

Internet Engineering Task Force (IETF)
Request for Comments: 7987
Category: Standards Track
ISSN: 2070-1721

L. Ginsberg
P. Wells
Cisco Systems
B. Decraene
Orange
T. Przygienda
Juniper
H. Gredler
RtBrick Inc.
October 2016

IS-IS Minimum Remaining Lifetime

Abstract

Corruption of the Remaining Lifetime field in a Link State Protocol Data Unit (LSP) can go undetected. In certain scenarios, this may cause or exacerbate flooding storms. It is also a possible denial-of-service attack vector. This document defines a backwards-compatible solution to this problem.

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 7841.

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/rfc7987>.

Copyright Notice

Copyright (c) 2016 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.

Table of Contents

1. Problem Statement	3
1.1. Requirements Language	4
2. Solution	4
3. Deployment Considerations	5
3.1. Inconsistent Values for MaxAge	5
3.2. Reporting Corrupted Lifetime	6
3.3. Impact of Delayed LSP Purging	7
4. Security Considerations	7
5. References	7
5.1. Normative References	7
5.2. Informative References	8
Acknowledgements	8
Contributors	8
Authors' Addresses	9

1. Problem Statement

[ISO10589] defines the format of a Link State PDU (LSP) that includes a Remaining Lifetime field. This field is set by the originator based on local configuration and then decremented by all systems once the entry is stored in their LSP Database (LSPDB) consistent with the passing of time. This allows all Intermediate Systems (ISs) to age out the LSP at approximately the same time.

Each LSP also has a checksum field to allow receiving systems to detect errors that may have occurred during transmission. [ISO10589] mandates that the checksum is calculated by the originator of the LSP and cannot be modified by other routers. Therefore, the Remaining Lifetime is deliberately excluded from the checksum calculation. In cases where cryptographic authentication is included in an LSP ([RFC5304] or [RFC5310]), the Remaining Lifetime field is also excluded from the hash calculation. If the Remaining Lifetime field gets corrupted during flooding, this corruption is therefore undetectable. The consequences of such corruption depend upon how the Remaining Lifetime is altered.

In cases where the Remaining Lifetime becomes larger than the originator intended, the impact is benign. As the originator is responsible for refreshing the LSP before it ages out, a new version of the LSP will be generated before the LSP ages out, so no harm is done.

In cases where the Remaining Lifetime field becomes smaller than the originator intended, the LSP may age out prematurely (i.e., before the originator refreshes the LSP). This has two negative consequences:

1. The LSP will be purged by an IS when the Remaining Lifetime expires. This will cause a temporary loss of reachability to destinations impacted by the content of that LSP.
2. Unnecessary LSP flooding will occur as a result of the premature purge and subsequent regeneration/flooding of a new version of the LSP by the originator.

If the corrupted Remaining Lifetime is only modestly shorter than the lifetime assigned by the originator, the negative impacts are also modest. If, however, the corrupted Remaining Lifetime becomes very small, then the negative impacts can become significant, especially in cases where the cause of the corruption is persistent so that the cycle repeats itself frequently.

A backwards-compatible solution to this problem is defined in the following sections.

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].

2. Solution

As discussed in the previous section, the problematic case is when the Remaining Lifetime is corrupted and becomes much smaller than it should be. The goal of the solution is then to prevent premature purging.

Under normal circumstances, updates to an LSP -- including purging, if appropriate -- are the responsibility of the originator of the LSP. There is a maximum time between generations of a given LSP. Once this time has expired, it is the responsibility of the originator to refresh the LSP (i.e., issue a new version with a higher sequence number) even if the contents of the LSP have not changed. [ISO10589] defines `maximumLSPGenerationInterval` to be sufficiently less than the maximum lifetime of an LSP so that the new version can be flooded network wide before the old version ages out on any IS.

[ISO10589] defines two cases where a system other than the originator of an LSP is allowed to purge an LSP:

1. The LSP ages out. This should only occur if the originating IS is no longer reachable and therefore is unable to update the LSP.
2. There is a Designated Intermediate System (DIS) change on a LAN. The pseudonode LSPs generated by the previous DIS are no longer required and may be purged by the new DIS.

In both of these cases, purging is not necessary for correct operation of the protocol. It is provided as an optimization to remove stale entries from the LSPDB.

In cases where the Remaining Lifetime in a received LSP has been corrupted and is smaller than the Remaining Lifetime at the originating node, when the Remaining Lifetime expires on the receiving node, it can appear as if the originating IS has failed to regenerate the LSP before it ages out (case #1 above). To prevent this from having a negative impact, a modest change to the storage of "new" LSPs in the LSPDB is specified.

Section 7.3.16 of [ISO10589] defines the rules to determine whether a received LSP is older, the same, or newer than the copy of the same LSP in the receiver's LSPDB. The key elements are:

- o Higher sequence numbers are newer.
- o If sequence numbers are the same, an LSP with a zero Remaining Lifetime (a purged LSP) is newer than the same LSP with a non-zero Remaining Lifetime.
- o If both the received and local copy of the LSP have a non-zero Remaining Lifetime, they are considered the same even if the Remaining Lifetimes differ.

Section 7.3.15.1.e(1) of [ISO10589] defines the actions to take on receipt of an LSP generated by another IS that is newer than the local copy and has a non-zero Remaining Lifetime. An additional action is defined by this document:

- vi. If the Remaining Lifetime of the new LSP is less than MaxAge, it is set to MaxAge.

This additional action ensures that no matter what value of Remaining Lifetime is received, a system other than the originator of an LSP will never purge the LSP until the LSP has existed in the database for at least MaxAge.

