

Internet Engineering Task Force
Internet-Draft
Intended status: Standards Track
Expires: August, 2014

Q. Zhao
Huawei Technology
K. Raza
C. Zhou
Cisco Systems
L. Fang
Microsoft
L. Li
China Mobile
D. King
Old Dog Consulting
February 14, 2014

LDP Extensions for Multi Topology
draft-ietf-mpls-ldp-multi-topology-10.txt

Abstract

Multi-Topology (MT) routing is supported in IP networks with the use of MT aware IGP. In order to provide MT routing within Multiprotocol Label Switching (MPLS) Label Distribution Protocol (LDP) networks new extensions are required.

This document describes the LDP protocol extensions required to support MT routing in an MPLS environment.

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."

This Internet-Draft will expire on August, 2014.

Copyright Notice

Copyright (c) 2014 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.

This document may contain material from IETF Documents or IETF Contributions published or made publicly available before November 10, 2008. The person(s) controlling the copyright in some of this material may not have granted the IETF Trust the right to allow modifications of such material outside the IETF Standards Process. Without obtaining an adequate license from the person(s) controlling the copyright in such materials, this document may not be modified outside the IETF Standards Process, and derivative works of it may not be created outside the IETF Standards Process, except to format it for publication as an RFC or to translate it into languages other than English.

Table of Contents

1.	Introduction	.3
2.	Terminology	.4
3.	Signaling Extensions	.4
3.1.	Topology-Scoped Forwarding Equivalence Class (FEC)	.4
3.2.	New Address Families: MT IP	.4
3.3.	LDP FEC Elements with MT IP AF	.5
3.4.	IGP MT-ID Mapping and Translation	.6
3.5.	LDP MT Capability Advertisement	.6
3.5.1.	Protocol Extension	.6
3.5.2.	Procedures	.7
3.6.	LDP Sessions	.8
3.7.	Reserved MT ID Values	.8
4.	MT Applicability on FEC-based features	.9
4.1.	Typed Wildcard FEC Element	.9
4.2.	End-of-LIB	.9
4.3.	LSP Ping	.9
4.3.1.	New FEC Sub-Types	.10
4.3.2.	MT LDP IPv4 FEC Sub-TLV	.10
4.3.3.	MT LDP IPv6 FEC Sub-TLV	.11
4.3.4.	Operation Considerations	.11
5.	Error Handling	.11
5.1.	MT Error Notification for Invalid Topology ID	.12
6.	Backwards Compatibility	.12

7.	MPLS Forwarding in MT	12
--------------------	---------------------------------	--------------------

8.	Security Consideration	12
9.	IANA Considerations	13
10.	Manageability Considerations	14
10.1.	Control of Function and Policy	14
10.2.	Information and Data Models	14
10.3.	Liveness Detection and Monitoring	14
10.4.	Verify Correct Operations	14
10.5.	Requirements On Other Protocols	15
10.6.	Impact On Network Operations	15
11.	Contributors	15
12.	Acknowledgement	15
13.	References	15
13.1.	Normative References	15
13.2.	Informative References	16
	Authors' Addresses	16

[1.](#) Introduction

Multi-Topology (MT) routing is supported in IP networks with the use of MT aware IGPs. It would be advantageous for communications Service Providers (CSP) to support Multiple Topologies (MT) within MPLS environments (MPLS-MT). The benefits of MPLS-MT enabled networks include:

- o A CSP may want to assign varying Quality of Service (QoS) profiles to traffic, based on a specific MT.
- o Separate routing and MPLS domains may be used to isolate multicast and IPv6 islands within the backbone network.
- o Specific IP address space could be routed across an MT based on security or operational isolation requirements.
- o Low latency links could be assigned to an MT for delay sensitive traffic.
- o Management traffic could be separated from customer traffic using multiple MTs, where the management traffic MT does not use links that carry customer traffic.

This document describes the Label Distribution Protocol (LDP) procedures and protocol extensions required to support MT routing in an MPLS environment.

This document also updates [RFC4379](#) by defining two new FEC types for Label Switched Path (LSP) ping.

2. Terminology

This document uses MPLS terminology defined in [[RFC5036](#)]. Additional terms are defined below:

- o MT-ID: A 16 bit value used to represent the Multi-Topology ID.
- o Default MT Topology: A topology that is built using the MT-ID default value of 0.
- o MT Topology: A topology that is built using the corresponding MT-ID.

