

Internet Engineering Task Force (IETF)
Request for Comments: 6822
Category: Standards Track
ISSN: 2070-1721

S. Previdi, Ed.
L. Ginsberg
Cisco Systems
M. Shand

A. Roy
D. Ward
Cisco Systems
December 2012

IS-IS Multi-Instance

Abstract

This document describes a mechanism that allows a single router to share one or more circuits among multiple Intermediate System to Intermediate System (IS-IS) routing protocol instances.

Multiple instances allow the isolation of resources associated with each instance. Routers will form instance-specific adjacencies. Each instance can support multiple topologies. Each topology has a unique Link State Database (LSDB). Each Protocol Data Unit (PDU) will contain a new Type-Length-Value (TLV) identifying the instance and the topology (or topologies) to which the PDU belongs.

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/rfc6822>.

Copyright Notice

Copyright (c) 2012 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.

This document may contain material from IETF Documents or IETF Contributions published or made publicly available before November 10, 2008. The person(s) controlling the copyright in some of this material may not have granted the IETF Trust the right to allow modifications of such material outside the IETF Standards Process. Without obtaining an adequate license from the person(s) controlling the copyright in such materials, this document may not be modified outside the IETF Standards Process, and derivative works of it may not be created outside the IETF Standards Process, except to format it for publication as an RFC or to translate it into languages other than English.

Table of Contents

1. Introduction	3
1.1. Requirements Language	4
2. Elements Of Procedure	4
2.1. Instance Identifier TLV	5
2.2. Instance Membership	6
2.3. Use of Authentication	6
2.4. Adjacency Establishment	6
2.4.1. Point-to-Point Adjacencies	7
2.4.2. Multi-Access Adjacencies	7
2.5. Update Process Operation	7
2.5.1. Update Process Operation on Point-to-Point Circuits	7
2.5.2. Update Process Operation on Broadcast Circuits	7
2.6. Interoperability Considerations	8
2.6.1. Interoperability Issues on Broadcast Circuits	8
2.6.2. Interoperability Using Point-to-Point Circuits	9
3. Usage Guidelines	9
3.1. One-to-One Mapping between Topologies and Instances	10
3.2. Many-to-One Mapping between Topologies and Instances	10
3.3. Considerations for the Number of Instances	11
4. Relationship to M-ISIS	11
5. Graceful Restart Interactions	11
6. IANA Considerations	12
7. Security Considerations	12
8. Acknowledgements	12
9. References	12
9.1. Normative References	13
9.2. Informative References	14

1. Introduction

An existing limitation of the protocol defined by [ISO10589] is that only one instance of the protocol can operate on a given circuit. This document defines an extension to IS-IS to remove this restriction. The extension is referred to as "Multi-Instance IS-IS" (MI-IS-IS).

Routers that support this extension are referred to as "Multi-Instance-capable routers" (MI-RTR).

The use of multiple instances enhances the ability to isolate the resources associated with a given instance both within a router and across the network. Instance-specific prioritization for processing PDUs and performing routing calculations within a router may be specified. Instance-specific flooding parameters may also be defined so as to allow different instances to consume network-wide resources at different rates.

Another existing protocol limitation is that a given instance supports a single Update Process operating on a single Link State Database (LSDB). This document defines an extension to IS-IS to allow non-zero instances of the protocol to support multiple Update Processes. Each Update Process is associated with a topology and a unique topology specific LSDB. Non-zero instances of the protocol are only supported by MI-RTRs. Legacy routers support the standard or zero instance of the protocol. The behavior of the standard instance is not changed in any way by the extensions defined in this document.

MI-IS-IS might be used to support topology-specific routing. When used for this purpose, it is an alternative to Multi-Topology IS-IS [RFC5120].

MI-IS-IS might also be used to support advertisement of information on behalf of applications [RFC6823]. The advertisement of information not directly related to the operation of the IS-IS protocol can therefore be done in a manner that minimizes its impact on the operation of routing.

The above are examples of how MI-IS-IS might be used. The specification of uses of MI-IS-IS is outside the scope of this document.

