Internet Engineering Task Force (IETF)
Request for Comments: 6952
Category: Informational
ISSN: 2070-1721

M. Jethanandani
Ciena Corporation
K. Patel
Cisco Systems, Inc
L. Zheng
Huawei Technologies

May 2013

Analysis of BGP, LDP, PCEP, and MSDP Issues According to the Keying and Authentication for Routing Protocols (KARP) Design Guide

Abstract

This document analyzes TCP-based routing protocols, the Border Gateway Protocol (BGP), the Label Distribution Protocol (LDP), the Path Computation Element Communication Protocol (PCEP), and the Multicast Source Distribution Protocol (MSDP), according to guidelines set forth in Section 4.2 of "Keying and Authentication for Routing Protocols Design Guidelines", RFC 6518.

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/rfc6952.

Table of Contents

1. Introduction 3
1.1. Abbreviations 4
2. Current Assessment of BGP, LDP, PCEP, and MSDP 5
2.1. Transport Layer 5
2.2. Keying Mechanisms 6
2.3. BGP 7
2.4. LDP 7
2.4.1. Spoofing Attacks 7
2.4.2. Denial-of-Service Attacks 8
2.5. PCEP 8
2.6. MSDP 10
3. Optimal State for BGP, LDP, PCEP, and MSDP 10
3.1. LDP 10
4. Gap Analysis for BGP, LDP, PCEP, and MSDP 11
4.1. LDP 12
4.2. PCEP 13
5. Transition and Deployment Considerations 13
6. Security Considerations 13
7. Acknowledgements 14
8. References 14
8.1. Normative References 14
8.2. Informative References 14

1. Introduction

In their "Report from the IAB Workshop on Unwanted Traffic March 9-10, 2006" [RFC4948], the Internet Architecture Board (IAB) described an attack on core routing infrastructure as an ideal attack that would inflict the greatest amount of damage and suggested steps to tighten the infrastructure against the attack. Four main steps were identified for that tightening:

1. Create secure mechanisms and practices for operating routers.

2. Clean up the Internet Routing Registry (IRR) repository, and secure both the database and the access, so that it can be used for routing verifications.

3. Create specifications for cryptographic validation of routing message content.

4. Secure the routing protocols' packets on the wire.

In order to secure the routing protocols, this document performs an initial analysis of the current state of four TCP-based protocols -- BGP [RFC4271], LDP [RFC5036], PCEP [RFC5440], and MSDP [RFC3618] -- according to the requirements of the KARP Design Guidelines [RFC6518]. Section 4.2 of that document uses the term "state", which will be referred to as the "state of the security method". Thus, a term like "Define Optimal State" would be referred to as "Define Optimal State of the Security Method".

This document builds on several previous efforts into routing security:

- "Issues with Existing Cryptographic Protection Methods for Routing Protocols" [RFC6039], describes issues with existing cryptographic protection methods for routing protocols.

- Analysis of OSPF Security According to the KARP Design Guide [RFC6863] analyzes Open Shortest Path First (OSPF) security according to the KARP Design Guide.

Section 2 of this document looks at the current state of the security method for the four routing protocols: BGP, LDP, PCEP, and MSDP. Section 3 examines what the optimal state of the security method would be for the four routing protocols according to the KARP Design Guidelines [RFC6518], and Section 4 does an analysis of the gap between the existing state of the security method and the optimal state of the security method for the protocols and suggests some areas where improvement is needed.

