

Network Working Group
Request for Comments: 3373
Category: Informational

D. Katz
Juniper Networks, Inc.
R. Saluja
Tenet Technologies
September 2002

Three-Way Handshake for
Intermediate System to Intermediate System (IS-IS)
Point-to-Point Adjacencies

Status of this Memo

This memo provides information for the Internet community. It does not specify an Internet standard of any kind. Distribution of this memo is unlimited.

Copyright Notice

Copyright (C) The Internet Society (2002). All Rights Reserved.

Abstract

The IS-IS routing protocol (ISO 10589) requires reliable protocols at the link layer for point-to-point links. As a result, it does not use a three-way handshake when establishing adjacencies on point-to-point media. This paper defines a backward-compatible extension to the protocol that provides for a three-way handshake. It is fully interoperable with systems that do not support the extension.

Additionally, the extension allows the robust operation of more than 256 point-to-point links on a single router.

This extension has been implemented by multiple router vendors; this paper is provided as information to the Internet community in order to allow interoperable implementations to be built by other vendors.

1. Terms

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 BCP 14, RFC 2119.

2. Introduction

The IS-IS protocol [1] assumes certain requirements stated in ISO 10589 (section 6.7.2) for the operation of IS-IS over point-to-point links and hence provides only a two-way handshake when establishing

adjacencies on point-to-point links. The protocol does not operate correctly if these subnetwork requirements for point-to-point links are not met. The basic mechanism defined in the standard is that each side declares the other side to be reachable if a Hello packet is heard from it. Once this occurs, each side then sends a Complete Sequence Number PDU (CSNP) to trigger database synchronization.

Three failure modes are known. First, if the link goes down and then comes back up, or one of the systems restarts, and the CSNP packet is lost, and the network has a cut set of one through the link, the link state databases on either side of the link will not synchronize for a full LSP refresh period (up to eighteen hours).

A second, more serious failure, is that if the link fails in only one direction, the failure will only be detected by one of the systems. Normally only one of the two systems will announce the adjacency in its link state packets, and the SPF algorithm will thus ignore the link. However, if there are two parallel links between systems and one of them fails in one direction, SPF will still calculate paths between the two systems, and the system that does not notice the failure will attempt to pass traffic down the failed link (in the direction that does not work).

The third issue is that on some physical layers, the interconnectivity between endpoints can change without causing a link-layer-down condition. In this case, a system may receive packets that are actually destined for a different system (or a different link on the same system). The receiving system may end up thinking that it has an adjacency with the remote system when in fact the remote system is adjacent with a third system.

The solution proposed here ensures correct operation of the protocol over unreliable point-to-point links. As part of the solution to the three-way handshaking issue, a method is defined to remove the limitation of 255 point-to-point interfaces imposed by IS-IS [1]. This method is more robust than the ad hoc methods currently in use.

3. Overview of Extensions

3.1 Handshaking

The intent is to provide a three-way handshake for point-to-point adjacency establishment in a backward compatible fashion. This is done by providing an optional mechanism that allows each system to report its adjacency three-way state; this allows a system to only declare an adjacency to be up if it knows that the other system is receiving its IS-IS Hello (IIH) packets.

The adjacency three-way state can be one of the following types:

Down

This is the initial point-to-point adjacency three-way state. The system has not received any IIH packet containing the three-way handshake option on this point-to-point circuit.

Initializing

The system has received IIH packet containing the three-way handshake option from a neighbor but does not know whether the neighbor is receiving its IIH packet.

Up

The system knows that the neighbor is receiving its IIH packets.

The adjacency three-way state that is reported by this mechanism is not equal or equivalent to the adjacency state that is described in ISO 10589 [1]. If this mechanism is supported then an adjacency may have two states, its state as defined in ISO 10589 [1], and its three-way state. For example according to ISO 10589 [1] receipt of an ISH will cause an adjacency to go to Initializing state; however receipt of an ISH will have no effect on the three-way state of an adjacency, which remains firmly Down until it receives an IIH from a neighbor that contains the three-way handshaking option.

In addition, the neighbor's system ID and (newly-defined) extended circuit ID are reported in order to detect the case where the same stream is being received by multiple systems (only one of which can talk back).

The mechanism is quite similar to the one defined in the Netware Link Services Protocol (NLSP) [2], a variant of IS-IS used for routing IPX traffic. The difference between this mechanism and the one used in NLSP is the location where the information is carried (NLSP uses two of the reserved bits in the IIH header, whereas this solution adds a separate option to the IIH), and the presence of the neighbor's system ID and circuit ID. In theory, using the reserved header bits should be backward compatible, since systems are supposed to ignore them. However, it was felt that this was risky, as the use of untested mechanisms such as this have led to problems in the past in other protocols. New option codes, on the other hand, have been demonstrated to work properly, as the deployment of Integrated IS-IS for IP [3] has done exactly this.