It is important to note that no change is proposed for handling the receipt of purged LSPs. The rules specified in Section 7.3.15.1.b of [ISO10589] still apply, i.e., an LSP received with a zero Remaining Lifetime is still considered newer than a matching LSP with a non-zero Remaining Lifetime. Therefore, the changes proposed here will not result in LSPDB inconsistency among routers in the network.

3. Deployment Considerations

This section discusses some possible deployment issues for this protocol extension.

3.1. Inconsistent Values for MaxAge

[ISO10589] defines MaxAge (the maximum value for the Remaining Lifetime in an LSP) as an architectural constant of 20 minutes (1200 seconds). However, in practice, implementations have supported allowing this value to be configurable. The common intent of a configurable value is to support longer lifetimes than the default, thus reducing the periodic regeneration of LSPs in the absence of topology changes. See a discussion of this point in [RFC3719]. It

is therefore possible for the value of MaxAge on the IS that originates an LSP to be higher or lower than the value of MaxAge on the ISs that receive the LSP.

If the value of MaxAge of the IS that originated the LSP is smaller than the value of MaxAge of the receiver of an LSP, then setting the Remaining Lifetime of the received LSP to the local value of MaxAge will ensure that it is not purged prematurely. However, if the value of MaxAge on the receiver is less than that of the originator, then it is still possible to have an LSP purged prematurely when using the extension defined in the previous section. Implementors of this extension may wish to protect against this case by making the value to which the Remaining Lifetime is set under the conditions described in the previous section configurable. If that is done, the configured value MUST be greater than or equal to the locally configured value of MaxAge.

3.2. Reporting Corrupted Lifetime

Reporting reception of an LSP with a possible corrupt Remaining Lifetime field can be useful in identifying a problem in the network. In order to minimize the reports of false positives, the following algorithm SHOULD be used in determining whether the Remaining Lifetime in the received LSP is possibly corrupt:

- o The LSP has passed all acceptance tests as specified in Section 7.3.15.1 of [ISO10589].
- o The LSP is newer than the copy in the local LSPDB (including the case of not being present in the LSPDB).
- o The Remaining Lifetime in the received LSP is less than ZeroAgeLifetime.
- o The adjacency to the neighbor from which the LSP is received has been up for a minimum of ZeroAgeLifetime.

In such a case, an IS SHOULD generate a CorruptRemainingLifetime event.

Note that it is not possible to guarantee that all cases of a corrupt Remaining Lifetime will be detected using the above algorithm. It is also not possible to guarantee that all CorruptRemainingLifetime events reported using the above algorithm are valid. As a diagnostic aid, an implementation MAY wish to retain the value of the Remaining Lifetime received when the LSP was added to the LSPDB.

3.3. Impact of Delayed LSP Purging

The extensions defined in this document may result in retaining an LSP longer than its original lifetime. In order for this to occur, the scheduled refresh of the LSP by the originator of the LSP must fail to occur -- this implies that the originator is no longer reachable. In such a case, its neighbors will update their own LSPs to report the loss of connectivity to the originator. [ISO10589] specifies that LSPs from a node that is unreachable (failure of the two-way connectivity check) will not be used. Note that this behavior applies to ALL information in the set of LSPs from such a node.

Retention of stale LSPs therefore has no negative side effects other than requiring additional memory for the LSPDB. In networks where a combination of pathological behaviors (e.g., LSP corruption and frequent resetting of nodes in the network) is seen, this could lead to a large number of stale LSPs being retained, but such networks are already compromised.

4. Security Considerations

The ability to introduce corrupt LSPs is not altered by the rules defined in this document. Use of authentication as defined in [RFC5304] and [RFC5310] prevents such LSPs from being intentionally introduced. A man-in-the-middle attack that modifies an existing LSP by changing the Remaining Lifetime to a small value can cause premature purges even in the presence of cryptographic authentication. The mechanisms defined in this document prevent such an attack from being effective.

5. References

5.1. Normative References

- [ISO10589] International Organization for Standardization, "Information technology -- Telecommunications and information exchange between systems -- Intermediate System to Intermediate System intra-domain routing 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, November 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>>.

- [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>>.

5.2. Informative References

- [PROB-STATEMENT] Decraene, B. and C. Schmitz, "IS-IS LSP lifetime corruption - Problem Statement", Work in Progress, draft-decraene-isis-lsp-lifetime-problem-statement-02, July 2016.
- [RFC3719] Parker, J., Ed., "Recommendations for Interoperable Networks using Intermediate System to Intermediate System (IS-IS)", RFC 3719, DOI 10.17487/RFC3719, February 2004, <<http://www.rfc-editor.org/info/rfc3719>>.

Acknowledgements

The problem statement in [PROB-STATEMENT] motivated this work.

Contributors

The following individual contributed substantially to the content of this document and should be considered a co-author:

Stefano Previdi
Cisco Systems
Email: sprevidi@cisco.com

Authors' Addresses

Les Ginsberg
Cisco Systems
510 McCarthy Blvd.
Milpitas, CA 95035
United States of America

Email: ginsberg@cisco.com

Paul Wells
Cisco Systems
170 W. Tasman Dr.
San Jose, CA 95035
United States of America

Email: pauwells@cisco.com

Bruno Decraene
Orange
44 avenue de la Republique, CS 50010
92326 Chatillon Cedex 92794
France

Email: bruno.decraene@orange.com

Tony Przygienda
Juniper
1137 Innovation Way
Sunnyvale, CA 94089
United States of America

Email: prz@juniper.net

Hannes Gredler
RtBrick Inc.

Email: hannes@rtbrick.com