

Network Working Group
Request for Comments: 4971
Category: Standards Track

JP. Vasseur, Ed.
N. Shen, Ed.
Cisco Systems, Inc.
R. Aggarwal, Ed.
Juniper Networks
July 2007

Intermediate System to Intermediate System (IS-IS) Extensions for Advertising Router Information

Status of This Memo

This document specifies an Internet standards track protocol for the Internet community, and requests discussion and suggestions for improvements. Please refer to the current edition of the "Internet Official Protocol Standards" (STD 1) for the standardization state and status of this protocol. Distribution of this memo is unlimited.

Copyright Notice

Copyright (C) The IETF Trust (2007).

Abstract

This document defines a new optional Intermediate System to Intermediate System (IS-IS) TLV named CAPABILITY, formed of multiple sub-TLVs, which allows a router to announce its capabilities within an IS-IS level or the entire routing domain.

Table of Contents

1. Introduction	2
1.1. Conventions Used in This Document	2
2. IS-IS Router CAPABILITY TLV	3
3. Elements of Procedure	4
4. Interoperability with Routers Not Supporting the Capability TLV	5
5. Security Considerations	6
6. IANA Considerations	6
7. Acknowledgment	6
8. References	6
8.1. Normative References	6
8.2. Informative References	8

1. Introduction

There are several situations where it is useful for the IS-IS [IS-IS] [IS-IS-IP] routers to learn the capabilities of the other routers of their IS-IS level, area, or routing domain. For the sake of illustration, three examples related to MPLS Traffic Engineering (TE) are described here:

1. Mesh-group: the setting up of a mesh of TE Label Switched Paths (LSPs) [IS-IS-TE] requires some significant configuration effort. [AUTOMESH] proposes an auto-discovery mechanism whereby every Label Switching Router (LSR) of a mesh advertises its mesh-group membership by means of IS-IS extensions.
2. Point to Multipoint TE LSP (P2MP LSP). A specific sub-TLV ([TE-NODE-CAP]) allows an LSR to advertise its Point To Multipoint capabilities ([P2MP] and [P2MP-REQS]).
3. Inter-area traffic engineering: Advertisement of the IPv4 and/or the IPv6 Traffic Engineering Router IDs.

The use of IS-IS for Path Computation Element (PCE) discovery may also be considered and will be discussed in the PCE WG.

The capabilities mentioned above require the specification of new sub-TLVs carried within the CAPABILITY TLV defined in this document.

Note that the examples above are provided for the sake of illustration. This document proposes a generic capability advertising mechanism that is not limited to MPLS Traffic Engineering.

This document defines a new optional IS-IS TLV named CAPABILITY, formed of multiple sub-TLVs, which allows a router to announce its capabilities within an IS-IS level or the entire routing domain. The applications mentioned above require the specification of new sub-TLVs carried within the CAPABILITY TLV defined in this document.

Definition of these sub-TLVs is outside the scope of this document.

1.1. Conventions Used in This Document

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 [RFC-2119].

2. IS-IS Router CAPABILITY TLV

The IS-IS Router CAPABILITY TLV is composed of 1 octet for the type, 1 octet that specifies the number of bytes in the value field, and a variable length value field that starts with 4 octets of Router ID, indicating the source of the TLV, and followed by 1 octet of flags.

A set of optional sub-TLVs may follow the flag field. Sub-TLVs are formatted as described in RFC 3784 [IS-IS-TE].

TYPE: 242

LENGTH: from 5 to 255

VALUE:

Router ID (4 octets)

Flags (1 octet)

Set of optional sub-TLVs (0-250 octets)

Flags

```

0 1 2 3 4 5 6 7
+-----+-----+
| Reserved | D | S |
+-----+-----+

```

Currently two bit flags are defined.

S bit (0x01): If the S bit is set(1), the IS-IS Router CAPABILITY TLV MUST be flooded across the entire routing domain. If the S bit is not set(0), the TLV MUST NOT be leaked between levels. This bit MUST NOT be altered during the TLV leaking.

D bit (0x02): When the IS-IS Router CAPABILITY TLV is leaked from level-2 to level-1, the D bit MUST be set. Otherwise, this bit MUST be clear. IS-IS Router capability TLVs with the D bit set MUST NOT be leaked from level-1 to level-2. This is to prevent TLV looping.

The Router CAPABILITY TLV is OPTIONAL. As specified in Section 3, more than one Router CAPABILITY TLV from the same source MAY be present.

This document does not specify how an application may use the Router Capability TLV and such specification is outside the scope of this document.

