Internet Engineering Task Force (IETF)
Request for Comments: 7602
Category: Standards Track
ISSN: 2070-1721
U. Chunduri
W. Lu
A. Tian
Ericsson Inc.
N. Shen
Cisco Systems, Inc.
July 2015

IS-IS Extended Sequence Number TLV

Abstract

This document defines the Extended Sequence Number TLV to protect Intermediate System to Intermediate System (IS-IS) PDUs from replay attacks.

Status of This Memo

This is an Internet Standards Track document.

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). Further information on Internet Standards is available in 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/rfc7602.

Copyright Notice

Table of Contents

1. Introduction

Intermediate System to Intermediate System (IS-IS) [ISO10589] has been adopted widely in various Layer 2 / Layer 3 routing and switching deployments of data centers and for critical business operations. Its flexibility and scalability make it well suited for the rapid development of new data center infrastructures. Also, while technologies such as Software-Defined Networking (SDN) may improve network management and enable new applications, their use has an effect on the security requirements of the routing infrastructure.

A replayed IS-IS PDU can potentially cause many problems in IS-IS networks, including bouncing adjacencies, blackholing, and even some form of Denial-of-Service (DoS) attacks as explained in Section 2. This problem is also discussed in the Security Considerations section, in the context of cryptographic authentication work as described in [RFC5304] and [RFC5310].

Currently, there is no mechanism to protect IS-IS Hello (IIH) PDUs and Sequence Number PDUs (SNPs) from replay attacks. However, Link State PDUs (LSPs) have a sequence number in the LSP header as defined in [ISO10589], with which they can effectively mitigate intra-session replay attacks. But, LSPs are still susceptible to inter-session replay attacks.

This document defines the Extended Sequence Number (ESN) TLV to protect IS-IS PDUs from replay attacks.

The new ESN TLV defined here thwarts these threats and can be deployed with the authentication mechanisms specified in [RFC5304] and [RFC5310] for a more secure network.

Replay attacks can be effectively mitigated by deploying a group key management protocol (being developed as defined in [GROUP-IKEv2] and [MRKMP]) with a frequent key change policy. Currently, there is no such mechanism defined for IS-IS. Even if such a mechanism is defined, usage of this TLV can be helpful to avoid replays before the keys are changed.

Also, it is believed that, even when such a key management system is deployed, there always will be some systems based on manual keying that coexist with systems based on key management protocols. The ESN TLV defined in this document is helpful for such deployments.

1.1. Requirements Language

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. Acronyms

CSNP - Complete Sequence Number PDU

ESN - Extended Sequence Number

IIH - IS-IS Hello

IS - Intermediate System

LSP - IS-IS Link State PDU

PDU - Protocol Data Unit

PSNP - Partial Sequence Number PDU

SNP - Sequence Number PDU

2. Replay Attacks and Impact on IS-IS Networks

Replaying a captured protocol packet to cause damage is a common threat for any protocol. Securing the packet with cryptographic authentication information alone cannot mitigate this threat completely. This section explains the replay attacks and their applicability to each IS-IS PDU.

2.1. IIHs

When an adjacency is brought up, an IS sends an IIH packet with an empty neighbor list (TLV 6); it can be sent with or without authentication information. Packets can be replayed later on the broadcast network, and this may cause all ISs to bounce the adjacency, thus churning the network. Note that mitigating replay is only possible when authentication information is present.

2.2. LSPs

Normal operation of the IS-IS update process as specified in [ISO10589] provides timely recovery from all LSP replay attacks. Therefore, the use of the extensions defined in this document is prohibited in LSPs. Further discussion of the vulnerability of LSPs to replay attacks can be found in [ISIS-ANALYSIS].

2.3. SNPs

A replayed CSNP can result in the sending of unnecessary PSNPs on a given link. A replayed CSNP or PSNP can result in unnecessary LSP flooding on the link.

