

Internet Engineering Task Force (IETF)
Request for Comments: 5920
Category: Informational
ISSN: 2070-1721

L. Fang, Ed.
Cisco Systems, Inc.
July 2010

Security Framework for MPLS and GMPLS Networks

Abstract

This document provides a security framework for Multiprotocol Label Switching (MPLS) and Generalized Multiprotocol Label Switching (GMPLS) Networks. This document addresses the security aspects that are relevant in the context of MPLS and GMPLS. It describes the security threats, the related defensive techniques, and the mechanisms for detection and reporting. This document emphasizes RSVP-TE and LDP security considerations, as well as inter-AS and inter-provider security considerations for building and maintaining MPLS and GMPLS networks across different domains or different Service Providers.

Status of This Memo

This document is not an Internet Standards Track specification; it is published for informational purposes.

This document is a product of the Internet Engineering Task Force (IETF). It represents the consensus of the IETF community. It has received public review and has been approved for publication by the Internet Engineering Steering Group (IESG). Not all documents approved by the IESG are a candidate for any level of Internet Standard; see Section 2 of RFC 5741.

Information about the current status of this document, any errata, and how to provide feedback on it may be obtained at <http://www.rfc-editor.org/info/rfc5920>.

Copyright Notice

Copyright (c) 2010 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (<http://trustee.ietf.org/license-info>) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

This document may contain material from IETF Documents or IETF Contributions published or made publicly available before November 10, 2008. The person(s) controlling the copyright in some of this material may not have granted the IETF Trust the right to allow modifications of such material outside the IETF Standards Process. Without obtaining an adequate license from the person(s) controlling the copyright in such materials, this document may not be modified outside the IETF Standards Process, and derivative works of it may not be created outside the IETF Standards Process, except to format it for publication as an RFC or to translate it into languages other than English.

Table of Contents

1. Introduction	4
2. Terminology	5
2.1. Acronyms and Abbreviations	5
2.2. MPLS and GMPLS Terminology	6
3. Security Reference Models	8
4. Security Threats	10
4.1. Attacks on the Control Plane	12
4.2. Attacks on the Data Plane	15
4.3. Attacks on Operation and Management Plane	17
4.4. Insider Attacks Considerations	19
5. Defensive Techniques for MPLS/GMPLS Networks	19
5.1. Authentication	20
5.2. Cryptographic Techniques	22
5.3. Access Control Techniques	33
5.4. Use of Isolated Infrastructure	38
5.5. Use of Aggregated Infrastructure	38
5.6. Service Provider Quality Control Processes	39
5.7. Deployment of Testable MPLS/GMPLS Service	39
5.8. Verification of Connectivity	40
6. Monitoring, Detection, and Reporting of Security Attacks	40
7. Service Provider General Security Requirements	42
7.1. Protection within the Core Network	42
7.2. Protection on the User Access Link	46
7.3. General User Requirements for MPLS/GMPLS Providers	48
8. Inter-Provider Security Requirements	48
8.1. Control-Plane Protection	49
8.2. Data-Plane Protection	53
9. Summary of MPLS and GMPLS Security	54
9.1. MPLS and GMPLS Specific Security Threats	55
9.2. Defense Techniques	56
9.3. Service Provider MPLS and GMPLS Best-Practice Outlines	57
10. Security Considerations	59
11. References	59
11.1. Normative References	59
11.2. Informative References	62
12. Acknowledgements	64
13. Contributors' Contact Information	65

1. Introduction

Security is an important aspect of all networks, MPLS and GMPLS networks being no exception.

MPLS and GMPLS are described in [RFC3031] and [RFC3945]. Various security considerations have been addressed in each of the many RFCs on MPLS and GMPLS technologies, but no single document covers general security considerations. The motivation for creating this document is to provide a comprehensive and consistent security framework for MPLS and GMPLS networks. Each individual document may point to this document for general security considerations in addition to providing security considerations specific to the particular technologies the document is describing.

In this document, we first describe the security threats relevant in the context of MPLS and GMPLS and the defensive techniques to combat those threats. We consider security issues resulting both from malicious or incorrect behavior of users and other parties and from negligent or incorrect behavior of providers. An important part of security defense is the detection and reporting of a security attack, which is also addressed in this document.

We then discuss possible service provider security requirements in an MPLS or GMPLS environment. Users have expectations for the security characteristics of MPLS or GMPLS networks. These include security requirements for equipment supporting MPLS and GMPLS and operational security requirements for providers. Service providers must protect their network infrastructure and make it secure to the level required to provide services over their MPLS or GMPLS networks.

Inter-AS and inter-provider security are discussed with special emphasis, because the security risk factors are higher with inter-provider connections. Note that inter-carrier MPLS security is also considered in [MFA-MPLS-ICI].

Depending on different MPLS or GMPLS techniques used, the degree of risk and the mitigation methodologies vary. This document discusses the security aspects and requirements for certain basic MPLS and GMPLS techniques and interconnection models. This document does not attempt to cover all current and future MPLS and GMPLS technologies, as it is not within the scope of this document to analyze the security properties of specific technologies.

It is important to clarify that, in this document, we limit ourselves to describing the providers' security requirements that pertain to MPLS and GMPLS networks, not including the connected user sites. Readers may refer to the "Security Best Practices Efforts and

Documents" [OPSEC-EFFORTS] and "Security Mechanisms for the Internet" [RFC3631] for general network operation security considerations. It is not our intention, however, to formulate precise "requirements" for each specific technology in terms of defining the mechanisms and techniques that must be implemented to satisfy such security requirements.

2. Terminology

2.1. Acronyms and Abbreviations

AS	Autonomous System
ASBR	Autonomous System Border Router
ATM	Asynchronous Transfer Mode
BGP	Border Gateway Protocol
BFD	Bidirectional Forwarding Detection
CE	Customer-Edge device
CoS	Class of Service
CPU	Central Processing Unit
DNS	Domain Name System
DoS	Denial of Service
ESP	Encapsulating Security Payload
FEC	Forwarding Equivalence Class
GMPLS	Generalized Multi-Protocol Label Switching
GCM	Galois Counter Mode
GRE	Generic Routing Encapsulation
ICI	InterCarrier Interconnect
ICMP	Internet Control Message Protocol
ICMPv6	ICMP in IP Version 6
IGP	Interior Gateway Protocol
IKE	Internet Key Exchange
IP	Internet Protocol
IPsec	IP Security
IPVPN	IP-based VPN
LDP	Label Distribution Protocol
L2TP	Layer 2 Tunneling Protocol
LMP	Link Management Protocol
LSP	Label Switched Path
LSR	Label Switching Router
MD5	Message Digest Algorithm
MPLS	Multiprotocol Label Switching
MP-BGP	Multiprotocol BGP
NTP	Network Time Protocol
OAM	Operations, Administration, and Maintenance
PCE	Path Computation Element
PE	Provider-Edge device
PPVPN	Provider-Provisioned Virtual Private Network
PSN	Packet-Switched Network

PW	Pseudowire
QoS	Quality of Service
RR	Route Reflector
RSVP	Resource Reservation Protocol
RSVP-TE	Resource Reservation Protocol with Traffic Engineering Extensions
SLA	Service Level Agreement
SNMP	Simple Network Management Protocol
SP	Service Provider
SSH	Secure Shell
SSL	Secure Sockets Layer
SYN	Synchronize packet in TCP
TCP	Transmission Control Protocol
TDM	Time Division Multiplexing
TE	Traffic Engineering
TLS	Transport Layer Security
ToS	Type of Service
TTL	Time-To-Live
UDP	User Datagram Protocol
VC	Virtual Circuit
VPN	Virtual Private Network
WG	Working Group of IETF
WSS	Web Services Security

2.2. MPLS and GMPLS Terminology

This document uses MPLS- and GMPLS-specific terminology. Definitions and details about MPLS and GMPLS terminology can be found in [RFC3031] and [RFC3945]. The most important definitions are repeated in this section; for other definitions, the reader is referred to [RFC3031] and [RFC3945].

Core network: An MPLS/GMPLS core network is defined as the central network infrastructure that consists of P and PE routers. An MPLS/GMPLS core network may consist of one or more networks belonging to a single SP.

Customer Edge (CE) device: A Customer Edge device is a router or a switch in the customer's network interfacing with the Service Provider's network.

Forwarding Equivalence Class (FEC): A group of IP packets that are forwarded in the same manner (e.g., over the same path, with the same forwarding treatment).

Label: A short, fixed length, physically contiguous identifier, usually of local significance.

Label merging: the replacement of multiple incoming labels for a particular FEC with a single outgoing label.

Label Switched Hop: A hop between two MPLS nodes, on which forwarding is done using labels.

Label Switched Path (LSP): The path through one or more LSRs at one level of the hierarchy followed by packets in a particular FEC.

Label Switching Routers (LSRs): An MPLS/GMPLS node assumed to have a forwarding plane that is capable of (a) recognizing either packet or cell boundaries, and (b) being able to process either packet headers or cell headers.

Loop Detection: A method of dealing with loops in which loops are allowed to be set up, and data may be transmitted over the loop, but the loop is later detected.

Loop Prevention: A method of dealing with loops in which data is never transmitted over a loop.

Label Stack: An ordered set of labels.

Merge Point: A node at which label merging is done.

MPLS Domain: A contiguous set of nodes that perform MPLS routing and forwarding and are also in one Routing or Administrative Domain.

MPLS Edge Node: An MPLS node that connects an MPLS domain with a node outside of the domain, either because it does not run MPLS, or because it is in a different domain. Note that if an LSR has a neighboring host not running MPLS, then that LSR is an MPLS edge node.

MPLS Egress Node: An MPLS edge node in its role in handling traffic as it leaves an MPLS domain.

MPLS Ingress Node: A MPLS edge node in its role in handling traffic as it enters a MPLS domain.

MPLS Label: A label carried in a packet header, which represents the packet's FEC.

MPLS Node: A node running MPLS. An MPLS node is aware of MPLS control protocols, runs one or more routing protocols, and is capable of forwarding packets based on labels. An MPLS node may optionally be also capable of forwarding native IP packets.

Multiprotocol Label Switching (MPLS): MPLS is an architecture for efficient data packet switching and routing. MPLS assigns data packets with labels. Instead of performing the longest match for each packet's destination as in conventional IP forwarding, MPLS makes the packet-forwarding decisions solely on the contents of the label without examining the packet itself. This allows the creation of end-to-end circuits across any type of transport medium, using any protocols.

P: Provider Router. A Provider Router is a router in the Service Provider's core network that does not have interfaces directly towards the customer. A P router is used to interconnect the PE routers and/or other P routers within the core network.

PE: Provider Edge device. A Provider Edge device is the equipment in the Service Provider's network that interfaces with the equipment in the customer's network.

PPVPN: Provider-Provisioned Virtual Private Network, including Layer 2 VPNs and Layer 3 VPNs.

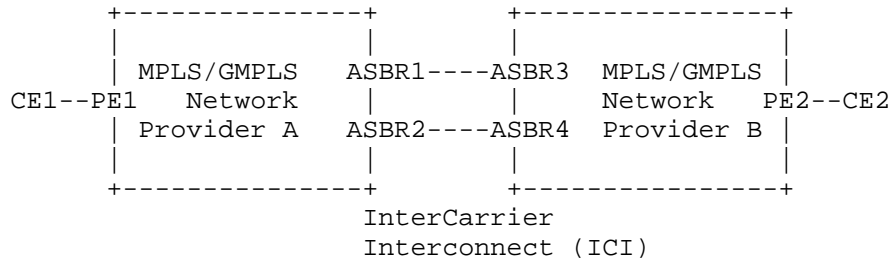
VPN: Virtual Private Network, which restricts communication between a set of sites, making use of an IP backbone shared by traffic not going to or not coming from those sites [RFC4110].

3. Security Reference Models

This section defines a reference model for security in MPLS/GMPLS networks.

This document defines each MPLS/GMPLS core in a single domain to be a trusted zone. A primary concern is about security aspects that relate to breaches of security from the "outside" of a trusted zone to the "inside" of this zone. Figure 1 depicts the concept of trusted zones within the MPLS/GMPLS framework.

inter-AS LSPs, the other provider is considered an untrusted but authorized neighbor.



For Provider A:

Trusted Zone: Provider A MPLS/GMPLS network
 Authorized but untrusted neighbor: provider B
 Unauthorized neighbors: CE1, CE2

Figure 2: MPLS/GMPLS Trusted Zone and Authorized Neighbor

All aspects of network security independent of whether a network is an MPLS/GMPLS network, are out of scope. For example, attacks from the Internet to a user's web-server connected through the MPLS/GMPLS network are not considered here, unless the way the MPLS/GMPLS network is provisioned could make a difference to the security of this user's server.

4. Security Threats

This section discusses the various network security threats that may endanger MPLS/GMPLS networks. RFC 4778 [RFC4778] provided the best current operational security practices in Internet Service Provider environments.

A successful attack on a particular MPLS/GMPLS network or on an SP's MPLS/GMPLS infrastructure may cause one or more of the following ill effects:

- Observation, modification, or deletion of a provider's or user's data.
- Replay of a provider's or user's data.
- Injection of inauthentic data into a provider's or user's traffic stream.
- Traffic pattern analysis on a provider's or user's traffic.
- Disruption of a provider's or user's connectivity.