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. Elements Of Procedure

An Instance Identifier (IID) is introduced to uniquely identify an IS-IS instance. The protocol extension includes a new TLV (IID-TLV) in each IS-IS PDU originated by an MI-RTR except as noted in this document. The IID-TLV identifies the unique instance as well as the topology/topologies to which the PDU applies. Each IS-IS PDU is associated with only one IS-IS instance.

MI-RTRs form instance-specific adjacencies. The IID-TLV included in IS-IS Hellos (IIH) includes the IID and the set of Instance-Specific Topology Identifiers (ITIDs) that the sending IS supports. When multiple instances share the same circuit, each instance will have a separate set of adjacencies.

MI-RTRs support the exchange of topology-specific Link State PDUs for the IID/ITID pairs that each neighbor supports. A unique IS-IS

Update Process (see [ISO10589]) operates for each IID/ITID pair. This MAY also imply IID/ITID-specific routing calculations and IID/ITID-specific routing and forwarding tables. However, this aspect is outside the scope of this specification.

The mechanisms used to implement support of the separation of IS-IS instances and topology-specific Update Processes within a router are outside the scope of this specification.

2.1. Instance Identifier TLV

A new TLV is defined in order to convey the IID and ITIDs supported. The IID-TLV associates a PDU with an IS-IS instance using a unique 16-bit number. The IID-TLV is carried in all IS-IS PDUs that are associated with a non-zero instance; this includes IIHs, Sequence Number PDUs (SNPs), and Link State PDUs (LSPs).

Multiple instances of IS-IS may coexist on the same circuit and on the same physical router. IIDs MUST be unique within the same routing domain.

IID #0 is reserved for the standard instance supported by legacy systems. IS-IS PDUs associated with the standard instance MUST NOT include an IID-TLV except where noted in this document.

The IID-TLV MAY include one or more ITIDs. An ITID is a 16-bit identifier where all values (0 - 65535) are valid.

The following format is used for the IID-TLV:

Type: 7
Length: 2 - 254
Value:

		No. of octets
+-----+		
IID (0 - 65535)		2
+-----+		
Supported ITID		2
+-----+		
:	:	
+-----+		
Supported ITID		2
+-----+		

When the IID = 0, the list of supported ITIDs MUST NOT be present.

An IID-TLV with IID = 0 MUST NOT appear in an SNP or LSP. When the TLV appears (with a non-zero IID) in an SNP or LSP, exactly one ITID

MUST be present indicating the topology with which the PDU is associated. If no ITIDs or multiple ITIDs are present or the IID is zero, then the PDU MUST be ignored.

When the IID is non-zero and the TLV appears in an IIH, the set of ITIDs supported on the circuit over which the IIH is sent is included. There MUST be at least one ITID present.

Multiple IID-TLVs MAY appear in IIHs. If multiple IID-TLVs are present and the IID value in all IID-TLVs is not the same, then the PDU MUST be ignored.

A single IID-TLV will support advertisement of up to 126 ITIDs. If multiple IID-TLVs are present in an IIH PDU, the supported set of ITIDs is the union of all ITIDs present in all IID-TLVs.

When an LSP purge is initiated, the IID-TLV MUST be retained, but the remainder of the body of the LSP SHOULD be removed. The purge procedure is described in [RFC6233] and [RFC6232].

A PDU without an IID-TLV belongs to the standard instance.

2.2. Instance Membership

Each MI-RTR is configured to be participating in one or more instances of IS-IS. For each non-zero instance in which it participates, an MI-RTR marks IS-IS PDUs (IIHs, LSPs, or SNPs) generated that pertain to that instance by including the IID-TLV with the appropriate instance identifier.

2.3. Use of Authentication

When authentication is in use, the IID, if present, is first used to select the authentication configuration that is applicable. The authentication check is then performed as normal. When multiple ITIDs are supported, ITID-specific authentication MAY be used in SNPs and LSPs.

2.4. Adjacency Establishment

In order to establish adjacencies, IS-IS routers exchange IIH PDUs. Two types of adjacencies exist in IS-IS: point-to-point and broadcast. The following subsections describe the additional rules an MI-RTR MUST follow when establishing adjacencies.

