

Internet Engineering Task Force (IETF)
Request for Comments: 6213
Category: Standards Track
ISSN: 2070-1721

C. Hopps
L. Ginsberg
Cisco Systems
April 2011

IS-IS BFD-Enabled TLV

Abstract

This document describes a type-length-value (TLV) for use in the IS-IS routing protocol that allows for the proper use of the Bidirectional Forwarding Detection (BFD) protocol. There exist certain scenarios in which IS-IS will not react appropriately to a BFD-detected forwarding plane failure without use of either this TLV or some other method.

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/rfc6213>.

Copyright Notice

Copyright (c) 2011 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. Introduction	2
1.1. Requirements Language	2
2. The Problem	2
3. The Solution	3
3.1. State Definitions	3
3.2. Adjacency Establishment and Maintenance	4
3.3. Advertisement of Topology-Specific IS Neighbors	4
4. Transition	4
5. Graceful Restart	5
6. The BFD-Enabled TLV	5
7. Security Considerations	6
8. IANA Considerations	6
9. Acknowledgements	6
10. Normative References	7

1. Introduction

The Bidirectional Forwarding Detection (BFD) protocol [RFC5880] is a protocol that allows for detection of a forwarding plane failure between two routers. A router can use [RFC5880] to validate that a peer router's forwarding ability is functioning.

One specific application of BFD as described in [RFC5882] is to verify the forwarding ability of an IS-IS [RFC1195] router's adjacencies; however, the method described in [RFC5882] does not allow for certain failure scenarios. We will define a TLV that will allow for proper response to the detection of all forwarding failures where the use of BFD is employed with IS-IS.

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. The Problem

We observe that, in order to allow for mixed use (i.e., some routers running BFD and some not), [RFC5882] does not require a BFD session be established prior to the establishment of an IS-IS adjacency. Thus, if a router A has neighbors B and C, and B does not support BFD, A would still form adjacencies with B and C, and it would only establish a BFD session with C.

The problem with this solution is that it assumes that the transmission and receipt of IS-IS Hellos (IIHs) shares fate with forwarded data packets. This is not a fair assumption to make given that the primary use of BFD is to protect IPv4 (and IPv6) forwarding, and IS-IS does not utilize IPv4 or IPv6 for sending or receiving its hellos.

Thus, if we consider our previous example, and if C is currently experiencing an IPv4 forwarding failure that allows for IIHs to be sent and received, when A first starts (or restarts), A will assume that C simply does not support BFD, will form an adjacency with C, and may incorrectly forward IPv4 traffic through C.

3. The Solution

A simple solution to this problem is for an IS-IS router to advertise that it has BFD enabled on a given interface. It can do this through the inclusion of a TLV in its IIHs as described in this document.

When sending an IIH on a BFD enabled interface, a router that supports this extension MUST include the BFD-enabled TLV in its IIH. The contents of the TLV MUST indicate what topologies/protocols [RFC5120] have been enabled for BFD by including the appropriate Multi-Topology Identifier (MTID)/ Network Layer Protocol Identifier (NLPID) pairs.

When sending an IIH on an interface on which BFD is NOT enabled, a router MUST NOT include the BFD-enabled TLV.

3.1. State Definitions

The following definitions apply to each IS-IS neighbor:

For each locally supported MTID/NLPID pair, an "ISIS_TOPO_NLPID_BFD_REQUIRED" variable is assigned. If BFD is supported by both the local system and the neighbor of the MTID/NLPID, this variable is set to "TRUE". Otherwise, the variable is set to "FALSE".

For each locally supported MTID, an "ISIS_TOPO_BFD_REQUIRED" variable is set to the logical "OR" of all "ISIS_TOPO_NLPID_BFD_REQUIRED" variables associated with that MTID.

An "ISIS_BFD_REQUIRED" variable is set to the logical "AND" of all "ISIS_TOPO_BFD_REQUIRED" variables.

For each locally supported MTID/NLPID pair, an "ISIS_TOPO_NLPID_STATE" variable is assigned. If "ISIS_TOPO_NLPID_BFD_REQUIRED" is "TRUE", this variable follows the BFD session state for that MTID/NLPID ("UP == TRUE"). Otherwise, the variable is set to "TRUE".

For each locally supported topology (MTID), an "ISIS_TOPO_USEABLE" variable is set to the logical "AND" of the set of "ISIS_TOPO_NLPID_STATE" variables associated with that MTID.

An "ISIS_NEIGHBOR_USEABLE" variable is set to the logical "OR" of all "ISIS_TOPO_USEABLE" variables.

3.2. Adjacency Establishment and Maintenance

Whenever "ISIS_BFD_REQUIRED" is "TRUE", the following extensions to the rules for adjacency establishment and maintenance MUST apply:

- o "ISIS_NEIGHBOR_USEABLE" MUST be "TRUE" before the adjacency can transition from "INIT" to "UP" state.
- o When the IS-IS adjacency is "UP" and "ISIS_NEIGHBOR_USEABLE" becomes "FALSE", the IS-IS adjacency MUST transition to "DOWN".
- o On a Point-to-Point circuit whenever "ISIS_NEIGHBOR_USEABLE" is "FALSE", the Three-Way adjacency state MUST be set to "DOWN" in the Point-to-Point Three-Way Adjacency TLV [RFC5303] in all transmitted IIHs.
- o On a LAN circuit whenever "ISIS_NEIGHBOR_USEABLE" is "FALSE", the IS Neighbors TLV advertising the Media Access Control (MAC) address of the neighbor MUST be omitted in all transmitted IIHs.

