C. Hopps
L. Ginsberg
Cisco Systems

November 18, 2007

IS-IS BFD Enabled TLV
draft-chopps-isis-bfd-tlv-01

Status of this Memo

Copyright Notice

Abstract

This document describes a TLV for use in the IS-IS routing protocol that allows for the proper use of the Bidirectional Forwarding Detection protocol (BFD). 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.

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].

Table of Contents

1. Introduction 3
2. The Problem 3
3. The Solution 3
3.1. Determining Local Significance 4
3.2. Adjacency Establishment and Maintenance 4
3.3. Advertisement of Topology Specific IS Neighbors 5
4. Transition 5
5. Graceful Restart 6
6. The BFD Enabled TLV 6
7. Security Considerations 7
8. IANA Considerations 7
9. Acknowledgements 7
10. References 7
10.1. Normative References 7
10.2. Informative References 8
Authors' Addresses 8
Intellectual Property and Copyright Statements 9

1. Introduction

The Bidirectional Forwarding Detection protocol [I-D.ietf-bfd-base] is a protocol that allows for detection of a forwarding plane failure between two routers. A router can use [I-D.ietf-bfd-base] to validate that a peer router's forwarding ability is functioning.

One specific application of BFD as described in [I-D.ietf-bfd-generic] is to verify the forwarding ability of an IS-IS [RFC1195] router's adjacencies; however, the method described in [I-D.ietf-bfd-generic] 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.

2. The Problem

We observe that to allow for mixed use (i.e., some routers running BFD and some not) [I-D.ietf-bfd-generic] 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 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 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 IS-IS 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, and indeed that is our proposal.

When sending an IIH on a BFD enabled interface, a router which supports this extension MUST include the BFD enabled TLV in its IIH. The contents of the TLV MUST indicate what protocols have been enabled for BFD by including the appropriate NLPID(s).

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

3.1. Determining Local Significance

When receiving an IIH from a neighbor on an interface with BFD enabled, if the IIH contains the BFD enabled TLV the contents of the BFD TLV are examined to determine if they are of local significance. The logic used to determine local significance is impacted by the combination of topologies and NLPIDs supported on each topology by the local system.[I-D.ietf-isis-wg-multi-topology]. We introduce the following definitions:

NLPID_LOCAL_BFD - The set of NLPIDs for which BFD has been locally enabled on an interface.

NLPID_LOCAL_TOPO - The set of NLPIDs supported on a given topology.

NLPID_BFD_TLV - The set of NLPIDs advertised in the BFD TLV in a received IIH.

NLPID_BFD_TOPO - The set of NLPIDs which are common to (NLPID_LOCAL_BFD and NLPID_LOCAL_TOPO and NLPID_BFD_TLV).

IIH_BFD_LSIG - A boolean which is TRUE when there exists at least one topology which is supported by both the local system and the neighbor where NLPID_BFD_TOPO is not empty.

IS-IS_BFD_TOPO_UP - A per topology boolean whose value is TRUE when IIH_BFD_LSIG is TRUE, the topology is supported by both the local system and the neighbor, and either BFD session state for all NLPIDs in the corresponding NLPID_BFD_TOPO set is UP or the NLPID_BFD_TOPO set is empty for that topology.

IS-IS_BFD_UP - A boolean whose value is TRUE when IIH_BFD_LSIG is TRUE and there is at least one topology supported by the local system and the neighbor which has an IS-IS_BFD_TOPO_UP value which is TRUE.

If IIH_BFD_LSIG is FALSE then the contents of the corresponding received BFD TLV are ignored. Note that this includes the case where BFD is not locally enabled on an interface for any NLPID.

3.2. Adjacency Establishment and Maintenance

When IIH_BFD_LSIG is TRUE, the following extensions to the rules for adjacency establishment and maintenance MUST apply:

- IS-IS_BFD_UP state MUST be TRUE before the adjacency can transition from INIT to UP state

- When the IS-IS adjacency is UP and IS-IS_BFD_UP becomes FALSE the IS-IS adjacency MUST transition to DOWN.

- On a Point-to-Point circuit whenever IS-IS_BFD_UP is FALSE, the Three-Way adjacency state MUST be set to DOWN in the Point-to-Point Three Way Adjacency TLV[RFC3373] in all transmitted IIHs.