2.4.1. Point-to-Point Adjacencies

MI-RTRs include the IID-TLV in the point-to-point Hello PDUs they originate. Upon reception of an IIH, an MI-RTR inspects the received IID-TLV and if the IID matches any of the IIDs that the router supports on that circuit, normal adjacency establishment procedures are used to establish an instance-specific adjacency. Note that the absence of the IID TLV implies IID #0. For instances other than IID #0, an adjacency SHOULD NOT be established unless there is at least one ITID in common.

This extension allows an MI-RTR to establish multiple adjacencies to the same physical neighbor over a point-to-point circuit. However, as the instances are logically independent, the normal expectation of at most one neighbor on a given point-to-point circuit still applies.

2.4.2. Multi-Access Adjacencies

Multi-Access (broadcast) circuits behave differently than point-to-point in that PDUs sent by one router are visible to all routers and all routers must agree on the election of a Designated Intermediate System (DIS) independent of the set of ITIDs supported.

MI-RTRs will establish adjacencies and elect a DIS per IS-IS instance. Each MI-RTR will form adjacencies only with routers that advertise support for the instances that the local router has been configured to support on that circuit. Since an MI-RTR is not required to support all possible instances on a LAN, it's possible to elect a different DIS for different instances.

2.5. Update Process Operation

For non-zero instances, a unique Update Process exists for each supported ITID.

2.5.1. Update Process Operation on Point-to-Point Circuits

On Point-to-Point circuits -- including Point-to-Point Operation over LAN [RFC5309] -- the ITID-specific Update Process only operates on that circuit for those ITIDs that are supported by both ISs operating on the circuit.

2.5.2. Update Process Operation on Broadcast Circuits

On broadcast circuits, a single DIS is elected for each supported IID independent of the set of ITIDs advertised in LAN IIHs. This requires that the DIS generate pseudo-node LSPs for all supported ITIDs and that the Update Process for all supported ITIDs operate on

the broadcast circuit. Among MI-RTRs operating on a broadcast circuit, if the set of supported ITIDs for a given non-zero IID is inconsistent, connectivity for the topology (or topologies) associated with the ITIDs not supported by some MI-RTRs can be compromised.

2.6. Interoperability Considerations

[ISO10589] requires that any TLV that is not understood is silently ignored without compromising the processing of the whole IS-IS PDU (IIH, LSP, SNP).

To a router not implementing this extension, all IS-IS PDUs received will appear to be associated with the standard instance regardless of whether an IID TLV is present in those PDUs. This can cause interoperability issues unless the mechanisms and procedures discussed below are followed.

2.6.1. Interoperability Issues on Broadcast Circuits

In order for routers to correctly interoperate with routers not implementing this extension and in order not to cause disruption, a specific and dedicated Media Access Control (MAC) address is used for multicasting IS-IS PDUs with any non-zero IID. Each level will use a specific layer 2 multicast address. Such an address allows MI-RTRs to exchange IS-IS PDUs with non-zero IIDs without these PDUs being processed by legacy routers, and therefore no disruption is caused.

An MI-RTR will use the AllL1IS or AllL2IS ISIS MAC-layer address (as defined in [ISO10589]) as the destination address when sending an IS-IS PDU for the standard instance. An MI-RTR will use one of two new dedicated layer 2 multicast addresses (AllL1MI-ISs or AllL2MI-ISs) as the destination address when sending an IS-IS PDU for any non-zero IID. These addresses are specified in Section 6. If operating in point-to-point mode on a broadcast circuit [RFC5309], an MI-RTR MUST use one of the two new multicast addresses as the destination address when sending point-to-point IIHs associated with a non-zero instance. (Either address will do.)

MI-RTRs MUST discard IS-IS PDUs received if either of the following is true:

- o The destination multicast address is AllL1IS or AllL2IS and the PDU contains an IID-TLV.
- o The destination multicast address is one of the two new addresses, and the PDU contains an IID-TLV with a zero value for the IID or has no IID-TLV.

NOTE: If the multicast addresses AllL1IS and/or AllL2IS are improperly used to send IS-IS PDUs for non-zero IIDs, legacy systems will interpret these PDUs as being associated with IID #0. This will cause inconsistencies in the LSDB in those routers, may incorrectly maintain adjacencies, and may lead to inconsistent DIS election.