The new mechanism only comes into play when the remote system includes the new option in its IIH packet; if the option is not present, it is assumed that the system does not support the new mechanism, and so the old procedures are used.

3.2 More Than 256 Interfaces

The IS-IS specification has an implicit limit of 256 interfaces, as constrained by the eight bit Circuit ID field carried in various packets. Moderately clever implementors have realized that the only true constraint is that of 256 LAN interfaces, and for that matter only 256 LAN interfaces for which a system is the Designated IS. This is because the only place that the circuit ID is advertised in LSPs is in the pseudonode LSP ID.

Implementors have treated the point-to-point Circuit ID number space as being independent from that of the LAN interfaces, since these Circuit IDs appear only in IIH PDUs and are only used for detection of a change in identity at the other end of a link. More than 256 point-to-point interfaces have been supported by sending the same circuit ID on multiple interfaces. This reduces the robustness of the ID change detection algorithm, since it would then be possible to switch links between interfaces on a system without detecting the change.

Since the Circuit ID is an integral part of the new handshaking mechanism, a backward compatible mechanism for expanding the circuit ID number space is included in this specification.

4. Details

4.1 Syntax

A new IS-IS Option type, "Point-to-Point Three-Way Adjacency", is defined:

x Type - 0xF0 (decimal 240)

x Length - total length of the value field (1 to 17 octets)

x Value -

| | No. of Octets |
|---|---------------|
| +-----+ Adjacency Three-Way State | 1 |
| +-----+ Extended Local Circuit ID | 4 |
| +-----+ Neighbor System ID | ID Length |
| +-----+ Neighbor Extended Local Circuit ID | 4 |
| +-----+ | |

Adjacency Three-Way State

The adjacency three-way state of the point-to-point adjacency. The following values are defined:

- 0 - Up
- 1 - Initializing
- 2 - Down

Extended Local Circuit ID

Unique ID assigned to this circuit when it is created by this Intermediate system.

Neighbor System ID

System ID of neighbor Intermediate system if known. The length of this field is equal to "ID Length" of IIH PDU described in section "Point-to-point IS to IS hello PDU" (section 9.7 of [1]).

Neighbor Extended Local Circuit ID

Extended Local Circuit ID of the other end of the point-to-point adjacency if known.

Any system that supports this mechanism SHALL include this option in its Point-to-Point IIH packets.

Any system that does not understand this option SHALL ignore it, and (of course) SHALL NOT include it in its own IIH packets.

Any system that supports this mechanism MUST include Adjacency Three-Way State field in this option. The other fields in this option SHOULD be included as explained below in section 3.2.

Any system that is able to process this option SHALL follow the procedures below.

4.2 Elements of Procedure

The new handshake procedure is added to the IS-IS point-to-point IIH state machine after the PDU acceptance tests have been performed.

Although the extended circuit ID is only used in the context of the three-way handshake, it is worth noting that it effectively protects against the unlikely event where a link is moved to another interface on a system that has the same local circuit ID, as the received PDUs will be ignored (via the checks defined below) and the existing adjacency will fail.

Add a clause e) to the end of section "Receiving ISH PDUs by an intermediate system" (section 8.2.2 of [1]):

Set the state to be reported in the Adjacency Three-Way State field of the Point-to-Point Three-Way Adjacency option to Down.

Add a clause e) to the end of section "Sending point-to-point IIH PDUs" (section 8.2.3 of [1]):

The IS SHALL include the Point-to-Point Three-Way Adjacency option in the transmitted Point-to-Point IIH PDU. The current three-way state of the adjacency with its neighbor on the link (as defined in new section 8.2.4.1.1 introduced later in the document) SHALL be reported in the Adjacency Three-Way State field. If no adjacency exists, the state SHALL be reported as Down.

The Extended Local Circuit ID field SHALL contain a value assigned by this IS when the circuit is created. This value SHALL be unique among all the circuits of this Intermediate System. The value is not necessarily related to that carried in the Local Circuit ID field of the IIH PDU.

If the system ID and Extended Local Circuit ID of the neighboring system are known (in adjacency three-way state Initializing or Up), the neighbor's system ID SHALL be reported in the Neighbor System ID field, and the neighbor's Extended Local Circuit ID SHALL be reported in the Neighbor Extended Local Circuit ID field.