1.1. Abbreviations

AES - Advanced Encryption Standard

AO - Authentication Option

AS - Autonomous System

BGP - Border Gateway Protocol

CMAC - Cipher-Based Message Authentication Code

DoS - Denial of Service

GTSM - Generalized Time-to-Live (TTL) Security Mechanism

HMAC - Hash-Based MAC

KARP - Key and Authentication for Routing Protocols

KDF - Key Derivation Function

KEK - Key Encrypting Key

KMP - Key Management Protocol

LDP - Label Distribution Protocol

LSR - Label Switching Routers

MAC - Message Authentication Code

MKT - Master Key Table

MSDP - Multicast Source Distribution Protocol

MD5 - Message Digest Algorithm 5

OSPF - Open Shortest Path First

PCEP - Path Computation Element Communication Protocol

PCC - Path Computation Client

PCE - Path Computation Element

SHA - Secure Hash Algorithm

TCP - Transmission Control Protocol

TTL - Time-to-Live

UDP - User Datagram Protocol

WG - Working Group

2. Current Assessment of BGP, LDP, PCEP, and MSDP

This section assesses the transport protocols for any authentication or integrity mechanisms used by the protocol. It describes the current security mechanisms, if any, used by BGP, LDP, PCEP, and MSDP.

2.1. Transport Layer

At the transport layer, routing protocols are subject to a variety of DoS attacks, as outlined in "Internet Denial-of-Service Considerations" [RFC4732]. Such attacks can cause the routing protocol to become congested, resulting in the routing updates being supplied too slowly to be useful. In extreme cases, these attacks prevent routers from converging after a change.

Routing protocols use several methods to protect themselves. Those that use TCP as a transport protocol use access lists to accept packets only from known sources. These access lists also help protect edge routers from attacks originating outside the protected domain. In addition, for edge routers running the External Border Gateway Protocol (eBGP), TCP LISTEN is run only on interfaces on which its peers have been discovered or via which routing sessions are expected (as specified in router configuration databases).

"Generalized TTL Security Mechanism (GTSM)" [RFC5082] describes a generalized Time-to-Live (TTL) security mechanism to protect a protocol stack from CPU-utilization-based attacks. TCP Robustness [RFC5961] recommends some TCP-level mitigations against spoofing attacks targeted towards long-lived routing protocol sessions.

Even when BGP, LDP, PCEP, and MSDP sessions use access lists, they are vulnerable to spoofing and man-in-the-middle attacks. Authentication and integrity checks allow the receiver of a routing protocol update to know that the message genuinely comes from the node that claims to have sent it and to know whether the message has been modified. Sometimes routers can be subjected to a large number of authentication and integrity requests, exhausting connection resources on the router in a way that could lead to the denial of genuine requests.

TCP MD5 [RFC2385] has been obsoleted by TCP-AO [RFC5925]. However, it is still widely used to authenticate TCP-based routing protocols such as BGP. It provides a way for carrying a MD5 digest in a TCP segment. This digest is computed using information known only to the endpoints, and this ensures authenticity and integrity of messages. The MD5 key used to compute the digest is stored locally on the router. This option is used by routing protocols to provide for session-level protection against the introduction of spoofed TCP segments into any existing TCP streams, in particular, TCP Reset segments. TCP MD5 does not provide a generic mechanism to support key rollover. It also does not support algorithm agility.

The Message Authentication Codes (MACs) used by TCP MD5 are considered too weak both because of the use of the hash function and because of the way the secret key used by TCP MD5 is managed. Furthermore, TCP MD5 does not support any algorithm agility. TCP-AO [RFC5925] and its companion document Cryptographic Algorithms for TCP-AO [RFC5926], describe steps towards correcting both the MAC weakness and the management of secret keys. Those steps require that two MAC algorithms be supported. They are HMAC-SHA-1-96, as specified in HMAC [RFC2104], and AES-128-CMAC-96, as specified in NIST-SP800-38B [NIST-SP800-38B]. Cryptographic research suggests that both these MAC algorithms are fairly secure. By supporting multiple MAC algorithms, TCP-AO supports algorithm agility. TCP-AO also allows additional MACs to be added in the future.

2.2. Keying Mechanisms

For TCP-AO [RFC5925], there is no Key Management Protocol (KMP) used to manage the keys that are employed to generate the MAC. TCP-AO talks about coordinating keys derived from the Master Key Table (MKT) between endpoints and allows for a master key to be configured manually or for it to be managed via an out-of-band mechanism.