- Degradation of a provider's service quality.
- Probing a provider's network to determine its configuration, capacity, or usage.

It is useful to consider that threats, whether malicious or accidental, may come from different categories of sources. For example, they may come from:

- Other users whose services are provided by the same MPLS/GMPLS core.
- The MPLS/GMPLS SP or persons working for it.
- Other persons who obtain physical access to an MPLS/GMPLS SP's site.
- Other persons who use social engineering methods to influence the behavior of an SP's personnel.
- Users of the MPLS/GMPLS network itself, e.g., intra-VPN threats. (Such threats are beyond the scope of this document.)
- Others, e.g., attackers from the Internet at large.
- Other SPs in the case of MPLS/GMPLS inter-provider connection. The core of the other provider may or may not be using MPLS/GMPLS.
- Those who create, deliver, install, and maintain software for network equipment.

Given that security is generally a tradeoff between expense and risk, it is also useful to consider the likelihood of different attacks occurring. There is at least a perceived difference in the likelihood of most types of attacks being successfully mounted in different environments, such as:

- An MPLS/GMPLS core interconnecting with another provider's core.
- An MPLS/GMPLS configuration transiting the public Internet.

Most types of attacks become easier to mount and hence more likely as the shared infrastructure via which service is provided expands from a single SP to multiple cooperating SPs to the global Internet. Attacks that may not be of sufficient likeliness to warrant concern in a closely controlled environment often merit defensive measures in broader, more open environments. In closed communities, it is often

practical to deal with misbehavior after the fact: an employee can be disciplined, for example.

The following sections discuss specific types of exploits that threaten MPLS/GMPLS networks.

4.1. Attacks on the Control Plane

This category encompasses attacks on the control structures operated by the SP with MPLS/GMPLS cores.

It should be noted that while connectivity in the MPLS control plane uses the same links and network resources as are used by the data plane, the GMPLS control plane may be provided by separate resources from those used in the data plane. That is, the GMPLS control plane may be physically separate from the data plane.

The different cases of physically congruent and physically separate control/data planes lead to slightly different possibilities of attack, although most of the cases are the same. Note that, for example, the data plane cannot be directly congested by an attack on a physically separate control plane as it could be if the control and data planes shared network resources. Note also that if the control plane uses diverse resources from the data plane, no assumptions should be made about the security of the control plane based on the security of the data plane resources.

This section is focused the outsider attack. The insider attack is discussed in Section 4.4.

4.1.1. LSP Creation by an Unauthorized Element

The unauthorized element can be a local CE or a router in another domain. An unauthorized element can generate MPLS signaling messages. At the least, this can result in extra control plane and forwarding state, and if successful, network bandwidth could be reserved unnecessarily. This may also result in theft of service or even compromise the entire network.

4.1.2. LSP Message Interception

This threat might be accomplished by monitoring network traffic, for example, after a physical intrusion. Without physical intrusion, it could be accomplished with an unauthorized software modification. Also, many technologies such as terrestrial microwave, satellite, or free-space optical could be intercepted without physical intrusion. If successful, it could provide information leading to label spoofing attacks. It also raises confidentiality issues.

4.1.3. Attacks against RSVP-TE

RSVP-TE, described in [RFC3209], is the control protocol used to set up GMPLS and traffic engineered MPLS tunnels.

There are two major types of denial-of-service (DoS) attacks against an MPLS domain based on RSVP-TE. The attacker may set up numerous unauthorized LSPs or may send a storm of RSVP messages. It has been demonstrated that unprotected routers running RSVP can be effectively disabled by both types of DoS attacks.

These attacks may even be combined, by using the unauthorized LSPs to transport additional RSVP (or other) messages across routers where they might otherwise be filtered out. RSVP attacks can be launched against adjacent routers at the border with the attacker, or against non-adjacent routers within the MPLS domain, if there is no effective mechanism to filter them out.

4.1.4. Attacks against LDP

LDP, described in [RFC5036], is the control protocol used to set up MPLS tunnels without TE.

There are two significant types of attack against LDP. An unauthorized network element can establish an LDP session by sending LDP Hello and LDP Init messages, leading to the potential setup of an LSP, as well as accompanying LDP state table consumption. Even without successfully establishing LSPs, an attacker can launch a DoS attack in the form of a storm of LDP Hello messages or LDP TCP SYN messages, leading to high CPU utilization or table space exhaustion on the target router.

4.1.5. Denial-of-Service Attacks on the Network Infrastructure

DoS attacks could be accomplished through an MPLS signaling storm, resulting in high CPU utilization and possibly leading to control-plane resource starvation.

Control-plane DoS attacks can be mounted specifically against the mechanisms the SP uses to provide various services, or against the general infrastructure of the service provider, e.g., P routers or shared aspects of PE routers. (An attack against the general infrastructure is within the scope of this document only if the attack can occur in relation with the MPLS/GMPLS infrastructure; otherwise, it is not an MPLS/GMPLS-specific issue.)

The attacks described in the following sections may each have denial of service as one of their effects. Other DoS attacks are also possible.

4.1.6. Attacks on the SP's MPLS/GMPLS Equipment via Management Interfaces

This includes unauthorized access to an SP's infrastructure equipment, for example, to reconfigure the equipment or to extract information (statistics, topology, etc.) pertaining to the network.

4.1.7. Cross-Connection of Traffic between Users

This refers to the event in which expected isolation between separate users (who may be VPN users) is breached. This includes cases such as:

- A site being connected into the "wrong" VPN.
- Traffic being replicated and sent to an unauthorized user.
- Two or more VPNs being improperly merged together.
- A point-to-point VPN connecting the wrong two points.
- Any packet or frame being improperly delivered outside the VPN to which it belongs

Misconnection or cross-connection of VPNs may be caused by service provider or equipment vendor error, or by the malicious action of an attacker. The breach may be physical (e.g., PE-CE links misconnected) or logical (e.g., improper device configuration).

Anecdotal evidence suggests that the cross-connection threat is one of the largest security concerns of users (or would-be users).

4.1.8. Attacks against Routing Protocols

This encompasses attacks against underlying routing protocols that are run by the SP and that directly support the MPLS/GMPLS core. (Attacks against the use of routing protocols for the distribution of backbone routes are beyond the scope of this document.) Specific attacks against popular routing protocols have been widely studied and are described in [RFC4593].

4.1.9. Other Attacks on Control Traffic

Besides routing and management protocols (covered separately in the previous sections), a number of other control protocols may be directly involved in delivering services by the MPLS/GMPLS core. These include but may not be limited to:

- MPLS signaling (LDP, RSVP-TE) discussed above in subsections 4.1.4 and 4.1.3
- PCE signaling
- IPsec signaling (IKE and IKEv2)
- ICMP and ICMPv6
- L2TP
- BGP-based membership discovery
- Database-based membership discovery (e.g., RADIUS)
- Other protocols that may be important to the control infrastructure, e.g., DNS, LMP, NTP, SNMP, and GRE.

Attacks might subvert or disrupt the activities of these protocols, for example via impersonation or DoS.

Note that all of the data-plane attacks can also be carried out against the packets of the control and management planes: insertion, spoofing, replay, deletion, pattern analysis, and other attacks mentioned above.

4.2. Attacks on the Data Plane

This category encompasses attacks on the provider's or end-user's data. Note that from the MPLS/GMPLS network end user's point of view, some of this might be control-plane traffic, e.g., routing protocols running from user site A to user site B via IP or non-IP connections, which may be some type of VPN.

4.2.1. Unauthorized Observation of Data Traffic

This refers to "sniffing" provider or end user packets and examining their contents. This can result in exposure of confidential information. It can also be a first step in other attacks (described below) in which the recorded data is modified and re-inserted, or simply replayed later.

4.2.2. Modification of Data Traffic

This refers to modifying the contents of packets as they traverse the MPLS/GMPLS core.

4.2.3. Insertion of Inauthentic Data Traffic: Spoofing and Replay

Spoofing refers to sending a user packets or inserting packets into a data stream that do not belong, with the objective of having them accepted by the recipient as legitimate. Also included in this category is the insertion of copies of once-legitimate packets that have been recorded and replayed.

4.2.4. Unauthorized Deletion of Data Traffic

This refers to causing packets to be discarded as they traverse the MPLS/GMPLS networks. This is a specific type of denial-of-service attack.

4.2.5. Unauthorized Traffic Pattern Analysis

This refers to "sniffing" provider or user packets and examining aspects or meta-aspects of them that may be visible even when the packets themselves are encrypted. An attacker might gain useful information based on the amount and timing of traffic, packet sizes, source and destination addresses, etc. For most users, this type of attack is generally considered to be significantly less of a concern than the other types discussed in this section.

4.2.6. Denial-of-Service Attacks

Denial-of-service (DoS) attacks are those in which an attacker attempts to disrupt or prevent the use of a service by its legitimate users. Taking network devices out of service, modifying their configuration, or overwhelming them with requests for service are several of the possible avenues for DoS attack.

Overwhelming the network with requests for service, otherwise known as a "resource exhaustion" DoS attack, may target any resource in the network, e.g., link bandwidth, packet forwarding capacity, session capacity for various protocols, CPU power, table size, storage overflows, and so on.

DoS attacks of the resource exhaustion type can be mounted against the data plane of a particular provider or end user by attempting to insert (spoofing) an overwhelming quantity of inauthentic data into the provider or end-user's network from outside of the trusted zone. Potential results might be to exhaust the bandwidth available to that

provider or end user, or to overwhelm the cryptographic authentication mechanisms of the provider or end user.

Data-plane resource exhaustion attacks can also be mounted by overwhelming the service provider's general (MPLS/GMPLS-independent) infrastructure with traffic. These attacks on the general infrastructure are not usually an MPLS/GMPLS-specific issue, unless the attack is mounted by another MPLS/GMPLS network user from a privileged position. (For example, an MPLS/GMPLS network user might be able to monopolize network data-plane resources and thus disrupt other users.)

Many DoS attacks use amplification, whereby the attacker co-opts otherwise innocent parties to increase the effect of the attack. The attacker may, for example, send packets to a broadcast or multicast address with the spoofed source address of the victim, and all of the recipients may then respond to the victim.

4.2.7. Misconnection

Misconnection may arise through deliberate attack, or through misconfiguration or misconnection of the network resources. The result is likely to be delivery of data to the wrong destination or black-holing of the data.

In GMPLS with physically diverse control and data planes, it may be possible for data-plane misconnection to go undetected by the control plane.

In optical networks under GMPLS control, misconnection may give rise to physical safety risks as unprotected lasers may be activated without warning.

4.3. Attacks on Operation and Management Plane

Attacks on the Operation and Management plane have been discussed extensively as general network security issues over the last 20 years. RFC 4778 [RFC4778] may serve as the best current operational security practices in Internet Service Provider environments. RFC 4377 [RFC4377] provided Operations and Management Requirements for MPLS networks. See also the Security Considerations of RFC 4377 and Section 7 of RFC 4378 [RFC4378].

Operation and Management across the MPLS-ICI could also be the source of security threats on the provider infrastructure as well as the service offered over the MPLS-ICI. A large volume of Operation and Management messages could overwhelm the processing capabilities of an ASBR if the ASBR is not properly protected. Maliciously generated

Operation and Management messages could also be used to bring down an otherwise healthy service (e.g., MPLS Pseudowire), and therefore affect service security. LSP ping does not support authentication today, and that support should be a subject for future considerations. Bidirectional Forwarding Detection (BFD), however, does have support for carrying an authentication object. It also supports Time-To-Live (TTL) processing as an anti-replay measure. Implementations conformant with this MPLS-ICI should support BFD authentication and must support the procedures for TTL processing.

Regarding GMPLS Operation and Management considerations in optical interworking, there is a good discussion on security for management interfaces to Network Elements [OIF-Sec-Mag].

Network elements typically have one or more (in some cases many) Operation and Management interfaces used for network management, billing and accounting, configuration, maintenance, and other administrative activities.

Remote access to a network element through these Operation and Management interfaces is frequently a requirement. Securing the control protocols while leaving these Operation and Management interfaces unprotected opens up a huge security vulnerability. Network elements are an attractive target for intruders who want to disrupt or gain free access to telecommunications facilities. Much has been written about this subject since the 1980s. In the 1990s, telecommunications facilities were identified in the U.S. and other countries as part of the "critical infrastructure", and increased emphasis was placed on thwarting such attacks from a wider range of potentially well-funded and determined adversaries.

At one time, careful access controls and password management were a sufficient defense, but are no longer. Networks using the TCP/IP protocol suite are vulnerable to forged source addresses, recording and later replay, packet sniffers picking up passwords, re-routing of traffic to facilitate eavesdropping or tampering, active hijacking attacks of TCP connections, and a variety of denial-of-service attacks. The ease of forging TCP/IP packets is the main reason network management protocols lacking strong security have not been used to configure network elements (e.g., with the SNMP SET command).

Readily available hacking tools exist that let an eavesdropper on a LAN take over one end of any TCP connection, so that the legitimate party is cut off. In addition, enterprises and Service Providers in some jurisdictions need to safeguard data about their users and network configurations from prying. An attacker could eavesdrop and

observe traffic to analyze usage patterns and map a network configuration; an attacker could also gain access to systems and manipulate configuration data or send malicious commands.