Add a section 8.2.4.1.1, "Three-Way Handshake", immediately prior to section "IIH PDU Processing" (section 8.2.4.2 of [1]):

A received Point-to-Point IIH PDU may or may not contain the Point-to-Point Three-Way Adjacency option. If it does not, the link is assumed to be functional in both directions, and the procedures described in section 8.2.4.2 are followed.

If the option is present and contains invalid Adjacency Three-Way State, the PDU SHALL be discarded and no further action is taken.

If the option with a valid Adjacency Three-Way State is present, the Neighbor System ID and Neighbor Extended Local Circuit ID fields, if present, SHALL be examined. If they are present, and the Neighbor System ID contained therein does not match the local system's ID, or the Neighbor Extended Local Circuit ID does not match the local system's extended circuit ID, the PDU SHALL be discarded and no further action is taken.

If the Neighbor System ID and Neighbor Extended Local Circuit ID fields match those of the local system, or are not present, the procedures described in section 8.2.4.2 are followed with following changes:

- a) In section 8.2.4.2 a and b, the action "Up" from state tables 5, 6, 7 and 8 may create a new adjacency but the three-way state of the adjacency SHALL be Down.
- b) If the action taken from section 8.2.4.2 a or b is "Up" or "Accept", the IS SHALL perform the action indicated by the new adjacency three-way state table below, based on the current adjacency three-way state and the received Adjacency Three-Way State value from the option. (Note that the procedure works properly if neither field is ever included. This provides backward compatibility to an earlier version of this option.)

| | | Received Adjacency Three-Way State | | |
|-------------------------------|--------------|------------------------------------|--------------|--------|
| | | Down | Initializing | Up |
| adj three -way state | Down | Initialize | Up | Down |
| | Initializing | Initialize | Up | Up |
| | Up | Initialize | Accept | Accept |

Adjacency Three-Way State Table

If the new action is "Down", an adjacencyStateChange(Down) event is generated with the reason "Neighbor restarted" and the adjacency SHALL be deleted.

If the new action is "Initialize", no event is generated and the adjacency three-way state SHALL be set to "Initializing".

If the new action is "Up", an adjacencyStateChange(Up) event is generated.

- c) Skip section 8.2.4.2 c and d.
- d) If the new action is "Initialize", "Up" or "Accept", follow section 8.2.4.2 e.

5. Security Considerations

This document raises no new security issues for IS-IS.

6. References

- [1] ISO, "Intermediate system to Intermediate system routing information exchange protocol for use in conjunction with the Protocol for providing the Connectionless-mode Network Service (ISO 8473)", ISO/IEC 10589:1992.
- [2] "Netware Link Services Protocol Specification, Version 1.0", Novell, Inc., February 1994.
- [3] Callon, R., "OSI IS-IS for IP and Dual Environment", RFC 1195, December 1990.

7. Acknowledgements

The authors would like to thank Tony Li, Henk Smit, Naiming Shen, Dave Ward, Jeff Learman, Les Ginsberg and Philip Christian for their contributions to this document.

8. Authors' Addresses

Dave Katz
Juniper Networks
1194 N. Mathilda Ave.
Sunnyvale, CA 94089

Phone: (408) 745-2073
EMail: dkatz@juniper.net

Rajesh Saluja
Tenet Technologies
30/31, 100 Feet Road, Madiwala
Bangalore - 560 068 INDIA

Phone: +91 80 552 2215
EMail: rajesh.saluja@tenetindia.com

9. Full Copyright Statement

Copyright (C) The Internet Society (2002). All Rights Reserved.

This document and translations of it may be copied and furnished to others, and derivative works that comment on or otherwise explain it or assist in its implementation may be prepared, copied, published and distributed, in whole or in part, without restriction of any kind, provided that the above copyright notice and this paragraph are included on all such copies and derivative works. However, this document itself may not be modified in any way, such as by removing the copyright notice or references to the Internet Society or other Internet organizations, except as needed for the purpose of developing Internet standards in which case the procedures for copyrights defined in the Internet Standards process must be followed, or as required to translate it into languages other than English.

The limited permissions granted above are perpetual and will not be revoked by the Internet Society or its successors or assigns.

This document and the information contained herein is provided on an "AS IS" basis and THE INTERNET SOCIETY AND THE INTERNET ENGINEERING TASK FORCE DISCLAIMS ALL WARRANTIES, EXPRESS OR IMPLIED, INCLUDING BUT NOT LIMITED TO ANY WARRANTY THAT THE USE OF THE INFORMATION HEREIN WILL NOT INFRINGE ANY RIGHTS OR ANY IMPLIED WARRANTIES OF MERCHANTABILITY OR FITNESS FOR A PARTICULAR PURPOSE.

Acknowledgement

Funding for the RFC Editor function is currently provided by the Internet Society.

