

Internet Engineering Task Force (IETF)
Request for Comments: 7981
Obsoletes: 4971
Category: Standards Track
ISSN: 2070-1721

L. Ginsberg
S. Previdi
Cisco Systems
M. Chen
Huawei Technologies Co., Ltd
October 2016

IS-IS Extensions for Advertising Router Information

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. This document obsoletes RFC 4971.

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 7841.

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/rfc7981>.

Copyright Notice

Copyright (c) 2016 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.

Table of Contents

1. Introduction	2
1.1. Requirements Language	3
2. IS-IS Router CAPABILITY TLV	3
3. Elements of Procedure	4
4. Interoperability with Routers Not Supporting the IS-IS Router CAPABILITY TLV	6
5. Security Considerations	6
6. IANA Considerations	7
7. References	7
7.1. Normative References	7
7.2. Informative References	8
Appendix A. Changes to RFC 4971	9
Acknowledgements	10
Authors' Addresses	10

1. Introduction

There are several situations where it is useful for the IS-IS [ISO10589] [RFC1195] 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) [RFC5305] requires some significant configuration effort. [RFC4972] 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 (RFC4875): A specific sub-TLV [RFC5073] allows an LSR to advertise its Point-to-Multipoint capabilities ([RFC4875] and [RFC4461]).
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 IS-IS Router 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 IS-IS Router CAPABILITY TLV defined in this document.

Definition of these sub-TLVs 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. 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, followed by 1 octet of flags.

A set of optional sub-TLVs may follow the flag field. Sub-TLVs are formatted as described in [RFC5305].

TYPE: 242
LENGTH: from 5 to 255
VALUE:
Router ID (4 octets)
Flags (1 octet)
Set of optional sub-TLVs (0-250 octets)

Flags



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 (L2) to Level 1 (L1), 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 IS-IS Router CAPABILITY TLV is OPTIONAL. As specified in Section 3, more than one IS-IS Router CAPABILITY TLV from the same source MAY be present.

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

3. Elements of Procedure

The Router ID SHOULD be identical to the value advertised in the Traffic Engineering Router ID TLV [RFC5305]. If no Traffic Engineering Router ID is assigned, the Router ID SHOULD be identical to an IP Interface Address [RFC1195] advertised by the originating IS. If the originating node does not support IPv4, then the reserved value 0.0.0.0 MUST be used in the Router ID field, and the IPv6 TE Router ID sub-TLV [RFC5316] MUST be present in the TLV. IS-IS Router CAPABILITY TLVs that have a Router ID of 0.0.0.0 and do NOT have the IPv6 TE Router ID sub-TLV present MUST NOT be used.

When advertising capabilities with different flooding scopes, a router MUST originate a minimum of two IS-IS 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 IS-IS Router CAPABILITY TLVs:

- o One IS-IS Router CAPABILITY TLV with the S flag cleared, carrying the sub-TLV(s) relative to C1. This IS-IS Router CAPABILITY TLV will not be leaked into another level.

- o One IS-IS Router CAPABILITY TLV with the S flag set, carrying the sub-TLV(s) relative to C2. This IS-IS 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 IS-IS Router CAPABILITY TLVs, a system MUST NOT use an IS-IS Router 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 IS-IS Router CAPABILITY TLV.

When an IS-IS Router CAPABILITY TLV is not used, either due to a lack of reachability to the originating router or due to an unusable Router ID, note that leaking the IS-IS Router CAPABILITY TLV is one of the uses that is prohibited under these conditions.

Example: If Level 1 router A generates an IS-IS Router 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 without the above prohibition, 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 IS-IS Router CAPABILITY TLV or S's copy of A's IS-IS Router CAPABILITY TLV, and they have no reliable way to choose. By making sure that T stops leaking A's information, the possibility that other routers will use stale information from A is eliminated.

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 an IS-IS Router CAPABILITY 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 IS-IS Router 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 an IS-IS Router CAPABILITY TLV from the same system that have conflicting information for a given sub-TLV, the procedure used to choose which copy shall be used is undefined.

4. Interoperability with Routers Not Supporting the IS-IS Router CAPABILITY TLV

Routers that do not support the IS-IS 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 an IS-IS Router CAPABILITY TLV MUST silently ignore the unsupported sub-TLVs and continue processing those sub-TLVs that are supported in the IS-IS Router CAPABILITY TLV. How partial support may impact the operation of the capabilities advertised within the IS-IS Router CAPABILITY TLV is outside the scope of this document.

In order for IS-IS 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 IS-IS Router CAPABILITY TLV is required, the entire CAPABILITY TLV MUST be leaked into another level without change (except for changes to the TLV flags as noted in Section 2) even though it may contain some sub-TLVs that are unsupported by the router doing the leaking.

