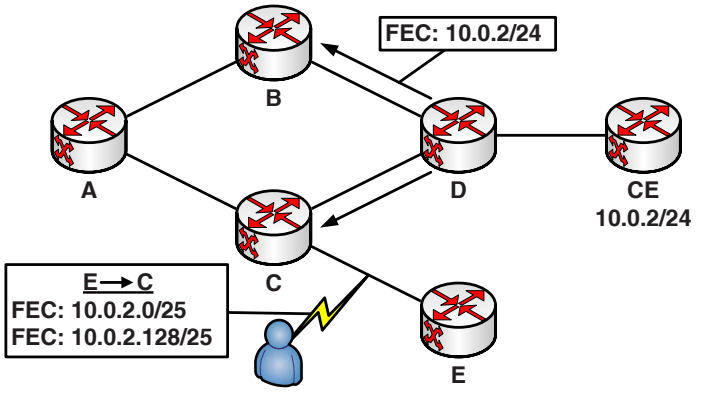

Figure 6. Route modification by creating more specific FECs.

through a compromised link); affect accounting (e.g., trigger automatic financial transactions among cooperating providers); or route traffic across domains (e.g., send one customer's traffic to another customer's network). Our analysis has revealed four route modification attacks.

- **Exploiting FEC Specificity:** This attack takes advantage of the "most specific" or "longest match" rule applied by ingress routers to incoming IP packets. An attacker needs access to a link or a connection to an interface to establish an LDP session. The attacker identifies a target FEC and advertises label bindings for more specific FECs. LSRs that receive the label mappings distribute them throughout the network, thereby building new LSPs toward the compromised link.

  For example, in Figure 6 the attacker targets FEC 10.0.2/24 by distributing mappings for 10.0.2.0/25 and 10.0.2.128/25. When ingress LER A sees a packet for 10.0.2.1, it selects FEC 10.0.2.0/25 and forwards the traffic to the compromised link. The attacker may now read, modify and/or forward this packet to its original destination. The attacker may also be very specific by sending label bindings for a single host such as FEC 10.0.2.1.

- **Fabricating Label Mapping Messages:** This attack reroutes traffic or creates loops by modifying the labels in Label Mapping messages. The attacker needs knowledge of downstream labels, which can be obtained by listening on the compromised link. The attacker may either modify a message in transit or fabricate a Label Mapping message. The message is sent to the upstream router causing it to adjust its LIB. When the upstream LSR receives a packet for the target FEC, it applies the incorrect label, which causes the downstream router to mistakenly recognize the packet as belonging to a different FEC. The packet is then forwarded along the desired LSP. A nefarious attacker can exploit this vulnerability to forward all targeted traffic to a different domain.

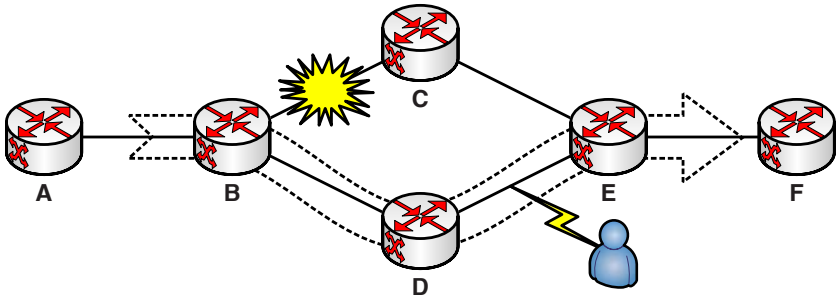

Figure 7. DoS-based route modification attacks.

- **Fabricating Address Messages:** This attack reroutes traffic or creates loops by manipulating the "least cost" mechanism used to select the next hop. Traffic can be redirected using access to a compromised link adjacent to an LSR along a selected LSP (the attacker does not require access to the link carrying the targeted traffic). The attacker crafts an LDP Address message that spoofs the address of the IP next hop. The fabricated message causes the LSR to adjust its LIB and generate a Label Request message. Thus, a new LSP is constructed that forces the targeted traffic along the compromised link.

- **Strategic Placement of DoS Attacks:** As shown in Figure 7, an attacker may execute DoS attacks that force the network to reroute traffic. These attacks change traffic flow within an MPLS network; however, they lack the varying degrees of granularity provided by the other route modification attacks. Nevertheless, the attacks are quite effective and their strategic placement can disable large portions of the network and force traffic through desired paths.

6. Mitigation Strategies

As in the case of traditional networks, most security mechanisms are applied at the perimeter of MPLS networks. However, many of the attacks discussed above occur from within administrative domains. Therefore, it is essential to apply security mechanisms that protect the internal operations of MPLS networks.

Many vulnerabilities in LDP stem from the lack of authentication, integrity and confidentiality mechanisms. LDP messages are sent in the clear, which enables an attacker to gather valuable network information, identify important targets and perform insidious attacks. Without integrity or authentication checks, LSRs are unable to discern the source of a message or verify that a message has not been modified or replayed.