Therefore, in addition to authenticating the human user, more sophisticated protocol security is needed for Operation and Management interfaces, especially when they are configured over TCP/IP stacks. Finally, relying on a perimeter defense, such as firewalls, is insufficient protection against "insider attacks" or against penetrations that compromise a system inside the firewall as a launching pad to attack network elements. The insider attack is discussed in the following session.

4.4. Insider Attacks Considerations

The chain of trust model means that MPLS and GMPLS networks are particularly vulnerable to insider attacks. These can be launched by any malign person with access to any LSR in the trust domain. Insider attacks could also be launched by compromised software within the trust domain. Such attacks could, for example, advertise non-existent resources, modify advertisements from other routers, request unwanted LSPs that use network resources, or deny or modify legitimate LSP requests.

Protection against insider attacks is largely for future study in MPLS and GMPLS networks. Some protection can be obtained by providing strict security for software upgrades and tight OAM access control procedures. Further protection can be achieved by strict control of user (i.e., operator) access to LSRs. Software change management and change tracking (e.g., CVS diffs from text-based configuration files) helps in spotting irregularities and human errors. In some cases, configuration change approval processes may also be warranted. Software tools could be used to check configurations for consistency and compliance. Software tools may also be used to monitor and report network behavior and activity in order to quickly spot any irregularities that may be the result of an insider attack.

5. Defensive Techniques for MPLS/GMPLS Networks

The defensive techniques discussed in this document are intended to describe methods by which some security threats can be addressed. They are not intended as requirements for all MPLS/GMPLS implementations. The MPLS/GMPLS provider should determine the applicability of these techniques to the provider's specific service offerings, and the end user may wish to assess the value of these techniques to the user's service requirements. The operational environment determines the security requirements. Therefore,

protocol designers need to provide a full set of security services, which can be used where appropriate.

The techniques discussed here include encryption, authentication, filtering, firewalls, access control, isolation, aggregation, and others.

Often, security is achieved by careful protocol design, rather than by adding a security method. For example, one method of mitigating DoS attacks is to make sure that innocent parties cannot be used to amplify the attack. Security works better when it is "designed in" rather than "added on".

Nothing is ever 100% secure. Defense therefore involves protecting against those attacks that are most likely to occur or that have the most direct consequences if successful. For those attacks that are protected against, absolute protection is seldom achievable; more often it is sufficient just to make the cost of a successful attack greater than what the adversary will be willing or able to expend.

Successfully defending against an attack does not necessarily mean the attack must be prevented from happening or from reaching its target. In many cases, the network can instead be designed to withstand the attack. For example, the introduction of inauthentic packets could be defended against by preventing their introduction in the first place, or by making it possible to identify and eliminate them before delivery to the MPLS/GMPLS user's system. The latter is frequently a much easier task.

5.1. Authentication

To prevent security issues arising from some DoS attacks or from malicious or accidental misconfiguration, it is critical that devices in the MPLS/GMPLS should only accept connections or control messages from valid sources. Authentication refers to methods to ensure that message sources are properly identified by the MPLS/GMPLS devices with which they communicate. This section focuses on identifying the scenarios in which sender authentication is required and recommends authentication mechanisms for these scenarios.

Cryptographic techniques (authentication, integrity, and encryption) do not protect against some types of denial-of-service attacks, specifically resource exhaustion attacks based on CPU or bandwidth exhaustion. In fact, the software-based cryptographic processing required to decrypt or check authentication may in some cases increase the effect of these resource exhaustion attacks. With a hardware cryptographic accelerator, attack packets can be dropped at line speed without a cost to software cycles. Cryptographic

techniques may, however, be useful against resource exhaustion attacks based on the exhaustion of state information (e.g., TCP SYN attacks).

The MPLS data plane, as presently defined, is not amenable to source authentication, as there are no source identifiers in the MPLS packet to authenticate. The MPLS label is only locally meaningful. It may be assigned by a downstream node or upstream node for multicast support.

When the MPLS payload carries identifiers that may be authenticated (e.g., IP packets), authentication may be carried out at the client level, but this does not help the MPLS SP, as these client identifiers belong to an external, untrusted network.

5.1.1. Management System Authentication

Management system authentication includes the authentication of a PE to a centrally managed network management or directory server when directory-based "auto-discovery" is used. It also includes authentication of a CE to the configuration server, when a configuration server system is used.

Authentication should be bidirectional, including PE or CE to configuration server authentication for the PE or CE to be certain it is communicating with the right server.

5.1.2. Peer-to-Peer Authentication

Peer-to-peer authentication includes peer authentication for network control protocols (e.g., LDP, BGP, etc.) and other peer authentication (i.e., authentication of one IPsec security gateway by another).

Authentication should be bidirectional, including PE or CE to configuration server authentication for the PE or CE to be certain it is communicating with the right server.

As indicated in Section 5.1.1, authentication should be bidirectional.

5.1.3. Cryptographic Techniques for Authenticating Identity

Cryptographic techniques offer several mechanisms for authenticating the identity of devices or individuals. These include the use of shared secret keys, one-time keys generated by accessory devices or software, user-ID and password pairs, and a range of public-private

key systems. Another approach is to use a hierarchical Certification Authority system to provide digital certificates.

This section describes or provides references to the specific cryptographic approaches for authenticating identity. These approaches provide secure mechanisms for most of the authentication scenarios required in securing an MPLS/GMPLS network.

5.2. Cryptographic Techniques

MPLS/GMPLS defenses against a wide variety of attacks can be enhanced by the proper application of cryptographic techniques. These same cryptographic techniques are applicable to general network communications and can provide confidentiality (encryption) of communication between devices, authenticate the identities of the devices, and detect whether the data being communicated has been changed during transit or replayed from previous messages.

Several aspects of authentication are addressed in some detail in a separate "Authentication" section (Section 5.1).

Cryptographic methods add complexity to a service and thus, for a few reasons, may not be the most practical solution in every case. Cryptography adds an additional computational burden to devices, which may reduce the number of user connections that can be handled on a device or otherwise reduce the capacity of the device, potentially driving up the provider's costs. Typically, configuring encryption services on devices adds to the complexity of their configuration and adds labor cost. Some key management system is usually needed. Packet sizes are typically increased when the packets are encrypted or have integrity checks or replay counters added, increasing the network traffic load and adding to the likelihood of packet fragmentation with its increased overhead. (This packet length increase can often be mitigated to some extent by data compression techniques, but at the expense of additional computational burden.) Finally, some providers may employ enough other defensive techniques, such as physical isolation or filtering and firewall techniques, that they may not perceive additional benefit from encryption techniques.

Users may wish to provide confidentiality end to end. Generally, encrypting for confidentiality must be accompanied with cryptographic integrity checks to prevent certain active attacks against the encrypted communications. On today's processors, encryption and integrity checks run extremely quickly, but key management may be more demanding in terms of both computational and administrative overhead.

The trust model among the MPLS/GMPLS user, the MPLS/GMPLS provider, and other parts of the network is a major element in determining the applicability of cryptographic protection for any specific MPLS/GMPLS implementation. In particular, it determines where cryptographic protection should be applied:

- If the data path between the user's site and the provider's PE is not trusted, then it may be used on the PE-CE link.
- If some part of the backbone network is not trusted, particularly in implementations where traffic may travel across the Internet or multiple providers' networks, then the PE-PE traffic may be cryptographically protected. One also should consider cases where L1 technology may be vulnerable to eavesdropping.
- If the user does not trust any zone outside of its premises, it may require end-to-end or CE-CE cryptographic protection. This fits within the scope of this MPLS/GMPLS security framework when the CE is provisioned by the MPLS/GMPLS provider.
- If the user requires remote access to its site from a system at a location that is not a customer location (for example, access by a traveler), there may be a requirement for cryptographically protecting the traffic between that system and an access point or a customer's site. If the MPLS/GMPLS provider supplies the access point, then the customer must cooperate with the provider to handle the access control services for the remote users. These access control services are usually protected cryptographically, as well.

Access control usually starts with authentication of the entity. If cryptographic services are part of the scenario, then it is important to bind the authentication to the key management. Otherwise, the protocol is vulnerable to being hijacked between the authentication and key management.

Although CE-CE cryptographic protection can provide integrity and confidentiality against third parties, if the MPLS/GMPLS provider has complete management control over the CE (encryption) devices, then it may be possible for the provider to gain access to the user's traffic or internal network. Encryption devices could potentially be reconfigured to use null encryption, bypass cryptographic processing altogether, reveal internal configuration, or provide some means of sniffing or diverting unencrypted traffic. Thus an implementation using CE-CE encryption needs to consider the trust relationship between the MPLS/GMPLS user and provider. MPLS/GMPLS users and providers may wish to negotiate a service level agreement (SLA) for CE-CE encryption that provides an acceptable demarcation of

responsibilities for management of cryptographic protection on the CE devices. The demarcation may also be affected by the capabilities of the CE devices. For example, the CE might support some partitioning of management, a configuration lock-down ability, or shared capability to verify the configuration. In general, the MPLS/GMPLS user needs to have a fairly high level of trust that the MPLS/GMPLS provider will properly provision and manage the CE devices, if the managed CE-CE model is used.

5.2.1. IPsec in MPLS/GMPLS

IPsec [RFC4301] [RFC4302] [RFC4835] [RFC4306] [RFC4309] [RFC2411] [IPSECME-ROADMAP] is the security protocol of choice for protection at the IP layer. IPsec provides robust security for IP traffic between pairs of devices. Non-IP traffic, such as IS-IS routing, must be converted to IP (e.g., by encapsulation) in order to use IPsec. When the MPLS is encapsulating IP traffic, then IPsec covers the encryption of the IP client layer; for non-IP client traffic, see Section 5.2.4 (MPLS PWs).

In the MPLS/GMPLS model, IPsec can be employed to protect IP traffic between PEs, between a PE and a CE, or from CE to CE. CE-to-CE IPsec may be employed in either a provider-provisioned or a user-provisioned model. Likewise, IPsec protection of data performed within the user's site is outside the scope of this document, because it is simply handled as user data by the MPLS/GMPLS core. However, if the SP performs compression, pre-encryption will have a major effect on that operation.

IPsec does not itself specify cryptographic algorithms. It can use a variety of integrity or confidentiality algorithms (or even combined integrity and confidentiality algorithms) with various key lengths, such as AES encryption or AES message integrity checks. There are trade-offs between key length, computational burden, and the level of security of the encryption. A full discussion of these trade-offs is beyond the scope of this document. In practice, any currently recommended IPsec protection offers enough security to reduce the likelihood of its being directly targeted by an attacker substantially; other weaker links in the chain of security are likely to be attacked first. MPLS/GMPLS users may wish to use a Service Level Agreement (SLA) specifying the SP's responsibility for ensuring data integrity and confidentiality, rather than analyzing the specific encryption techniques used in the MPLS/GMPLS service.

Encryption algorithms generally come with two parameters: mode such as Cipher Block Chaining and key length such as AES-192. (This should not be confused with two other senses in which the word "mode" is used: IPsec itself can be used in Tunnel Mode or Transport Mode,

and IKE [version 1] uses Main Mode, Aggressive Mode, or Quick Mode). It should be stressed that IPsec encryption without an integrity check is a state of sin.

For many of the MPLS/GMPLS provider's network control messages and some user requirements, cryptographic authentication of messages without encryption of the contents of the message may provide appropriate security. Using IPsec, authentication of messages is provided by the Authentication Header (AH) or through the use of the Encapsulating Security Protocol (ESP) with NULL encryption. Where control messages require integrity but do not use IPsec, other cryptographic authentication methods are often available. Message authentication methods currently considered to be secure are based on hashed message authentication codes (HMAC) [RFC2104] implemented with a secure hash algorithm such as Secure Hash Algorithm 1 (SHA-1) [RFC3174]. No attacks against HMAC SHA-1 are likely to play out in the near future, but it is possible that people will soon find SHA-1 collisions. Thus, it is important that mechanisms be designed to be flexible about the choice of hash functions and message integrity checks. Also, many of these mechanisms do not include a convenient way to manage and update keys.

A mechanism to provide a combination of confidentiality, data-origin authentication, and connectionless integrity is the use of AES in GCM (Counter with CBC-MAC) mode (RFC 4106) [RFC4106].

5.2.2. MPLS / GMPLS Diffserv and IPsec

MPLS and GMPLS, which provide differentiated services based on traffic type, may encounter some conflicts with IPsec encryption of traffic. Because encryption hides the content of the packets, it may not be possible to differentiate the encrypted traffic in the same manner as unencrypted traffic. Although Diffserv markings are copied to the IPsec header and can provide some differentiation, not all traffic types can be accommodated by this mechanism. Using IPsec without IKE or IKEv2 (the better choice) is not advisable. IKEv2 provides IPsec Security Association creation and management, entity authentication, key agreement, and key update. It works with a variety of authentication methods including pre-shared keys, public key certificates, and EAP. If DoS attacks against IKEv2 are considered an important threat to mitigate, the cookie-based anti-spoofing feature of IKEv2 should be used. IKEv2 has its own set of cryptographic methods, but any of the default suites specified in [RFC4308] or [RFC4869] provides more than adequate security.

5.2.3. Encryption for Device Configuration and Management

For configuration and management of MPLS/GMPLS devices, encryption and authentication of the management connection at a level comparable to that provided by IPsec is desirable.

