IS-IS WG
Internet-Draft
Intended status: Standards Track
Expires: March 19, 2018

S. Hegde
C. Bowers
Juniper Networks
P. Mattes
M. Nanduri
S. Giacalone
Microsoft
I. Mohammad
Arista Networks

September 15, 2017

Advertising TE protocols in IS-IS
draft-hegde-isis-advertising-te-protocols-03

Abstract

This document defines a mechanism to indicate which traffic engineering protocols are enabled on a link in IS-IS. It does so by introducing a new traffic-engineering protocol sub-TLV for TLV-22. This document also describes mechanisms to address backward compatibility issues for implementations that have not yet been upgraded to software that understands this new sub-TLV.

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].

Status of This Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at https://datatracker.ietf.org/drafts/current/.

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."

This Internet-Draft will expire on March 19, 2018.

Table of Contents

1. Introduction

IS-IS extensions for traffic engineering are specified in [RFC5305]. [RFC5305] defines several link attributes such as administrative group, maximum link bandwidth, and shared risk link groups (SRLGs) which can be used by traffic engineering applications. Additional link attributes for traffic engineering have subsequently been defined in other documents as well. Most recently [RFC7810] defined link attributes for delay, loss, and measured bandwidth utilization.

The primary consumers of these traffic engineering link attributes have been RSVP-based applications that use the advertised link attributes to compute paths which will subsequently be signalled using RSVP-TE. However, these traffic engineering link attributes have also been used by other applications, such as IP/LDP fast-reroute using loop-free alternates as described in [RFC7916]. In the future, it is likely that traffic engineering applications based on Segment Routing [I-D.ietf-spring-segment-routing] will also use these link attributes.

Existing IS-IS standards do not provide a mechanism to explicitly indicate whether or not RSVP has been enabled on a link. Instead, different RSVP-TE implementations have used the presence of certain traffic engineering sub-TLVs in IS-IS to infer that RSVP signalling is enabled on a given link. A study was conducted with various vendor implementations to determine which traffic engineering sub-TLVs cause an implementation to infer that RSVP signalling is enabled on a link. The results are shown in Figure 1.

TLV/ sub-TLV	Sub-TLV name	Implementation		
		X	Y	Z
22	Extended IS Reachability TLV	N	N	N
22/3	Administrative group (color)	N	Y	Y
22/4	Link Local/Remote ID	N	N	N
22/6	IPV4 Interface Address	N	N	N
22/8	IPV4 Neighbor Address	N	N	N
22/9	Max Link Bandwidth	N	Y	Y
22/10	Max Reservable Link Bandwidth	N	Y	Y
22/11	Unreserved Bandwidth	Y	Y	Y
22/14	Extended Admin Group	N	Y	N
22/18	TE Default Metric	N	N	N
22/20	Link Protection Type	N	Y	Y
22/21	Interface Switching Capability	N	Y	Y
22/22	TE Bandwidth Constraints	N	Y	Y
22/33-39	TE Metric Extensions(RFC7180)	N	N	N
138	SRLG TLV	N	Y	Y

Figure 1: Traffic engineering Sub-TLVs that cause implementation X, Y, or Z to infer that RSVP signalling is enabled on a link

The study indicates that the different implementations use the presence of different sub-TLVs under TLV 22 (or the presence of TLV 138) to infer that RSVP signalling is enabled on a link. It is

possible that other implementations may use other sub-TLVs to infer that RSVP is enabled on a link.

This document defines a standard way to indicate whether or not RSVP, segment routing, or another future protocol is enabled on a link. In this way, implementations will not have to infer whether or not RSVP is enabled based on the presence of different sub-TLVs, but can use the explicit indication. When network operators want to use a non-RSVP traffic engineering application (such as IP/LDP FRR or segment routing), they will be able to advertise traffic engineering sub-TLVs and explicitly indicate what traffic engineering protocols are enabled on a link.

2. Goals

1. The solution should allow the TE protocol enabled on a link to be communicated unambiguously.

2. The solution should decouple the advertisement of which TE protocols are enabled on a link from the advertisement of other TE attributes.

3. The solution should be backward compatible so that nodes that do not understand the new advertisement do not cause issues for existing RSVP deployments.