Adequate authentication and integrity mechanisms would mitigate the majority of attacks discussed above. However, implementing these mechanisms

requires significant effort and overhead for key management. According to RFC 3562 [11], keys should be changed at least every 90 days. Additionally, the Internet Engineering Task Force (IETF) suggests strict guidelines for key distribution. Unfortunately, a manageable implementation scheme has yet to be demonstrated. Similar problems surface when using pre-shared keys to encrypt traffic for protecting messaging confidentiality.

In addition to authentication, integrity and confidentiality, simple filtering techniques can be applied to protect LDP from exploitation. For example, an LSR should not accept a Link Hello (used in direct peer discovery) unless the packet is addressed to the link multicast address and the source address is on the same subnet [1]. Without this restriction, it may be possible for an attacker to create LDP adjacencies by addressing Link Hellos directly to a target LSR. To prevent the abuse of Extended Hellos (used in extended peer discovery), each LSR should be configured with an access control list that specifies authorized remote peers. Extended Hello messages should also be filtered at the ingress; unless the source and destination addresses identify an authorized external LDP adjacency, the message should be discarded.

To mitigate memory exhaustion attacks, LSRs should favor existing label bindings over new label bindings. LSRs in the liberal retention mode are susceptible to memory exhaustion because they maintain all label bindings advertised by their peers. LSRs in the conservative retention mode, however, are not susceptible because they release bindings that do not correspond to the IP next hop. Unfortunately, configuring all LSRs for the conservative mode comes at the cost of increased time required to recover from network failures. Alternatively, LSRs in the liberal retention mode should not discard bindings corresponding to the IP next hop when limited by memory constraints. LSRs may also prioritize label bindings based on recent use to protect the most common alternate routes.

7. Conclusions

MPLS has emerged as a mainstay for transporting large volumes of traffic over a wide array of networks. Indeed, much of the world's enterprise traffic already depends on MPLS-based infrastructures to deliver reliable voice, video and application services. A persistent attack on the MPLS infrastructure could cripple corporate, national and even global operations.

LDP, a critical component for discovering and constructing MPLS routes, is vulnerable to several types of attacks. An attacker with internal link access can disable portions of a network or modify traffic flow. Therefore, mitigation strategies should focus on internal operations as well as external operations. We hope that this work prompts a more thorough analysis of security for LDP and related MPLS protocols.

References

[1] L. Anderson, P. Doolan, N. Feldman, A. Fredette and B. Thomas, LDP Specification, RFC 3036, 2001.

[2] AT&T, AT&T wins Frost & Sullivan 2009 North American Market Leadership Award in MPLS/IP VPN services, Dallas, Texas (att.centralcast.net /rss/feeds.aspx), July 20, 2009.

[3] D. Awduche, J. Malcolm, J. Agogbua, M. O'Dell and J. McManus, Requirements for Traffic Engineering over MPLS, RFC 2702, 1999.

[4] BT, Delivering the future: BT's 21 Century Network, London, United Kingdom (www.btplc.com/21CN/index.htm).

[5] Congressional Budget Office, Issues Associated with the Global Information Grid Bandwidth Expansion, Washington, DC (www.cbo.gov/doc .cfm?index=6132&type=0), February 28, 2005.

[6] B. Davie and Y. Rekhter, *MPLS: Technology and Applications*, Morgan Kaufmann, San Francisco, California, 2000.

[7] L. Ghein, *MPLS Fundamentals*, Cisco Press, Indianapolis, Indiana, 2007.

[8] E. Gray, *MPLS: Implementing the Technology*, Addison-Wesley, Reading, Massachusetts, 2001.

[9] D. Grayson, D. Guernsey, J. Butts, M. Spainhower and S. Shenoi, Analysis of security threats to MPLS virtual private networks, *International Journal of Critical Infrastructure Protection*, vol. 2(4), pp. 146–153, 2009.

[10] Joint Interoperability Test Command, JITC DISN OT&E Support, Fort Huachuca, Arizona (jitc.fhu.disa.mil/ot&e/gigbe.htm), June 26, 2002.

[11] M. Leech, Key Management Considerations for the TCP MD5 Signature Option, RFC 3562, 2003.

[12] L. Peterson and B. Davie, *Computer Networks: A Systems Approach*, Morgan Kaufmann, San Francisco, California, 2007.

[13] B. Reed, What's next for MPLS? *Network World*, December 21, 2009.

[14] E. Rosen, A. Viswanathan and R. Callon, Multiprotocol Label Switching Architecture, RFC 3031, 2001.

[15] M. Spainhower, J. Butts, D. Guernsey and S. Shenoi, Security analysis of RSVP-TE signaling in MPLS networks, *International Journal of Critical Infrastructure Protection*, vol. 1, pp. 68–74, 2008.

[16] TelecomWeb, Verizon Business Promises Aggressive 2009 Network Investment, Parsippany, New Jersey (www.telecomweb.com/international/262 500.html), February 10, 2009.