Several methods of transporting MPLS/GMPLS device management traffic offer authentication, integrity, and confidentiality.

- Secure Shell (SSH) offers protection for TELNET [STD8] or terminal-like connections to allow device configuration.
- SNMPv3 [STD62] provides encrypted and authenticated protection for SNMP-managed devices.
- Transport Layer Security (TLS) [RFC5246] and the closely-related Secure Sockets Layer (SSL) are widely used for securing HTTP-based communication, and thus can provide support for most XML- and SOAP-based device management approaches.
- Since 2004, there has been extensive work proceeding in several organizations (OASIS, W3C, WS-I, and others) on securing device management traffic within a "Web Services" framework, using a wide variety of security models, and providing support for multiple security token formats, multiple trust domains, multiple signature formats, and multiple encryption technologies.
- IPsec provides security services including integrity and confidentiality at the network layer. With regards to device management, its current use is primarily focused on in-band management of user-managed IPsec gateway devices.
- There is recent work in the ISMS WG (Integrated Security Model for SNMP Working Group) to define how to use SSH to secure SNMP, due to the limited deployment of SNMPv3, and the possibility of using Kerberos, particularly for interfaces like TELNET, where client code exists.

5.2.4. Security Considerations for MPLS Pseudowires

In addition to IP traffic, MPLS networks may be used to transport other services such as Ethernet, ATM, Frame Relay, and TDM. This is done by setting up pseudowires (PWs) that tunnel the native service through the MPLS core by encapsulating at the edges. The PWE architecture is defined in [RFC3985].

PW tunnels may be set up using the PWE control protocol based on LDP [RFC4447], and thus security considerations for LDP will most likely be applicable to the PWE3 control protocol as well.

PW user packets contain at least one MPLS label (the PW label) and may contain one or more MPLS tunnel labels. After the label stack, there is a four-byte control word (which is optional for some PW types), followed by the native service payload. It must be stressed that encapsulation of MPLS PW packets in IP for the purpose of enabling use of IPsec mechanisms is not a valid option.

The following is a non-exhaustive list of PW-specific threats:

- Unauthorized setup of a PW (e.g., to gain access to a customer network)
- Unauthorized teardown of a PW (thus causing denial of service)
- Malicious reroute of a PW
- Unauthorized observation of PW packets
- Traffic analysis of PW connectivity
- Unauthorized insertion of PW packets
- Unauthorized modification of PW packets
- Unauthorized deletion of PW packets replay of PW packets
- Denial of service or significant impact on PW service quality

These threats are not mutually exclusive, for example, rerouting can be used for snooping or insertion/deletion/replay, etc. Multisegment PWs introduce additional weaknesses at their stitching points.

The PW user plane suffers from the following inherent security weaknesses:

- Since the PW label is the only identifier in the packet, there is no authenticatable source address.
- Since guessing a valid PW label is not difficult, it is relatively easy to introduce seemingly valid foreign packets.
- Since the PW packet is not self-describing, minor modification of control-plane packets renders the data-plane traffic useless.

- The control-word sequence number processing algorithm is susceptible to a DoS attack.

The PWE control protocol introduces its own weaknesses:

- No (secure) peer autodiscovery technique has been standardized .
- PE authentication is not mandated, so an intruder can potentially impersonate a PE; after impersonating a PE, unauthorized PWs may be set up, consuming resources and perhaps allowing access to user networks.
- Alternately, desired PWs may be torn down, giving rise to denial of service.

The following characteristics of PWs can be considered security strengths:

- The most obvious attacks require compromising edge or core routers (although not necessarily those along the PW path).
- Adequate protection of the control-plane messaging is sufficient to rule out many types of attacks.
- PEs are usually configured to reject MPLS packets from outside the service provider network, thus ruling out insertion of PW packets from the outside (since IP packets cannot masquerade as PW packets).

5.2.5. End-to-End versus Hop-by-Hop Protection Tradeoffs in MPLS/GMPLS

In MPLS/GMPLS, cryptographic protection could potentially be applied to the MPLS/GMPLS traffic at several different places. This section discusses some of the tradeoffs in implementing encryption in several different connection topologies among different devices within an MPLS/GMPLS network.

Cryptographic protection typically involves a pair of devices that protect the traffic passing between them. The devices may be directly connected (over a single "hop"), or intervening devices may transport the protected traffic between the pair of devices. The extreme cases involve using protection between every adjacent pair of devices along a given path (hop-by-hop), or using protection only between the end devices along a given path (end-to-end). To keep this discussion within the scope of this document, the latter ("end-to-end") case considered here is CE-to-CE rather than fully end-to-end.

Figure 3 depicts a simplified topology showing the Customer Edge (CE) devices, the Provider Edge (PE) devices, and a variable number (three are shown) of Provider core (P) devices, which might be present along the path between two sites in a single VPN operated by a single service provider (SP).

Site_1---CE---PE---P---P---P---PE---CE---Site_2

Figure 3: Simplified Topology Traversing through MPLS/GMPLS Core

Within this simplified topology, and assuming that the P devices are not involved with cryptographic protection, four basic, feasible configurations exist for protecting connections among the devices:

- 1) Site-to-site (CE-to-CE) - Apply confidentiality or integrity services between the two CE devices, so that traffic will be protected throughout the SP's network.
- 2) Provider edge-to-edge (PE-to-PE) - Apply confidentiality or integrity services between the two PE devices. Unprotected traffic is received at one PE from the customer's CE, then it is protected for transmission through the SP's network to the other PE, and finally it is decrypted or checked for integrity and sent to the other CE.
- 3) Access link (CE-to-PE) - Apply confidentiality or integrity services between the CE and PE on each side or on only one side.
- 4) Configurations 2 and 3 above can also be combined, with confidentiality or integrity running from CE to PE, then PE to PE, and then PE to CE.

Among the four feasible configurations, key tradeoffs in considering encryption include:

- Vulnerability to link eavesdropping or tampering - assuming an attacker can observe or modify data in transit on the links, would it be protected by encryption?
- Vulnerability to device compromise - assuming an attacker can get access to a device (or freely alter its configuration), would the data be protected?
- Complexity of device configuration and management - given the number of sites per VPN customer as N_{ce} and the number of PEs participating in a given VPN as N_{pe} , how many device configurations need to be created or maintained, and how do those configurations scale?

- Processing load on devices - how many cryptographic operations must be performed given N packets? - This raises considerations of device capacity and perhaps end-to-end delay.
- Ability of the SP to provide enhanced services (QoS, firewall, intrusion detection, etc.) - Can the SP inspect the data to provide these services?

These tradeoffs are discussed for each configuration, below:

1) Site-to-site (CE-to-CE)

Link eavesdropping or tampering - protected on all links. Device compromise - vulnerable to CE compromise.

Complexity - single administration, responsible for one device per site (N_{ce} devices), but overall configuration per VPN scales as N_{ce}^2 .

Though the complexity may be reduced: 1) In practice, as N_{ce} grows, the number of VPNs falls off from being a full clique; 2) If the CEs run an automated key management protocol, then they should be able to set up and tear down secured VPNs without any intervention.

Processing load - on each of the two CEs, each packet is cryptographically processed (2P), though the protection may be "integrity check only" or "integrity check plus encryption."

Enhanced services - severely limited; typically only Diffserv markings are visible to the SP, allowing some QoS services. The CEs could also use the IPv6 Flow Label to identify traffic classes.

2) Provider Edge-to-Edge (PE-to-PE)

Link eavesdropping or tampering - vulnerable on CE-PE links; protected on SP's network links.

Device compromise - vulnerable to CE or PE compromise.

Complexity - single administration, N_{pe} devices to configure. (Multiple sites may share a PE device so N_{pe} is typically much smaller than N_{ce} .) Scalability of the overall configuration depends on the PPVPN type: if the cryptographic protection is separate per VPN context, it scales as N_{pe}^2 per customer VPN. If it is per-PE, it scales as N_{pe}^2 for all customer VPNs combined.

Processing load - on each of the two PEs, each packet is cryptographically processed (2P).

Enhanced services - full; SP can apply any enhancements based on detailed view of traffic.

3) Access Link (CE-to-PE)

Link eavesdropping or tampering - protected on CE-PE link; vulnerable on SP's network links.

Device compromise - vulnerable to CE or PE compromise.

Complexity - two administrations (customer and SP) with device configuration on each side (Nce + Npe devices to configure), but because there is no mesh, the overall configuration scales as Nce.

Processing load - on each of the two CEs, each packet is cryptographically processed, plus on each of the two PEs, each packet is cryptographically processed (4P).

Enhanced services - full; SP can apply any enhancements based on a detailed view of traffic.

4) Combined Access link and PE-to-PE (essentially hop-by-hop).

Link eavesdropping or tampering - protected on all links.

Device compromise - vulnerable to CE or PE compromise.

Complexity - two administrations (customer and SP) with device configuration on each side (Nce + Npe devices to configure). Scalability of the overall configuration depends on the PPVPN type: If the cryptographic processing is separate per VPN context, it scales as Npe^{**2} per customer VPN. If it is per-PE, it scales as Npe^{**2} for all customer VPNs combined.

Processing load - on each of the two CEs, each packet is cryptographically processed, plus on each of the two PEs, each packet is cryptographically processed twice (6P).

Enhanced services - full; SP can apply any enhancements based on a detailed view of traffic.

Given the tradeoffs discussed above, a few conclusions can be drawn:

- Configurations 2 and 3 are subsets of 4 that may be appropriate alternatives to 4 under certain threat models; the remainder of these conclusions compare 1 (CE-to-CE) versus 4 (combined access links and PE-to-PE).
- If protection from link eavesdropping or tampering is all that is important, then configurations 1 and 4 are equivalent.
- If protection from device compromise is most important and the threat is to the CE devices, both cases are equivalent; if the threat is to the PE devices, configuration 1 is better.
- If reducing complexity is most important, and the size of the network is small, configuration 1 is better. Otherwise, configuration 4 is better because rather than a mesh of CE devices, it requires a smaller mesh of PE devices. Also, under some PPVPN approaches, the scaling of 4 is further improved by sharing the same PE-PE mesh across all VPN contexts. The scaling advantage of 4 may be increased or decreased in any given situation if the CE devices are simpler to configure than the PE devices, or vice-versa.
- If the overall processing load is a key factor, then 1 is better, unless the PEs come with a hardware encryption accelerator and the CEs do not.
- If the availability of enhanced services support from the SP is most important, then 4 is best.
- If users are concerned with having their VPNs misconnected with other users' VPNs, then encryption with 1 can provide protection.

As a quick overall conclusion, CE-to-CE protection is better against device compromise, but this comes at the cost of enhanced services and at the cost of operational complexity due to the $O(n^2)$ scaling of a larger mesh.

This analysis of site-to-site vs. hop-by-hop tradeoffs does not explicitly include cases of multiple providers cooperating to provide a PPVPN service, public Internet VPN connectivity, or remote access VPN service, but many of the tradeoffs are similar.

In addition to the simplified models, the following should also be considered:

- There are reasons, perhaps, to protect a specific P-to-P or PE-to-P.
- There may be reasons to do multiple encryptions over certain segments. One may be using an encrypted wireless link under our IPsec VPN to access an SSL-secured web site to download encrypted email attachments: four layers.)
- It may be appropriate that, for example, cryptographic integrity checks are applied end to end, and confidentiality is applied over a shorter span.
- Different cryptographic protection may be required for control protocols and data traffic.
- Attention needs to be given to how auxiliary traffic is protected, e.g., the ICMPv6 packets that flow back during PMTU discovery, among other examples.

5.3. Access Control Techniques

Access control techniques include packet-by-packet or packet-flow-by-packet-flow access control by means of filters and firewalls on IPv4/IPv6 packets, as well as by means of admitting a "session" for a control, signaling, or management protocol. Enforcement of access control by isolated infrastructure addresses is discussed in Section 5.4 of this document.

In this document, we distinguish between filtering and firewalls based primarily on the direction of traffic flow. We define filtering as being applicable to unidirectional traffic, while a firewall can analyze and control both sides of a conversation.

The definition has two significant corollaries:

- Routing or traffic flow symmetry: A firewall typically requires routing symmetry, which is usually enforced by locating a firewall where the network topology assures that both sides of a conversation will pass through the firewall. A filter can operate upon traffic flowing in one direction, without considering traffic in the reverse direction. Beware that this concept could result in a single point of failure.

- Statefulness: Because it receives both sides of a conversation, a firewall may be able to interpret a significant amount of information concerning the state of that conversation and use this information to control access. A filter can maintain some limited state information on a unidirectional flow of packets, but cannot determine the state of the bidirectional conversation as precisely as a firewall.

For a general description on filtering and rate limiting for IP networks, please also see [OPSEC-FILTER].

5.3.1. Filtering

It is relatively common for routers to filter packets. That is, routers can look for particular values in certain fields of the IP or higher-level (e.g., TCP or UDP) headers. Packets matching the criteria associated with a particular filter may either be discarded or given special treatment. Today, not only routers, but most end hosts have filters, and every instance of IPsec is also a filter [RFC4301].

In discussing filters, it is useful to separate the filter characteristics that may be used to determine whether a packet matches a filter from the packet actions applied to those packets matching a particular filter.

o Filter Characteristics

Filter characteristics or rules are used to determine whether a particular packet or set of packets matches a particular filter.

