

Network Working Group
Internet Draft
Proposed status: Standard
Expires: August 2007

Jean-Philippe Vasseur(Ed)
Naiming Shen (Ed)
Cisco Systems, Inc.
Rahul Aggarwal(Ed)
Juniper Networks

February 2007

IS-IS Extensions for Advertising Router Information

[draft-ietf-isis-caps-07.txt](#)

Status of this Memo

By submitting this Internet-Draft, each author represents that any applicable patent or other IPR claims of which he or she is aware have been or will be disclosed, and any of which he or she becomes aware will be disclosed, in accordance with [Section 6 of BCP 79](#).

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/lid-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 August 15, 2007.

Abstract

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.

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](#)].

Table of Contents

1.	Introduction.....	2
2.	IS-IS Router CAPABILITY TLV.....	3
3.	Element of procedure.....	3
4.	Interoperability with routers not supporting the capability TLV.5	
5.	Security considerations.....	5
6.	Acknowledgment.....	6
7.	Intellectual Property Considerations.....	6
8.	References.....	6
	Normative references.....	6
	Informative references.....	6
9.	Author's Addresses.....	7

[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, two examples related to MPLS Traffic Engineering are described here:

1. Mesh-group: the setting up of a mesh of TE LSPs [[IS-IS-TE](#)] requires some significant configuration effort. [[AUTOMESH](#)] proposes an auto-discovery mechanism whereby every LSR of a mesh advertises its mesh-group membership by means of IS-IS extensions.
2. Point to Multi-point 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 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.

Vasseur et al.

[Page 2]

[draft-ietf-isis-caps-07.txt](#)

February 2007

2. IS-IS Router CAPABILITY TLV

The IS-IS Router CAPABILITY TLV is composed of 1 octet for the type, 1 octet specifying the number of bytes in the value field, and a variable length value field, starting 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 (To be assigned by IANA)

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 TLVs 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 which generates a CAPABILITY TLV MUST have a Router ID which 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 information A system MUST NOT use a Capability TLV present in an LSP of a system which 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 the sending system is the originator of the Capabilities TLV or not. Note that leaking a

Capabilities TLV is one of the uses which 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 which allow multiple entries to appear in the value field (e.g. IS-neighbors) - 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 which had the old advertisement and the fragment which 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 which it replaces. Where this is not possible, the two affected LSP fragments should be flooded as an atomic action.

Systems which 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 which 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 which do not support the Router CAPABILITY TLV MUST silently ignore the TLV(s) and continue processing other TLVs in the same LSP. Routers which 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 in the Router CAPABILITY TLV which are supported. 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 CAP TLV is required, the entire CAP 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 one defined in [RFC3567](#) or [draft-ietf-isis-hmac-sha](#), should be applied if there is high risk resulting from modification of capability information.

6. IANA considerations

IANA will assign 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. Suggested value is 242 (to be assigned by IANA).

7. Acknowledgment

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

8. Intellectual Property Considerations

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.

9. References

9.1 Normative references

[RFC-2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels," [RFC 2119](#).

[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.

[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] Li, T., Smit, H., "IS-IS extensions for Traffic Engineering", [RFC 3784](#), June 2004.

9.2 Informative references

[AUTOMESH] JP Vasseur, JL. Le Roux et al, "Routing extensions for discovery of Multiprotocol (MPLS) Label Switch Router (LSR) Traffic Engineering (TE) mesh membership", [draft-ietf-ccamp-automesh](#), work in progress.

[TE-NODE-CAP] JP Vasseur, JL. Le Roux et al, "Routing extensions for discovery of Traffic Engineering Node Capabilities", [draft-ietf-ccamp-te-node-cap](#), work in progress.

[P2MP] R. Aggarwal, D. Papadimitriou, S. Yasukawa, et. al. "Extensions to RSVP-TE for Point To Multipoint TE LSPs", [draft-ietf-mpls-rsvp-te-p2mp](#), work in progress.

[P2MP-REQS] S. Yasukawa et al. "Requirements for point to multipoint extension to RSVP", [draft-ietf-mpls-p2mp-sig-requirement](#), work in progress.

10. 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

Vasseur et al.

[Page 7]

[draft-ietf-isis-caps-07.txt](#)

February 2007

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

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

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

Scott Shaffer
e-mail: 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.

Vasseur et al.

[Page 8]