2.6.2. Interoperability Using Point-to-Point Circuits

In order for an MI-RTR to interoperate over a point-to-point circuit with a router that does NOT support this extension, the MI-RTR MUST NOT send IS-IS PDUs for instances other than IID #0 over the point-to-point circuit as these PDUs may affect the state of IID #0 in the neighbor.

The presence or absence of the IID-TLV in an IIH indicates that the neighbor does or does not support this extension, respectively. Therefore, all IIHs sent on a point-to-point circuit by an MI-RTR MUST include an IID-TLV. This includes IIHs associated with IID #0. Once it is determined that the neighbor does not support this extension, an MI-RTR MUST NOT send PDUs (including IIHs) for instances other than IID #0.

Until an IIH is received from a neighbor, an MI-RTR MAY send IIHs for a non-zero instance. However, once an IIH with no IID TLV has been received -- indicating that the neighbor is not an MI-RTR -- the MI-RTR MUST NOT send IIHs for a non-zero instance. The temporary relaxation of the restriction on sending IIHs for non-zero instances allows a non-zero instance adjacency to be established on an interface on which an MI-RTR does NOT support the standard instance.

Point-to-point adjacency setup MUST be done through the use of the three-way handshaking procedure as defined in [RFC5303] in order to prevent a non-MI capable neighbor from bringing up an adjacency prematurely based on reception of an IIH with an IID-TLV for a non-zero instance.

3. Usage Guidelines

As discussed above, MI-IS-IS extends IS-IS to support multiple instances on a given circuit. Each instance is uniquely identified by the IID and forms instance-specific adjacencies. Each instance supports one or more topologies as represented by the ITIDs. All topologies associated with a given instance share the instance-specific adjacencies. The set of topologies supported by a given IID MAY differ from circuit to circuit. Each topology has its own set of LSPs and runs a topology-specific Update Process. Flooding of topology-specific LSPs is only performed on circuits on which both the local router and the neighbor(s) support a given topology (i.e.,

advertise the same ITID in the set of supported ITIDs sent in the IID-TLV included in IIHs).

The following subsections provide some guidelines for usage of instances and topologies within each instance. While this represents examples based on the intent of the authors, implementors are not constrained by the examples.

3.1. One-to-One Mapping between Topologies and Instances

When the set of information to be flooded in LSPs is intended to be flooded to all MI-RTRs supporting a given IID, a single topology MAY be used. The information contained in the single LSDB MAY still contain information associated with multiple applications as the GENINFO TLV for each application has an application-specific ID that identifies the application to which the TLV applies [RFC6823].

3.2. Many-to-One Mapping between Topologies and Instances

When the set of information to be flooded in LSPs includes subsets that are of interest to a subset of the MI-RTRs supporting a given IID, support of multiple ITIDs allows each subset to be flooded only to those MI-RTRs that are interested in that subset. In the simplest case, a one-to-one mapping between a given application and an ITID allows the information associated with that application to be flooded only to MI-RTRs that support that application -- but a many-to-one mapping between applications and a given ITID is also possible. When the set of application-specific information is large, the use of multiple ITIDs provides significantly greater efficiencies, as MI-RTRs only need to maintain the LSDB for applications of interest and that information only needs to be flooded over a topology defined by the MI-RTRs who support a given ITID.

The use of multiple ITIDs also allows the dedication of a full LSP set (256 LSPs at each level) for the use of a given (set of) applications, thereby minimizing the possibility of exceeding the carrying capacity of an LSP set. Such a possibility might arise if information for all applications were to be included in a single LSP set.

Note that the topology associated with each ITID MUST be fully connected in order for ITID-specific LSPs to be successfully flooded to all MI-RTRs that support that ITID.

3.3. Considerations for the Number of Instances

The support of multiple topologies within the context of a single instance provides better scalability in support of multiple applications both in terms of the number of adjacencies that are required and in the flooding of topology-specific LSDB. In many cases, the use of a single non-zero instance would be sufficient and optimal. However, in cases where the set of topologies desired in support of a set of applications is largely disjoint from the set of topologies desired in support of a second set of applications, it could make sense to use multiple instances.