In many cases, filter characteristics may be stateless. A stateless filter determines whether a particular packet matches a filter based solely on the filter definition, normal forwarding information (such as the next hop for a packet), the interface on which a packet arrived, and the contents of that individual packet. Typically, stateless filters may consider the incoming and outgoing logical or physical interface, information in the IP header, and information in higher-layer headers such as the TCP or UDP header. Information in the IP header to be considered may for example include source and destination IP addresses; Protocol field, Fragment Offset, and TOS field in IPv4; or Next Header, Extension Headers, Flow label, etc. in IPv6. Filters also may consider fields in the TCP or UDP header such as the Port numbers, the SYN field in the TCP header, as well as ICMP and ICMPv6 type.

Stateful filtering maintains packet-specific state information to aid in determining whether a filter rule has been met. For example, a device might apply stateless filtering to the first fragment of a fragmented IPv4 packet. If the filter matches, then the data unit ID may be remembered and other fragments of the same packet may then be considered to match the same filter. Stateful filtering is more commonly done in firewalls, although firewall technology may be added to routers. The data unit ID can also be a Fragment Extension Header Identification field in IPv6.

o Actions based on Filter Results

If a packet, or a series of packets, matches a specific filter, then a variety of actions may be taken based on that match. Examples of such actions include:

- Discard

In many cases, filters are set to catch certain undesirable packets. Examples may include packets with forged or invalid source addresses, packets that are part of a DoS or Distributed DoS (DDoS) attack, or packets trying to access unallowed resources (such as network management packets from an unauthorized source). Where such filters are activated, it is common to discard the packet or set of packets matching the filter silently. The discarded packets may of course also be counted or logged.

- Set CoS

A filter may be used to set the class of service associated with the packet.

- Count packets or bytes

- Rate Limit

In some cases, the set of packets matching a particular filter may be limited to a specified bandwidth. In this case, packets or bytes would be counted, and would be forwarded normally up to the specified limit. Excess packets may be discarded or may be marked (for example, by setting a "discard eligible" bit in the IPv4 ToS field, or changing the EXP value to identify traffic as being out of contract).

- Forward and Copy

It is useful in some cases to forward some set of packets normally, but also to send a copy to a specified other address or interface. For example, this may be used to implement a lawful intercept capability or to feed selected packets to an Intrusion Detection System.

o Other Packet Filters Issues

Filtering performance may vary widely according to implementation and the types and number of rules. Without acceptable performance, filtering is not useful.

The precise definition of "acceptable" may vary from SP to SP, and may depend upon the intended use of the filters. For example, for some uses, a filter may be turned on all the time to set CoS, to prevent an attack, or to mitigate the effect of a possible future attack. In this case, it is likely that the SP will want the filter to have minimal or no impact on performance. In other cases, a filter may be turned on only in response to a major attack (such as a major DDoS attack). In this case, a greater performance impact may be acceptable to some service providers.

A key consideration with the use of packet filters is that they can provide few options for filtering packets carrying encrypted data. Because the data itself is not accessible, only packet header information or other unencrypted fields can be used for filtering.

5.3.2. Firewalls

Firewalls provide a mechanism for controlling traffic passing between different trusted zones in the MPLS/GMPLS model or between a trusted zone and an untrusted zone. Firewalls typically provide much more functionality than filters, because they may be able to apply detailed analysis and logical functions to flows, and not just to individual packets. They may offer a variety of complex services, such as threshold-driven DoS attack protection, virus scanning, acting as a TCP connection proxy, etc.

As with other access control techniques, the value of firewalls depends on a clear understanding of the topologies of the MPLS/GMPLS core network, the user networks, and the threat model. Their effectiveness depends on a topology with a clearly defined inside (secure) and outside (not secure).

Firewalls may be applied to help protect MPLS/GMPLS core network functions from attacks originating from the Internet or from

MPLS/GMPLS user sites, but typically other defensive techniques will be used for this purpose.

Where firewalls are employed as a service to protect user VPN sites from the Internet, different VPN users, and even different sites of a single VPN user, may have varying firewall requirements. The overall PPVPN logical and physical topology, along with the capabilities of the devices implementing the firewall services, has a significant effect on the feasibility and manageability of such varied firewall service offerings.

Another consideration with the use of firewalls is that they can provide few options for handling packets carrying encrypted data. Because the data itself is not accessible, only packet header information, other unencrypted fields, or analysis of the flow of encrypted packets can be used for making decisions on accepting or rejecting encrypted traffic.

Two approaches of using firewalls are to move the firewall outside of the encrypted part of the path or to register and pre-approve the encrypted session with the firewall.

Handling DoS attacks has become increasingly important. Useful guidelines include the following:

1. Perform ingress filtering everywhere.
2. Be able to filter DoS attack packets at line speed.
3. Do not allow oneself to amplify attacks.
4. Continue processing legitimate traffic. Over provide for heavy loads. Use diverse locations, technologies, etc.

5.3.3. Access Control to Management Interfaces

Most of the security issues related to management interfaces can be addressed through the use of authentication techniques as described in the section on authentication (Section 5.1). However, additional security may be provided by controlling access to management interfaces in other ways.

The Optical Internetworking Forum has done relevant work on protecting such interfaces with TLS, SSH, Kerberos, IPsec, WSS, etc. See "Security for Management Interfaces to Network Elements" [OIF-SMI-01.0] and "Addendum to the Security for Management Interfaces to Network Elements" [OIF-SMI-02.1]. See also the work in the ISMS WG (<http://datatracker.ietf.org/wg/isms/charter/>).

Management interfaces, especially console ports on MPLS/GMPLS devices, may be configured so they are only accessible out-of-band, through a system that is physically or logically separated from the rest of the MPLS/GMPLS infrastructure.

Where management interfaces are accessible in-band within the MPLS/GMPLS domain, filtering or firewalling techniques can be used to restrict unauthorized in-band traffic from having access to management interfaces. Depending on device capabilities, these filtering or firewalling techniques can be configured either on other devices through which the traffic might pass, or on the individual MPLS/GMPLS devices themselves.

5.4. Use of Isolated Infrastructure

One way to protect the infrastructure used for support of MPLS/GMPLS is to separate the resources for support of MPLS/GMPLS services from the resources used for other purposes (such as support of Internet services). In some cases, this may involve using physically separate equipment for VPN services, or even a physically separate network.

For example, PE-based IPVPNs may be run on a separate backbone not connected to the Internet, or may use separate edge routers from those supporting Internet service. Private IPv4 addresses (local to the provider and non-routable over the Internet) are sometimes used to provide additional separation. For a discussion of comparable techniques for IPv6, see "Local Network Protection for IPv6," RFC 4864 [RFC4864].

In a GMPLS network, it is possible to operate the control plane using physically separate resources from those used for the data plane. This means that the data-plane resources can be physically protected and isolated from other equipment to protect users' data while the control and management traffic uses network resources that can be accessed by operators to configure the network. Conversely, the separation of control and data traffic may lead the operator to consider that the network is secure because the data-plane resources are physically secure. However, this is not the case if the control plane can be attacked through a shared or open network, and control-plane protection techniques must still be applied.

5.5. Use of Aggregated Infrastructure

In general, it is not feasible to use a completely separate set of resources for support of each service. In fact, one of the main reasons for MPLS/GMPLS enabled services is to allow sharing of resources between multiple services and multiple users. Thus, even if certain services use a separate network from Internet services,

nonetheless there will still be multiple MPLS/GMPLS users sharing the same network resources. In some cases, MPLS/GMPLS services will share network resources with Internet services or other services.

It is therefore important for MPLS/GMPLS services to provide protection between resources used by different parties. Thus, a well-behaved MPLS/GMPLS user should be protected from possible misbehavior by other users. This requires several security measurements to be implemented. Resource limits can be placed on a per service and per user basis. Possibilities include, for example, using a virtual router or logical router to define hardware or software resource limits per service or per individual user; using rate limiting per Virtual Routing and Forwarding (VRF) or per Internet connection to provide bandwidth protection; or using resource reservation for control-plane traffic. In addition to bandwidth protection, separate resource allocation can be used to limit security attacks only to directly impacted service(s) or customer(s). Strict, separate, and clearly defined engineering rules and provisioning procedures can reduce the risks of network-wide impact of a control-plane attack, DoS attack, or misconfiguration.

In general, the use of aggregated infrastructure allows the service provider to benefit from stochastic multiplexing of multiple bursty flows, and also may in some cases thwart traffic pattern analysis by combining the data from multiple users. However, service providers must minimize security risks introduced from any individual service or individual users.

5.6. Service Provider Quality Control Processes

Deployment of provider-provisioned VPN services in general requires a relatively large amount of configuration by the SP. For example, the SP needs to configure which VPN each site belongs to, as well as QoS and SLA guarantees. This large amount of required configuration leads to the possibility of misconfiguration.

It is important for the SP to have operational processes in place to reduce the potential impact of misconfiguration. CE-to-CE authentication may also be used to detect misconfiguration when it occurs. CE-to-CE encryption may also limit the damage when misconfiguration occurs.

5.7. Deployment of Testable MPLS/GMPLS Service

This refers to solutions that can be readily tested to make sure they are configured correctly. For example, for a point-to-point connection, checking that the intended connectivity is working pretty

much ensures that there is no unintended connectivity to some other site.

5.8. Verification of Connectivity

In order to protect against deliberate or accidental misconnection, mechanisms can be put in place to verify both end-to-end connectivity and hop-by-hop resources. These mechanisms can trace the routes of LSPs in both the control plane and the data plane.

It should be noted that if there is an attack on the control plane, data-plane connectivity test mechanisms that rely on the control plane can also be attacked. This may hide faults through false positives or disrupt functioning services through false negatives.

6. Monitoring, Detection, and Reporting of Security Attacks

MPLS/GMPLS network and service may be subject to attacks from a variety of security threats. Many threats are described in Section 4 of this document. Many of the defensive techniques described in this document and elsewhere provide significant levels of protection from a variety of threats. However, in addition to employing defensive techniques silently to protect against attacks, MPLS/GMPLS services can also add value for both providers and customers by implementing security monitoring systems to detect and report on any security attacks, regardless of whether the attacks are effective.

Attackers often begin by probing and analyzing defenses, so systems that can detect and properly report these early stages of attacks can provide significant benefits.

Information concerning attack incidents, especially if available quickly, can be useful in defending against further attacks. It can be used to help identify attackers or their specific targets at an early stage. This knowledge about attackers and targets can be used to strengthen defenses against specific attacks or attackers, or to improve the defenses for specific targets on an as-needed basis. Information collected on attacks may also be useful in identifying and developing defenses against novel attack types.

Monitoring systems used to detect security attacks in MPLS/GMPLS typically operate by collecting information from the Provider Edge (PE), Customer Edge (CE), and/or Provider backbone (P) devices. Security monitoring systems should have the ability to actively retrieve information from devices (e.g., SNMP get) or to passively receive reports from devices (e.g., SNMP notifications). The systems may actively retrieve information from devices (e.g., SNMP get) or passively receive reports from devices (e.g., SNMP notifications).

The specific information exchanged depends on the capabilities of the devices and on the type of VPN technology. Particular care should be given to securing the communications channel between the monitoring systems and the MPLS/GMPLS devices.

The CE, PE, and P devices should employ efficient methods to acquire and communicate the information needed by the security monitoring systems. It is important that the communication method between MPLS/GMPLS devices and security monitoring systems be designed so that it will not disrupt network operations. As an example, multiple attack events may be reported through a single message, rather than allowing each attack event to trigger a separate message, which might result in a flood of messages, essentially becoming a DoS attack against the monitoring system or the network.

The mechanisms for reporting security attacks should be flexible enough to meet the needs of MPLS/GMPLS service providers, MPLS/GMPLS customers, and regulatory agencies, if applicable. The specific reports should depend on the capabilities of the devices, the security monitoring system, the type of VPN, and the service level agreements between the provider and customer.

While SNMP/syslog type monitoring and detection mechanisms can detect some attacks (usually resulting from flapping protocol adjacencies, CPU overload scenarios, etc.), other techniques, such as netflow-based traffic fingerprinting, are needed for more detailed detection and reporting.

With netflow-based traffic fingerprinting, each packet that is forwarded within a device is examined for a set of IP packet attributes. These attributes are the IP packet identity or fingerprint of the packet and determine if the packet is unique or similar to other packets.

The flow information is extremely useful for understanding network behavior, and detecting and reporting security attacks:

- Source address allows the understanding of who is originating the traffic.
- Destination address tells who is receiving the traffic.
- Ports characterize the application utilizing the traffic.
- Class of service examines the priority of the traffic.

- The device interface tells how traffic is being utilized by the network device.
- Tallied packets and bytes show the amount of traffic.
- Flow timestamps allow the understanding of the life of a flow; timestamps are useful for calculating packets and bytes per second.
- Next-hop IP addresses including BGP routing Autonomous Systems (ASes).
- Subnet mask for the source and destination addresses are for calculating prefixes.
- TCP flags are useful for examining TCP handshakes.

7. Service Provider General Security Requirements

This section covers security requirements the provider may have for securing its MPLS/GMPLS network infrastructure including LDP and RSVP-TE-specific requirements.