It should be noted that most routers configured with static keys have not seen the key changed ever. The common reason given for not changing the key is the difficulty in coordinating the change between pairs of routers when using TCP MD5. It is well known that the longer the same key is used, the greater the chance that it can be guessed or exposed, e.g., when an administrator with knowledge of the keys leaves the company.

For point-to-point key management, the IKEv2 protocol [RFC5996] provides for automated key exchange under a Security Association (SA) and can be used for a comprehensive KMP solution for routers. IKEv2 can be used for both IPsec SAs [RFC4301] and other types of SAs. For example, Fibre Channel SAs [RFC4595] are currently negotiated with IKEv2. Using IKEv2 to negotiate TCP-AO is a possible option.

2.3. BGP

All BGP communications take place over TCP. Therefore, all security vulnerabilities for BGP can be categorized as relating to the security of the transport protocol itself, or to the compromising of individual routers and the data they handle. This document examines the issues for the transport protocol, while the SIDR Working Group (WG) looks at ways to sign and secure the data exchanged in BGP as described in "An Infrastructure to Support Secure Internet Protocol" [RFC6480].

2.4. LDP

"Security Framework for MPLS and GMPLS Networks" [RFC5920] outlines security aspects that are relevant in the context of MPLS and GMPLS. It describes the security threats, the related defensive techniques, and the mechanism for detection and reporting.

Section 5 of LDP [RFC5036] states that LDP is subject to two different types of attacks: spoofing and denial-of-service attacks.

2.4.1. Spoofing Attacks

A spoofing attack against LDP can occur both during the discovery phase and during the session communication phase.

2.4.1.1. Discovery Exchanges using UDP

Label Switching Routers (LSRs) indicate their willingness to establish and maintain LDP sessions by periodically sending Hello messages. Reception of a Hello message serves to create a new "Hello adjacency", if one does not already exist, or to refresh an existing one.

There are two variants of the discovery mechanism. A Basic Discovery mechanism is used to discover LSR neighbors that are directly connected at the link level, and an Extended Discovery mechanism is used by LSRs that are more than one hop away.

Unlike all other LDP messages, the Hello messages are sent using UDP. This means that they cannot benefit from the security mechanisms available with TCP. LDP [RFC5036] does not provide any security mechanisms for use with Hello messages except for some configuration that may help protect against bogus discovery events. These configurations include directly connected links and interfaces. Routers that do not use directly connected links have to use the Extended Discovery mechanism and will not be able to use configuration to protect against bogus discovery events.

Spoofing a Hello packet for an existing adjacency can cause the adjacency to time out and result in termination of the associated session. This can occur when the spoofed Hello message specifies a small Hold Time, causing the receiver to expect Hello messages within this interval, while the true neighbor continues sending Hello messages at the lower, previously agreed to frequency.

Spoofing a Hello packet can also cause the LDP session to be terminated. This can occur when the spoofed Hello specifies a different Transport Address from the previously agreed one between neighbors. Spoofed Hello messages are observed and reported as a real problem in production networks.

2.4.1.2. Session Communication using TCP

LDP, like other TCP-based routing protocols, specifies use of the TCP MD5 Signature Option to provide for the authenticity and integrity of session messages. As stated in Section 2.1 of this document and in Section 2.9 of LDP [RFC5036], MD5 authentication is considered too weak for this application as outlined in MD5 and HMAC-MD5 Security Considerations [RFC6151]. It also does not support algorithm agility. A stronger hashing algorithm, e.g., SHA1, which is supported by TCP-AO [RFC5925], could be deployed to take care of the weakness.

Alternatively, one could move to using TCP-AO, which provides for stronger MAC algorithms, makes it easier to set up manual keys, and protects against replay attacks.

2.4.2. Denial-of-Service Attacks

LDP is subject to Denial-of-Service (DoS) attacks both in discovery mode and session mode. The potential targets are documented in Section 5.3 of LDP [RFC5036].