4. The solution should be extensible for new protocols.

5. The solution should try to limit any increases to the quantity and size of link state advertisements.

2.1. Explicit and unambiguous indication of TE protocol

Communicating unambiguously which TE protocol is enabled on a link is important to be able to share this information with other consumers through other protocols, aside from just the IGP. For example, for a network running both RSVP-TE and SR, it will be useful to communicate which TE protocols are enabled on which links via BGP-LS [RFC7752] to a central controller. Typically, BGP-LS relies on the IGP to distribute IGP topology and traffic engineering information so that only a few BGP-LS sessions with the central controller are needed. In order for a router running a BGP-LS session to a central controller to correctly communicate what TE protocols are enabled on the links in the IGP domain, that information first needs to be communicated unambiguously within the IGP itself. As Figure 1 illustrates, that is currently not the case.

3. Solution

3.1. Traffic-engineering protocol sub-TLV

A new Traffic-engineering protocol sub-TLV is added in the TLV 22 [RFC5305] or TLV 222 to indicate the protocols enabled on the link. The sub-TLV has flags in the value field to indicate the protocol enabled on the link. The length field is variable to allow the flags field to grow for future requirements.

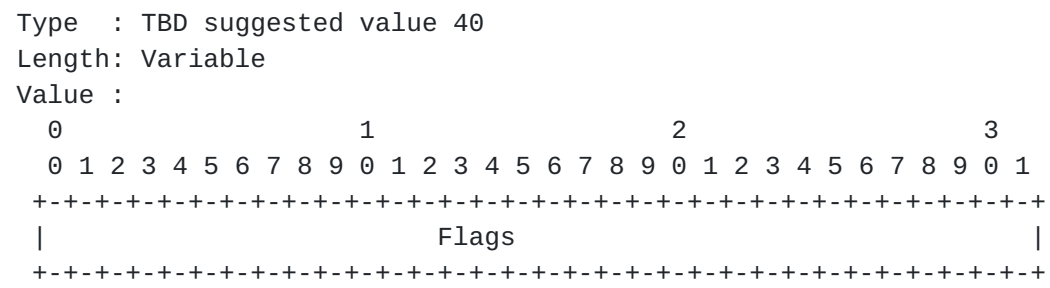

Figure 2: Traffic-Engineering Protocol sub-TLV

Type : TBA (suggested value 40)

Length: variable (in bytes)

Value: The value field consists of bits indicating the protocols enabled on the link. This document defines the two protocol values below.

Value	Protocol Name
0x01	RSVP
0x02	Segment Routing

Figure 3: Flags for the protocols

The RSVP flag is set to one to indicate that RSVP-TE is enabled on a link. The RSVP flag is set to zero to indicate that RSVP-TE is not enabled on a link.

The Segment Routing flag is set to one to indicate that Segment Routing is enabled on a link. The Segment Routing flag is set to zero to indicate that Segment Routing is not enabled on a link.

All undefined flags MUST be set to zero on transmit and ignored on receipt.

An implementation that supports the TE protocol sub-TLV and sends TLV 22 MUST advertise the TE protocol sub-TLV in TLV 22 for that link, even when both the RSVP and SR flags are set to zero. In other words, whenever the TE protocol sub-TLV is supported, it MUST be sent, even if no TE protocols are enabled on the link. This allows a receiving router to determine whether or not the sending router is capable of sending the TE protocol sub-TLV.

A router supporting the TE protocol sub-TLV which receives an advertisement for a link containing TLV 22 with the TE protocol sub-TLV present SHOULD respect the values of the flags in the TE protocol sub-TLV. The receiving router SHOULD only consider links with a given TE protocol enabled for inclusion in a path using that TE protocol. Conversely, links for which the TE protocol sub-TLV is present, but for which the TE protocol flag is not set to one, SHOULD NOT be included in any TE CSPF computations on the receiving router for the protocol in question.

The ability for a receiving router to determine whether or not the sending router is capable of sending the TE protocol sub-TLV is also used for backward compatibility as described in Section 4.

An implementation that supports the TE protocol sub-TLV SHOULD be able to advertise TE sub-TLVs without enabling RSVP-TE signalling on the link.

