Networking Working Group
Internet-Draft
Intended status: Standards Track
Expires: February 1, 2016

L. Ginsberg
P. Wells
S. Previdi
Cisco Systems
B. Decraene
Orange
T. Przygienda
Ericsson
H. Gredler
Juniper Networks, Inc

July 31, 2015

IS-IS Minimum Remaining Lifetime
draft-ginsberg-isis-remaining-lifetime-00.txt

Abstract

Corruption of the Remainining Lifetime Field in a Link State PDU 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.

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].

Status of This Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at http://datatracker.ietf.org/drafts/current/.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on February 1, 2016.

Table of Contents

1. Problem Statement

Each Link State PDU (LSP) 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 Link

State PDU 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 which may have occurred during transmission. As the Remaining Lifetime field changes as it is flooded and as the checksum field MUST NOT be altered by receiving ISs 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.

2. Solution

As discussed in the previous section, the problematic case is when 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 higher sequence number) even if the contents of the LSP have not changed. [ISO10589] specifies that maximumLSPGenerationInterval MUST 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.

There are 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 pseudo-node 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 RemainingLifetime expires on the receiving node it can appear as if the originating IS has failed to regenerate the LSP (case #1 above) when in fact the LSP still has significant lifetime remaining. To prevent this from having a negative impact a modest change to the storage of "new" LSPs in the LSPDB is specified.

[ISO10589] Section 7.3.16 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:

- Higher sequence numbers are newer

- If sequence numbers are the same, an LSP with zero RemainingLifetime (a purged LSP) is newer than the same LSP w non-zero RemainingLifetime

- o If both the received and local copy of the LSP have non-zero RemainingLifetime they are considered the same even if the RemainingLifetimes differ

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

vi. If the RemainingLifetime of the new LSP is less than MaxAge it is set to MaxAge

This additional action insures 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 [ISO10589] Section 7.3.15.1b still apply i.e., an LSP received with zero RemainingLifetime is still considered newer than a matching LSP with non-zero RemainingLifetime. Therefore the changes proposed here will not result in LSPDB inconsistency among routers in the newtork.

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 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 which originates an LSP to be higher or lower than the value of MaxAge on the ISs which receive the LSP.

If the value of MaxAge of the IS which originated the LSP is smaller than the value of MaxAge of the receiver of an LSP, then setting the RemainingLifetime of the received LSP to the local value of MaxAge will insure 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 when using the extension defined in the previous

section to have an LSP purged prematurely. Implementors of this extension MAY wish to protect against this case by making the value to which RemainingLifetime 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

It may be useful for an IS to report reception of an LSP with a possible corrupt RemainingLifetime field. In order to minimize the reports of false positives the following algorithm SHOULD be used in determining whether the RemainingLifetime in the received LSP is possibly corrupt:

- The LSP has passed all acceptance tests as specified in [ISO10589] Section 7.3.15.1
- The LSP is newer than the copy in the local LSPDB (including the case of not being present in the LSPDB)
- RemainingLifetime in the received LSP is less than ZeroAgeLifetime
- 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 MAY generate a CorruptRemainingLifetime event.

Note that it is not possible to guarantee that all cases of corrupt RemainingLifetime 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 RemainingLifetime 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 - which implies the originator is no longer reachable. In such a case its neighbors will update their own LSPs reporting the loss of connectivity to the originator. LSPs from a node which is unreachable (failure of the two-way-connectivity check) MUST NOT be used. Note 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 (LSP corruption, 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. IANA Considerations

None.

5. 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 which 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.

6. Acknowledgements

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

7. References

7.1. Normative References

[ISO10589] International Organization for Standardization, "Intermediate system to Intermediate system intra-domain routeing 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>.

[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>.

7.2. Informational References

[LIFE-PROB] "IS-IS LSP lifetime corruption - Problem Statement, draft-decraene-isis-lsp-lifetime-problem-statement-00(work in progress)", July 2015.

[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>.

Authors' Addresses

Les Ginsberg
Cisco Systems
510 McCarthy Blvd.
Milpitas, CA 95035
USA

Email: ginsberg@cisco.com

Paul Wells
Cisco Systems
170 W Tasman Dr
San Jose, Ca 95035
USA

Email: pauwells@cisco.com

Stefano Previdi
Cisco Systems
Via Del Serafico 200
Rome 0144
Italy

Email: sprevidi@cisco.com

Bruno Decraene
Orange
38 rue du General Leclerc
Issy Moulineaux cedex 9 92794
France

Email: bruno.decraene@orange.com

Tony Przygienda
Ericsson
300 Holger Way
San Jose, Ca 95134
USA

Email: antoni.przygienda@ericsson.com

Hannes Gredler
Juniper Networks, Inc
1194 N. Matilda Ave
Sunnyvale, Ca 94089
USA

Email: hannes@juniper.net