Workgroup: Networking Working Group
Internet-Draft:
draft-ppsenak-lsr-igp-pfx-reach-loss-00
Published: 7 March 2022
Intended Status: Informational
Expires: 8 September 2022
Authors: P. Psenak, Ed. L. Ginsberg D. Voyer
Cisco Systems Cisco Systems Bell Canada

IGP Prefix Reachability Loss Anouncement

Abstract

In the presence of summarization, there is a need to signal loss of reachability to an individual prefix covered by the summary in order to enable fast convergence away from paths to the node which owns the prefix which is no longer reachable. This document describes how to use existing protocol mechanisms in IS-IS and OSPF to advertise such prefix reachability loss.

Requirements Language

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "NOT RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in BCP 14 [RFC2119][RFC8174] when, and only when, they appear in all capitals, as shown here.

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 https://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 8 September 2022.

Copyright Notice

Table of Contents

1. Introduction
2. Supporting PRLA in IS-IS
   2.1. Advertisement of PRLA in IS-IS
   2.2. Propagation of PRLA in IS-IS
3. Supporting PRLA in OSPF
   3.1. Advertisement of PRLA in OSPF
   3.2. Propagation of PRLA in OSPF
4. Deployment Considerations for PRLA
5. IANA Considerations
6. Security Considerations
7. Acknowledgements
8. Normative References
Authors' Addresses

1. Introduction

Link-state IGP protocols like IS-IS and OSPF are primarily used to distribute routing information between routers belonging to a single Autonomous System (AS) and to calculate the reachability for IPv4 or IPv6 prefixes advertised by the individual nodes inside the AS. Each node advertises the state of its local adjacencies, connected prefixes, capabilities, etc. The collection of these states from all the routers inside the area form a link-state database (LSDB) that describes the topology of the area and holds additional state information about the prefixes, router capabilities, etc.

The growth of networks running a link-state routing protocol results in the addition of more state which leads to scalability and convergence challenges. The organization of networks into levels/areas and IGP domains helps limit the scope of link-state information within certain boundaries. However, the state related to prefix reachability often requires propagation across a multi-area/level and/or multi-domain IGP network. Techniques such as summarization have been used traditionally to address the scale challenges associated with advertising prefix state outside of the local area/domain. However, this results in suppression of the individual prefix state that is useful for triggering fast-

convergence mechanisms outside of the IGPs - e.g., BGP PIC Edge [I-D.ietf-rtgwg-bgp-pic].

This document describes how the use of existing protocol mechanisms can support the necessary functionality without the need for any protocol extensions. The functionality being described is called Prefix Reachability Loss Announcement (PRLA).

2. Supporting PRLA in IS-IS

[RFC5305] defines the encoding for advertising IPv4 prefixes using 4 octets of metric information. Section 4 specifies:

"If a prefix is advertised with a metric larger then MAX_PATH_METRIC (0xFE000000, see paragraph 3.0), this prefix MUST NOT be considered during the normal SPF computation. This allows advertisement of a prefix for purposes other than building the normal IP routing table. "

Similarly, [RFC5308] defines the encoding for advertising IPv6 prefixes using 4 octets of metric information. Section 2 states:

"...if a prefix is advertised with a metric larger than MAX_V6_PATH_METRIC (0xFE000000), this prefix MUST NOT be considered during the normal Shortest Path First (SPF) computation. This will allow advertisement of a prefix for purposes other than building the normal IPv6 routing table."

This functionality can be used to advertise a prefix (IPv4 or IPv6) in a manner which indicates that reachability has been lost - and to do so without requiring all nodes in the network to be upgraded to support the functionality.

2.1. Advertisement of PRLA in IS-IS

Existing nodes in a network which receive PRLA advertisements will ignore them. This allows flooding of such advertisements to occur without the need to upgrade all nodes in a network.

Recognition of the advertisement as PRLA is only required on routers which have a use case for this information. Area Border Routers (ABRs), which would be responsible for propagating PRLA advertisements into other areas would need to recognize such advertisements.

As per the definitions referenced in the preceding section, any prefix advertisement with a metric value greater than 0xFE000000 can be used for purposes other than normal routing calculations. Such an advertisement can be interpreted by the receiver as a PRLA.

Optionally, an implementation may use local configuration to limit the set of metric values which will be interpreted as PRLA. The only restriction is that such values MUST be greater than 0xFE000000.

2.2. Propagation of PRLA in IS-IS

ISIS L1/L2 routers may wish to advertise received PRLAs into other areas (upwards and/or downwards). When propagating PRLAs the original metric value MUST be preserved. The cost to reach the originator of the received PRLA MUST NOT be considered when readvertising the PRLA.

3. Supporting PRLA in OSPF

[RFC2328] Appendix B defines the following architectural constant for OSPF:

"LSInfinity The metric value indicating that the destination described by an LSA is unreachable. Used in summary-LSAs and AS-external-LSAs as an alternative to premature aging (see Section 14.1). It is defined to be the 24-bit binary value of all ones: 0xffffff."