3. Extended Sequence Number TLV

The Extended Sequence Number (ESN) TLV is composed of 1 octet for the Type, 1 octet that specifies the number of bytes in the Value field, and a 12-byte Value field. This TLV is defined only for IIH and SNP PDUs.

Code - 11.

Length - total length of the value field, which is 12 bytes.

Value - 64-bit Extended Session Sequence Number (ESSN), which is followed by a 32-bit, monotonically increasing, per-packet sequence number.

```
 0                   1                   2                   3
 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+-+-+-+-+-+-+-+-+
|    Type       |
+-+-+-+-+-+-+-+-+
|    Length     |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|    Extended Session Sequence Number (High-Order 32 Bits)      |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|    Extended Session Sequence Number (Low-Order 32 Bits)       |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|             Packet Sequence Number (32 Bits)                  |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

Figure 1: Extended Sequence Number (ESN) TLV

The ESN TLV defined here is optional. Though this is an optional TLV, it can be mandatory on a link when 'verify' mode is enabled as specified in Section 5.1. The ESN TLV MAY be present only in IIH PDUs and SNPs. A PDU with multiple ESN TLVs is invalid and MUST be discarded on receipt.

The 64-bit ESSN MUST be nonzero and MUST contain a number that is increased whenever it is changed due any situation, as specified in Section 3.1. Encoding the 64-bit unsigned integer ESSN value is a local matter, and implementations MAY use one of the alternatives provided in Appendix A. Effectively, for each PDU that contains the ESN TLV, the 96-bit unsigned integer value consisting of the 64-bit ESSN and 32-bit Packet Sequence Number (PSN) -- where the ESSN is the higher-order 64 bits -- MUST be greater than the most recently received value in a PDU of the same type originated by the same IS.

3.1. Sequence Number Wrap

If the 32-bit Packet Sequence Number in the ESN TLV wraps or the router performs a cold restart, the 64-bit ESSN value MUST be set higher than the previous value. IS-IS implementations MAY use the guidelines provided in Appendix A for accomplishing this.

4. Mechanism and Packet Encoding

The encoding of the ESN TLV in each applicable IS-IS PDU is detailed below. Please refer to Section 5 for appropriate operations on how to interoperate with legacy node(s) that do not support the extensions defined in this document. If the received PDU with the ESN TLV is accepted, then the stored value for the corresponding originator and PDU type MUST be updated with the latest value received. Please note that level information is included in the PDU type.

4.1. IIHs

ESN TLV information is maintained for each type of IIH PDU being sent on a given circuit. The procedures for encoding, verification, and sequence number wrapping are explained in Section 3.

4.2. SNPs

Separate CSNP/PSNP ESN TLV information is maintained per PDU type, per originator, and per link. The procedures for encoding, verification, and sequence number wrapping are explained in Section 3.

5. Backward Compatibility and Deployment

The implementation and deployment of the ESN TLV can be done to support backward compatibility and gradual deployment in the network without requiring a flag day. This feature can also be deployed for the links in a certain area of the network where the maximum security mechanism is needed, or it can be deployed for the entire network.

The implementation SHOULD allow the configuration of ESN TLV features on each IS-IS link level. The implementation SHOULD also allow operators to control the configuration of the 'send' and/or 'verify' feature of IS-IS PDUs for the links and for the node. In this document, the 'send' mode is to include the ESN TLV in its own IS-IS PDUs, and the 'verify' mode is to process the ESN TLV in the receiving IS-IS PDUs from neighbors.

