Routing Working Group
Internet-Draft
Intended status: Informational
Expires: April 21, 2013

M. Jethanandani
Ciena Corporation
K. Patel
Cisco Systems, Inc
L. Zheng
Huawei Technologies
October 18, 2012

Analysis of BGP, LDP, PCEP and MSDP Issues According to KARP Design Guide

draft-ietf-karp-routing-tcp-analysis-05.txt

Abstract

This document analyzes Border Gateway Protocol (BGP) [RFC4271], Label Distribution Protocol (LDP) [RFC5036], Path Computation Element Protocol (PCEP) [RFC5440] and Multicast Source Distribution Protocol (MSDP) [RFC3618] according to guidelines set forth in section 4.2 of Keying and Authentication for Routing Protocols Design Guidelines [RFC6518].

Table of Contents

1. Introduction

In March 2006 the Internet Architecture Board (IAB) in its "Unwanted Internet Traffic" workshop documented in Report from the IAB workshop on Unwanted Traffic March 9-10, 2006 [RFC4948] described an attack on core routing infrastructure as an ideal attack with the most amount of damage. 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 securing 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 BGP, LDP, PCEP and MSDP according to the requirements of KARP Design Guidelines [RFC6518]. Section 4.2 of the 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 analysis efforts into routing security. The OPSEC working group published Issues with existing Cryptographic Protection Methods for Routing Protocols [RFC6039] an analysis of cryptographic issues with routing protocols and Analysis of OSPF Security According to KARP Design Guide [draft-ietf-karp-ospf-analysis-03].

Section 2 of this document looks at the current state of 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 KARP Design Guidelines [RFC6518] and Section 4 does a analysis of the gap between the existing state of the security method and the optimal state of the security method for protocols and suggests some areas where improvement is needed.

1.1. Conventions Used in This Document

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in RFC 2119 [RFC2119].

1.2. Abbreviations

AS - Autonomous Systems

BGP - Border Gateway Protocol

DoS - Denial of Service

GTSM - Generalized TTL Security Mechanism

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 Switch Routers

MAC - Message Authentication Code

MKT - Master Key Tuple

MSDP - Multicast Source Distribution Protocol

MD5 - Message Digest algorithm 5

OSPF - OPen Shortest Path First

PCEP - Path Computation Element Protocol

TCP - Transmission Control Protocol

TTL - Time To Live

UDP - User Datagram Protocol

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 a 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 with the result that routing updates are 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 from outside the protected domain. In addition for edge routers running 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 purports 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 deny 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 acts like a signature for that segment, computed using information known only to the connection end points. 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 roll-over.

The Message Authentication Codes (MACs) used by the TCP MD5 option is 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. TCP-AO [RFC5925] and its companion document Crypto Algorithms for TCP-AO [RFC5926] describe steps towards correcting both the MAC weakness and the management of secret keys. For MAC it specifies two MAC algorithms that MUST 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 defined are fairly secure. TCP-AO 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 Message Authentication Code (MAC). TCP-AO allows for a master key to be configured manually or for it to be managed via a 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 IKEv2 [RFC5996] provides for automated key exchange under a SA and can be used for a comprehensive Key Management Protocol (KMP) solution.

2.3. LDP

Section 5 of LDP [RFC5036] states that LDP is subject to two different types of attacks: spoofing, and denial of service attacks. In addition, LDP distributes labels in the clear, enabling hackers to see what labels are being distributed. The attacker can use that information to spoof a connection and distribute a different set of labels causing traffic to be dropped.

2.3.1. Spoofing attacks

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

2.3.1.1. Discovery exchanges using UDP

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

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 which 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 Extended Hello messages.

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 real problem in production networks.

2.3.1.2. 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 above, MD5 authentication is considered too weak for this application. 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 MACs, makes it easier to setup manual keys and protects against replays.

2.3.2. Privacy Issues

LDP provides no mechanism for protecting the privacy of label distribution. The security requirements of label distribution are similar to other routing protocols that need to distribute routing

information.