The MPLS/GMPLS service provider's requirements defined here are for the MPLS/GMPLS core in the reference model. The core network can be implemented with different types of network technologies, and each core network may use different technologies to provide the various services to users with different levels of offered security. Therefore, an MPLS/GMPLS service provider may fulfill any number of the security requirements listed in this section. This document does not state that an MPLS/GMPLS network must fulfill all of these requirements to be secure.

These requirements are focused on: 1) how to protect the MPLS/GMPLS core from various attacks originating outside the core including those from network users, both accidentally and maliciously, and 2) how to protect the end users.

7.1. Protection within the Core Network

7.1.1. Control-Plane Protection - General

- Filtering spoofed infrastructure IP addresses at edges

Many attacks on protocols running in a core involve spoofing a source IP address of a node in the core (e.g., TCP-RST attacks). It makes sense to apply anti-spoofing filtering at edges, e.g., using strict unicast reverse path forwarding (uRPF) [RFC3704] and/or by preventing

the use of infrastructure addresses as source. If this is done comprehensively, the need to cryptographically secure these protocols is smaller. See [BACKBONE-ATTKS] for more elaborate description.

- Protocol authentication within the core

The network infrastructure must support mechanisms for authentication of the control-plane messages. If an MPLS/GMPLS core is used, LDP sessions may be authenticated with TCP MD5. In addition, IGP and BGP authentication should be considered. For a core providing various IP, VPN, or transport services, PE-to-PE authentication may also be performed via IPsec. See the above discussion of protocol security services: authentication, integrity (with replay detection), and confidentiality. Protocols need to provide a complete set of security services from which the SP can choose. Also, the important but often more difficult part is key management. Considerations, guidelines, and strategies regarding key management are discussed in [RFC3562], [RFC4107], [RFC4808].

With today's processors, applying cryptographic authentication to the control plane may not increase the cost of deployment for providers significantly, and will help to improve the security of the core. If the core is dedicated to MPLS/GMPLS enabled services without any interconnects to third parties, then this may reduce the requirement for authentication of the core control plane.

- Infrastructure Hiding

Here we discuss means to hide the provider's infrastructure nodes. An MPLS/GMPLS provider may make its infrastructure routers (P and PE) unreachable from outside users and unauthorized internal users. For example, separate address space may be used for the infrastructure loopbacks.

Normal TTL propagation may be altered to make the backbone look like one hop from the outside, but caution needs to be taken for loop prevention. This prevents the backbone addresses from being exposed through trace route; however, this must also be assessed against operational requirements for end-to-end fault tracing.

An Internet backbone core may be re-engineered to make Internet routing an edge function, for example, by using MPLS label switching for all traffic within the core and possibly making the Internet a VPN within the PPVPN core itself. This helps to detach Internet access from PPVPN services.

Separating control-plane, data-plane, and management-plane functionality in hardware and software may be implemented on the PE

devices to improve security. This may help to limit the problems when attacked in one particular area, and may allow each plane to implement additional security measures separately.

PEs are often more vulnerable to attack than P routers, because PEs cannot be made unreachable from outside users by their very nature. Access to core trunk resources can be controlled on a per-user basis by using of inbound rate limiting or traffic shaping; this can be further enhanced on a per-class-of-service basis (see Section 8.2.3)

In the PE, using separate routing processes for different services, for example, Internet and PPVPN service, may help to improve the PPVPN security and better protect VPN customers. Furthermore, if resources, such as CPU and memory, can be further separated based on applications, or even individual VPNs, it may help to provide improved security and reliability to individual VPN customers.

7.1.2. Control-Plane Protection with RSVP-TE

- General RSVP Security Tools

Isolation of the trusted domain is an important security mechanism for RSVP, to ensure that an untrusted element cannot access a router of the trusted domain. However, ASBR-ASBR communication for inter-AS LSPs needs to be secured specifically. Isolation mechanisms might also be bypassed by an IPv4 Router Alert or IPv6 using Next Header 0 packets. A solution could consist of disabling the processing of IP options. This drops or ignores all IP packets with IPv4 options, including the router alert option used by RSVP; however, this may have an impact on other protocols using IPv4 options. An alternative is to configure access-lists on all incoming interfaces dropping IPv4 protocol or IPv6 next header 46 (RSVP).

RSVP security can be strengthened by deactivating RSVP on interfaces with neighbors who are not authorized to use RSVP, to protect against adjacent CE-PE attacks. However, this does not really protect against DoS attacks or attacks on non-adjacent routers. It has been demonstrated that substantial CPU resources are consumed simply by processing received RSVP packets, even if the RSVP process is deactivated for the specific interface on which the RSVP packets are received.

RSVP neighbor filtering at the protocol level, to restrict the set of neighbors that can send RSVP messages to a given router, protects against non-adjacent attacks. However, this does not protect against DoS attacks and does not effectively protect against spoofing of the source address of RSVP packets, if the filter relies on the neighbor's address within the RSVP message.

RSVP neighbor filtering at the data-plane level, with an access list to accept IP packets with port 46 only for specific neighbors, requires Router Alert mode to be deactivated and does not protect against spoofing.

Another valuable tool is RSVP message pacing, to limit the number of RSVP messages sent to a given neighbor during a given period. This allows blocking DoS attack propagation.

- Another approach is to limit the impact of an attack on control-plane resources.

To ensure continued effective operation of the MPLS router even in the case of an attack that bypasses packet filtering mechanisms such as Access Control Lists in the data plane, it is important that routers have some mechanisms to limit the impact of the attack. There should be a mechanism to rate limit the amount of control-plane traffic addressed to the router, per interface. This should be configurable on a per-protocol basis, (and, ideally, on a per-sender basis) to avoid letting an attacked protocol or a given sender block all communications. This requires the ability to filter and limit the rate of incoming messages of particular protocols, such as RSVP (filtering at the IP protocol level), and particular senders. In addition, there should be a mechanism to limit CPU and memory capacity allocated to RSVP, so as to protect other control-plane elements. To limit memory allocation, it will probably be necessary to limit the number of LSPs that can be set up.

- Authentication for RSVP messages

RSVP message authentication is described in RFC 2747 [RFC2747] and RFC 3097 [RFC3097]. It is one of the most powerful tools for protection against RSVP-based attacks. It applies cryptographic authentication to RSVP messages based on a secure message hash using a key shared by RSVP neighbors. This protects against LSP creation attacks, at the expense of consuming significant CPU resources for digest computation. In addition, if the neighboring RSVP speaker is compromised, it could be used to launch attacks using authenticated RSVP messages. These methods, and certain other aspects of RSVP security, are explained in detail in RFC 4230 [RFC4230]. Key management must be implemented. Logging and auditing as well as multiple layers of cryptographic protection can help here. IPsec can also be used in some cases (see [RFC4230]).

One challenge using RSVP message authentication arises in many cases where non-RSVP nodes are present in the network. In such cases, the RSVP neighbor may not be known up front, thus neighbor-based keying approaches fail, unless the same key is used everywhere, which is not

recommended for security reasons. Group keying may help in such cases. The security properties of various keying approaches are discussed in detail in [RSVP-key].

7.1.3. Control-Plane Protection with LDP

The approaches to protect MPLS routers against LDP-based attacks are similar to those for RSVP, including isolation, protocol deactivation on specific interfaces, filtering of LDP neighbors at the protocol level, filtering of LDP neighbors at the data-plane level (with an access list that filters the TCP and UDP LDP ports), authentication with a message digest, rate limiting of LDP messages per protocol per sender, and limiting all resources allocated to LDP-related tasks. LDP protection could be considered easier in a certain sense. UDP port matching may be sufficient for LDP protection. Router alter options and beyond might be involved in RSVP protection.

7.1.4. Data-Plane Protection

IPsec can provide authentication, integrity, confidentiality, and replay detection for provider or user data. It also has an associated key management protocol.

In today's MPLS/GMPLS, ATM, or Frame Relay networks, encryption is not provided as a basic feature. Mechanisms described in Section 5 can be used to secure the MPLS data-plane traffic carried over an MPLS core. Both the Frame Relay Forum and the ATM Forum standardized cryptographic security services in the late 1990s, but these standards are not widely implemented.

7.2. Protection on the User Access Link

Peer or neighbor protocol authentication may be used to enhance security. For example, BGP MD5 authentication may be used to enhance security on PE-CE links using eBGP. In the case of inter-provider connections, cryptographic protection mechanisms, such as IPsec, may be used between ASes.

If multiple services are provided on the same PE platform, different WAN address spaces may be used for different services (e.g., VPN and non-VPN) to enhance isolation.

Firewall and Filtering: access control mechanisms can be used to filter any packets destined for the service provider's infrastructure prefix or eliminate routes identified as illegitimate. Filtering should also be applied to prevent sourcing packets with infrastructure IP addresses from outside.

Rate limiting may be applied to the user interface/logical interfaces as a defense against DDoS bandwidth attack. This is helpful when the PE device is supporting both multiple services, especially VPN and Internet Services, on the same physical interfaces through different logical interfaces.

7.2.1. Link Authentication

Authentication can be used to validate site access to the network via fixed or logical connections, e.g., L2TP or IPsec, respectively. If the user wishes to hold the authentication credentials for access, then provider solutions require the flexibility for either direct authentication by the PE itself or interaction with a customer authentication server. Mechanisms are required in the latter case to ensure that the interaction between the PE and the customer authentication server is appropriately secured.

7.2.2. Access Routing Control

Choice of routing protocols, e.g., RIP, OSPF, or BGP, may be used to provide control access between a CE and a PE. Per-neighbor and per-VPN routing policies may be established to enhance security and reduce the impact of a malicious or non-malicious attack on the PE; the following mechanisms, in particular, should be considered:

- Limiting the number of prefixes that may be advertised on a per-access basis into the PE. Appropriate action may be taken should a limit be exceeded, e.g., the PE shutting down the peer session to the CE
- Applying route dampening at the PE on received routing updates
- Definition of a per-VPN prefix limit after which additional prefixes will not be added to the VPN routing table.

In the case of inter-provider connection, access protection, link authentication, and routing policies as described above may be applied. Both inbound and outbound firewall or filtering mechanisms between ASes may be applied. Proper security procedures must be implemented in inter-provider interconnection to protect the providers' network infrastructure and their customers. This may be custom designed for each inter-provider peering connection, and must be agreed upon by both providers.

7.2.3. Access QoS

MPLS/GMPLS providers offering QoS-enabled services require mechanisms to ensure that individual accesses are validated against their subscribed QoS profile and as such gain access to core resources that match their service profile. Mechanisms such as per-class-of-service rate limiting or traffic shaping on ingress to the MPLS/GMPLS core are two options for providing this level of control. Such mechanisms may require the per-class-of-service profile to be enforced either by marking, remarking, or discarding of traffic outside of the profile.

7.2.4. Customer Service Monitoring Tools

End users needing specific statistics on the core, e.g., routing table, interface status, or QoS statistics, place requirements on mechanisms at the PE both to validate the incoming user and limit the views available to that particular user. Mechanisms should also be considered to ensure that such access cannot be used as means to construct a DoS attack (either maliciously or accidentally) on the PE itself. This could be accomplished either through separation of these resources within the PE itself or via the capability to rate limiting, which is performed on the basis of each physical interface or each logical connection.

7.3. General User Requirements for MPLS/GMPLS Providers

MPLS/GMPLS providers must support end users' security requirements. Depending on the technologies used, these requirements may include:

- User control plane separation through routing isolation when applicable, for example, in the case of MPLS VPNs.
- Protection against intrusion, DoS attacks, and spoofing
- Access Authentication
- Techniques highlighted throughout this document that identify methodologies for the protection of resources and the MPLS/GMPLS infrastructure.

Hardware or software errors in equipment leading to breaches in security are not within the scope of this document.

8. Inter-Provider Security Requirements

This section discusses security capabilities that are important at the MPLS/GMPLS inter-provider connections and at devices (including ASBR routers) supporting these connections. The security

capabilities stated in this section should be considered as complementary to security considerations addressed in individual protocol specifications or security frameworks.

Security vulnerabilities and exposures may be propagated across multiple networks because of security vulnerabilities arising in one peer's network. Threats to security originate from accidental, administrative, and intentional sources. Intentional threats include events such as spoofing and denial-of-service (DoS) attacks.

The level and nature of threats, as well as security and availability requirements, may vary over time and from network to network. This section, therefore, discusses capabilities that need to be available in equipment deployed for support of the MPLS InterCarrier Interconnect (MPLS-ICI). Whether any particular capability is used in any one specific instance of the ICI is up to the service providers managing the PE equipment offering or using the ICI services.

8.1. Control-Plane Protection

This section discusses capabilities for control-plane protection, including protection of routing, signaling, and OAM capabilities.

8.1.1. Authentication of Signaling Sessions

Authentication may be needed for signaling sessions (i.e., BGP, LDP, and RSVP-TE) and routing sessions (e.g., BGP), as well as OAM sessions across domain boundaries. Equipment must be able to support the exchange of all protocol messages over IPsec ESP, with NULL encryption and authentication, between the peering ASBRs. Support for message authentication for LDP, BGP, and RSVP-TE authentication must also be provided. Manual keying of IPsec should not be used. IKEv2 with pre-shared secrets or public key methods should be used. Replay detection should be used.

Mechanisms to authenticate and validate a dynamic setup request must be available. For instance, if dynamic signaling of a TE-LSP or PW is crossing a domain boundary, there must be a way to detect whether the LSP source is who it claims to be and that it is allowed to connect to the destination.