3.2. Segment Routing flag considerations

The Segment Routing (SR) architecture assumes that the SR topology is congruent with the IGP topology. The path described by a prefix segment is computed using the SPF algorithm applied to the IGP topology, which is the same as the SR topology. Therefore, the presence or absence of the Segment Routing flag MUST NOT be interpreted as modifying the SR topology, which is always congruent with the IGP topology.

It is however useful for a centralized application (or an ingress router) to know whether or not it should expect to be able to forward traffic over a given link using labels distributed via SR. If a link is advertised with the TE protocol sub-TLV and the SR flag set to zero, then a centralized application can assume that traffic sent

with a prefix segment whose path crosses that link is unlikely to be forwarded across that link. With this information, a centralized application can decide to use a different path for that traffic by using a different label stack.

4. Backward compatibility

Routers running older software that do not understand the new Traffic-Engineering protocol sub-TLV will continue to interpret the presence of some sub-TLVs in TLV 22 or the presence of TLV 138 as meaning that RSVP is enabled a link. A network operator may not want to or be able to upgrade all routers in the domain at the same time. There are two backward compatibility scenarios to consider depending on whether the router that doesn't understand the new TE protocol sub-TLV is an RSVP-TE ingress router or an RSVP-TE transit router.

4.1. Scenario with upgraded RSVP-TE transit router but RSVP-TE ingress router not upgraded

An upgraded RSVP-TE transit router is able to explicitly indicate that RSVP is not enabled on a link by advertising the TE protocol sub-TLV with the RSVP flag set to zero. However, an RSVP-TE ingress router that has not been upgraded to understand the new TE protocol sub-TLV will not understand that RSVP-TE is not enabled on the link, and may include the link on a path computed for RSVP-TE. When the network tries to signal an explicit path LSP using RSVP-TE through that link, it will fail. In order to avoid this scenario, an operator can use the mechanism described below.

For this scenario, the basic idea is to use the existing administrative group link attribute as a means of preventing existing RSVP implementations from using a link. The network operator defines an administrative group to mean that RSVP is not enabled on a link. We call this admin group the RSVP-not-enabled admin group. If the operator needs to advertise a TE sub-TLV (maximum link bandwidth, for example) on a link, but doesn't want to enable RSVP on that link, then the operator also advertises the RSVP-not-enabled admin group on that link. The operator can then use existing mechanisms to exclude links advertising the RSVP-not-enabled admin group from the constrained shortest path first (CSPF) computation used by RSVP. This will prevent RSVP implementations from attempting to signal RSVP-TE LSPs across links that do not have RSVP enabled. Once the entire network domain is upgraded to understand the TE protocol sub-TLV in this draft, the configuration involving the RSVP-not-enabled admin group is no longer needed for this network.

4.2. Scenario with upgraded RSVP-TE ingress router but RSVP-TE transit router not upgraded

The other scenario to consider is when the RSVP-TE ingress router has been upgraded to understand the TE protocol sub-TLV, but the RSVP-TE transit router has not. In this case, the transit router has not been upgraded, so it is not yet capable of sending the TE protocol sub-TLV. If the transit router has RSVP-TE enabled on a link, we would like for the RSVP-TE ingress router to still be able to use the link for RSVP-TE paths. While it is possible to describe a solution for this scenario that makes use of administrative groups, we describe a simpler solution below.

The solution for this scenario relies on the following observation. If the RSVP-TE ingress router can understand that the transit router is not capable of sending the TE protocol sub-TLV, then it can continue inferring whether or not RSVP-TE is enabled on the transit router links based on the presence of TE sub-TLVs, just as it does today.

To accomplish this, we require an upgraded router to send the TE protocol sub-TLV if it sends TLV 22, even when both the RSVP and SR flags are set to zero. In other words, whenever the TE protocol sub-TLV is supported, it MUST be sent, even if no TE protocols are enabled on the link. see Section 3. This allows the receiving router to interpret the absence of the TE-protocol sub-TLV together with presence of TLV 22 to mean that the sending router has not been upgraded. This allows the upgraded RSVP-TE ingress router to distinguish between transit routers that have been upgraded and those that haven't. When the transit router has been upgraded, then the RSVP-TE ingress router uses the information in the TE protocol sub-TLV. When the transit router has not been upgraded, then RSVP-TE ingress router contines to infer whether or not RSVP-TE is enabled on the transit router links based on the presence of TE sub-TLVs, just as it does today. The solution for this scenario requires no configuration on the part of network operators.