When an adversary is actively attacking, it is possible to have inconsistent data views in the network, if there is a considerable delay in enabling the 'verify' mode where nodes were configured to the 'send' mode, e.g., from the first to the last node or all nodes of a particular LAN segment. This happens primarily because replay PDUs can potentially be accepted by the nodes where the 'verify' mode is still not provisioned at the time of the attack. To minimize such a window, it is recommended that provisioning of 'verify' SHOULD be done in a timely fashion by the network operators.

5.1. IIHs and SNPs

On the link level, the ESN TLV involves the IIH PDUs and SNPs (both CSNP and PSNP). The 'send' and 'verify' modes described above can be set independently on each link and, in the case of a broadcast network, independently on each level.

To introduce ESN support without disrupting operations, ISs on a given interface are first configured to operate in 'send' mode. Once all routers operating on an interface are operating in 'send' mode, 'verify' mode can be enabled on each IS. Once 'verify' mode is set for an interface, all the IIH PDUs and SNPs being sent on that interface MUST contain the ESN TLV. Any such PDU received without an ESN TLV MUST be discarded when 'verify' mode is enabled. Similarly, to safely disable ESN support on a link, 'verify' mode is disabled on all ISs on the link. Once 'verify' mode is disabled on all routers operating on an interface, 'send' mode can be disabled on each IS. Please refer to Section 5 for considerations on enabling or disabling 'verify' mode on all ISs on a link.

6. IANA Considerations

A new TLV codepoint, as defined in this document, has been assigned by IANA from the "IS-IS TLV Codepoints" registry. It is referred to as the Extended Sequence Number TLV and has the following attributes:

Value	Name	IIH	LSP	SNP	Purge
11	ESN TLV	y	n	y	n

7. Security Considerations

This document describes a mechanism to mitigate the replay attack threat as discussed in the Security Considerations sections of [RFC5304] and [RFC5310]. If an adversary interferes either by not forwarding packets or by delaying messages as described in Section 3.3 of [RFC6862], the mechanism specified in this document cannot mitigate those threats. Also, some of the threats described in Section 2.3 of [ISIS-ANALYSIS] are not addressable with the ESN TLV as specified in this document. This document does not introduce any new security concerns to IS-IS or any other specifications referenced.

8. References

8.1. Normative References

[ISO10589] International Organization for Standardization, "Intermediate system to intermediate system intra-domain-routing routine information exchange protocol for use in conjunction with the protocol for providing the connectionless-mode Network Service (ISO 8473)", ISO/IEC 10589:2002, Second Edition, Nov. 2002.

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <http://www.rfc-editor.org/info/rfc2119>.

[RFC5905] Mills, D., Martin, J., Ed., Burbank, J., and W. Kasch, "Network Time Protocol Version 4: Protocol and Algorithms Specification", RFC 5905, DOI 10.17487/RFC5905, June 2010, <http://www.rfc-editor.org/info/rfc5905>.

8.2. Informative References

[MRKMP] Hartman, S., Zhang, D., and G. Lebovitz, "Multicast Router Key Management Protocol (MaRK)", Work in Progress, draft-hartman-karp-mrkmp-05, September 2012.

[ISIS-ANALYSIS] Chunduri, U., Tian, A., and W. Lu, "KARP IS-IS security analysis", Work in Progress, draft-ietf-karp-isis-analysis-07, July 2015.

[GROUP-IKEv2] Rowles, S., Yeung, A., Ed., Tran, P., and Y. Nir, "Group Key Management using IKEv2", Work in Progress, draft-yeung-g-ikev2-08, October 2014.

[RFC5304] Li, T. and R. Atkinson, "IS-IS Cryptographic Authentication", RFC 5304, DOI 10.17487/RFC5304, October 2008, <http://www.rfc-editor.org/info/rfc5304>.

[RFC5310] Bhatia, M., Manral, V., Li, T., Atkinson, R., White, R., and M. Fanto, "IS-IS Generic Cryptographic Authentication", RFC 5310, DOI 10.17487/RFC5310, February 2009, <http://www.rfc-editor.org/info/rfc5310>.

[RFC6862] Lebovitz, G., Bhatia, M., and B. Weis, "Keying and Authentication for Routing Protocols (KARP) Overview, Threats, and Requirements", RFC 6862, DOI 10.17487/RFC6862, March 2013, <http://www.rfc-editor.org/info/rfc6862>.