2.3.3. Denial of Service Attacks

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

2.4. PCEP

Attacks on PCEP [RFC5440] may result in damage to active networks. These include computation responses, which if changed can cause protocols like LDP to setup sub-optimal or inappropriate LSPs. In addition, PCE itself can be attacked by 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.

As RFC 5440 states, PCEP could be the target of the following attacks.

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

In inter-Autonomous Systems (AS) scenarios where PCE-to-PCE communication is required, attacks may be particularly significant with commercial 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 end-points do not reside in the same AS. An attacker that intercepts a PCE message could obtain sensitive information related to computed paths and resources.

2.5. MSDP

Similar to BGP and LDP, 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.

MSDP also advocates imposing a limit on number of source address and group addresses (S,G) that can be cached within the protocol and thereby 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 are when they can withstand any of the known types of attacks.

Additionally, Key Management Protocol (KMP) for the routing sessions should help negotiate unique, pair wise random keys without administrator involvement. It should also negotiate Security Association (SA) parameter required for the session connection, including key life times. 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 discovery phase of LDP. TCP and UDP should be able to withstand any of 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.

3.1. LDP

To harden LDP against its 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.

There is currently no requirement to protect the privacy of label distribution as labels are carried in the clear like other routing information.

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 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 protocol. It covers issues that are common to the four protocols in this section, 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 Improving TCP's Robustness to Blind In-Window Attacks. [RFC5961] should also be followed and implemented to strengthen TCP.

Routers lack comprehensive key management and keys derived from it that they can use to authenticate data. As an example TCP-AO [RFC5925], talks about coordinating keys derived from 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, tamper protection, and encryption all require the use of keys by sender and receiver. An automated KMP therefore has to include a way to distribute MKT between two end points with little or no administration 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 link layer 2 header.

There are two methods of automatic key rollover. Implicit key rollover can be initiated after 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

Crypto Algorithms for TCP-AO [RFC5926] suggest that two MAC algorithms that MUST be supported are HMAC-SHA-1-96 as specified in HMAC [RFC2104] and AES-128-CMAC-96 as specified in NIST-SP800-38B [NIST-SP800-38B].

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 SIDR WG. Similar mechanisms could be used for intra-domain protocols.

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 an access lists. However, performing the filtering using access lists requires LSR resource, 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 [draft-zheng-mpls-ldp-hello-crypto-auth-04] 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 its RFC [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 KARP Design Guidelines [RFC6518] it is imperative that the new authentication and security mechanisms defined 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 fallback to the mechanisms that were already in place e.g. use static keys if that was the mechanism in place. 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 a manual coordination is not required.

If the automatic KMP requires use of public/private keys to exchange key material, the required CA root certificates may need to be installed to verify authenticity of requests initiated by a peer. Such a step does not require coordination with the peer except to decide what 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.

Once keys are used, mechanisms are required to support key rollover. This 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.

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.

7. Acknowledgements

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

8. References

8.1. Normative References

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

[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.

[draft-ietf-karp-threats-reqs] Lebovitz, G. and M. Bhatia, "KARP Threats and Requirements", March 2012.

8.2. Informative References

[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.

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.

[RFC2409] Harkins, D. and D. Carrel, "The Internet Key Exchange (IKE)", RFC 2409, November 1998.

[RFC3547] Baugher, M., Weis, B., Hardjono, T., and H. Harney, "The Group Domain of Interpretation", RFC 3547, July 2003.

[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.

[RFC4732] Handley, M., Rescorla, E., and 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.

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

[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.

[draft-ietf-karp-ospf-analysis-03] Hartman, S., "Analysis of OSPF Security According to KARP Design Guide", March 2012.

[draft-zheng-mpls-ldp-hello-crypto-auth-04] Zheng, "LDP Hello Cryptographic Authentication", May 2012.

Authors' Addresses

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

Phone: + (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
Fax:
Email: vero.zheng@huawei.com
URI: