

Network Working Group
Internet-Draft
Intended status: Standards Track
Expires: June 14, 2014

L. Jin
F. Jounay
France Telecom
I. Wijnands
Cisco Systems, Inc
N. Leymann
Deutsche Telekom AG
December 11, 2013

LDP Extensions for Hub & Spoke Multipoint Label Switched Path
draft-ietf-mpls-mldp-hsmp-05.txt

Abstract

This draft introduces a hub & spoke multipoint (HSMP) Label Switched Path (LSP), which allows traffic both from root to leaf through point-to-multipoint (P2MP) LSP and also leaf to root along the reverse path. That means traffic entering the HSMP LSP from application/customer at the root node travels downstream to each leaf node, exactly as if it is travelling downstream along a P2MP LSP to each leaf node. Upstream traffic entering the HSMP LSP at any leaf node travels upstream along the tree to the root, as if it is unicast to the root. Direct communication among the leaf nodes is not allowed.

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](#) [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 <http://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."

Internet-Draft

[draft-ietf-mpls-mldp-hsmp](#)

December 2013

This Internet-Draft will expire on June 14, 2014.

Copyright Notice

Copyright (c) 2013 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.

Internet-Draft

[draft-ietf-mpls-mldp-hsmp](#)

December 2013

Table of Contents

1.	Introduction	4
2.	Terminology	4
3.	Setting up HSMP LSP with LDP	4
3.1.	Support for HSMP LSP Setup with LDP	5
3.2.	HSMP FEC Elements	6
3.3.	Using the HSMP FEC Elements	6
3.4.	HSMP LSP Label Map	7
3.4.1.	HSMP LSP Leaf Node Operation	8
3.4.2.	HSMP LSP Transit Node Operation	8
3.4.3.	HSMP LSP Root Node Operation	9
3.5.	HSMP LSP Label Withdraw	10
3.5.1.	HSMP Leaf Operation	10
3.5.2.	HSMP Transit Node Operation	10
3.5.3.	HSMP Root Node Operation	10
3.6.	HSMP LSP Upstream LSR Change	11
3.7.	Determining Forwarding Interface	11
4.	HSMP LSP on a LAN	11
5.	Redundancy Considerations	12
6.	Failure Detection of HSMP LSP	12
7.	Security Considerations	13
8.	IANA Considerations	13
8.1.	New LDP FEC Element types	13
8.2.	HSMP LSP capability TLV	13
8.3.	New sub-TLVs for the Target Stack TLV	14
9.	Acknowledgement	14
10.	References	14
10.1.	Normative references	14
10.2.	Informative References	15
	Authors' Addresses	15

1. Introduction

The point-to-multipoint (P2MP) Label Switched Path (LSP) defined in [[RFC6388](#)] allows traffic to transmit from root to several leaf nodes, and multipoint-to-multipoint (MP2MP) LSP allows traffic from every node to transmit to every other node. This draft introduces a hub & spoke multipoint (HSMP) LSP, which has one root node and one or more leaf nodes. HSMP LSP allows traffic both from root to leaf through downstream LSP and also leaf to root along the upstream LSP. That means traffic entering the HSMP LSP at the root node travels along downstream LSP, exactly as if it is travelling along a P2MP LSP, and traffic entering the HSMP LSP at any other leaf nodes travels along upstream LSP toward only the root node. The upstream LSP should be thought of unicast LSP to the root node, except that it follows the reverse direction of the downstream LSP, rather than routing protocol based unicast path. The combination of upstream LSPs initiated from all leaf nodes forms a multipoint-to-point LSP.

2. Terminology

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](#)].

This document uses some terms and acronyms as follows:

mLDP: Multipoint extensions for Label Distribution Protocol (LDP) defined in [[RFC6388](#)].

P2MP LSP: point-to-multipoint Label Switched Path. An LSP that has one Ingress Label Switching Router (LSR) and one or more Egress LSRs.

MP2MP LSP: multipoint-to-multipoint Label Switched Path. An LSP that connects a set of nodes, such that traffic sent by any node in the LSP is delivered to all others.

HSMP LSP: hub & spoke multipoint Label Switched Path. An LSP that has one root node and one or more leaf nodes, allows traffic from root to all leaf nodes along downstream P2MP LSP and also leaf to root node along the upstream unicast LSP.

3. Setting up HSMP LSP with LDP

HSMP LSP is similar to MP2MP LSP described in [\[RFC6388\]](#), with the difference that, when the leaf LSRs send traffic on the LSP, the

traffic is first delivered only to the root node and follows the upstream path from the leaf node to the root node. The root node then distributes the traffic on the P2MP tree to all of the leaf nodes.

HSMP LSP consists of a downstream path and upstream path. The downstream path is same as P2MP LSP, while the upstream path is only from leaf to root node, without communication between leaf and leaf nodes. The transmission of packets from the root node of an HSMP LSP to the receivers (the leaf nodes) is identical to that of a P2MP LSP. Traffic from a leaf node to the root follows the upstream path that is the reverse of the path from the root to the leaf. Unlike an MP2MP LSP, traffic from a leaf node does not branch toward other leaf nodes, but is sent direct to the root where it is placed on the P2MP path and distributed to all leaf nodes including the original sender.

To set up the upstream path of an HSMP LSP, ordered mode MUST be used. Ordered mode can guarantee a leaf to start sending packets to root immediately after the upstream path is installed, without being dropped due to an incomplete LSP.

3.1. Support for HSMP LSP Setup with LDP

HSMP LSP requires the LDP capabilities [[RFC5561](#)] for nodes to indicate that they support setup of HSMP LSPs. An implementation supporting the HSMP LSP procedures specified in this document MUST implement the procedures for Capability Parameters in Initialization Messages. Advertisement of the HSMP LSP Capability indicates support of the procedures for HSMP LSP setup.

A new Capability Parameter TLV is defined, the HSMP LSP Capability Parameter. Following is the format of the HSMP LSP Capability Parameter.




Figure 1. HSMP LSP Capability Parameter encoding

U-bit: Unknown TLV bit, as described in [[RFC5036](#)]. The value MUST be 1. The unknown TLV MUST be silently ignored and the rest of the message processed as if the unknown TLV did not exist.

F-bit: Forward unknown TLV bit, as described in [[RFC5036](#)]. The value of this bit MUST be 0 since a Capability Parameter TLV is sent only in Initialization and Capability messages, which are not forwarded.

The length SHOULD be 1, and the S bit and reserved bits are defined in [[RFC5561](#)] [section 3](#).

The HSMP LSP Capability Parameter type is to be assigned by IANA.

If the peer has not advertised the corresponding capability, then label messages using the HSMP FEC Element SHOULD NOT be sent to the peer. Since ordered mode is applied for HSMP LSP signalling, the label message break would ensure that the initiating leaf node is unable to establish the upstream path to root node.

[3.2.](#) HSMP FEC Elements

Similar as MP2MP LSP, we define two new protocol entities, the HSMP Downstream FEC Element and Upstream FEC Element. If a FEC TLV contains one of the HSMP FEC Elements, the HSMP FEC Element MUST be the only FEC Element in the FEC TLV. The structure, encoding and error handling for the HSMP Downstream FEC Element and Upstream FEC Element are the same as for the P2MP FEC Element described in [\[RFC6388\] Section 2.2](#). The difference is that two additional new FEC types are defined: HSMP Downstream FEC (to be assigned by IANA) and HSMP Upstream FEC (to be assigned by IANA).

[3.3.](#) Using the HSMP FEC Elements

In order to describe the message processing clearly, the entries in the list below define the processing of the HSMP FEC Elements. Additionally, the entries defined in [\[RFC6388\] section 3.3](#) are also reused in the following sections.

1. HSMP downstream LSP <X, Y> (or simply downstream <X, Y>): an HSMP LSP downstream path with root node address X and opaque value Y.
2. HSMP upstream LSP <X, Y> (or simply upstream <X, Y>): an HSMP LSP upstream path for root node address X and opaque value Y which will be used by any of downstream node to send traffic upstream to root node.
3. HSMP downstream FEC Element <X, Y>: a FEC Element with root node address X and opaque value Y used for a downstream HSMP LSP.
4. HSMP upstream FEC Element <X, Y>: a FEC Element with root node address X and opaque value Y used for an upstream HSMP LSP.

5. HSMP-D Label Mapping <X, Y, L>: A Label Mapping message with a single HSMP downstream FEC Element <X, Y> and label TLV with label L. Label L MUST be allocated from the per-platform label space of the LSR sending the Label Mapping Message.

6. HSMP-U Label Mapping <X, Y, Lu>: A Label Mapping message with a single HSMP upstream FEC Element <X, Y> and label TLV with label Lu. Label Lu MUST be allocated from the per-platform label space of the

LSR sending the Label Mapping Message.

7. HSMP-D Label Withdraw $\langle X, Y, L \rangle$: a Label Withdraw message with a FEC TLV with a single HSMP downstream FEC Element $\langle X, Y \rangle$ and label TLV with label L.

8. HSMP-U Label Withdraw $\langle X, Y, Lu \rangle$: a Label Withdraw message with a FEC TLV with a single HSMP upstream FEC Element $\langle X, Y \rangle$ and label TLV with label Lu.

9. HSMP-D Label Release $\langle X, Y, L \rangle$: a Label Release message with a FEC TLV with a single HSMP downstream FEC Element $\langle X, Y \rangle$ and Label TLV with label L.

10. HSMP-U Label Release $\langle X, Y, Lu \rangle$: a Label Release message with a FEC TLV with a single HSMP upstream FEC Element $\langle X, Y \rangle$ and label TLV with label Lu.

3.4. HSMP LSP Label Map

This section specifies the procedures for originating HSMP Label Mapping messages and processing received HSMP Label Mapping messages for a particular HSMP LSP. The procedure of downstream HSMP LSP is similar as that of downstream MP2MP LSP described in [RFC6388]. When LDP operates in Ordered Label Distribution Control mode [RFC5036], the upstream LSP will be set up by sending HSMP LSP LDP Label Mapping message with a label which is allocated by upstream LSR to its downstream LSR hop by hop from root to leaf node, installing the upstream forwarding table by every node along the LSP. The detail procedure of setting up upstream HSMP LSP is different with that of upstream MP2MP LSP, and is specified in below section.

All labels discussed here are downstream-assigned [RFC5332] except those which are assigned using the procedures described in [Section 4](#).

Determining the upstream LSR for the HSMP LSP $\langle X, Y \rangle$ follows the procedure for a P2MP LSP described in [RFC6388] [Section 2.4.1.1](#). That is, a node Z that wants to join an HSMP LSP $\langle X, Y \rangle$ determines the LDP peer U that is Z's next-hop on the best path from Z to the root node X. If there are multiple upstream LSRs, local algorithm

should be applied to ensure that there is a single upstream LSRs for

a particular LSP.

To determining one's HSMP downstream LSR, an upstream LDP peer which receives a Label Mapping with HSMP downstream FEC Element from an LDP peer D will treat D as HSMP downstream LDP peer.

3.4.1. HSMP LSP Leaf Node Operation

The leaf node operation is much the same as the operation of MP2MP LSP defined in [\[RFC6388\] Section 3.3.1.4](#). The only difference is the FEC elements as specified below.

A leaf node Z of an HSMP LSP $\langle X, Y \rangle$ determines its upstream LSR U for $\langle X, Y \rangle$ as per [Section 3.3](#), allocates a label L, and sends an HSMP-D Label Mapping $\langle X, Y, L \rangle$ to U. Leaf node Z expects an HSMP-U Label Mapping $\langle X, Y, Lu \rangle$ from node U and checks whether it already has forwarding state for upstream $\langle X, Y \rangle$. If not, Z creates forwarding state to push label Lu onto the traffic that Z wants to forward over the HSMP LSP. How it determines what traffic to forward on this HSMP LSP is outside the scope of this document.

3.4.2. HSMP LSP Transit Node Operation

The procedure of HSMP-D Label Mapping message is much the same as processing MP2MP-D Label Mapping message defined in [\[RFC6388\] Section 3.3.1.5](#). The processing of HSMP-U Label Mapping message is different with that of MP2MP-U Label Mapping message as specified below.

Suppose node Z receives an HSMP-D Label Mapping $\langle X, Y, L \rangle$ from LSR D. Z checks whether it has forwarding state for downstream $\langle X, Y \rangle$. If not, Z determines its upstream LSR U for $\langle X, Y \rangle$ as per [Section 3.3](#). Using this Label Mapping to update the label forwarding table MUST NOT be done as long as LSR D is equal to LSR U. If LSR U is different from LSR D, Z will allocate a label L' and install downstream forwarding state to swap label L' with label L over interface I associated with LSR D and send an HSMP-D Label Mapping $\langle X, Y, L' \rangle$ to U. Interface I is determined via the procedures in [Section 3.7](#).

If Z already has forwarding state for downstream $\langle X, Y \rangle$, all that Z needs to do in this case is check that LSR D is not equal to the upstream LSR of $\langle X, Y \rangle$ and update its forwarding state. Assuming its old forwarding state was $L' \rightarrow \{ \langle I1, L1 \rangle \langle I2, L2 \rangle \dots, \langle In, Ln \rangle \}$, its new forwarding state becomes $L' \rightarrow \{ \langle I1, L1 \rangle \langle I2, L2 \rangle \dots, \langle In, Ln \rangle, \langle I, L \rangle \}$. If the LSR D is equal to the installed upstream LSR, the Label Mapping from LSR D MUST be retained and MUST NOT update the label forwarding table.

Node Z checks if upstream LSR U already has assigned a label Lu to upstream <X, Y>. If not, transit node Z waits until it receives an HSMP-U Label Mapping <X, Y, Lu> from LSR U. Once the HSMP-U Label Mapping is received from LSR U, node Z checks whether it already has forwarding state upstream <X, Y> with incoming label Lu' and outgoing label Lu. If it does not, it allocates a label Lu' and creates a new label swap for Lu' with Label Lu over interface Iu. Interface Iu is determined via the procedures in [Section 3.7](#). Node Z determines the downstream HSMP LSR as per [Section 4.3.1](#), and sends an HSMP-U Label Mapping <X, Y, Lu'> to node D.

Since a packet from any downstream node is forwarded only to the upstream node, the same label (representing the upstream path) SHOULD be distributed to all downstream nodes. This differs from the procedures for MP2MP LSPs [[RFC6388](#)], where a distinct label must be distributed to each downstream node. The forwarding state upstream <X, Y> on node Z will be like this {<Lu'>, <Iu Lu>}. Iu means the upstream interface over which Z receives HSMP-U Label Map <X, Y, Lu> from LSR U. Packets from any downstream interface over which Z sends HSMP-U Label Map <X, Y, Lu'> with label Lu' will be forwarded to Iu with label Lu' swap to Lu.

[3.4.3](#). HSMP LSP Root Node Operation

The procedure of HSMP-D Label Mapping message is much the same as processing MP2MP-D Label Mapping message defined in [[RFC6388](#)] [Section 3.3.1.6](#). The processing of HSMP-U Label Mapping message is different with that of MP2MP-U Label Mapping message as specified below.

Suppose the root node Z receives an HSMP-D Label Mapping <X, Y, L> from node D. Z checks whether it already has forwarding state for downstream <X, Y>. If not, Z creates downstream forwarding state and installs a outgoing label L over interface I. Interface I is determined via the procedures in [Section 3.7](#). If Z already has forwarding state for downstream <X, Y>, then Z will add label L over interface I to the existing state.