Message authentication support for all TCP-based protocols within the scope of the MPLS-ICI (i.e., LDP signaling and BGP routing) and Message authentication with the RSVP-TE Integrity Object must be provided to interoperate with current practices. Equipment should be able to support the exchange of all signaling and routing (LDP, RSVP-TE, and BGP) protocol messages over a single IPsec association pair

in tunnel or transport mode with authentication but with NULL encryption, between the peering ASBRs. IPsec, if supported, must be supported with HMAC-SHA-1 and alternatively with HMAC-SHA-2 and optionally SHA-1. It is expected that authentication algorithms will evolve over time and support can be updated as needed.

OAM operations across the MPLS-ICI could also be the source of security threats on the provider infrastructure as well as the service offered over the MPLS-ICI. A large volume of OAM messages could overwhelm the processing capabilities of an ASBR if the ASBR is not properly protected. Maliciously generated OAM messages could also be used to bring down an otherwise healthy service (e.g., MPLS Pseudowire), and therefore affect service security. LSP ping does not support authentication today, and that support should be a subject for future consideration. Bidirectional Forwarding Detection (BFD), however, does have support for carrying an authentication object. It also supports Time-To-Live (TTL) processing as an anti-replay measure. Implementations conformant with this MPLS-ICI should support BFD authentication and must support the procedures for TTL processing.

8.1.2. Protection Against DoS Attacks in the Control Plane

Implementations must have the ability to prevent signaling and routing DoS attacks on the control plane per interface and provider. Such prevention may be provided by rate limiting signaling and routing messages that can be sent by a peer provider according to a traffic profile and by guarding against malformed packets.

Equipment must provide the ability to filter signaling, routing, and OAM packets destined for the device, and must provide the ability to rate limit such packets. Packet filters should be capable of being separately applied per interface, and should have minimal or no performance impact. For example, this allows an operator to filter or rate limit signaling, routing, and OAM messages that can be sent by a peer provider and limit such traffic to a given profile.

During a control-plane DoS attack against an ASBR, the router should guarantee sufficient resources to allow network operators to execute network management commands to take corrective action, such as turning on additional filters or disconnecting an interface under attack. DoS attacks on the control plane should not adversely affect data-plane performance.

Equipment running BGP must support the ability to limit the number of BGP routes received from any particular peer. Furthermore, in the case of IPVPN, a router must be able to limit the number of routes

learned from a BGP peer per IPVPN. In the case that a device has multiple BGP peers, it should be possible for the limit to vary between peers.

8.1.3. Protection against Malformed Packets

Equipment should be robust in the presence of malformed protocol packets. For example, malformed routing, signaling, and OAM packets should be treated in accordance with the relevant protocol specification.

8.1.4. Ability to Enable/Disable Specific Protocols

Equipment must have the ability to drop any signaling or routing protocol messages when these messages are to be processed by the ASBR but the corresponding protocol is not enabled on that interface.

Equipment must allow an administrator to enable or disable a protocol (by default, the protocol is disabled unless administratively enabled) on an interface basis.

Equipment must be able to drop any signaling or routing protocol messages when these messages are to be processed by the ASBR but the corresponding protocol is not enabled on that interface. This dropping should not adversely affect data-plane or control-plane performance.

8.1.5. Protection against Incorrect Cross Connection

The capability to detect and locate faults in an LSP cross-connect must be provided. Such faults may cause security violations as they result in directing traffic to the wrong destinations. This capability may rely on OAM functions. Equipment must support MPLS LSP ping [RFC4379]. This may be used to verify end-to-end connectivity for the LSP (e.g., PW, TE Tunnel, VPN LSP, etc.), and to verify PE-to-PE connectivity for IPVPN services.

When routing information is advertised from one domain to the other, operators must be able to guard against situations that result in traffic hijacking, black-holing, resource stealing (e.g., number of routes), etc. For instance, in the IPVPN case, an operator must be able to block routes based on associated route target attributes. In addition, mechanisms to defend against routing protocol attack must exist to verify whether a route advertised by a peer for a given VPN is actually a valid route and whether the VPN has a site attached to or reachable through that domain.

Equipment (ASBRs and Route Reflectors (RRs)) supporting operation of BGP must be able to restrict which route target attributes are sent to and accepted from a BGP peer across an ICI. Equipment (ASBRs, RR) should also be able to inform the peer regarding which route target attributes it will accept from a peer, because sending an incorrect route target can result in an incorrect cross-connection of VPNs. Also, sending inappropriate route targets to a peer may disclose confidential information. This is another example of defense against routing protocol attacks.

8.1.6. Protection against Spoofed Updates and Route Advertisements

Equipment must support route filtering of routes received via a BGP peer session by applying policies that include one or more of the following: AS path, BGP next hop, standard community, or extended community.

8.1.7. Protection of Confidential Information

The ability to identify and block messages with confidential information from leaving the trusted domain that can reveal confidential information about network operation (e.g., performance OAM messages or LSP ping messages) is required. SPs must have the flexibility to handle these messages at the ASBR.

Equipment should be able to identify and restrict where it sends messages that can reveal confidential information about network operation (e.g., performance OAM messages, LSP Traceroute messages). Service Providers must have the flexibility to handle these messages at the ASBR. For example, equipment supporting LSP Traceroute may limit to which addresses replies can be sent. Note that this capability should be used with care. For example, if an SP chooses to prohibit the exchange of LSP ping messages at the ICI, it may make it more difficult to debug incorrect cross-connection of LSPs or other problems.

An SP may decide to progress these messages if they arrive from a trusted provider and are targeted to specific, agreed-on addresses. Another provider may decide to traffic police, reject, or apply other policies to these messages. Solutions must enable providers to control the information that is relayed to another provider about the path that an LSP takes. For example, when using the RSVP-TE record route object or LSP ping / trace, a provider must be able to control the information contained in corresponding messages when sent to another provider.

8.1.8. Protection against Over-provisioned Number of RSVP-TE LSPs and Bandwidth Reservation

In addition to the control-plane protection mechanisms listed in the previous section on control-plane protection with RSVP-TE, the ASBR must be able both to limit the number of LSPs that can be set up by other domains and to limit the amount of bandwidth that can be reserved. A provider's ASBR may deny an LSP setup request or a bandwidth reservation request sent by another provider's whose limits have been reached.

8.2. Data-Plane Protection

8.2.1. Protection against DoS in the Data Plane

This is described in Section 5 of this document.

8.2.2. Protection against Label Spoofing

Equipment must be able to verify that a label received across an interconnect was actually assigned to an LSP arriving across that interconnect. If a label not assigned to an LSP arrives at this router from the correct neighboring provider, the packet must be dropped. This verification can be applied to the top label only. The top label is the received top label and every label that is exposed by label popping is to be used for forwarding decisions.

Equipment must provide the capability to drop MPLS-labeled packets if all labels in the stack are not processed. This lets SPs guarantee that every label that enters its domain from another carrier is actually assigned to that carrier.

The following requirements are not directly reflected in this document but must be used as guidance for addressing further work.

Solutions must NOT force operators to reveal reachability information to routers within their domains. Note that it is believed that this requirement is met via other requirements specified in this section plus the normal operation of IP routing, which does not reveal individual hosts.

Mechanisms to authenticate and validate a dynamic setup request must be available. For instance, if dynamic signaling of a TE-LSP or PW is crossing a domain boundary, there must be a way to detect whether the LSP source is who it claims to be and that it is allowed to connect to the destination.

8.2.3. Protection Using Ingress Traffic Policing and Enforcement

The following simple diagram illustrates a potential security issue on the data plane across an MPLS interconnect:

```

SP2 - ASBR2 - labeled path - ASBR1 - P1 - SP1's PSN - P2 - PE1
|         |         |         |         |         |         |
|<  AS2  >|<MPLS interconnect>|<           AS1           >|

```

Traffic flow direction is from SP2 to SP1

In the case of downstream label assignment, the transit label used by ASBR2 is allocated by ASBR1, which in turn advertises it to ASBR2 (downstream unsolicited or on-demand); this label is used for a service context (VPN label, PW VC label, etc.), and this LSP is normally terminated at a forwarding table belonging to the service instance on PE (PE1) in SP1.

In the example above, ASBR1 would not know whether the label of an incoming packet from ASBR2 over the interconnect is a VPN label or PSN label for AS1. So it is possible (though unlikely) that ASBR2 can be accidentally or intentionally configured such that the incoming label could match a PSN label (e.g., LDP) in AS1. Then, this LSP would end up on the global plane of an infrastructure router (P or PE1), and this could invite a unidirectional attack on that P or PE1 where the LSP terminates.

To mitigate this threat, implementations should be able to do a forwarding path look-up for the label on an incoming packet from an interconnect in a Label Forwarding Information Base (LFIB) space that is only intended for its own service context or provide a mechanism on the data plane that would ensure the incoming labels are what ASBR1 has allocated and advertised.

A similar concept has been proposed in "Requirements for Multi-Segment Pseudowire Emulation Edge-to-Edge (PWE3)" [RFC5254].

When using upstream label assignment, the upstream source must be identified and authenticated so the labels can be accepted as from a trusted source.

9. Summary of MPLS and GMPLS Security

The following summary provides a quick checklist of MPLS and GMPLS security threats, defense techniques, and the best-practice outlines for MPLS and GMPLS deployment.

9.1. MPLS and GMPLS Specific Security Threats

9.1.1. Control-Plane Attacks

Types of attacks on the control plane:

- Unauthorized LSP creation
- LSP message interception

Attacks against RSVP-TE: DoS attacks that set up unauthorized LSP and/or LSP messages.

Attacks against LDP: DoS attack with storms of LDP Hello messages or LDP TCP SYN messages.

Attacks may be launched from external or internal sources, or through an SP's management systems.

Attacks may be targeted at the SP's routing protocols or infrastructure elements.

In general, control protocols may be attacked by:

- MPLS signaling (LDP, RSVP-TE)
- PCE signaling
- IPsec signaling (IKE and IKEv2)
- ICMP and ICMPv6
- L2TP
- BGP-based membership discovery
- Database-based membership discovery (e.g., RADIUS)
- OAM and diagnostic protocols such as LSP ping and LMP
- Other protocols that may be important to the control infrastructure, e.g., DNS, LMP, NTP, SNMP, and GRE

9.1.2. Data-Plane Attacks

- Unauthorized observation of data traffic
- Data-traffic modification
- Spoofing and replay
- Unauthorized deletion
- Unauthorized traffic-pattern analysis
- Denial of Service

9.2. Defense Techniques

1) Authentication:

- Bidirectional authentication
- Key management
- Management system authentication
- Peer-to-peer authentication

2) Cryptographic techniques

3) Use of IPsec in MPLS/GMPLS networks

4) Encryption for device configuration and management

5) Cryptographic techniques for MPLS pseudowires

6) End-to-End versus Hop-by-Hop protection (CE-CE, PE-PE, PE-CE)

7) Access control techniques

- Filtering
- Firewalls
- Access Control to management interfaces

8) Infrastructure isolation

9) Use of aggregated infrastructure

- 10) Quality control processes
- 11) Testable MPLS/GMPLS service
- 12) End-to-end connectivity verification
- 13) Hop-by-hop resource configuration verification and discovery

9.3. Service Provider MPLS and GMPLS Best-Practice Outlines

9.3.1. SP Infrastructure Protection

- 1) General control-plane protection
 - Filtering out infrastructure source addresses at edges
 - Protocol authentication within the core
 - Infrastructure hiding (e.g., disable TTL propagation)
- 2) RSVP control-plane protection
 - RSVP security tools
 - Isolation of the trusted domain
 - Deactivating RSVP on interfaces with neighbors who are not authorized to use RSVP
 - RSVP neighbor filtering at the protocol level and data-plane level
 - Authentication for RSVP messages
 - RSVP message pacing
- 3) LDP control-plane protection (similar techniques as for RSVP)
- 4) Data-plane protection
 - User access link protection
 - Link authentication
 - Access routing control (e.g., prefix limits, route dampening, routing table limits (such as VRF limits))
 - Access QoS control

- Customer service monitoring tools
- Use of LSP ping (with its own control-plane security) to verify end-to-end connectivity of MPLS LSPs
- LMP (with its own security) to verify hop-by-hop connectivity.

9.3.2. Inter-Provider Security

Inter-provider connections are high security risk areas. Similar techniques and procedures as described for SP's general core protection are listed below for inter-provider connections.

1) Control-plane protection at inter-provider connections

- Authentication of signaling sessions
- Protection against DoS attacks in the control plane
- Protection against malformed packets
- Ability to enable/disable specific protocols
- Protection against incorrect cross connection
- Protection against spoofed updates and route advertisements
- Protection of confidential information
- Protection against an over-provisioned number of RSVP-TE LSPs and bandwidth reservation

2) Data-plane protection at the inter-provider connections

- Protection against DoS in the data plane
- Protection against label spoofing

For MPLS VPN interconnections [RFC4364], in practice, inter-AS option a), VRF-to-VRF connections at the AS (Autonomous System) border, is commonly used for inter-provider connections. Option c), Multi-hop EBGP redistribution of labeled VPN-IPv4 routes between source and destination ASes with EBGP redistribution of labeled IPv4 routes from AS to a neighboring AS, on the other hand, is not normally used for inter-provider connections due to higher security risks. For more details, please see [RFC4111].