2.5. PCEP

For effective selection by Path Computation Clients (PCCs), a PCC needs to know the location of Path Computation Elements (PCEs) in its domain along with some information relevant for PCE selection. Such PCE information could be learned through manual configuration, on each PCC, along with the capabilities of the PCE or automatically through a PCE discovery mechanism as outlined in Requirements for PCE Discovery [RFC4674].

Attacks on PCEP [RFC5440] may result in damage to active networks. These include computation responses, which if changed can cause protocols like RSVP-TE [RFC3209] to set up suboptimal or

inappropriate LSPs. In addition, PCE itself can be a target for a variety of DoS attacks. Such attacks can cause path computations to be supplied too slowly to be of any value, particularly as it relates to recovery or establishment of LSPs.

Finally, PCE discovery, as outlined in OSPF Protocol Extensions for PCE Discovery [RFC5088] and IS-IS Protocol Extensions for PCE Discovery [RFC5089], is a significant feature for the successful deployment of PCEP in large networks. These mechanisms allow PCC to discover the existence of PCEs within the network. If the discovery mechanism is compromised, it will impair the ability of the nodes to function as described below.

As RFC 5440 states, PCEP (which makes use of TCP as a transport) could be the target of the following attacks:

- Spoofing (PCC or PCE implementation)
- Snooping (message interception)
- Falsification
- Denial of Service

In inter-Autonomous System (inter-AS) scenarios where PCE-to-PCE communication is required, attacks may be particularly significant with commercial implications as well as service-level agreement implications.

Additionally, snooping of PCEP requests and responses may give an attacker information about the operation of the network. By viewing the PCEP messages, an attacker can determine the pattern of service establishment in the network and can know where traffic is being routed, thereby making the network susceptible to targeted attacks and the data within specific LSPs vulnerable.

Ensuring PCEP communication privacy is of key importance, especially in an inter-AS context, where PCEP communication endpoints do not reside in the same AS. An attacker that intercepts a PCE message could obtain sensitive information related to computed paths and resources.

At the time PCEP was documented in [RFC5440], TCP-AO had not been fully specified. Therefore, [RFC5440] mandates that PCEP implementations include support for TCP MD5 and that use of the function should be configurable by the operator. [RFC5440] also describes the vulnerabilities and weaknesses of TCP MD5 as noted in this document. [RFC5440] goes on to state that PCEP implementations

should include support for TCP-AO as soon as that specification is complete. Since TCP-AO [RFC5925] has now been published, new PCEP implementations should fully support TCP-AO.

2.6. MSDP

Similar to BGP and LDP, the Multicast Source Distribution Protocol (MSDP) uses TCP MD5 [RFC2385] to protect TCP sessions via the TCP MD5 option. But with a weak MD5 authentication, TCP MD5 is not considered strong enough for this application. It also does not support algorithm agility.

MSDP advocates imposing a limit on the number of source address and group addresses (S,G) that can be cached within the protocol in order to mitigate state explosion due to any denial of service and other attacks.

3. Optimal State for BGP, LDP, PCEP, and MSDP

The ideal state of the security method for BGP, LDP, PCEP, and MSDP protocols is when they can withstand any of the known types of attacks. The protocols also need to support algorithm agility, i.e., they must not hardwire themselves to one algorithm.

Additionally, the KMP for the routing sessions should help negotiate unique, pair-wise random keys without administrator involvement. It should also negotiate Security Association (SA) parameters required for the session connection, including key lifetimes. It should keep track of those lifetimes and negotiate new keys and parameters before they expire and do so without administrator involvement. In the event of a breach, including when an administrator with knowledge of the keys leaves the company, the keys should be changed immediately.

The DoS attacks for BGP, LDP, PCEP, and MSDP are attacks to the transport protocol -- TCP for the most part, and UDP in case of the discovery phase of LDP. TCP and UDP should be able to withstand any of the DoS scenarios by dropping packets that are attack packets in a way that does not impact legitimate packets.

The routing protocols should provide a mechanism to authenticate the routing information carried within the payload, and administrators should enable it.