- On a LAN circuit whenever IS-IS_BFD_UP is FALSE, the IS Neighbors TLV advertising the MAC address of the neighbor MUST be omitted in all transmitted IIHs.

3.3. Advertisement of Topology Specific IS Neighbors

When IIH_BFD_LSIG is TRUE for a given neighbor, 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 IS-IS_BFD_TOPO_UP for each topology. If IS-IS_BFD_TOPO_UP is TRUE then the topology specific neighbor is advertised. If IS-IS_BFD_TOPO_UP 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 IIH_BFD_LSIG becomes TRUE as a result of the introduction of the BFD TLV or the modification (by adding a new supported 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 IS-IS_BFD_UP becomes TRUE.

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

If a BFD session is administratively shut down [I-D.ietf-bfd-base] and the BFD session state change would impact the value of IS-IS_BFD_UP, then IS-IS SHOULD allow time for the corresponding NLPID to be removed from the neighbor's BFD TLV by not updating the adjacency hold time until IIH_BFD_LSIG 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 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

It is worth considering what if anything should be done when IS-IS is gracefully restarting [RFC3847].

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. How this is implemented is beyond the scope of this document. Consult [I-D.ietf-bfd-generic] for further details.

6. The BFD Enabled TLV

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

```
Type   139 (suggested - to be assigned by IANA)
Length # of octets in the value field (1 to 255)
Value  one octet NLPID for each data protocol
       for which BFD support is enabled
```

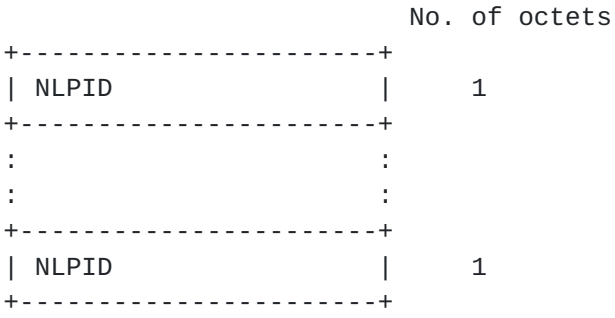

7. Security Considerations

The TLV defined within this document describes an addition to the IS-IS Hello protocol and does not impact the security mechanism of the IS-IS protocol.

8. IANA Considerations

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

Name	Value	IIH	LSP	SNP
BFD Enabled TLV	139	y	n	n

Please update the IS-IS TLV Codepoint Registry accordingly.

Note to RFC Editor: this section may be removed on publication as an RFC.

9. Acknowledgements

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

10. References

10.1. Normative References

[ISO9577] International Organization for Standardization, "Protocol identification in the network layer(ISO/IEC 9577)", ISO/IEC 9577:1999, Fourth Edition, Dec 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.

[RFC3373] Katz, D. and R. Saluja, "Three-Way Handshake for Intermediate System to Intermediate System (IS-IS) Point-to-Point Adjacencies", RFC 3373, September 2002.

10.2. Informative References

[I-D.ietf-bfd-base]
Katz, D. and D. Ward, "Bidirectional Forwarding Detection", draft-ietf-bfd-base-06 (work in progress), March 2007.

[I-D.ietf-bfd-generic]
Katz, D. and D. Ward, "Generic Application of BFD", draft-ietf-bfd-generic-03 (work in progress), March 2007.

[I-D.ietf-bfd-v4v6-1hop]
Katz, D. and D. Ward, "BFD for IPv4 and IPv6 (Single Hop)", draft-ietf-bfd-v4v6-1hop-06 (work in progress), March 2007.

[I-D.ietf-isis-wg-multi-topology]
Przygienda, T., "M-ISIS: Multi Topology (MT) Routing in IS-IS", draft-ietf-isis-wg-multi-topology-12 (work in progress), November 2007.

[RFC3847] Shand, M. and L. Ginsberg, "Restart Signaling for Intermediate System to Intermediate System (IS-IS)", RFC 3847, July 2004.

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, Ca. 95035
USA

Email: ginsberg@cisco.com

Full Copyright Statement

Intellectual Property

Acknowledgment

Funding for the RFC Editor function is provided by the IETF Administrative Support Activity (IASA).