4.3. Need for a long term solution

The use of the adminstrative group link attribute to prevent an RSVP-TE ingress router from computing a path using a given link is an effective short term workaround to allow networks to incrementally upgrade the routers to software that understands the new TE-protocol sub-TLV. One might also consider a long term solution based solely on the use of operator-defined adminstrative groups to communicate the TE protocol enabled on a link. However, we do not consider this workaround to be an effective long term solution because it relies on operator configuration that would have to be maintained in the long

term. As discussed in Section 2, continuing to have to infer which TE protocol is enabled on a link also limits our ability to communicate this information unambiguously in an interoperable manner for use by other applications such as central controllers.

5. Security Considerations

This document does not introduce any further security issues other than those discussed in [RFC1195] and [RFC5305].

6. IANA Considerations

This specification updates one IS-IS registry:

The extended IS reachability TLV Registry

i) Traffic-engineering Protocol sub-tlv = Suggested value 40

7. Acknowledgements

The authors thank Alia Atlas, Les Ginsberg, and Peter Psenak for helpful discussions on the topic of this draft.

8. References

8.1. Normative References

[I-D.ietf-spring-segment-routing]
Filsfils, C., Previdi, S., Decraene, B., Litkowski, S., and R. Shakir, "Segment Routing Architecture", draft-ietf-spring-segment-routing-09 (work in progress), July 2016.

[RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <https://www.rfc-editor.org/info/rfc2119>.

[RFC5305] Li, T. and H. Smit, "IS-IS Extensions for Traffic Engineering", RFC 5305, DOI 10.17487/RFC5305, October 2008, <https://www.rfc-editor.org/info/rfc5305>.

[RFC7810] Previdi, S., Ed., Giacalone, S., Ward, D., Drake, J., and Q. Wu, "IS-IS Traffic Engineering (TE) Metric Extensions", RFC 7810, DOI 10.17487/RFC7810, May 2016, <https://www.rfc-editor.org/info/rfc7810>.

8.2. Informative References

[RFC1195] Callon, R., "Use of OSI IS-IS for routing in TCP/IP and dual environments", RFC 1195, DOI 10.17487/RFC1195, December 1990, <https://www.rfc-editor.org/info/rfc1195>.

[RFC7752] Gredler, H., Ed., Medved, J., Previdi, S., Farrel, A., and S. Ray, "North-Bound Distribution of Link-State and Traffic Engineering (TE) Information Using BGP", RFC 7752, DOI 10.17487/RFC7752, March 2016, <https://www.rfc-editor.org/info/rfc7752>.

[RFC7916] Litkowski, S., Ed., Decraene, B., Filsfils, C., Raza, K., Horneffer, M., and P. Sarkar, "Operational Management of Loop-Free Alternates", RFC 7916, DOI 10.17487/RFC7916, July 2016, <https://www.rfc-editor.org/info/rfc7916>.

Authors' Addresses

Shraddha Hegde
Juniper Networks
Embassy Business Park
Bangalore, KA 560093
India

Email: shraddha@juniper.net

Chris Bowers
Juniper Networks
1194 N. Mathilda Ave.
Sunnyvale, CA 94089
US

Email: cbowers@juniper.net

Paul Mattes
Microsoft
One Microsoft Way
Redmond, WA 98052
US

Email: pamattes@microsoft.com

Mohan Nanduri
Microsoft
One Microsoft Way
Redmond, WA 98052
US

Email: mnanduri@microsoft.com

Spencer Giacalone
Microsoft
One Microsoft Way
Redmond, WA 98052
US

Email: Spencer.Giacalone@microsoft.com

Imtiyaz Mohammad
Arista Networks
Global Tech Park
Bangalore, KA 560103
India

Email: imtiyaz@arista.com