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](#)].

3. Signaling Extensions

3.1. Topology-Scoped Forwarding Equivalence Class (FEC)

LDP assigns and binds a label to a Forwarding Equivalence Class (FEC), where a FEC is a list of one or more FEC elements. To setup LSPs for unicast IP routing paths, LDP assigns local labels for IP prefixes, and advertises these labels to its peers so that an LSP is setup along the routing path. To setup MT LSPs for IP prefixes under a given topology scope, the LDP "prefix-related" FEC element must be extended to include topology information. This implies that MT-ID becomes an attribute of Prefix-related FEC element, and all FEC-Label binding operations are performed under the context of given topology (MT-ID).

The following Sub[section 3.2](#)(New Address Families (AF): MT IP) defines the extension required to bind "prefix-related" FEC to a topology.

3.2. New Address Families: MT IP

The LDP base specification [[RFC5036](#)] ([Section 4.1](#)) defines the "Prefix" FEC Element. The "Prefix" encoding is defined for a given "Address Family" (AF), and has length (in bits) specified by the "PreLen" field.

To extend IP address families for MT, two new Address Families named "MT IP" and "MT IPv6" are used to specify IPv4 and IPv6 prefixes within a topology scope.

The format of data associated with these new Address Families is

described below:

Zhao, et al.

Expires August, 2014

[Page 4]

The MT Typed Wildcard FEC element encoding is as follows:


```

      0              1              2              3
0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|Typed Wcard (5)|   FEC Type   |   Len = 6   |   AF = MT IP ..|
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|... or MT IPv6 |           MT ID           |
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+

```

Figure 3: MT Typed Wildcard FEC Element

The above format can be used for any LDP FEC Element that allows use of IP/IPv6 address family. In the scope of this document, the allowed "FEC Type" in a MT Typed Wildcard FEC Element is "Prefix" FEC element.

3.4. IGP MT-ID Mapping and Translation

The non-reserved non-special IGP MT-ID values can be used and carried in LDP without the need for translation. However, there is a need for translating reserved or special IGP MT-ID values to corresponding LDP MT-IDs. The assigned, unassigned and special LDP MT-ID values are requested In [Section 9](#). (IANA Considerations).

How future LDP MT-ID values are allocated are out of scope of this document. Instead a new Internet-Draft will be created to document the allocation policy and process for requesting new MT-ID values.

3.5. LDP MT Capability Advertisement

3.5.1. Protocol Extension

We specify a new LDP capability, named "Multi-Topology (MT)", which is defined in accordance with LDP Capability definition guidelines [[RFC5561](#)]. The LDP "MT" capability can be advertised by an LDP speaker to its peers either during the LDP session initialization or after the LDP session is setup to announce LSR capability to support MT for the given IP address family. An LDP speaker MUST NOT send messages containing MT FEC elements unless the peer has said it can handle it.

The MT capability is specified using "Multi-Topology Capability" TLV. The "Multi-Topology Capability" TLV format is in accordance with LDP capability guidelines as defined in [[RFC5561](#)]. To be able to specify IP address family, the capability specific data (i.e. "Capability Data" field of Capability TLV) is populated using "Typed Wildcard FEC Element" as defined in [[RFC5918](#)].

The format of "Multi-Topology Capability" TLV is as follows:

- o The capability for supporting multi-topology in LDP can be

advertised during LDP session initialization stage by including the LDP MT capability TLV in LDP Initialization message. After an LDP session is established, the MT capability can also be advertised or withdrawn using Capability message (only if "Dynamic Announcement" capability [[RFC5561](#)] has already been successfully negotiated).

- o If an LSR has not advertised MT capability, its peer MUST NOT send any LDP messages with FEC elements that include MT identifier to this LSR.
- o If an LSR is changed from non-MT capable to MT capable, it sets the S bit in MT capability TLV and advertises via the Capability message (if it supports Dynamic Announcement Capability). The existing LSP is treated as LSP for default MT (ID 0).
- o o If an LSR is changed from LDP-MT capable to non-MT capable, it initiates withdraw of all label mapping for existing LSPs of all non-default MTs. It also cleans up all the LSPs of all non-default MTs locally. Then it clears the S bit in MT capability TLV and advertises via the Capability message (if it supports Dynamic Announcement Capability). When an LSR knows the peer node is changed from LDP-MT capable to non-MT capable, it cleanup all the LSPs of all non-default MTs locally and initiate withdraw of all label mapping for existing LSPs of all non-default MTs. Both sides of the nodes send label release to its peer once they receive the label release messages even both sides have already cleaned up all the LSPs locally.
- o If an LSR does not support "Dynamic Announcement Capability", it MUST reset session with its peer whenever LSR changes its local capability with regards to supporting LDP MT.
- o If an LSR is changed from IGP-MT capable to non-MT capable, it may wait until the routes update to withdraw FEC and release the label mapping for existing LSPs of specific MT.

[3.6.](#) LDP Sessions

Since using different label spaces for different topologies would imply significant changes to the data plane, a single global label space is supported in this solution. There will be one session supported between a pair of peers, even if there are multiple topologies supported between these two peers.

[3.7.](#) Reserved MT ID Values

Certain MT topologies are assigned to serve predetermined purposes.

In [Section 9](#). (IANA Considerations), this document defines a new IANA registry "LDP Multi-Topology ID Name Space" under IANA "LDP Parameter" namespace to keep an LDP MT-ID reserved value.

If an LSR receives a FEC element with an "MT-ID" value that is "Reserved" for future use (and not IANA allocated yet), the LSR MUST abort the processing of the FEC element, and SHOULD send a notification message with status code "Invalid Topology ID" to the sender.

[4.](#) MT Applicability on FEC-based features

[4.1.](#) Typed Wildcard FEC Element

[RFC5918] extends base LDP and defines Typed Wildcard FEC Element framework. Typed Wildcard FEC element can be used in any LDP message to specify a wildcard operation/action for given type of FEC.

The MT extensions defined in document do not require any extension to procedures for Typed Wildcard FEC element, and these procedures apply as-is to MT wildcarding. The MT extensions, though, allow use of "MT IP" or "MT IPv6" in the Address Family field of the Typed Wildcard FEC element in order to use wildcard operations in the context of a given topology. The use of MT-scoped address family also allows us to specify MT-ID in these operations.

The defined format in [Section 3.3](#) (Typed Wildcard FEC Element) allows an LSR to perform wildcard FEC operations under the scope of a topology. If an LSR wishes to perform wildcard operation that applies to all topologies, it can use a "Wildcard Topology" MT-ID. For example, upon local de-configuration of a topology "x", an LSR may send a typed wildcard label withdraw message with MT-ID "x" to withdraw all its labels from the peer that advertised under the scope of topology "x". Additionally, upon a global configuration change, an LSR may send a typed wildcard label withdraw message with the MT-ID set to "Wildcard Topology" to withdraw all its labels under all topologies from the peer.

[4.2.](#) End-of-LIB

[RFC5919] specifies extensions and procedures for an LDP speaker to signal its convergence for a given FEC type towards a peer. The procedures defined in [\[RFC5919\]](#) applies as-is to an MT FEC element. This allows an LDP speaker to signal its IP convergence using Typed Wildcard FEC element, and its MT IP convergence per topology using a MT Typed Wildcard FEC element.

[4.3.](#) LSP Ping

[RFC4379] defines procedures to detect data-plane failures in MPLS LSPs via LSP ping. That specification defines a "Target FEC Stack" TLV that describes the FEC stack being tested. This TLV is sent in an MPLS echo request message towards LSPs egress LSR, and is forwarded along the same data path as other packets belonging to the FEC.

"Target FEC Stack" TLV contains one or more sub-TLVs pertaining to different FEC types. [Section 3.2 of \[RFC4379\]](#) defines Sub-Types and format for the FEC. To support LSP ping for MT LDP LSPs, this document defines following extensions to [\[RFC4379\]](#).

[4.3.1.](#) New FEC Sub-Types

We define two new FEC types for LSP ping:

- o MT LDP IPv4 FEC
- o MT LDP IPv6 FEC

We also define following new sub-types for sub-TLVs to specify these FECs in the "Target FEC Stack" TLV of [\[RFC4379\]](#):

Sub-Type	Length	Value Field
-----	-----	-----
TBA5	8	MT LDP IPv4 prefix
TBA6	20	MT LDP IPv6 prefix

Figure 5: new sub-types for sub-TLVs

The rules and procedures of using these sub-TLVs in an MPLS echo request message are same as defined for LDP IPv4/IPv6 FEC sub-TLV types in [\[RFC4379\]](#).