Node Z checks if it has forwarding state for upstream <X, Y>. If not, Z creates a forwarding state for incoming label Lu' that indicates that Z is the HSMP LSP egress LER. E.g., the forwarding state might specify that the label stack is popped and the packet passed to some specific application. Node Z determines the downstream HSMP LSR as per [Section 3.3](#), and sends an HSMP-U Label Map <X, Y, Lu'> to node D.

Since Z is the root of the tree, Z will not send an HSMP-D Label Map

and will not receive an HSMP-U Label Mapping.

Root node could also be a leaf node, and it is able to determine what traffic to forward on this HSMP LSP which is outside the scope of this document.

[3.5.](#) HSMP LSP Label Withdraw

[3.5.1.](#) HSMP Leaf Operation

If a leaf node Z discovers that it has no need to be an Egress LSR for that LSP (by means outside the scope of this document), then it SHOULD send an HSMP-D Label Withdraw $\langle X, Y, L \rangle$ to its upstream LSR U for $\langle X, Y \rangle$, where L is the label it had previously advertised to U for $\langle X, Y \rangle$. Leaf node Z will also send an unsolicited HSMP-U Label Release $\langle X, Y, Lu \rangle$ to U to indicate that the upstream path is no longer used and that label Lu can be removed.

Leaf node Z expects the upstream router U to respond by sending a downstream Label Release for L.

[3.5.2.](#) HSMP Transit Node Operation

If a transit node Z receives an HSMP-D Label Withdraw message $\langle X, Y, L \rangle$ from node D, it deletes label L from its forwarding state downstream $\langle X, Y \rangle$. Node Z sends an HSMP-D Label Release message with label L to D. There is no need to send an HSMP-U Label Withdraw $\langle X, Y, Lu \rangle$ to D because node D already removed Lu and sent a label release for Lu to Z.

If deleting L from Z's forwarding state for downstream $\langle X, Y \rangle$ results in no state remaining for $\langle X, Y \rangle$, then Z propagates the HSMP-D Label Withdraw $\langle X, Y, L \rangle$ to its upstream node U for $\langle X, Y \rangle$. Z should also check if there are any incoming interface in forwarding state upstream $\langle X, Y \rangle$. If all downstream nodes are released and there is no incoming interface, Z should delete the forwarding state upstream $\langle X, Y \rangle$ and send HSMP-U Label Release message to its upstream node. Otherwise, no HSMP-U Label Release message will be sent to the upstream node.

[3.5.3.](#) HSMP Root Node Operation

When the root node of an HSMP LSP receives an HSMP-D Label Withdraw and HSMP-U Label Release message, the procedure is the same as that for transit nodes, except that the root node will not propagate the Label Withdraw and Label Release upstream (as it has no upstream).

[3.6.](#) HSMP LSP Upstream LSR Change

The procedure for changing the upstream LSR is the same as defined in [\[RFC6388\] Section 2.4.3](#), only with different processing FEC Element.

When the upstream LSR changes from U to U', node Z should set up the HSMP LSP <X, Y> to U' by applying procedures in [Section 3.4](#). Z will also remove HSMP LSP <X, Y> to U by applying procedure in [Section 3.5](#).

To set up HSMP LSP to U' before/after removing HSMP LSP to U is a local matter, and the recommended default behavior is to remove before adding.

[3.7.](#) Determining Forwarding Interface

The co-routed path between upstream and downstream LSP would be achieved for HSMP LSP. Both LSR U and LSR D would ensure the same interface to send traffic by applying some procedures. For a network with symmetric IGP cost configuration, the following procedure MAY be used. To determine the downstream interface, LSR U MUST do a lookup in the unicast routing table to find the best interface and next-hop to reach LSR D. If the next-hop and interface are also advertised by LSR D via the LDP session, it should be used to transmit the packet to LSR D. Determine the upstream interface mechanism is same as determining the downstream interface by exchanging the role of LSR U and LSR D. If symmetric IGP cost could not be ensured, static route configuration on LSR U and D could also be a possible way to ensure co-routed path.

If co-routed is not required for HSMP LSP, the procedure defined in [\[RFC6388\] Section 2.4.1.2](#) could be applied. LSR U is free to

transmit the packet on any of the interfaces to LSR D. The algorithm it uses to choose a particular interface is a local matter. Determine the upstream interface mechanism is the same as determining the downstream interface.

4. HSMP LSP on a LAN

The procedure to process the downstream HSMP LSP on a LAN is much the same as downstream MP2MP LSP described in [\[RFC6388\] section 6.1.1](#).

When establishing the downstream path of an HSMP LSP, as defined in [\[RFC6389\]](#), a Label Request message for an LSP label is sent to the upstream LSR. The upstream LSR should send Label Mapping message that contains the LSP label for the downstream HSMP FEC and the upstream LSR context label defined in [\[RFC5331\]](#). When the LSR

forwards a packet downstream on one of those LSPs, the packet's top label must be the "upstream LSR context label", and the packet's second label is "LSP label". The HSMP downstream path will be installed in the context-specific forwarding table corresponding to the upstream LSR label. Packets sent by the upstream LSR can be forwarded downstream using this forwarding state based on a two-label lookup.

The upstream path of an HSMP LSP on a LAN is the same as the one on other kind of links. That is, the upstream LSR must send Label Mapping message that contains the LSP label for upstream HSMP FEC to downstream node. Packets travelling upstream need to be forwarded in the direction of the root by using the label allocated for upstream HSMP FEC.

5. Redundancy Considerations

In some scenarios, it is necessary to provide two root nodes for redundancy purpose. One way to implement this is to use two independent HSMP LSPs acting as active/standby. At one time, only one HSMP LSP will be active, and the other will be standby. After detecting the failure of active HSMP LSP, the root and leaf nodes will switch the traffic to the standby HSMP LSP which takes on the role as active HSMP LSP. The detail of redundancy mechanism is out

of the scope.

6. Failure Detection of HSMP LSP

The idea of LSP ping for HSMP LSPs could be expressed as an intention to test the LSP Ping Echo Request packets that enter at the root along a particular downstream path of HSMP LSP, and end their MPLS path on the leaf. The leaf node then sends the LSP Ping Echo Reply along the upstream path of HSMP LSP, and end on the root that are the (intended) root node.

New sub-TLVs are required to be assigned by IANA in Target FEC Stack TLV and Reverse-path Target FEC Stack TLV to define the corresponding HSMP-downstream FEC type and HSMP-upstream FEC type. In order to ensure the leaf node to send the LSP Ping Echo Reply along the HSMP upstream path, the R bit (Validate Reverse Path) in Global Flags Field defined in [[RFC6426](#)] is reused here.

The node processing mechanism of LSP Ping Echo Request and Echo Reply for HSMP LSP is inherited from [[RFC6425](#)] and [[RFC6426](#)] [Section 3.4](#), except the following:

1. The root node sending LSP Ping Echo Request message for HSMP LSP MUST attach Target FEC Stack with HSMP downstream FEC, and set R bit to '1' in Global Flags Field.
2. When the leaf node receiving the LSP Ping Echo Request, it MUST send the LSP Ping Echo Reply to the associated HSMP upstream path. The Reverse-path Target FEC Stack TLV attached by leaf node in Echo Reply message SHOULD contain the sub-TLV of associated HSMP upstream FEC.

7. Security Considerations

The same security considerations apply as for the MP2MP LSP described in [[RFC6388](#)] and [[RFC6425](#)].

Although this document introduces new FEC Elements and corresponding procedures, the protocol does not bring any new security issues

compared to [[RFC6388](#)] and [[RFC6425](#)].

[8.](#) IANA Considerations

[8.1.](#) New LDP FEC Element types

This document requires allocation of two new LDP FEC Element types from the "Label Distribution Protocol (LDP) Parameters registry" the "Forwarding Equivalence Class (FEC) Type Name Space":

1. the HSMP-upstream FEC type - requested value TBD
2. the HSMP-downstream FEC type - requested value TBD

The values should be allocated using the lowest free values from the "IETF Consensus"-range (0-127).

[8.2.](#) HSMP LSP capability TLV

This document requires allocation of one new code points for the HSMP LSP capability TLV from "Label Distribution Protocol (LDP) Parameters registry" the "TLV Type Name Space":

HSMP LSP Capability Parameter - requested value TBD

The value should be allocated from the range 0x0901-0x3DFF (IETF Consensus) using the first free value within this range.

[8.3.](#) New sub-TLVs for the Target Stack TLV

This document requires allocation of two new sub-TLV types for inclusion within the LSP ping [[RFC4379](#)] Target FEC Stack TLV (TLV type 1) and Reverse-path Target FEC Stack TLV (TLV type 16).

1. the HSMP-upstream LDP FEC Stack - requested value TBD
2. the HSMP-downstream LDP FEC Stack - requested value TBD

The value should be allocated from the IETF Standards Action range

(0-16383) that is used for mandatory and optional sub-TLVs that requires a response if not understood. The value should be allocated using the lowest free value within this range.

9. Acknowledgement

The author would like to thank Eric Rosen, Sebastien Jobert, Fei Su, Edward, Mach Chen, Thomas Morin, Loa Andersson for their valuable comments.

10. References

10.1. Normative references

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), March 1997.
- [RFC5331] Aggarwal, R., Rekhter, Y., and E. Rosen, "MPLS Upstream Label Assignment and Context-Specific Label Space", [RFC 5331](#), August 2008.
- [RFC5332] Eckert, T., Rosen, E., Aggarwal, R., and Y. Rekhter, "MPLS Multicast Encapsulations", [RFC 5332](#), August 2008.
- [RFC5561] Thomas, B., Raza, K., Aggarwal, S., Aggarwal, R., and JL. Le Roux, "LDP Capabilities", [RFC 5561](#), July 2009.
- [RFC6388] Wijnands, IJ., Minei, I., Kompella, K., and B. Thomas, "Label Distribution Protocol Extensions for Point-to-Multipoint and Multipoint-to-Multipoint Label Switched Paths", [RFC 6388](#), November 2011.
- [RFC6389] Aggarwal, R. and JL. Le Roux, "MPLS Upstream Label Assignment for LDP", [RFC 6389](#), November 2011.

- [RFC6425] Saxena, S., Swallow, G., Ali, Z., Farrel, A., Yasukawa, S., and T. Nadeau, "Detecting Data-Plane Failures in Point-to-Multipoint MPLS - Extensions to LSP Ping", [RFC 6425](#), November 2011.

- [RFC6426] Gray, E., Bahadur, N., Boutros, S., and R. Aggarwal, "MPLS On-Demand Connectivity Verification and Route Tracing", [RFC 6426](#), November 2011.

10.2. Informative References

- [RFC4379] Kompella, K. and G. Swallow, "Detecting Multi-Protocol Label Switched (MPLS) Data Plane Failures", [RFC 4379](#), February 2006.
- [RFC5036] Andersson, L., Minei, I., and B. Thomas, "LDP Specification", [RFC 5036](#), October 2007.

Authors' Addresses

Lizhong Jin
Shanghai, China

Email: lizho.jin@gmail.com

Frederic Jounay
France Telecom
2, avenue Pierre-Marzin
22307 Lannion Cedex, FRANCE

Email: frederic.jounay@orange.ch

IJsbrand Wijnands
Cisco Systems, Inc
De kleetlaan 6a
Diegem 1831, Belgium

Email: ice@cisco.com

Internet-Draft

[draft-ietf-mpls-mldp-hsmp](#)

December 2013

Nicolai Leymann
Deutsche Telekom AG
Winterfeldtstrasse 21
Berlin 10781

Email: N.Leymann@telekom.de