3.1. LDP

To mitigate LDP’s current vulnerability to spoofing attacks, LDP needs to be upgraded such that an implementation is able to determine the authenticity of the neighbors sending the Hello message.

Labels are similar to routing information, which is distributed in the clear. However, there is currently no requirement that the labels be encrypted. Such a requirement is out of scope for this document.

Similarly, it is important to ensure that routers exchanging labels are mutually authenticated, and that there are no rogue peers or unauthenticated peers that can compromise the stability of the network.

4. Gap Analysis for BGP, LDP, PCEP, and MSDP

This section outlines the differences between the current state of the security methods for routing protocols and the desired state of the security methods as outlined in Section 4.2 of the KARP Design Guidelines [RFC6518]. As that document states, these routing protocols fall into the category of one-to-one peering messages and will use peer keying protocols. This section covers issues that are common to the four protocols, leaving protocol-specific issues to sub-sections.

At a transport level, these routing protocols are subject to some of the same attacks that TCP applications are subject to. These include DoS and spoofing attacks. "Internet Denial-of-Service Considerations" [RFC4732] outlines some solutions. "Defending TCP Against Spoofing Attacks" [RFC4953] recommends ways to prevent spoofing attacks. In addition, the recommendations in [RFC5961] should also be followed and implemented to strengthen TCP.

Routers lack comprehensive key management and keys derived that they can use to authenticate data. As an example, TCP-AO [RFC5925], talks about coordinating keys derived from the Master Key Table (MKT) between endpoints, but the MKT itself has to be configured manually or through an out-of-band mechanism. Also, TCP-AO does not address the issue of connectionless reset, as it applies to routers that do not store MKT across reboots.

Authentication, integrity protection, and encryption all require the use of keys by sender and receiver. An automated KMP, therefore has to include a way to distribute key material between two endpoints with little or no administrative overhead. It has to cover automatic key rollover. It is expected that authentication will cover the packet, i.e., the payload and the TCP header, and will not cover the frame, i.e., the layer 2 header.

There are two methods of automatic key rollover. Implicit key rollover can be initiated after a certain volume of data gets exchanged or when a certain time has elapsed. This does not require

explicit signaling nor should it result in a reset of the TCP connection in a way that the links/adjacencies are affected. On the other hand, explicit key rollover requires an out-of-band key signaling mechanism. It can be triggered by either side and can be done anytime a security parameter changes, e.g., an attack has happened, or a system administrator with access to the keys has left the company. An example of this is IKEv2 [RFC5996], but it could be any other new mechanisms also.

As stated earlier, TCP-AO [RFC5925] and its accompanying document, Cryptographic Algorithms for TCP-AO [RFC5926], require that two MAC algorithms be supported, and they are HMAC-SHA-1-96, as specified in HMAC [RFC2104], and AES-128-CMAC-96, as specified in NIST-SP800-38B [NIST-SP800-38B]. Therefore, TCP-AO meets the algorithm agility requirement.

There is a need to protect authenticity and validity of the routing/label information that is carried in the payload of the sessions. However, that is outside the scope of this document and is being addressed by the SIDR WG. Similar mechanisms could be used for intra-domain protocols.

Finally, replay protection is required. The replay mechanism needs to be sufficient to prevent an attacker from creating a denial of service or disrupting the integrity of the routing protocol by replaying packets. It is important that an attacker not be able to disrupt service by capturing packets and waiting for replay state to be lost.

4.1. LDP

As described in LDP [RFC5036], the threat of spoofed Basic Hellos can be reduced by only accepting Basic Hellos on interfaces that LSRs trust, employing GTSM [RFC5082], and ignoring Basic Hellos not addressed to the "all routers on this subnet" multicast group. Spoofing attacks via Targeted Hellos are potentially a more serious threat. An LSR can reduce the threat of spoofed Extended Hellos by filtering them and accepting Hellos from sources permitted by access lists. However, performing the filtering using access lists requires LSR resources, and the LSR is still vulnerable to the IP source address spoofing. Spoofing attacks can be solved by being able to authenticate the Hello messages, and an LSR can be configured to only accept Hello messages from specific peers when authentication is in use.