3.3. Advertisement of Topology-Specific IS Neighbors

The advertisement of a topology-specific IS neighbor (as well as the use of the neighbor in the topology-specific decision process) is determined by the value of "ISIS_TOPO_USEABLE" for each topology. If "ISIS_TOPO_USEABLE" is "TRUE", then the topology-specific neighbor is advertised. If "ISIS_TOPO_USEABLE" is "FALSE", then the topology-specific neighbor is not advertised.

4. Transition

To allow for a non-disruptive transition to the use of BFD, some amount of time should be allowed before bringing down an "UP" adjacency on a BFD enabled interface when the value of "ISIS_BFD_REQUIRED" becomes "TRUE" as a result of the introduction of

the BFD TLV or the modification (by adding a new supported MTID/NLPID) of an existing BFD TLV in a neighbor's IIH. A simple way to do this is to not update the adjacency hold time when receiving such an IIH from a neighbor with whom we have an "UP" adjacency until "ISIS_NEIGHBOR_USEABLE" becomes "TRUE".

If the value of "ISIS_BFD_REQUIRED" becomes "FALSE" as a result of the removal the BFD TLV or the modification (by removing a supported MTID/NLPID) of an existing BFD TLV in a neighbor's IIH, then BFD session establishment is no longer required to maintain the adjacency or transition the adjacency to the "UP" state.

If a BFD session is administratively shut down [RFC5880] and the BFD session state change impacts the value of "ISIS_NEIGHBOR_USEABLE", then IS-IS SHOULD allow time for the corresponding MTID/NLPID to be removed from the neighbor's BFD TLV by not updating the adjacency hold time until "ISIS_BFD_REQUIRED" becomes "FALSE". Note that while this allows a non-disruptive transition, it still enforces consistency between the administrative state of the BFD session and the MTID/NLPID(s) advertised in the BFD TLV. This is necessary to provide consistent behavior regardless of whether the BFD AdminDown state is introduced before or after an IS-IS adjacency "UP" state has been achieved.

5. Graceful Restart

This section describes IS-IS implementation considerations when both IS-IS graceful restart [RFC5306] and BFD are co-deployed.

In cases where BFD shares fate with the control plane, it can be expected that BFD session failure may occur in conjunction with the control-plane restart. In such cases, premature abort of IS-IS graceful restart as a result of BFD session failure is undesirable. Therefore, some mechanism to ignore the BFD session failure for a limited period of time would be beneficial. The issue of the interaction between graceful restart and BFD is described at length in RFC 5882. The implementation of this interaction is outside the scope of this document.

6. The BFD-Enabled TLV

The BFD-enabled TLV is formatted as shown below. The TLV SHALL only be included in an IIH and only when BFD is enabled for one or more supported MTID/protocols on the interface over which the IIH is being sent. The NLPIDs encoded in the TLV are defined in [ISO9577].

Type 148
 Length # of octets in the value field (3 to 255)
 Value 3 octets specifying the MTID/NLPID for each topology/data protocol for which BFD support is enabled

	No. of octets
+-----+ R R R R MTID	2
+-----+ NLPID	1
+-----+ : : :	:
+-----+ R R R R MTID	2
+-----+ NLPID	1
+-----+	

7. Security Considerations

The TLV defined within this document describes an addition to the IS-IS Hello protocol. Inappropriate use of this TLV could prevent an IS-IS adjacency from forming or lead to failure to detect bidirectional forwarding failures -- each of which is a form of denial of service. However, a party who can manipulate the contents of this TLV is already in a position to create such a denial of service by disrupting IS-IS routing in other ways.

Note that the introduction of this TLV has no impact on the use/non-use of authentication either by IS-IS or by BFD.

8. IANA Considerations

The following IS-IS TLV type is defined by this document.

Name	Value	IIH	LSP	SNP	Purge
-----	-----	---	---	---	-----
BFD-Enabled TLV	148	y	n	n	n

The IS-IS TLV Codepoint registry has been updated accordingly.

9. Acknowledgements

The authors wish to thank Jeffrey Haas, Matthew Jones, Dave Katz, Jonathan Moon, Stefano Previdi, Mike Shand, Michael Shiplett, and David Ward for various input on this document.

10. Normative References

- [ISO9577] International Organization for Standardization, "Protocol identification in the network layer(ISO/IEC 9577)", ISO/IEC 9577:1999, Fourth Edition, December 1999.
- [RFC1195] Callon, R., "Use of OSI IS-IS for routing in TCP/IP and dual environments", RFC 1195, December 1990.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [RFC5120] Przygienda, T., Shen, N., and N. Sheth, "M-ISIS: Multi Topology (MT) Routing in Intermediate System to Intermediate Systems (IS-ISs)", RFC 5120, February 2008.
- [RFC5303] Katz, D., Saluja, R., and D. Eastlake, "Three-Way Handshake for IS-IS Point-to-Point Adjacencies", RFC 5303, October 2008.
- [RFC5306] Shand, M. and L. Ginsberg, "Restart Signaling for IS-IS", RFC 5306, October 2008.
- [RFC5880] Katz, D. and D. Ward, "Bidirectional Forwarding Detection (BFD)", RFC 5880, June 2010.
- [RFC5882] Katz, D. and D. Ward, "Generic Application of Bidirectional Forwarding Detection (BFD)", RFC 5882, June 2010.

Authors' Addresses

Christian E. Hopps
Cisco Systems
170 W. Tasman Dr.
San Jose, California 95134
USA
EMail: chopps@cisco.com

Les Ginsberg
Cisco Systems
510 McCarthy Blvd.
Milpitas, California 95035
USA
EMail: ginsberg@cisco.com