3. Elements of Procedure

A router that generates a CAPABILITY TLV MUST have a Router ID that is a 32-bit number. The ID MUST be unique within the IS-IS area. If the router generates any capability TLVs with domain flooding scope, then the ID MUST also be unique within the IS-IS routing domain.

When advertising capabilities with different flooding scopes, a router MUST originate a minimum of two Router CAPABILITY TLVs, each TLV carrying the set of sub-TLVs with the same flooding scope. For instance, if a router advertises two sets of capabilities, C1 and C2, with an area/level scope and routing domain scope respectively, C1 and C2 being specified by their respective sub-TLV(s), the router will originate two Router CAPABILITY TLVs:

- One Router CAPABILITY TLV with the S flag cleared, carrying the sub-TLV(s) relative to C1. This Router CAPABILITY TLV will not be leaked into another level.
- One Router CAPABILITY TLV with the S flag set, carrying the sub-TLV(s) relative to C2. This Router CAPABILITY TLV will be leaked into other IS-IS levels. When the TLV is leaked from level-2 to level-1, the D bit will be set in the level-1 LSP advertisement.

In order to prevent the use of stale capabilities, a system MUST NOT use a Capability TLV present in an LSP of a system that is not currently reachable via Level-x paths, where "x" is the level (1 or 2) in which the sending system advertised the TLV. This requirement applies regardless of whether or not the sending system is the originator of the Capabilities TLV. Note that leaking a Capabilities TLV is one of the uses that is prohibited under these conditions.

Example: If Level-1 router A generates a Capability TLV and floods it to two L1/L2 routers, S and T, they will flood it into the Level-2 domain. Now suppose the Level-1 area partitions, such that A and S are in one partition and T is in another. IP routing will still continue to work, but if A now issues a revised version of the CAP TLV, or decides to stop advertising it, S will follow suit, but T will continue to advertise the old version until the LSP times out.

Routers in other areas have to choose whether to trust T's copy of A's capabilities or S's copy of A's information and, they have no reliable way to choose. By making sure that T stops leaking A's information, this removes the possibility that other routers will use stale information from A.

In IS-IS, the atomic unit of the update process is a TLV -- or more precisely, in the case of TLVs that allow multiple entries to appear in the value field (e.g., IS-neighbors), the atomic unit is an entry in the value field of a TLV. If an update to an entry in a TLV is advertised in an LSP fragment different from the LSP fragment associated with the old advertisement, the possibility exists that other systems can temporarily have either 0 copies of a particular advertisement or 2 copies of a particular advertisement, depending on the order in which new copies of the LSP fragment that had the old advertisement and the fragment that has the new advertisement arrive at other systems.

Wherever possible, an implementation SHOULD advertise the update to a capabilities TLV in the same LSP fragment as the advertisement that it replaces. Where this is not possible, the two affected LSP fragments should be flooded as an atomic action.

Systems that receive an update to an existing capability TLV can minimize the potential disruption associated with the update by employing a holddown time prior to processing the update so as to allow for the receipt of multiple LSP fragments associated with the same update prior to beginning processing.

Where a receiving system has two copies of a capabilities TLV from the same system that have different settings for a given attribute, the procedure used to choose which copy shall be used is undefined.

4. Interoperability with Routers Not Supporting the Capability TLV

Routers that do not support the Router CAPABILITY TLV MUST silently ignore the TLV(s) and continue processing other TLVs in the same LSP. Routers that do not support specific sub-TLVs carried within a Router CAPABILITY TLV MUST silently ignore the unsupported sub-TLVs and continue processing those sub-TLVs that are supported in the Router CAPABILITY TLV. How partial support may impact the operation of the capabilities advertised within the Router CAPABILITY TLV is outside the scope of this document.

In order for Router CAPABILITY TLVs with domain-wide scope originated by L1 Routers to be flooded across the entire domain, at least one L1/L2 Router in every area of the domain MUST support the Router CAPABILITY TLV.

If leaking of the CAPABILITY TLV is required, the entire CAPABILITY TLV MUST be leaked into another level even though it may contain some of the unsupported sub-TLVs.

5. Security Considerations

Any new security issues raised by the procedures in this document depend upon the opportunity for LSPs to be snooped and modified, the ease/difficulty of which has not been altered. As the LSPs may now contain additional information regarding router capabilities, this new information would also become available to an attacker. Specifications based on this mechanism need to describe the security considerations around the disclosure and modification of their information. Note that an integrity mechanism, such as the one defined in [RFC-3567] or [IS-IS-HMAC], should be applied if there is high risk resulting from modification of capability information.

6. IANA Considerations

IANA assigned a new IS-IS TLV code-point for the newly defined IS-IS TLV type named the IS-IS Router CAPABILITY TLV and defined in this document. The assigned value is 242.