LDP Hello Cryptographic Authentication [HELLO-CRYPTO] suggest a new Cryptographic Authentication TLV that can be used as an authentication mechanism to secure Hello messages.

4.2. PCEP

Path Computation Element (PCE) discovery, according to [RFC5440], is a significant feature for the successful deployment of PCEP in large networks. This mechanism allows a Path Computation Client (PCC) to discover the existence of suitable PCEs within the network without the necessity of configuration. It should be obvious that, where PCEs are discovered and not configured, the PCC cannot know the correct key to use. There are different approaches to retain some aspect of security, but all of them require use of a keys and a keying mechanism, the need for which has been discussed above.

5. Transition and Deployment Considerations

As stated in the KARP Design Guidelines [RFC6518], it is imperative that the new authentication, security mechanisms, and key management protocol support incremental deployment, as it is not feasible to deploy the new routing protocol authentication mechanism overnight.

Typically, authentication and security in a peer-to-peer protocol requires that both parties agree to the mechanisms that will be used. If an agreement is not reached, the setup of the new mechanism will fail or will be deferred. Upon failure, the routing protocols can fall back to the mechanisms that were already in place, e.g., use static keys if that was the mechanism in place. The fallback should be configurable on a per-node or per-interface basis. It is usually not possible for one end to use the new mechanism while the other end uses the old. Policies can be put in place to retry upgrading after a said period of time, so that manual coordination is not required.

If the automatic KMP requires use of Public Key Infrastructure Certificates [RFC5280] to exchange key material, the required Certificate Authority (CA) root certificates may need to be installed to verify the authenticity of requests initiated by a peer. Such a step does not require coordination with the peer, except to decide which CA authority will be used.

6. Security Considerations

This section describes security considerations that BGP, LDP, PCEP, and MSDP should try to meet.

As with all routing protocols, they need protection from both on-path and off-path blind attacks. A better way to protect them would be with per-packet protection using a cryptographic MAC. In order to provide for the MAC, keys are needed.

The routing protocols need to support algorithm agility, i.e., they must not hardwire themselves to one algorithm.

Once keys are used, mechanisms are required to support key rollover. They should cover both manual and automatic key rollover. Multiple approaches could be used. However, since the existing mechanisms provide a protocol field to identify the key as well as management mechanisms to introduce and retire new keys, focusing on the existing mechanism as a starting point is prudent.

Furthermore, it is strongly suggested that these routing protocols support algorithm agility. It has been proven that algorithms weaken over time. Supporting algorithm agility assists in smooth transitions from old to new algorithms.

7. Acknowledgements

We would like to thank Brian Weis for encouraging us to write this document, and thanks to Anantha Ramaiah and Mach Chen for providing comments on it.

8. References

8.1. Normative References

[RFC5926] Lebovitz, G. and E. Rescorla, "Cryptographic Algorithms for the TCP Authentication Option (TCP-AO)", RFC 5926, June 2010.

[RFC6518] Lebovitz, G. and M. Bhatia, "Keying and Authentication for Routing Protocols (KARP) Design Guidelines", RFC 6518, February 2012.

[RFC6863] Hartman, S. and D. Zhang, "Analysis of OSPF Security According to the Keying and Authentication for Routing Protocols (KARP) Design Guide", RFC 6863, March 2013.

8.2. Informative References

[HELLO-CRYPTO]
Zheng, L., Chen, M., and M. Bhatia, "LDP Hello Cryptographic Authentication", Work in Progress, January 2013.

[NIST-SP800-38B]
Dworking, , "Recommendation for Block Cipher Modes of Operation: The CMAC Mode for Authentication", May 2005.

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

[RFC2385] Heffernan, A., "Protection of BGP Sessions via the TCP MD5 Signature Option", RFC 2385, August 1998.

[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.