[4.3.2.](#) MT LDP IPv4 FEC Sub-TLV

The format of "MT LDP IPv4 FEC" sub-TLV to be used in a "Target FEC Stack" [\[RFC4379\]](#) is:

```

      0               1               2               3
      0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|  Type = TBA5(MT LDP IPv4 FEC) |           Length = 8           |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
|                               IPv4 prefix                        |
+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+---+
| Prefix Length |           MBZ           |           MT-ID       |

```

[Page 10]

To detect data plane failures using LSP Ping for a specific topology, the router will initiate an LSP Ping request with the target FEC stack TLV containing LDP MT IP Prefix Sub-TLV in the Echo Request packet. The Echo Request packet is sent with the label bound to the IP Prefix in the topology. Once the echo request packet reaches the target router, it will process the packet and perform checks for the LDP MT IP Prefix sub-TLV present in the Target FEC Stack as described in [RFC4379] and respond according to [RFC4379] processing rules. For the case that the LSP ping with return path is not specified, the reply packet must go through the default topology instead of the topology where the Echo Request goes through.

5. Error Handling

Zhao, et al.

Expires August, 2014

[Page 11]

The extensions defined in this document utilize the existing LDP error handling defined in [[RFC5036](#)]. If an LSR receives an error notification from a peer for a session, it terminates the LDP session by closing the TCP transport connection for the session and discarding all multi-topology label mappings learned via the session.

5.1. MT Error Notification for Invalid Topology ID

An LSR should respond with an "Invalid Topology ID" status code in LDP Notification message when it receives an LDP message with a FEC element specifying an MT-ID which is not locally known or not supported. The LSR MUST also discard the entire message before sending the Notification.

6. Backwards Compatibility

The MPLS-MT solution is backwards compatible with existing LDP enhancements defined in [[RFC5036](#)], including message authenticity, integrity of message, and topology loop detection.

The legacy node which does not support MT should not receive any MT related LDP messages. In case the bad things happen, according to [[RFC5036](#)], processing of such messages should be aborted.

7. MPLS Forwarding in MT

Although forwarding is out of the scope of this draft, we include some forwarding consideration for informational purpose here.

The specified signaling mechanisms allow all the topologies to share the platform-specific label space, This feature allows the existing data plane techniques to be used. Also, there is no way for the data plane to associate a received packet with any one topology, meaning that topology-specific label spaces cannot be used.

8. Security Consideration

The use of MT over existing MPLS solutions does not offer any specific security benefit.

General LDP Communication security threats and how these may be mitigated are described in [[RFC5036](#)], these threats include:

- o Spoofing

- o Privacy
- o Denial of Service

For further discussion regarding possible LDP communication threats and mitigation techniques see [[RFC5920](#)].

9. IANA Considerations

The document introduces following new protocol elements that require IANA consideration and assignments:

- o New LDP Capability TLV: "Multi-Topology Capability" TLV (requested code point: TBA1 from LDP registry "TLV Type Name Space").
- o New Status Code: "Invalid Topology ID" (requested code point: TBA2 from LDP registry "Status Code Name Space").

Registry:	
Range/Value	Description
-----	-----
TBA1	Invalid Topology ID

Figure 8: New Code Points for LDP Multi Topology Extensions

- o New address families under IANA registry "Address Family Numbers":
 - MT IP: Multi-Topology IP version 4 (requested codepoint:26)
 - MT IPv6: Multi-Topology IP version 6 (requested codepoint:27)

Figure 9: Address Family Numbers

- o New registry "MPLS Multi-Topology Identifiers". The allocation policies for this registry are:

Range/Value	Purpose	Reference
-----	-----	-----
0	Default/standard topology	[This.I-D]
1	IPv4 in-band management	[This.I-D]
2	IPv6 routing topology	[This.I-D]
3	IPv4 multicast topology	[This.I-D]
4	IPv6 multicast topology	[This.I-D]
5	IPv6 in-band management	[This.I-D]
6-3995	Unassigned for future IGP topologies	[This.I-D]
	Assigned by Standards Action	[This.I-D]
3996-4095	Experimental	[This.I-D]

4096-65534	Unassigned for MPLS topologies	[This.I-D]
	Assigned by Standards Action	
65535	Wildcard Topology	[This.I-D]

Zhao, et al.	Expires August, 2014	[Page 13]
--------------	----------------------	-----------

Figure 10: MPLS Multi-Topology Identifier registry

- o New Sub-TLV Types for LSP ping: Following new sub-type values under TLV type 1 (Target FEC Stack) from "Multi-Protocol Label Switching (MPLS) Label Switched Paths (LSPs) Ping Parameters" registry, and "TLVs and sub-TLVs" sub-registry.