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 ones defined in [RFC5304] or [RFC5310], should be applied if there is high risk resulting from modification of capability information.

6. IANA Considerations

IANA originally assigned a TLV codepoint for the IS-IS Router CAPABILITY TLV (242) as described in RFC 4971. IANA has updated this entry in the "TLV Codepoints Registry" to refer to this document.

7. References

7.1. Normative References

- [ISO10589] International Organization for Standardization, "Information technology -- Telecommunications and information exchange between systems -- Intermediate System to Intermediate System intra-domain 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, November 2002.
- [RFC1195] Callon, R., "Use of OSI IS-IS for routing in TCP/IP and dual environments", RFC 1195, DOI 10.17487/RFC1195, December 1990, <<http://www.rfc-editor.org/info/rfc1195>>.
- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <<http://www.rfc-editor.org/info/rfc2119>>.
- [RFC5073] Vasseur, J., Ed. and J. Le Roux, Ed., "IGP Routing Protocol Extensions for Discovery of Traffic Engineering Node Capabilities", RFC 5073, DOI 10.17487/RFC5073, December 2007, <<http://www.rfc-editor.org/info/rfc5073>>.
- [RFC5304] Li, T. and R. Atkinson, "IS-IS Cryptographic Authentication", RFC 5304, DOI 10.17487/RFC5304, October 2008, <<http://www.rfc-editor.org/info/rfc5304>>.
- [RFC5305] Li, T. and H. Smit, "IS-IS Extensions for Traffic Engineering", RFC 5305, DOI 10.17487/RFC5305, October 2008, <<http://www.rfc-editor.org/info/rfc5305>>.

- [RFC5310] Bhatia, M., Manral, V., Li, T., Atkinson, R., White, R., and M. Fanto, "IS-IS Generic Cryptographic Authentication", RFC 5310, DOI 10.17487/RFC5310, February 2009, <<http://www.rfc-editor.org/info/rfc5310>>.
- [RFC5316] Chen, M., Zhang, R., and X. Duan, "ISIS Extensions in Support of Inter-Autonomous System (AS) MPLS and GMPLS Traffic Engineering", RFC 5316, DOI 10.17487/RFC5316, December 2008, <<http://www.rfc-editor.org/info/rfc5316>>.

7.2. Informative References

- [RFC4461] Yasukawa, S., Ed., "Signaling Requirements for Point-to-Multipoint Traffic-Engineered MPLS Label Switched Paths (LSPs)", RFC 4461, DOI 10.17487/RFC4461, April 2006, <<http://www.rfc-editor.org/info/rfc4461>>.
- [RFC4875] 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, DOI 10.17487/RFC4875, May 2007, <<http://www.rfc-editor.org/info/rfc4875>>.
- [RFC4972] Vasseur, JP., Ed., Leroux, 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, DOI 10.17487/RFC4972, July 2007, <<http://www.rfc-editor.org/info/rfc4972>>.

Appendix A. Changes to RFC 4971

This document makes the following changes to RFC 4971.

RFC 4971 only allowed a 32-bit Router ID in the fixed header of TLV 242. This is problematic in an IPv6-only deployment where an IPv4 address may not be available. This document specifies:

1. The Router ID SHOULD be identical to the value advertised in the Traffic Engineering Router ID TLV (134) if available.
2. If no Traffic Engineering Router ID is assigned, the Router ID SHOULD be identical to an IP Interface Address [RFC1195] advertised by the originating IS.
3. If the originating node does not support IPv4, then the reserved value 0.0.0.0 MUST be used in the Router ID field, and the IPv6 TE Router ID sub-TLV [RFC5316] MUST be present in the TLV.

In addition, some clarifying editorial changes have been made.

Acknowledgements

The authors of RFC 4971 thanked Jean-Louis Le Roux, Paul Mabey, Andrew Partan, and Adrian Farrel for their useful comments.

The authors of this document would like to thank Kris Michielsen for calling attention to the problem associated with an IPv6-only router.

Authors' Addresses

Les Ginsberg
Cisco Systems
510 McCarthy Blvd.
Milpitas, CA 95035
United States of America

Email: ginsberg@cisco.com

Stefano Previdi
Cisco Systems
Via Del Serafico 200
Rome 0144
Italy

Email: sprevidi@cisco.com

Mach(Guoyi) Chen
Huawei Technologies Co., Ltd
KuiKe Building, No. 9 Xinxu Rd. Hai-Dian District
Beijing 100085
China

Email: mach.chen@huawei.com