[RFC3618] Fenner, B. and D. Meyer, "Multicast Source Discovery Protocol (MSDP)", RFC 3618, October 2003.

[RFC4271] Rekhter, Y., Li, T., and S. Hares, "A Border Gateway Protocol 4 (BGP-4)", RFC 4271, January 2006.

[RFC4301] Kent, S. and K. Seo, "Security Architecture for the Internet Protocol", RFC 4301, December 2005.

[RFC4595] Maino, F. and D. Black, "Use of IKEv2 in the Fibre Channel Security Association Management Protocol", RFC 4595, July 2006.

[RFC4674] Le Roux, J.L., "Requirements for Path Computation Element (PCE) Discovery", RFC 4674, October 2006.

[RFC4732] Handley, M., Rescorla, E., IAB, "Internet Denial-of-Service Considerations", RFC 4732, December 2006.

[RFC4948] Andersson, L., Davies, E., and L. Zhang, "Report from the IAB workshop on Unwanted Traffic March 9-10, 2006", RFC 4948, August 2007.

[RFC4953] Touch, J., "Defending TCP Against Spoofing Attacks", RFC 4953, July 2007.

[RFC5036] Andersson, L., Minei, I., and B. Thomas, "LDP Specification", RFC 5036, October 2007.

[RFC5082] Gill, V., Heasley, J., Meyer, D., Savola, P., and C. Pignataro, "The Generalized TTL Security Mechanism (GTSM)", RFC 5082, October 2007.

[RFC5088] Le Roux, JL., Vasseur, JP., Ikejiri, Y., and R. Zhang, "OSPF Protocol Extensions for Path Computation Element (PCE) Discovery", RFC 5088, January 2008.

[RFC5089] Le Roux, JL., Vasseur, JP., Ikejiri, Y., and R. Zhang, "IS-IS Protocol Extensions for Path Computation Element (PCE) Discovery", RFC 5089, January 2008.

[RFC5280] Cooper, D., Santesson, S., Farrell, S., Boeyen, S., Housley, R., and W. Polk, "Internet X.509 Public Key Infrastructure Certificate and Certificate Revocation List (CRL) Profile", RFC 5280, May 2008.

[RFC5440] Vasseur, JP. and JL. Le Roux, "Path Computation Element (PCE) Communication Protocol (PCEP)", RFC 5440, March 2009.

[RFC5920] Fang, L., "Security Framework for MPLS and GMPLS Networks", RFC 5920, July 2010.

[RFC5925] Touch, J., Mankin, A., and R. Bonica, "The TCP Authentication Option", RFC 5925, June 2010.

[RFC5961] Ramaiah, A., Stewart, R., and M. Dalal, "Improving TCP's Robustness to Blind In-Window Attacks", RFC 5961, August 2010.

[RFC5996] Kaufman, C., Hoffman, P., Nir, Y., and P. Eronen, "Internet Key Exchange Protocol Version 2 (IKEv2)", RFC 5996, September 2010.

[RFC6039] Manral, V., Bhatia, M., Jaeggli, J., and R. White, "Issues with Existing Cryptographic Protection Methods for Routing Protocols", RFC 6039, October 2010.

[RFC6151] Turner, S. and L. Chen, "Updated Security Considerations for the MD5 Message-Digest and the HMAC-MD5 Algorithms", RFC 6151, March 2011.

[RFC6480] Lepinski, M. and S. Kent, "An Infrastructure to Support Secure Internet Routing", RFC 6480, February 2012.

Authors' Addresses

Mahesh Jethanandani
Ciena Corporation
1741 Technology Drive
San Jose, CA 95110
USA

Phone: +1 (408) 436-3313
EMail: mjethanandani@gmail.com

Keyur Patel
Cisco Systems, Inc
170 Tasman Drive
San Jose, CA 95134
USA

Phone: +1 (408) 526-7183
EMail: keyupate@cisco.com

Lianshu Zheng
Huawei Technologies
China

Phone: +86 (10) 82882008
EMail: vero.zheng@huawei.com