4. Relationship to M-ISIS

[RFC5120] defines support for multi-topology routing. In that document, 12-bit Multi-Topology Identifiers (MTIDs) are defined to identify the topologies that an IS-IS instance (a "standard instance" as defined by this document) supports. There is no relationship between the Multi-topology IDs defined in [RFC5120] and the ITIDs defined in this document.

If an MI-RTR uses the extensions in support of the BFD-Enabled TLV [RFC6213], the ITID SHOULD be used in place of the MTID, in which case all 16 bits of the identifier field are usable.

An MI-RTR MAY use the extensions defined in this document to support multiple topologies in the context of an instance with a non-zero IID. Each MI topology is associated with a unique LSDB identified by an ITID. An ITID-specific IS-IS Update Process operates on each topology. This differs from [RFC5120] where a single LSDB or single IS-IS Update Process is used in support of all topologies.

An MI-RTR MUST NOT support [RFC5120] multi-topology within a non-zero instance. The following TLVs MUST NOT be sent in an LSP associated with a non-zero instance and MUST be ignored when received:

- TLV 222 - MT IS Neighbors
- TLV 235 - MT IP Reachability
- TLV 237 - MT IPv6 Reachability

5. Graceful Restart Interactions

[RFC5306] defines protocol extensions in support of graceful restart of a routing instance. The extensions defined there apply to MI-RTRs with the notable addition that as there are topology-specific LSP databases all of the topology-specific LSP databases must be synchronized following restart in order for database synchronization

to be complete. This involves the use of additional T2 timers. See [RFC5306] for further details.

6. IANA Considerations

Per this document, IANA has registered a new IS-IS TLV, which is reflected in the "IS-IS TLV Codepoints" registry:

Type	Description	IIH	LSP	SNP	Purge
----	-----	---	---	---	-----
7	Instance Identifier	Y	Y	Y	Y

Per this document, IANA has registered two EUI-48 multicast addresses from the IANA-managed EUI address space as specified in [RFC5342]. The addresses are as follows:

```
01-00-5E-90-00-02 AllL1MI-ISs
01-00-5E-90-00-03 AllL2MI-ISs
```

7. Security Considerations

Security concerns for IS-IS are addressed in [ISO10589], [RFC5304], and [RFC5310].

8. Acknowledgements

The authors would like to acknowledge contributions made by Dino Farinacci and Tony Li.

9. References

9.1. 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, 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.
- [RFC5304] Li, T. and R. Atkinson, "IS-IS Cryptographic Authentication", RFC 5304, October 2008.
- [RFC5306] Shand, M. and L. Ginsberg, "Restart Signaling for IS-IS", RFC 5306, October 2008.
- [RFC5310] Bhatia, M., Manral, V., Li, T., Atkinson, R., White, R., and M. Fanto, "IS-IS Generic Cryptographic Authentication", RFC 5310, February 2009.
- [RFC6213] Hopps, C. and L. Ginsberg, "IS-IS BFD-Enabled TLV", RFC 6213, April 2011.
- [RFC6232] Wei, F., Qin, Y., Li, Z., Li, T., and J. Dong, "Purge Originator Identification TLV for IS-IS", RFC 6232, May 2011.
- [RFC6233] Li, T. and L. Ginsberg, "IS-IS Registry Extension for Purges", RFC 6233, May 2011.
- [RFC6823] Ginsberg, L., Previdi, S., and M. Shand, "Advertising Generic Information in IS-IS", RFC 6823, December 2012.

9.2. Informative References

- [RFC5309] Shen, N. and A. Zinin, "Point-to-Point Operation over LAN in Link State Routing Protocols", RFC 5309, October 2008.
- [RFC5342] Eastlake, D., "IANA Considerations and IETF Protocol Usage for IEEE 802 Parameters", BCP 141, RFC 5342, September 2008.

Authors' Addresses

Stefano Previdi (editor)
Cisco Systems
Via Del Serafico 200
Rome 0144
Italy

EMail: sprevidi@cisco.com

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

EMail: ginsberg@cisco.com

Mike Shand

EMail: imc.shand@gmail.com

Abhay Roy
Cisco Systems
170 W. Tasman Dr.
San Jose, CA 95134
USA

EMail: akr@cisco.com

Dave Ward
Cisco Systems
3700 Cisco Way
San Jose, CA 95134
USA

EMail: wardd@cisco.com