7. Acknowledgment

The authors would like to thank Jean-Louis Le Roux, Paul Mabey, Andrew Partan, and Adrian Farrel for their useful comments.

8. References

8.1. Normative References

- [RFC-2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, March 1997.
- [IS-IS] "Intermediate System to Intermediate System Intra-Domain Routeing Exchange Protocol for use in Conjunction with the Protocol for Providing the Connectionless-mode Network Service (ISO 8473)", ISO 10589.
- [RFC-3567] Li, T. and R. Atkinson, "Intermediate System to Intermediate System (IS-IS) Cryptographic Authentication", RFC 3567, July 2003.
- [IS-IS-IP] Callon, R., "Use of OSI IS-IS for routing in TCP/IP and dual environments", RFC 1195, December 1990.
- [IS-IS-TE] Smit, H. and T. Li, "Intermediate System to Intermediate System (IS-IS) Extensions for Traffic Engineering (TE)", RFC 3784, June 2004.

8.2. Informative References

- [AUTOMESH] Vasseur, JP., Ed., Le Roux, JL., Ed., Yasukawa, S., Previdi, S., Psenak, P., and P. Mabbey, "Routing extensions for Discovery of Multiprotocol (MPLS) Label Switch Router (LSR) Traffic Engineering (TE) Mesh Membership", RFC 4972, July 2007.
- [TE-NODE-CAP] Vasseur, JP., Ed., and J.L. Le Roux, "Routing Extensions for Discovery of Traffic Engineering Node Capabilities", Work in Progress, April 2007.
- [P2MP] Aggarwal, R., Ed., Papadimitriou, D., Ed., and S. Yasukawa, Ed., "Extensions to Resource Reservation Protocol - Traffic Engineering (RSVP-TE) for Point-to-Multipoint TE Label Switched Paths (LSPs)", RFC 4875, May 2007.
- [P2MP-REQS] Yasukawa, S., Ed., "Signaling Requirements for Point-to-Multipoint Traffic-Engineered MPLS Label Switched Paths (LSPs)", RFC 4461, April 2006.
- [IS-IS-HMAC] Bhatia, M., Ed. and V. Manral, Ed., "IS-IS Generic Cryptographic Authentication", Work in Progress, May 2007.

Authors' Addresses

Jean-Philippe Vasseur
CISCO Systems, Inc.
1414 Massachusetts Avenue
Boxborough, MA 01719
USA
EMail: jpv@cisco.com

Stefano Previdi
CISCO Systems, Inc.
Via Del Serafico 200
00142 - Roma
ITALY
EMail: sprevidi@cisco.com

Mike Shand
Cisco Systems
250 Longwater Avenue,
Reading,
Berkshire,
RG2 6GB
UK
EMail: mshand@cisco.com

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

Acee Lindem
Redback Networks
102 Carric Bend Court
Cary, NC 27519
USA
EMail: acee@redback.com

Naiming Shen
Cisco Systems
225 West Tasman Drive
San Jose, CA 95134
USA
EMail: naiming@cisco.com

Rahul Aggarwal
Juniper Networks
1194 N. Mathilda Avenue
San Jose, CA 94089
USA
EMail: rahul@juniper.net

Scott Shaffer
EMail: sshaffer@bridgeport-networks.com

Full Copyright Statement

Copyright (C) The IETF Trust (2007).

This document is subject to the rights, licenses and restrictions contained in BCP 78, and except as set forth therein, the authors retain all their rights.

This document and the information contained herein are provided on an "AS IS" basis and THE CONTRIBUTOR, THE ORGANIZATION HE/SHE REPRESENTS OR IS SPONSORED BY (IF ANY), THE INTERNET SOCIETY, THE IETF TRUST AND THE INTERNET ENGINEERING TASK FORCE DISCLAIM 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.

Intellectual Property

The IETF takes no position regarding the validity or scope of any Intellectual Property Rights or other rights that might be claimed to pertain to the implementation or use of the technology described in this document or the extent to which any license under such rights might or might not be available; nor does it represent that it has made any independent effort to identify any such rights. Information on the procedures with respect to rights in RFC documents can be found in BCP 78 and BCP 79.

Copies of IPR disclosures made to the IETF Secretariat and any assurances of licenses to be made available, or the result of an attempt made to obtain a general license or permission for the use of such proprietary rights by implementers or users of this specification can be obtained from the IETF on-line IPR repository at <http://www.ietf.org/ipr>.

The IETF invites any interested party to bring to its attention any copyrights, patents or patent applications, or other proprietary rights that may cover technology that may be required to implement this standard. Please address the information to the IETF at ietf-ipr@ietf.org.

Acknowledgement

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