[RFC7474] Bhatia, M., Hartman, S., Zhang, D., and A. Lindem, Ed., "Security Extension for OSPFv2 When Using Manual Key Management", RFC 7474, DOI 10.17487/RFC7474, April 2015, <http://www.rfc-editor.org/info/rfc7474>.

Appendix A. ESSN Encoding Mechanisms

IS-IS nodes implementing this specification SHOULD use available mechanisms to preserve the 64-bit Extended Session Sequence Number's strictly increasing property, whenever it is changed for the deployed life of the IS-IS node (including cold restarts).

This appendix provides guidelines for maintaining the strictly increasing property of the 64-bit ESSN in the ESN TLV, and implementations can resort to any similar method as long as this property is maintained.

A.1. Using Timestamps

One mechanism for accomplishing this is by encoding the 64-bit ESSN as the system time represented by a 64-bit unsigned integer value. This MAY be similar to the system timestamp encoding for the NTP long format as defined in Appendix A.4 of [RFC5905]. The new current time MAY be used when the IS-IS node loses its sequence number state including when the Packet Sequence Number wraps.

Implementations MUST make sure while encoding the 64-bit ESN value with the current system time that it does not default to any previous value or some default node time of the system, especially after cold restarts or any other similar events. In general, system time must be preserved across cold restarts in order for this mechanism to be feasible. One example of such implementation is to use a battery backed real-time clock (RTC).

A.2. Using Non-volatile Storage

One other mechanism for accomplishing this is similar to the one specified in [RFC7474] -- use the 64-bit ESSN as a wrap/boot count stored in non-volatile storage. This value is incremented anytime the IS-IS node loses its sequence number state, including when the Packet Sequence Number wraps.

There is a drawback to this approach, which is described as follows in Section 8 of [RFC7474]. It requires the IS-IS implementation to be able to save its boot count in non-volatile storage. If the non-volatile storage is ever repaired or router hardware is upgraded such that the contents are lost, keys MUST be changed to prevent replay attacks.

Appendix B. Operational/Implementation Considerations

Since the ESN is maintained per PDU type, per originator, and per link, this scheme can be useful for monitoring the health of the IS-IS adjacency. A Packet Sequence Number skip that occurs upon receiving an IIH can be recorded by the neighbors and can be used later to correlate adjacency state changes over the interface. For instance, in multi-access media, completely different issues on the network may be indicated when all neighbors record skips from the same IIH sender versus when only one neighbor records skips. For operational issues, effective usage of the TLV defined in this document MAY also need more system information before making concrete conclusions; defining all that information is beyond the scope of this document.

Acknowledgements

As some sort of sequence number mechanism to thwart protocol replays is a old concept, the authors of this document do not make any claims on the originality of the overall protection idea described. The authors are thankful for the review and the valuable feedback provided by Acee Lindem and Joel Halpern. Thanks to Alia Atlas, Chris Hopps, Nevil Brownlee, and Adam W. Montville for their reviews and suggestions during IESG directorate review. The authors also thank Christer Holmberg, Ben Campbell, Barry Leiba, Stephen Farrell, and Alvaro Retana for their reviews of this document.

Contributors

The authors would like to thank Les Ginsberg for his significant contribution in detailed reviews and suggestions.

Authors' Addresses

Uma Chunduri
Ericsson Inc.
300 Holger Way,
San Jose, California 95134
United States

Phone: 408 750-5678
Email: uma.chunduri@ericsson.com

Wenhu Lu
Ericsson Inc.
300 Holger Way,
San Jose, California 95134
United States

Email: wenhu.lu@ericsson.com

Albert Tian
Ericsson Inc.
300 Holger Way,
San Jose, California 95134
United States

Phone: 408 750-5210
Email: albert.tian@ericsson.com

Naiming Shen
Cisco Systems, Inc.
225 West Tasman Drive,
San Jose, California 95134
United States

Email: naiming@cisco.com