Network Working Group
Internet-Draft
Updates RFC 4379
Intended Status: Standards Track
Created: November 17, 2007
Expires: May 17, 2008

Adrian Farrel
Old Dog Consulting
George Swallow
Cisco

Future-Proof TLV Codepoint Space for LSP Ping
draft-farrel-mpls-lsp-ping-tlvs-01.txt

Status of this Memo

Abstract

Techniques for detecting Multi-Protocol Label Switched (MPLS) data plane failures are described in RFC 4379 and include the definition of a control protocol for testing and verifying Label Switched Path (LSP) connectivity that is known as LSP Ping. The protocol consists of a set of messages each of which contains data encoded as TLVs.

The LSP Ping protocol is being extended for several related uses. Each extension gives rise to the definition of new TLVs to be carried on the existing protocol messages.

The LSP Ping specification makes it mandatory that all TLVs are understood by each participating Label Switching Router (LSR) that receives an LSP Ping message. This makes future extensibility hard without upgrading all LSRs in the network.

This document defines ranges in the TLV codepoint space so that TLVs may be associated with different processing rules within an LSR if the TLV is unknown. This will make extensibility more simple.

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 [RFC2119].

1. Introduction

[RFC4379] defines the LSP Ping control protocol that can be used to detect Multi-Protocol Label Switched (MPLS) data plane failures. Specifically, LSP Ping is used to test and verify Label Switched Path (LSP) connectivity. The protocol consists of a set of messages each of which contains data encoded as TLVs.

The LSP Ping protocol is being extended for several related uses.

- [SELF-TEST] defines LSP Ping extensions to allow a Label Switching Router (LSR) to verify that its data plane is functioning for MPLS applications such as unicast forwarding and traffic engineering tunnels.

- [MPLS-BFD] makes use of the LSP Ping protocol exchanges to bootstrap an in-band bidirectional forwarding detection (BFD) mechanism for MPLS LSPs.

- [P2MP-PING] defines extensions to LSP Ping for use with point-to-multipoint LSPs.

Each extension gives rise to the definition of new TLVs to be carried on the existing protocol messages.

The LSP Ping specification makes it mandatory that all TLVs are understood by each participating Label Switching Router (LSR) that receives an LSP Ping message. This makes protocol extensions hard to deploy without upgrading all LSRs in the network. However, in many cases all that is actually required is that an LSR that does not recognise a TLV should ignore it and process the rest of the message as usual.

This document defines ranges in the TLV codepoint space so that TLVs may be associated with different processing rules within an LSR if the TLV is unknown.

2. Ranges and Rules for TLVs

2.1. Desired Function

Four desired functional behaviors are identified for an LSR that receives an LSP Ping message containing a TLV with an unrecognized Type value.

a. Support for the TLV is mandatory.

   If the TLV is not recognized the message MUST be dropped and an

   LSP Ping error MUST be reported according to the normal protocol procedures.

b. Support for the TLV is mandatory at an egress LSR, but is optional at a transit LSR.

   At an egress LSR, if the TLV is not recognized the message MUST be dropped and an LSP Ping error MUST be reported according to the normal protocol procedures.

   At a transit LSR, if the TLV is not recognized, the TLV SHOULD be ignored. The rest of the message MUST be processed as normal.

c. Support for the TLV is optional at all nodes.

   At a transit or egress LSR, if the TLV is not recognized, the TLV SHOULD be ignored. The rest of the message MUST be processed as normal.

d. Ignore the entire message if TLV is unrecognized.

   If the TLV is unrecognized the entire message SHOULD be ignored by the LSR. The LSR MUST NOT perform any processing associated with the message or the other TLVs on the message.

2.1.1. Extensibility Rules

The intention of the new rules and ranges defined in this document is that they are fully backward compatible with Type values already defined in existing RFCs and Internet-Drafts, and registered in the IANA registry.

2.2. Pre-Existing Ranges and Rules

The Type field of the TLV encoding for LSP Ping is a two byte (16 bit) field [RFC4379] so that the valid range for TLVs and sub-TLVs is 0 to 65535 inclusive.

The codespace has been managed by IANA according to the following rules using terminology from [RFC2434].

```
0x0000 - 0x3FFF Assignments via "Standards Action"
0x4000 - 0x7BFF Assignments via "Specification Required"
0x7C00 - 0x7FFF "Vendor Private Use"; MUST NOT be assigned.
0x8000 - 0xC009 Assignments via "Standards Action"
0xC00A - 0xFBFF Assignments via "Specification Required"
0xFC00 - 0xFFFF "Vendor Private Use"; MUST NOT be assigned.
```

As can be seen, these ranges are duplicated according to the setting of the most significant bit, but no use is made of this in [RFC4379]. All TLV Types currently registered are treated according to Rule a. in Section 2.1.

2.3. New Ranges and Rules

The rules set out in Section 2.1 require two bits to be fully encoded. In order to achieve backward compatibility, Rule a. must use the top two bits clear.

The following settings of the top two bits in the TLV Type field are defined:

```
00 Support for the TLV is mandatory (Rule a.)
01 Mandatory at egress, optional at transit (Rule b.)
10 Optional (Rule c.)
11 Ignore message (Rule d.)
```

This causes a change to the IANA registry as described in Section 4.

3. Ranges and Rules for Sub-TLVs

Each TLV defined for use in LSP Ping may carry Sub-TLVs. Sub-TLVs are formatted as TLVs, and have a 16-bit Type field.

3.1. Desired Function

Four desired functional behaviors are identified for an LSR that receives an LSP Ping message containing a TLV with a Sub-TLV with an unrecognized Type value.

e. Support for the Sub-TLV is mandatory.

If the Sub-TLV is not recognized the message MUST be dropped and

an LSP Ping error MUST be reported according to the normal protocol procedures.

f. Unrecognized Sub-TLV makes whole TLV unrecognized.

If a Sub-TLV is not recognized, the whole TLV must me treated as unrecognized and the whole TLV MUST be treated according to the rules set out in Section 2.1 as indicated by the top bit settings in the parent TLV Type.

g. Support for the Sub-TLV is optional.

Any LSR receiving an unrecognized Sub-TLV MUST ignore it and continue to process the rest of the TLV.

h. Ignore the entire TLV if the Sub-TLV is unrecognized.

If the Sub-TLV is unrecognized the entire TLV SHOULD be ignored by the LSR regardless of the parent TLV's type. The LSR MUST NOT perform any processing associated with the TLV, but SHOULD continue to process the other TLVs on the message.

3.1.1. Extensibility Rules

The intention of the new rules and ranges defined in this document is that they are fully backward compatible with Sub-TLV Type values already defined in existing RFCs and Internet-Drafts, and registered in the IANA registry.

3.2. Pre-Existing Ranges and Rules

The Type values defined for Sub-TLVs have the same assignment rules as the Type values for TLVs, but it should be noted that while the TLV registry has global scope, the Sub-TLV registry is scoped to the context of the parent TLV.

3.3. New Ranges and Rules

The rules set out in Section 3.1 require two bits to be fully encoded. In order to achieve backward compatibility, Rule e. must use the top two bits clear.

The following settings of the top two bits in the TLV Type field are defined:

00 Support for the Sub-TLV is mandatory (Rule e.)
01 Unrecognized Sub-TLV makes whole TLV unrecognized (Rule f.)
10 Support for the Sub-TLV is optional (Rule g.)
11 Ignore the entire TLV if the Sub-TLV is unrecognized (Rule h.)

This causes a change to the IANA registry as described in Section 4.

4. IANA Considerations

IANA maintains a registry entitled "Multi-Protocol Label Switching (MPLS) Label Switched Paths (LSPs) Parameters" that is used to assign and track codepoints for use with LSP Ping.

A subregistry called "TLVs and sub-TLVs" is used to assign and track Type values for use in TLVs and Sub-TLVs on LSP Ping messages.

This document requests IANA to modify the "TLVs and sub-TLVs" subregistry to further subdivide the codepoint space as described in this document. No changes to the values already assigned are requested, and IANA must not make such changes while applying the requests included in this document.

4.1. New Subregistry Preamble

The subregistry preamble previously read:

```
TLVs and sub-TLVs - per [RFC4379]
Registration Procedures:
0-16383 and 32768-49161 - Standards Action
16384-31743 and 49162-64511 - Specification Required (Experimental
                              RFC needed)
31744-32767 and 64512-65535 - Vendor Private Use, and MUST NOT be
                              allocated
```

IANA is requested to replace this with the following text:

```
TLVs and sub-TLVs - per [RFC4379] and [ID.thisdocument]

Each TLV and Sub-TLV in an LSP Ping message is identified by a
16-bit Type field. The Type field is divided into a 14-bit Type
value and a 2-bit action discriminator.

The TLV Types have global context, the Sub-TLV Types are scoped
only to their parent TLV.

Registration Procedures - as per [ID.thisdocument]

0 to 8191 are to be assigned by Standards Action
8192 to 12287 are to be assigned by Specification Required
12288 to 16383 are for Vendor Private Use and must not be assigned
```

4.2. New Registry Information

IANA is requested to modify the information stored in the registry to include the settings of the action discriminator bits for each Type value. The registry should look as follows:

Type	Sub-Type	Action	Value Field	Reference

So the existing registry should be updated as follows:

Type	Sub-Type	Action	Value Field	Reference
1		00	Target FEC Stack	[RFC4379]
	1	00	LDP IPv4 prefix	[RFC4379]
	2	00	LDP IPv6 prefix	[RFC4379]
	3	00	RSVP IPv4 LSP	[RFC4379]
	4	00	RSVP IPv6 LSP	[RFC4379]
	5	00	Not Assigned	[RFC4379]
	6	00	VPN IPv4 prefix	[RFC4379]
	7	00	VPN IPv6 prefix	[RFC4379]
	8	00	L2 VPN endpoint	[RFC4379]
	9	00	"FEC 128" Pseudowire (Deprecated)	[RFC4379]
	10	00	"FEC 128" Pseudowire	[RFC4379]
	11	00	"FEC 129" Pseudowire	[RFC4379]
	12	00	BGP labeled IPv4 prefix	[RFC4379]
	13	00	BGP labeled IPv6 prefix	[RFC4379]
	14	00	Generic IPv4 prefix	[RFC4379]
	15	00	Generic IPv6 prefix	[RFC4379]
	16	00	Nil FEC	[RFC4379]
2		00	Downstream Mapping	[RFC4379]
3		00	Pad	[RFC4379]
4		00	Not Assigned	[RFC4379]
5		00	Vendor Enterprise Number	[RFC4379]
6		00	Not Assigned	[RFC4379]
7		00	Interface and Label Stack	[RFC4379]
8		00	Not Assigned	[RFC4379]
9		00	Errored TLVs	[RFC4379]
	Any value	00	The type of the TLV not understood	[RFC4379]
10		00	Reply TOS Byte	[RFC4379]

4.3. A Note to IANA on Backward Compatibility

[RFC Editor - Please remove this section before publication.]

This note is provided to clarify backward compatibility in the modified codepoint space to reassure IANA that no existing allocations are impacted.

If the top two bits were included in the Type value, the new registry definition in Section 4.1 would give rise to the following twelve ranges. These can easily be seen to encompass all of the existing defined codepoints.

0 to 8191 are to be assigned by Standards Action
8192 to 12287 are to be assigned by Specification Required
12288 to 16383 are for Vendor Private Use and must not be assigned

16384 to 24575 are to be assigned by Standards Action
24576 to 28671 are to be assigned by Specification Required
28672 to 32767 are for Vendor Private Use and must not be assigned

32768 to 40959 are to be assigned by Standards Action
40960 to 45055 are to be assigned by Specification Required
45056 to 49151 are for Vendor Private Use and must not be assigned

49152 to 57343 are to be assigned by Standards Action
57344 to 61439 are to be assigned by Specification Required
61440 to 65535 are for Vendor Private Use and must not be assigned

5. Security Considerations

This document does not change the trust model for LSP Ping and so does not introduce security concerns over and above those described in [RFC4379].

It may be argued that defining these TLV action codes increases the scope for denial of service attacks by allowing spoof echo requests to include false TLVs that cause specific actions from the LSRs in the network. However, this is not the case since the existing default behavior is to generate an echo response indicating an error for any unknown TLV. Thus, the new behaviors introduced in this document do not exacerbate the scope for such attacks.

6. Acknowledgements

Thanks to Nitin Bahadur for discussions.

7. Intellectual Property Considerations

8. Normative References

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

[RFC2434]   Narten, T., and H. Alvestrand, "Guidelines for Writing an IANA Considerations Section in RFCs", BCP 26, RFC 2434, October 1998.

[RFC4379]   Kompella, K., and Swallow, G., "Detecting Multi-Protocol Label Switched (MPLS) Data Plane Failures", RFC 4379, February 2006.

9. Informative References

[SELF-TEST] Swallow, G., Kompella, K., and Tappan, D., "Label Switching Router Self-Test", draft-ietf-mpls-lsr-self-test, work in progress.

[MPLS-BFD] Aggarwal, R., Kompella, K., Nadeau, T., and Swallow, G., "BFD For MPLS LSPs", draft-ietf-bfd-mpls, work in progress.

[P2MP-PING] Yasukawa, S. (Ed.), and Farrel, A. (Ed.), "Detecting Data Plane Failures in Point-to-Multipoint Multiprotocol Label Switching (MPLS) - Extensions to LSP Ping", draft-ietf-mpls-p2mp-lsp-ping, work in progress.

10. Authors' Addresses

Adrian Farrel
Old Dog Consulting
EMail: adrian@olddog.co.uk

George Swallow
Cisco Systems, Inc.
1414 Massachusetts Ave
Boxborough, MA 01719
Email: swallow@cisco.com

11. Full Copyright Statement