Sub-Type	Value Field
-----	-----
TBA3	MT LDP IPv4 prefix
TBA4	MT LDP IPv6 prefix

Figure 11: New Sub-TLV Types for LSP ping

IANA should allocate the next available numbers for these TBAs.

As highlighted at the end of [Section 3.4](#) (IGP MT-ID Mapping and Translation), a new Internet-Draft will be created to document the policy and process for allocating new MT-ID values.

[10.](#) Manageability Considerations

[10.1.](#) Control of Function and Policy

There are capabilities that should be configurable to enable good manageability. One such example is to allow enable or disable LDP Multi-Topology capability. It is assumed that the mapping of the LDP MT ID and IGP MT ID is manually configured on every router by default. If an automatic mapping between IGP MT IDs and LDP MT IDs is needed, there must be explicit configuration to do so.

[10.2.](#) Information and Data Models

Any extensions that may be required for existing MIBs are beyond the scope of this document.

[10.3.](#) Liveness Detection and Monitoring

Mechanisms defined in this document do not imply any new liveness detection and monitoring requirements.

[10.4.](#) Verify Correct Operations

If an operator is trying to debug LDP MT enabled network and wants to make the association between the LDP label advertisement and the IGP routing advertisement, then the user MUST understand the mapping mechanism to convert the IGP MT ID to the LDP MT ID. This type of mapping mechanisms is out of the scope of this document.

10.5. Requirements On Other Protocols

If the LDP MT ID has an implicit dependency on IGP MT ID, then the corresponding IGP MT features will need to be supported.

10.6. Impact On Network Operations

Mechanisms defined in this document do not have any impact on network operations.

11. Contributors

12. Acknowledgement

The authors would like to thank Dan Tappan, Nabil Bitar, Huang Xin, Eric Rosen, IJsbrand Wijnands, Dimitri Papadimitriou, Yiqun Chai, Pranjal Dutta, George Swallow, Curtis Villamizar, Adrian Farrel, Alia Atlas, Loa Anderson, Joel Halpern and Kathleen Moriarty for their valuable comments on this draft.

13. References

13.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), March 1997.
- [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.
- [RFC5561] Thomas, B., Raza, K., Aggarwal, S., Aggarwal, R., and JL. Le Roux, "LDP Capabilities", [RFC 5561](#), July 2009.
- [RFC5918] Asati, R., Minei, I., and B. Thomas, "Label Distribution Protocol (LDP) 'Typed Wildcard' Forward Equivalence Class (FEC)", [RFC 5918](#), August 2010.

[RFC5919] Asati, R., Mohapatra, P., Chen, E., and B. Thomas,
"Signaling LDP Label Advertisement Completion", [RFC 5919](#),
August 2010.

13.2. Informative References

[RFC5920] Fang, L., "Security Framework for MPLS and GMPLS
Networks", [RFC 5920](#), July 2010.

Authors' Addresses

Quintin Zhao
Huawei Technology
125 Nagog Technology Park
Acton, MA 01719
US

Email: quintin.zhao@huawei.com

Kamran Raza
Cisco Systems
2000 Innovation Drive
Kanata, ON K2K-3E8, MA
Canada

Email: E-mail: skraza@cisco.com

Chao Zhou
Cisco Systems
300 Beaver Brook Road
Boxborough, MA 01719
US

Email: czhou@cisco.com

Luyuan Fang
Microsoft

Email: lufang@microsoft.com

Lianyuan Li
China Mobile
53A, Xibianmennei Ave.
Xunwu District, Beijing 01719
China

Email: lilianyuan@chinamobile.com

Daniel King
Old Dog Consulting

Email: lufang@microsoft.com