10. Security Considerations

Security considerations constitute the sole subject of this memo and hence are discussed throughout. Here we recap what has been presented and explain at a high level the role of each type of consideration in an overall secure MPLS/GMPLS system.

The document describes a number of potential security threats. Some of these threats have already been observed occurring in running networks; others are largely hypothetical at this time.

DoS attacks and intrusion attacks from the Internet against an SPs' infrastructure have been seen. DoS "attacks" (typically not malicious) have also been seen in which CE equipment overwhelms PE equipment with high quantities or rates of packet traffic or routing information. Operational or provisioning errors are cited by SPs as one of their prime concerns.

The document describes a variety of defensive techniques that may be used to counter the suspected threats. All of the techniques presented involve mature and widely implemented technologies that are practical to implement.

The document describes the importance of detecting, monitoring, and reporting attacks, both successful and unsuccessful. These activities are essential for "understanding one's enemy", mobilizing new defenses, and obtaining metrics about how secure the MPLS/GMPLS network is. As such, they are vital components of any complete PPVPN security system.

The document evaluates MPLS/GMPLS security requirements from a customer's perspective as well as from a service provider's perspective. These sections re-evaluate the identified threats from the perspectives of the various stakeholders and are meant to assist equipment vendors and service providers, who must ultimately decide what threats to protect against in any given configuration or service offering.

11. References

11.1. Normative References

- [RFC2747] Baker, F., Lindell, B., and M. Talwar, "RSVP Cryptographic Authentication", RFC 2747, January 2000.

- [RFC3031] Rosen, E., Viswanathan, A., and R. Callon, "Multiprotocol Label Switching Architecture", RFC 3031, January 2001.
- [RFC3097] Braden, R. and L. Zhang, "RSVP Cryptographic Authentication -- Updated Message Type Value", RFC 3097, April 2001.
- [RFC3209] Awduche, D., Berger, L., Gan, D., Li, T., Srinivasan, V., and G. Swallow, "RSVP-TE: Extensions to RSVP for LSP Tunnels", RFC 3209, December 2001.
- [RFC3945] Mannie, E., Ed., "Generalized Multi-Protocol Label Switching (GMPLS) Architecture", RFC 3945, October 2004.
- [RFC4106] Viega, J. and D. McGrew, "The Use of Galois/Counter Mode (GCM) in IPsec Encapsulating Security Payload (ESP)", RFC 4106, June 2005.
- [RFC4301] Kent, S. and K. Seo, "Security Architecture for the Internet Protocol", RFC 4301, December 2005.
- [RFC4302] Kent, S., "IP Authentication Header", RFC 4302, December 2005.
- [RFC4306] Kaufman, C., Ed., "Internet Key Exchange (IKEv2) Protocol", RFC 4306, December 2005.
- [RFC4309] Housley, R., "Using Advanced Encryption Standard (AES) CCM Mode with IPsec Encapsulating Security Payload (ESP)", RFC 4309, December 2005.
- [RFC4364] Rosen, E. and Y. Rekhter, "BGP/MPLS IP Virtual Private Networks (VPNs)", RFC 4364, February 2006.
- [RFC4379] Kompella, K. and G. Swallow, "Detecting Multi-Protocol Label Switched (MPLS) Data Plane Failures", RFC 4379, February 2006.
- [RFC4447] Martini, L., Ed., Rosen, E., El-Aawar, N., Smith, T., and G. Heron, "Pseudowire Setup and Maintenance Using the Label Distribution Protocol (LDP)", RFC 4447, April 2006.

- [RFC4835] Manral, V., "Cryptographic Algorithm Implementation Requirements for Encapsulating Security Payload (ESP) and Authentication Header (AH)", RFC 4835, April 2007.
- [RFC5246] Dierks, T. and E. Rescorla, "The Transport Layer Security (TLS) Protocol Version 1.2", RFC 5246, August 2008.
- [RFC5036] Andersson, L., Ed., Minei, I., Ed., and B. Thomas, Ed., "LDP Specification", RFC 5036, October 2007.
- [STD62] Harrington, D., Presuhn, R., and B. Wijnen, "An Architecture for Describing Simple Network Management Protocol (SNMP) Management Frameworks", STD 62, RFC 3411, December 2002.
- Case, J., Harrington, D., Presuhn, R., and B. Wijnen, "Message Processing and Dispatching for the Simple Network Management Protocol (SNMP)", STD 62, RFC 3412, December 2002.
- Levi, D., Meyer, P., and B. Stewart, "Simple Network Management Protocol (SNMP) Applications", STD 62, RFC 3413, December 2002.
- Blumenthal, U. and B. Wijnen, "User-based Security Model (USM) for version 3 of the Simple Network Management Protocol (SNMPv3)", STD 62, RFC 3414, December 2002.
- Wijnen, B., Presuhn, R., and K. McCloghrie, "View-based Access Control Model (VACM) for the Simple Network Management Protocol (SNMP)", STD 62, RFC 3415, December 2002.
- Presuhn, R., Ed., "Version 2 of the Protocol Operations for the Simple Network Management Protocol (SNMP)", STD 62, RFC 3416, December 2002.
- Presuhn, R., Ed., "Transport Mappings for the Simple Network Management Protocol (SNMP)", STD 62, RFC 3417, December 2002.
- Presuhn, R., Ed., "Management Information Base (MIB) for the Simple Network Management Protocol (SNMP)", STD 62, RFC 3418, December 2002.

[STD8] Postel, J. and J. Reynolds, "Telnet Protocol Specification", STD 8, RFC 854, May 1983.

Postel, J. and J. Reynolds, "Telnet Option Specifications", STD 8, RFC 855, May 1983.

11.2. Informative References

[OIF-SMI-01.0] Renee Esposito, "Security for Management Interfaces to Network Elements", Optical Internetworking Forum, Sept. 2003.

[OIF-SMI-02.1] Renee Esposito, "Addendum to the Security for Management Interfaces to Network Elements", Optical Internetworking Forum, March 2006.

[RFC2104] Krawczyk, H., Bellare, M., and R. Canetti, "HMAC: Keyed-Hashing for Message Authentication", RFC 2104, February 1997.

[RFC2411] Thayer, R., Doraswamy, N., and R. Glenn, "IP Security Document Roadmap", RFC 2411, November 1998.

[RFC3174] Eastlake 3rd, D. and P. Jones, "US Secure Hash Algorithm 1 (SHA1)", RFC 3174, September 2001.

[RFC3562] Leech, M., "Key Management Considerations for the TCP MD5 Signature Option", RFC 3562, July 2003.

[RFC3631] Bellovin, S., Ed., Schiller, J., Ed., and C. Kaufman, Ed., "Security Mechanisms for the Internet", RFC 3631, December 2003.

[RFC3704] Baker, F. and P. Savola, "Ingress Filtering for Multihomed Networks", BCP 84, RFC 3704, March 2004.

[RFC3985] Bryant, S., Ed., and P. Pate, Ed., "Pseudo Wire Emulation Edge-to-Edge (PWE3) Architecture", RFC 3985, March 2005.

[RFC4107] Bellovin, S. and R. Housley, "Guidelines for Cryptographic Key Management", BCP 107, RFC 4107, June 2005.

[RFC4110] Callon, R. and M. Suzuki, "A Framework for Layer 3 Provider-Provisioned Virtual Private Networks (PPVPNs)", RFC 4110, July 2005.

- [RFC4111] Fang, L., Ed., "Security Framework for Provider-Provisioned Virtual Private Networks (PPVPNs)", RFC 4111, July 2005.
- [RFC4230] Tschofenig, H. and R. Graveman, "RSVP Security Properties", RFC 4230, December 2005.
- [RFC4308] Hoffman, P., "Cryptographic Suites for IPsec", RFC 4308, December 2005.
- [RFC4377] Nadeau, T., Morrow, M., Swallow, G., Allan, D., and S. Matsushima, "Operations and Management (OAM) Requirements for Multi-Protocol Label Switched (MPLS) Networks", RFC 4377, February 2006.
- [RFC4378] Allan, D., Ed., and T. Nadeau, Ed., "A Framework for Multi-Protocol Label Switching (MPLS) Operations and Management (OAM)", RFC 4378, February 2006.
- [RFC4593] Barbir, A., Murphy, S., and Y. Yang, "Generic Threats to Routing Protocols", RFC 4593, October 2006.
- [RFC4778] Kaeo, M., "Operational Security Current Practices in Internet Service Provider Environments", RFC 4778, January 2007.
- [RFC4808] Bellovin, S., "Key Change Strategies for TCP-MD5", RFC 4808, March 2007.
- [RFC4864] Van de Velde, G., Hain, T., Droms, R., Carpenter, B., and E. Klein, "Local Network Protection for IPv6", RFC 4864, May 2007.
- [RFC4869] Law, L. and J. Solinas, "Suite B Cryptographic Suites for IPsec", RFC 4869, May 2007.
- [RFC5254] Bitar, N., Ed., Bocci, M., Ed., and L. Martini, Ed., "Requirements for Multi-Segment Pseudowire Emulation Edge-to-Edge (PWE3)", RFC 5254, October 2008.
- [MFA-MPLS-ICI] N. Bitar, "MPLS InterCarrier Interconnect Technical Specification," IP/MPLS Forum 19.0.0, April 2008.

- [OIF-Sec-Mag] R. Esposito, R. Graveman, and B. Hazzard, "Security for Management Interfaces to Network Elements," OIF-SMI-01.0, September 2003.
- [BACKBONE-ATTKS] Savola, P., "Backbone Infrastructure Attacks and Protections", Work in Progress, January 2007.
- [OPSEC-FILTER] Morrow, C., Jones, G., and V. Manral, "Filtering and Rate Limiting Capabilities for IP Network Infrastructure", Work in Progress, July 2007.
- [IPSECME-ROADMAP] Frankel, S. and S. Krishnan, "IP Security (IPsec) and Internet Key Exchange (IKE) Document Roadmap", Work in Progress, May 2010.
- [OPSEC-EFFORTS] Lonvick, C. and D. Spak, "Security Best Practices Efforts and Documents", Work in Progress, May 2010.
- [RSVP-key] Behringer, M. and F. Le Faucheur, "Applicability of Keying Methods for RSVP Security", Work in Progress, June 2009.

12. Acknowledgements

The authors and contributors would also like to acknowledge the helpful comments and suggestions from Sam Hartman, Dimitri Papadimitriou, Kannan Varadhan, Stephen Farrell, Mircea Pisica, Scott Brim in particular for his comments and discussion through GEN-ART review, as well as Suresh Krishnan for his GEN-ART review and comments. The authors would like to thank Sandra Murphy and Tim Polk for their comments and help through Security AD review, thank Pekka Savola for his comments through ops-dir review, and Amanda Baber for her IANA review.

This document has used relevant content from RFC 4111 "Security Framework of Provider Provisioned VPN for Provider-Provisioned Virtual Private Networks (PPVPNs)" [RFC4111]. We acknowledge the authors of RFC 4111 for the valuable information and text.

Authors:

Luyuan Fang, Ed., Cisco Systems, Inc.
Michael Behringer, Cisco Systems, Inc.
Ross Callon, Juniper Networks
Richard Graveman, RFG Security, LLC
J. L. Le Roux, France Telecom
Raymond Zhang, British Telecom
Paul Knight, Individual Contributor

Yaakov Stein, RAD Data Communications
Nabil Bitar, Verizon
Monique Morrow, Cisco Systems, Inc.
Adrian Farrel, Old Dog Consulting

As a design team member for the MPLS Security Framework, Jerry Ash also made significant contributions to this document.

13. Contributors' Contact Information

Michael Behringer
Cisco Systems, Inc.
Village d'Entreprises Green Side
400, Avenue Roumanille, Batiment T 3
06410 Biot, Sophia Antipolis
FRANCE
EMail: mbehring@cisco.com

Ross Callon
Juniper Networks
10 Technology Park Drive
Westford, MA 01886
USA
EMail: rcallon@juniper.net

Richard Graveman
RFG Security
15 Park Avenue
Morristown, NJ 07960
EMail: rfg@acm.org

Jean-Louis Le Roux
France Telecom
2, avenue Pierre-Marzin
22307 Lannion Cedex
FRANCE
EMail: jeanlouis.leroux@francetelecom.com

Raymond Zhang
British Telecom
BT Center
81 Newgate Street
London, EC1A 7AJ
United Kingdom
EMail: raymond.zhang@bt.com

Paul Knight
39 N. Hancock St.
Lexington, MA 02420
EMail: paul.the.knight@gmail.com

Yaakov (Jonathan) Stein
RAD Data Communications
24 Raoul Wallenberg St., Bldg C
Tel Aviv 69719
ISRAEL
EMail: yaakov_s@rad.com

Nabil Bitar
Verizon
40 Sylvan Road
Waltham, MA 02145
EMail: nabil.bitar@verizon.com

Monique Morrow
Glatt-com
CH-8301 Glattzentrum
Switzerland
EMail: mmorrow@cisco.com

Adrian Farrel
Old Dog Consulting
EMail: adrian@olddog.co.uk

Editor's Address

Luyuan Fang (editor)
Cisco Systems, Inc.
300 Beaver Brook Road
Boxborough, MA 01719
USA
EMail: lufang@cisco.com

