isis
Internet-Draft
Intended status: Standards Track
Expires: April 23, 2015

D. Lamparter
NetDEF

October 20, 2014

IS-IS Reachability with critical Sub-TLVs
draft-lamparter-isis-reachability-critical-subtlvs-00

Abstract

While previously existing TLVs for IP Reachability extensibly support Sub-TLVs, these cannot be marked as critical. This is required for extending router behaviour with additional qualifiers on routes, hence this document introduces new Reachability TLVs that support critical Sub-TLVs.

Table of Contents

1. Introduction 2
1.1. Requirements Language 3
2. Design considerations 3
3. SPF Functional specification 4
3.1. Simplified SPF 4
4. TLV formats 4
4.1. IP/IPv6 Reachability TLV 4
4.2. Reachability Critical Sub-TLVs Supported TLV 5
5. IANA Considerations 6
5.1. IS-IS TLV Codepoints 7
5.2. TLVs 135, 235, 236, 237 Sub-TLV Registry 7
5.3. TLVs TBD1, TBD2 critical Sub-TLV Registry 7
6. Security Considerations 7
7. Privacy Considerations 8
8. Acknowledgements 8
9. Change Log 8
10. Normative References 8
Author's Address 8

1. Introduction

IS-IS is very extensible by design; Newly defined Sub-TLVs can be added in many places. However, the behaviour for unknown Sub-TLVs is always assumed to be "ignore", there is currently no way to prescribe different behaviour. Therefore, a system that receives a Reachability TLV with a Sub-TLV it doesn't recognise will silently process the Reachability with a reduced set of specified information.

This is not desirable for situations where Sub-TLVs provide essential information for the reachability, in particular if that information restricts the usability of the reachability. At the time of writing, usage by extensions of the following types is envisioned:

- further qualifications for the route target, e.g. restricted source address or flowlabel. In this case the reachability information is incomplete (and the route does not match) without these critical fields.

- mandatory encapsulation specifications, e.g. routing headers or labels required for the egress router or systems outside the domain. Here, ignorance of that information would render these systems unable to apply correct forwarding decisions.

Other future developments may find even more use cases for this TLV. The functionality defined here could also have been used for M-ISIS

[RFC5120] reachabilities in order to hide them from non-M-ISIS routers without introducing a new TLV type.

Therefore, this document creates a new Reachability TLV with a critical Sub-TLV part, where the specified behavior on unrecognized Sub-TLVs is to ignore the entire Reachability TLV, not just the Sub-TLV.

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 [RFC2119].

2. Design considerations

This document specifies new Reachability TLVs for IPv4 and IPv6. These new TLVs have two Sub-TLV blocks: one critical and one optional. Sub-TLVs in the optional block behave exactly as Sub-TLVs in previous Reachability TLVs (135, 235, 236, and 237.) This includes application of the same TLV namespace, all TLVs defined for these four TLVs are also applicable in the optional part of the new two TLVs.

The critical Sub-TLV block constitutes a separate namespace. A system MUST keep these separated, and specifications MUST define to which part exactly they apply. Expected combinations are:

- 135, 235, 236, 237, TBD1 and TBD2 optional
- 135, 235, TBD1 optional (IPv4)
- 236, 237, TBD2 optional (IPv6)
- TBD1 and TBD2 critical
- TBD1 critical (IPv4)
- TBD2 critical (IPv6)

Though no such use is foreseen at this point, a specification MAY specify a TLV to be valid in either the optional or critical part. This TLV may end up with different codepoints in each of the namespaces.

A system MUST NOT originate these new TLVs with an empty critical part. Doing so would create an alternate encoding of the previous TLVs, breaking interoperability. Systems SHOULD process a new TLV with an empty critical block.

There is no need for non-MT variants of these TLVs. If a system does not implement M-ISIS, it MUST ignore all TLVs with a MT ID other than zero.

3. SPF Functional specification

This document assumes that all transit routers need to support processing of the feature associated with a respective critical Sub-TLV. Hence, on calculating a path for a reachability with critical Sub-TLV A, all intermediate systems that do not indicate support for Sub-TLV A must be excluded.

The logical result from this is essentially that separate SPF trees MUST be calculated for each set of critical Sub-TLVs.

Calculation of these extra trees can be optimized by sharing intermediate calculation results as far as critical Sub-TLV support is identical.

A system MUST NOT blindly use a "more Sub-TLVs supported" SPF calculation result for calculating paths that require only a subset of these Sub-TLVs. This would result in a disagreement on shortest path with other routers, which correctly used a SPF tree for the specific combination.

3.1. Simplified SPF

TBD: It is possible to construct a variant of this that doesn't implicitly work with multiple topologies, instead marking routes as unreachable if they transit over routers that do not support the critical TLVs. This may be useful for simpler implementations.

4. TLV formats

4.1. IP/IPv6 Reachability TLV

The encoding for TLVs TBD1 and TBD2 is modified from TLVs 235 and 237 by inserting a second length field for the critical Sub-TLV part before the existing length field for the optional Sub-TLV part. The critical Sub-TLV part follows after the length field, then the optional part.

This results in the following TLV structure:

```
(2/4 bytes TLV header)
2 octets of MT ID (12 bits, top 4 bits reserved)
-- multiple (n >= 1) occurences of the following:
4 octets of metric information
1 octet of control information, consisting of
     1 bit of up/down information
     1 bit indicating the presence of optional sub-TLVs
     6 bits of prefix length
     0-4/0-16 octets of IPv4/IPv6 prefix
4-n optional octets of sub-TLVs, if present consisting of
     1/2 octets of length of critical sub-TLVs
     2-n octets of critical sub-TLVs,
     -- depending on presence of optional sub-TLVs indication:
     0-2 octets of length of optional sub-TLVs
     0-n octets of optional sub-TLVs,
          where each sub-TLV (critical or optional) is a sequence of
               1/2 octets of sub-type
               1/2 octets of length of the value field of the sub-TLV
               0-n octets of value
```

Unlike MT Reachability TLVs, this TLV MUST NOT be ignored if the MT ID is zero. Instead, the information applies to the "standard" topology.

The size of offset and length fields depends on the PDU in which the TLV is found, as per [RFC7356].

The critical sub-TLV part MUST NOT be empty. Reachability TLVs without a critical Sub-TLV field MUST be used instead in this case.

As in TLVs 235 and 237, the optional sub-TLV length and data fields are only present if the "presence of optional sub-TLVs" bit is one.

4.2. Reachability Critical Sub-TLVs Supported TLV

Supported Critical Information Sub-TLV format

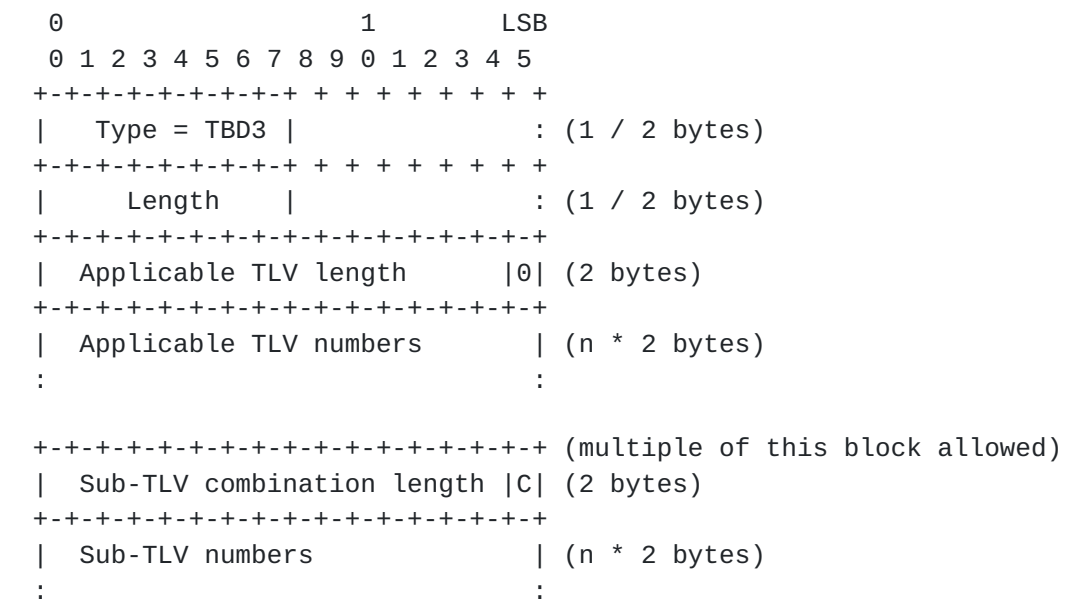

Applicable TLV length and numbers specify which (parent) TLVs this information applies to. Both fields are always in 2-octet units each, which means the length is even. Thus, the LSB of the length field MUST be set to 0 on TLV origination. Systems SHOULD ignore the entire TLV if the applicable TLV length field is not even. The same applies if the applicable TLV length is zero, systems SHOULD ignore the entire TLV.

Sub-TLV combination length and numbers specify supported Sub-TLVs for the TLVs with applicable TLV numbers listed before. As with Applicable TLVs, these are units of 2 octets each.

The LSB of the combination length is redefined to be the "Combinatorial" bit. Any mixture (present or not present) of Sub-TLVs listed with C=1, plus any Sub-TLVs present in at most one list with C=0 are understood to be supported by the router. A combination of Sub-TLVs present in two distinct lists with C=0 MUST NOT be assumed to be in a router's supported set.

The block of Sub-TLV combination length and numbers MAY occur multiple times, as MAY the entire TLV. The information MUST be merged.

Systems MUST process known TLVs even if unknown TLVs are present. The latter MUST be ignored.

5. IANA Considerations

5.1. IS-IS TLV Codepoints

This document requests the allocation of two codepoints from the IS-IS TLV Codepoints registry. Suggested values are 238 for TBD1 and 239 for TBD2.

Top-level codepoints

Value	Name	IIH	LSP	SNP	Purge
TBD1	MT IPv4 Reach with Critical Sub-TLVs	n	y	n	n
TBD2	MT IPv6 Reach with Critical Sub-TLVs	n	y	n	n

A codepoint from the Sub-TLVs for TLV 144 registry is also requested:

TLV 144 Sub-TLV codepoints

Value	Name
TBD3	Reachability Critical Sub-TLVs Supported

5.2. TLVs 135, 235, 236, 237 Sub-TLV Registry

The registry for Sub-TLVs below TLVs 135, 235, 236, and 237 is requested to be renamed to "Sub-TLVs for TLVs 135, 235, 236, 237, TBD1 (optional) and TBD2 (optional)". Two new columns are added to the table: "TBD1 (optional)" and "TBD2 (optional)". The value for preexisting entries is copied from 235 to TBD1 and from 237 to TBD2. This document is added as reference.

5.3. TLVs TBD1, TBD2 critical Sub-TLV Registry

This document requests creation of a new registry named "Sub-TLVs for TLVs TBD1 (critical), and TBD2 (critical)". Procedures and experts are inherited from the registry in the previous paragraph. The registry's table is initially empty and has a total of two applicability columns titled "TBD1 (critical)" and "TBD2 (critical)". The starting value for allocations is 1.

6. Security Considerations

The mechanism outlined in this document can be used to perform memory and processor resource exhaustion attacks against routers. By introducing reachabilities with different sets of critical Sub-TLVs present, participating routers are forced to calculate different SPF trees.

As a countermeasure, routers SHOULD:

- only calculate SPF trees for critical TLV combinations they support
- conflate SPF trees where logically correct, i.e. where routers' lists of critical TLV combinations overlap

7. Privacy Considerations

No privacy considerations apply to this document, as it only specifies routing control plane information.

8. Acknowledgements

This document is largely the result of discussions with Fred Baker.

9. Change Log

Initial Version: October 2014

10. Normative References

[IS-IS] ISO/IEC, "Intermediate System to Intermediate System Intra-Domain Routing Exchange Protocol for use in Conjunction with the Protocol for Providing the Connectionless-mode Network Service (ISO 8473)", ISO/IEC 10589:2002, Second Edition, 2002.

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

[RFC5120] Przygienda, T., Shen, N., and N. Sheth, "M-ISIS: Multi Topology (MT) Routing in Intermediate System to Intermediate Systems (IS-ISs)", RFC 5120, February 2008.

[RFC7356] Ginsberg, L., Previdi, S., and Y. Yang, "IS-IS Flooding Scope Link State PDUs (LSPs)", RFC 7356, September 2014.

Author's Address

David Lamparter
NetDEF
Leipzig 04103
Germany

Email: david@opensourcerouting.org