[RFC5340] Appendix B states:

"Architectural constants for the OSPF protocol are defined in Appendix B of OSPFV2."

indicating that these same constants are applicable to OSPFv3.

[RFC2328] section 14.1. also describes the usage of LSInfinity as a way to indicate loss of prefix reachability:

"Premature aging can also be used when, for example, one of the router's previously advertised external routes is no longer reachable. In this circumstance, the router can flush its AS-external-LSA from the routing domain via premature aging. This procedure is preferable to the alternative, which is to originate a new LSA for the destination specifying a metric of LSInfinity."

3.1. Advertisement of PRLA in OSPF

Using the existing mechanism already defined in the standards, as described in previous section, an advertisement of the inter-area or external prefix inside OSPF or OSPFv3 LSA that has the age set to value lower than MaxAge and metic set to LSInfinity can be interpreted by the receiver as a PRLA.

Existing nodes in a network which receive PRLA advertisements will propagate it following existing standard procedures defined by OSPF.

OSPF Area Border Routers (ABRs), which would be responsible for propagating PRLA advertisements into other areas would need to recognize such advertisements.

3.2. Propagation of PRLA in OSPF

OSPF ABRs may wish to advertise received PRLAs into other connected areas. When doing so, the original LSInfinity metric value in PRLA MUST be preserved. The cost to reach the originator of the received PRLA MUST NOT be considered when readvertising the PRLA to connected areas.

4. Deployment Considerations for PRLA

The economy provided by the use of summary advertisements diminishes in the presence of PRLA. It is therefore recommended that implementations limit the number of PRLA advertisements which can be originated at a given time. This implies that PRLA can be used to signal the loss of reachablity to a modest number of nodes - but it is not a good tool to signal the loss of many nodes simultaneously.

The intent of PRLA is to provide an event driven signal of the transition of a destination from reachable to unreachable. It is not intended to advertise a persistent state. PRLA advertisements should therefore be withdrawn after a modest amount of time, that would provides sufficient time for PRLA to be flooded network-wide and acted upon by receiving nodes, but limits the presence of PRLA in the network to a short time period. The time the PRLA is kept in the network SHOULD also reflect the intended use-case for which the PRLA was advertised.

As PRLA advertisements in ISIS are advertised in existing Link State PDUs (LSPs) and the unit of flooding in IS-IS is an LSP, it is recommended that, when possible, PRLAs are advertised in LSPs dedicated to this type of advertisement. This will minimize the number of LSPs which need to be updated when PRLAs are advertised and withdrawn.

In OSPF and OSPFv3, each inter-area and external prefix is advertised in it's own LSA, so the above optimisation does not apply to OSPF.

5. IANA Considerations

This document makes no requests to IANA.

6. Security Considerations

The use of PRLAs introduces the possibility that an attacker could inject a false, but apparently valid, PRLA. However, the risk of

this occurring is no greater than the risk today of an attacker injecting any other type of false advertisement .

The risks can be reduced by the use of existing security extensions as described in [RFC5304] and [RFC5310] for IS-IS, in [RFC2328][and [RFC7474] for OSPFv2, and in [RFC5340] and [RFC4552] for OSPFv3.

7. Acknowledgements

TBD

8. Normative References

[ISO10589] ISO, "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)", 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, <https://www.rfc-editor.org/info/rfc2119>.

[RFC2328] Moy, J., "OSPF Version 2", STD 54, RFC 2328, DOI 10.17487/RFC2328, April 1998, <https://www.rfc-editor.org/info/rfc2328>.

[RFC4552] Gupta, M. and N. Melam, "Authentication/Confidentiality for OSPFv3", RFC 4552, DOI 10.17487/RFC4552, June 2006, <https://www.rfc-editor.org/info/rfc4552>.

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

[RFC5305] Li, T. and H. Smit, "IS-IS Extensions for Traffic Engineering", RFC 5305, DOI 10.17487/RFC5305, October 2008, <https://www.rfc-editor.org/info/rfc5305>.

[RFC5308] Hopps, C., "Routing IPv6 with IS-IS", RFC 5308, DOI 10.17487/RFC5308, October 2008, <https://www.rfc-editor.org/info/rfc5308>.

[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, <https://www.rfc-editor.org/info/rfc5310>.

[RFC5340] Coltun, R., Ferguson, D., Moy, J., and A. Lindem, "OSPF for IPv6", RFC 5340, DOI 10.17487/RFC5340, July 2008, <https://www.rfc-editor.org/info/rfc5340>.

[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, <https://www.rfc-editor.org/info/rfc7474>.

[RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in RFC 2119 Key Words", BCP 14, RFC 8174, DOI 10.17487/RFC8174, May 2017, <https://www.rfc-editor.org/info/rfc8174>.

Authors' Addresses

Peter Psenak (editor)
Cisco Systems
Pribinova Street 10
Bratislava 81109
Slovakia

Email: ppsenak@cisco.com

Les Ginsberg
Cisco Systems
821 Alder Drive
Milpitas, CA 95035
United States of America

Email: ginsberg@cisco.com

Daniel Voyer
Bell Canada

Email: daniel.voyer@bell.ca