

Network Working Group S.  
Previdi  
Internet-Draft L.  
Ginsberg  
Intended status: Standards Track M.  
Shand  
Expires: April 18, 2010 D.  
Ward A.  
  
Roy  
  
Cisco  
Systems  
  
October 15,  
  
2009

**IS-IS Multi-Instance  
draft-ietf-isis-mi-02**

Status of this Memo

This Internet-Draft is submitted to IETF in full conformance with the provisions of [BCP 78](#) and [BCP 79](#). 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.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

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."

The list of current Internet-Drafts can be accessed at <http://www.ietf.org/ietf/1id-abstracts.txt>.

The list of Internet-Draft Shadow Directories can be accessed at <http://www.ietf.org/shadow.html>.

This Internet-Draft will expire on April 18, 2010.

Copyright Notice

Copyright (c) 2009 IETF Trust and the persons identified as the

Previdi, et al.  
1]

Expires April 18, 2010

[Page

document authors. All rights reserved.

This document is subject to [BCP 78](#) and the IETF Trust's Legal Provisions Relating to IETF Documents in effect on the date of publication of this document (<http://trustee.ietf.org/license-info>). Please review these documents carefully, as they describe your rights and restrictions with respect to this document.

#### Abstract

This draft describes a mechanism that allows a single router to share one or more links among multiple IS-IS routing protocol instances.

Multiple instances allow the isolation of resources associated with each instance. Routers will form instance specific adjacencies, exchange instance specific routing updates and compute paths utilizing instance specific LSDB information. Each PDU will contain a new TLV identifying the instance to which the PDU belongs. This allows a network operator to deploy multiple IS-IS instances in parallel, using the same set of links when required and still have the capability of computing instance specific paths. This draft does not address the forwarding paradigm that needs to be used in order to ensure data PDUs are forwarded according to the paths computed by a specific instance.

#### 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

<a href="#">1.</a>	Introduction . . . . .	
<a href="#">4</a>		
<a href="#">2.</a>	Elements Of Procedure . . . . .	
<a href="#">4</a>		
<a href="#">2.1.</a>	Instance Identifier . . . . .	
<a href="#">5</a>		
<a href="#">2.2.</a>	Instance Membership . . . . .	
<a href="#">5</a>		
<a href="#">2.3.</a>	Adjacency Establishment . . . . .	
<a href="#">5</a>		
<a href="#">2.3.1.</a>	Point-to-Point Adjacencies . . . . .	
<a href="#">5</a>		
<a href="#">2.3.2.</a>	Multi-Access Adjacencies . . . . .	
<a href="#">6</a>		
<a href="#">2.4.</a>	Interoperability Considerations . . . . .	
<a href="#">6</a>		
<a href="#">2.4.1.</a>	Interoperability Issues on Broadcast Networks . . . . .	
<a href="#">6</a>		
<a href="#">2.4.2.</a>	Interoperability using p2p networks . . . . .	
<a href="#">7</a>		
<a href="#">3.</a>	IANA Considerations . . . . .	
<a href="#">7</a>		
<a href="#">4.</a>	Security Considerations . . . . .	
<a href="#">8</a>		
<a href="#">5.</a>	Informational References . . . . .	
<a href="#">8</a>		
<a href="#">6.</a>	Acknowledgements . . . . .	
<a href="#">8</a>		
<a href="#">7.</a>	Normative References . . . . .	
<a href="#">8</a>		
<a href="#">8</a>	Authors' Addresses . . . . .	
<a href="#">8</a>		

Previdi, et al.  
3]

Expires April 18, 2010

[Page

## 1. Introduction

An existing limitation of the protocol defined by [\[IS-IS\]](#) is that only one instance of the protocol can operate on a given link. 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 which 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.

MI-IS-IS might be used to support IS-IS for multiple topologies. When used for this purpose it is an alternative to [\[MT-IS-IS\]](#).

MI-IS-IS might also be used to support an instance which advertises information on behalf of applications. The advertisement of information not directly related to the operation of the IS-IS protocol can therefore be done in a manner which 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.

## 2. Elements Of Procedure

The protocol extension uses a new TLV called the Instance Identifier (IID) that is included in each IS-IS PDU originated by an MI-RTR. MI-RTRs form instance specific adjacencies and exchange instance specific routing updates only for the instance IDs which are supported both by the MI-RTR and its neighbor.

This also implies an instance specific flooding scheme, instance specific LSDBs and instance specific routing calculations. It MAY also imply instance specific routing and forwarding tables.

However,

this aspect is outside the scope of this specification. When multiple instances share the same link each instance will have a separate set of adjacencies. Each IS-IS PDU is associated with only one IS-IS instance.





The mechanisms used to implement support for the separation of IS-IS instances within a router are outside the scope of this specification.

### **2.1. Instance Identifier**

A new TLV is defined in order to convey an instance identifier (IID).

The purpose of the IID is to identify the PDUs associated with each IS-IS instance using a unique 16-bit number. The IID TLV is carried in all IS-IS PDUs (IIHs, SNPs and LSPs) originated by the router.

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

Instance identifier #0 is reserved for the standard instance supported by legacy systems.

The following format is used for the IID:

Type: 7  
Length: 2  
Value: Instance Identifier (0 to 65535)

### **2.2. Instance Membership**

Each router is configured to be participating in one or more instances of IS-IS. For each instance in which it participates, a router marks all IS-IS PDUs (IIHs, LSPs or SNPs) generated pertaining

to that instance by including the IID TLV with the appropriate instance identifier. Note that this applies to the standard instance

(instance identifier #0). A PDU MUST NOT be generated with multiple IID TLVs. PDUs received with multiple IID TLVs MUST be ignored. A PDU without an IID TLV is assumed to belong to the standard instance (#0).

### **2.3. 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 sub-sections describe the additional rules an MI-RTR MUST follow when establishing adjacencies.

#### **2.3.1. Point-to-Point Adjacencies**

MI-RTRs include the IID TLV in the p2p hello PDUs they originate. Upon reception of an IIH, an MI-RTR inspects the received IID TLV and  
if it matches any of the IIDs which the router supports on that

link,

Previdi, et al.  
5]

Expires April 18, 2010

[Page

normal adjacency establishment procedures are used to establish an instance specific adjacency. Note that the absence of the IID TLV implies instance ID #0.

This extension allows an MI-RTR to establish multiple adjacencies to the same physical neighbor over a p2p link. However, as the instances are logically independent, the normal expectation of at most one neighbor on a given p2p link still applies.

### **2.3.2. Multi-Access Adjacencies**

Multi-Access (broadcast) networks behave differently than p2p in that

PDU's sent by one router are visible to all routers and all routers must agree on the election of a DIS.

MI-RTRs will establish adjacencies and elect a DIS per IS-IS instance. Each MI-RTR will form adjacencies only with routers which advertise support for the instances which the local router has been configured to support on that link. 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.

### **2.4. Interoperability Considerations**

[IS-IS] 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 PDU's received will appear to be associated with the standard instance regardless of whether an IID TLV is present in those PDU's. This can cause interoperability issues unless the mechanisms and procedures discussed below are followed.

#### **2.4.1. Interoperability Issues on Broadcast Networks**

In order for routers to correctly interoperate with routers not implementing this extension and in order not to cause disruption, a specific and dedicated MAC address is used for multicasting IS-IS PDU's 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 PDU's with non-zero IIDs without these PDU's being processed by legacy routers and therefore no disruption is caused.

An MI-RTR will use the AllL1IS and AllL2IS IS-IS mac layer addresses (as defined in [\[IS-IS\]](#)) when sending IS-IS PDU's for the standard instance (IID #0). An MI-RTR will use two new (TBD) dedicated layer 2 multicast addresses (one for each level) when sending IS-IS PDU's



for any non-zero IID.

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 with non-zero value
- o The destination multicast address is one of the two new addresses and the PDU contains an IID TLV with a zero value 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.

**2.4.2. Interoperability using p2p networks**

In order for an MI-RTR to interoperate over a p2p link with a router which does NOT support this extension, the MI-RTR MUST NOT send IS-IS PDUs for instances other than IID #0 over the p2p link as these PDUs may affect the state of IID #0 in the neighbor.

The presence/absence of the IID TLV in an IIH indicates that the neighbor does/does not support this extension. 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 instance #0.

**3. IANA Considerations**

This document requires the definition a new ISIS TLV that needs to be reflected in the ISIS TLV code-point registry:

Type	Description	IIH	LSP	SNP
----	-----	---	---	---
TBA	MI-MT IID	y	y	y



#### **4. Security Considerations**

Security concerns for IS-IS are addressed in the IS-IS specification [[IS-IS](#)], and accompanying specifications on [[HMAC-MD5](#)]. No additional considerations need to be made for the extension.

#### **5. Informational References**

[IS-IS] 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 10589:2002, Second Edition.

[HMAC-MD5] Li, T. and R. Atkinson, "Intermediate System to Intermediate System (IS-IS) Cryptographic Authentication", [RFC 3567](#), July 2003.

[MT-IS-IS] Pryzgienda, T., Shen, N., and Sheth, N., "Multi Topology (MT) Routing in IS-IS", [RFC5120](#), February 2008.

#### **6. Acknowledgements**

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

#### **7. Normative References**

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), March 1997.

#### Authors' Addresses

Stefano Previdi  
Cisco Systems  
Via Del Serafico 200  
Roma, RM 00142  
Italy

Email: [sprevidi@cisco.com](mailto:sprevidi@cisco.com)





Internet-Draft  
2009

IS-IS Multi-Instance

October

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

Email: [ginsberg@cisco.com](mailto:ginsberg@cisco.com)

Mike Shand  
Cisco Systems  
250 Longwater Avenue  
Reading, Berkshire RG2 6GB  
UK

Email: [mshand@cisco.com](mailto:mshand@cisco.com)

Dave Ward  
Cisco Systems  
170 W. Tasman Dr.  
San Jose, CA 95134  
USA

Email: [dward@cisco.com](mailto:dward@cisco.com)

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

Email: [akr@cisco.com](mailto:akr@cisco.com)

