

Internet Engineering Task Force
Internet-Draft
Intended status: Standards Track
Expires: May 24, 2012

Q. Zhao
Huawei Technology
L. Fang
C. Zhou
Cisco Systems
L. Li
China Mobile
N. So
Verizon Business
R. Torvi
Juniper Networks
November 21, 2011

LDP Extensions for Multi Topology Routing
draft-ietf-mpls-ldp-multi-topology-02.txt

Abstract

Multi-Topology (MT) routing is supported in IP through extension of IGP protocols, such as OSPF and IS-IS. It would be advantageous to extend Multiprotocol Label Switching (MPLS), using Label Distribution Protocol (LDP), to support multiple topologies. These LDP extensions, known as Multiple Topology Label Distribution Protocol (MT LDP), would allow the configuration of multiple topologies within an MPLS LDP enabled network.

This document describes the protocol extensions required to extend the existing MPLS LDP signalling protocol for creating and maintaining LSPs in an MT environment.

Status of this Memo

This Internet-Draft is submitted to IETF 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 May 24, 2012.

Copyright Notice

Copyright (c) 2011 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to [BCP 78](http://trustee.ietf.org/license-info) 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.

Table of Contents

1.	Terminology	4
1.1.	Requirements Language	4
2.	Introduction	4
3.	Requirements	5
3.1.	Application Scenarios	6
3.1.1.	Simplified Data-plane	6
3.1.2.	Using MT for p2p Protection	6
3.1.3.	Using MT for mLDP Protection	7
3.1.4.	Service Separation	7
3.1.5.	An Alternative inter-AS VPN Solution	7
3.2.	Signaling Extensions	7
3.2.1.	Topology-Scoped Prefix FEC	7
3.2.2.	LDP MT Capability Advertisement	9
3.2.3.	Procedures	11
3.2.4.	LDP Sessions	12
3.2.5.	Reserved MT ID Values	12
3.3.	MT Applicability on FEC-based features	13
3.3.1.	Typed Wildcard Prefix FEC Element	13
3.3.2.	End-of-LIB	13
3.4.	MPLS Forwarding in MT	14
3.5.	Security Consideration	14
3.6.	IANA Considerations	14
3.7.	Acknowledgement	15
4.	References	15
4.1.	Normative References	15
4.2.	Informative References	15
	Authors' Addresses	16

1. Terminology

Terminology used in this document

MT-ID: A 12 bit value to represent Multi-Topology ID.

Default Topology: A topology that is built using the MT-ID value 0.

MT topology: A topology that is built using the corresponding MT-ID.

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 [RFC 2119](#) [[RFC2119](#)].

2. Introduction

There are increasing requirements to support multi-topology in MPLS network. For example, service providers may want to assign different level of service(s) to different topologies so that the service separation can be achieved. It is also possible to have an in-band management network on top of the original MPLS topology, or maintain separate routing and MPLS domains for isolated multicast or IPv6 islands within the backbone, or force a subset of an address space to follow a different MPLS topology for the purpose of security, QoS, or simplified management and/or operations.

OSPF and IS-IS use MT-ID (Multi-Topology Identifier) to identify different topologies. For each topology identified by an MT-ID, IGP computes a separate SPF tree independently to find the best paths to the IP prefixes associated with this topology.

For IP Prefix FECs that are associated with a specific topology, this solution utilizes the use of MT-ID of the topology in LDP. Thus the LSP for the given Prefix FEC may be created and maintained along the IGP path in this topology if it is needed.

Maintaining multiple MTs for MPLS network in a backwards-compatible manner requires several extensions to the label signaling encoding and processing procedures. When a label is associated with an IP Prefix FEC, the corresponding FEC element includes both the destination prefix (IP address) and the topology it belongs to.

MT based MPLS in general can be used for a variety of purposes such

as service separation by assigning each service or a group of services to a topology, where the management, QoS, and security of the service or the group of the services can be simplified and guaranteed, in-band management network "on top" of the original MPLS topology, maintain separate routing and MPLS forwarding domains for isolated multicast or IPv6 islands within the backbone, or force a subset of an address space to follow a different MPLS topology for the purpose of security, QoS or simplified management and/or operations.

One of the use of the MT based MPLS is where one class of data requires low latency links, for example Voice over IP (VoIP) data. As a result such data may be sent preferably via physical landlines rather than, for example, high latency links such as satellite links. As a result an additional topology is defined as all low latency links on the network and VoIP data packets are assigned to the additional topology. Further possible examples are File Transfer Protocol (FTP) or Simple Mail Transfer Protocol (SMTP) traffic which can be assigned to additional topology comprising high latency links, and Internet Protocol version 4 (IPv4) versus Internet Protocol version 6 (IPv6) traffic which may be assigned to different topology or data to be distinguished by the Quality of Service (QoS) assigned to it.

3. Requirements

MPLS-MT may be used for a variety of purposes such as service separation by assigning each service or a group of services to a topology, where the management, QoS and security of the service or the group of the services can be simplified and guaranteed, in-band management network "on top" of the original MPLS topology, maintain separate routing and MPLS forwarding domains for isolated multicast or IPv6 islands within the backbone, or force a subset of an address space to follow a different MPLS topology for the purpose of security, QoS or simplified management and/or operations.

The following specific requirements and objectives have been defined in order to provide the functionality described above, and facilitate service provider configuration and operation.

- o Deployment of MPLS-MT within existing MPLS networks should be possible, with MPLS-MT non-capable nodes existing with MPLS-MT capable nodes.
- o Minimise configuration and operation complexity of MPLS-MT across the network.

- o The MPLS-MT solution SHOULD NOT require data-plane modification.
- o The MPLS-MT solution MUST support multiple topologies. Allowing a MPLS LSP to be established across a specific, or set of, multiple topologies.
- o Control and filtering of LSPs using explicitly including or excluding multiple topologies MUST be supported.
- o The MPLS-MT solution MUST be capable of supporting QoS mechanisms.

[Editors Note - We expect these base MPLS-MT protocol requirements to be evolved over the next few versions of this document. Note that all Editors notes will be deleted before publication of the document]

3.1. Application Scenarios

3.1.1. Simplified Data-plane

IGP-MT requires additional data-plane resources maintain multiple forwarding for each configured MT. On the other hand, MPLS-MT does not change the data-plane system architecture, if an IGP-MT is mapped to an MPLS-MT. In case MPLS-MT, incoming label value itself can determine an MT, and hence it requires a single NHLFE space. MPLS-MT requires only MT-RIBs in the control-plane, no need to have MT-FIBs. Forwarding IP packets over a particular MT requires either configuration or some external means at every node, to maps an attribute of incoming IP packet header to IGP-MT, which is additional overhead for network management. Whereas, MPLS-MT mapping is required only at the ingress-PE of an MPLS-MT LSP, because of each node identifies MPLS-MT LSP switching based on incoming label, hence no additional configuration is required at every node.

3.1.2. Using MT for p2p Protection

We know that [IP-FRR-MT] can be used for configuring alternate path via backup-mt, such that if primary link fails, then backup-MT can be used for forwarding. However, such techniques require special marking of IP packets that needs to be forwarded using backup-MT. MPLS-LDP-MT procedures simplify the forwarding of the MPLS packets over backup-MT, as MPLS-LDP-MT procedure distribute separate labels for each MT. How backup paths are computed depends on the implementation, and the algorithm. The MPLS-LDP-MT in conjunction with IGP-MT could be used to separate the primary traffic and backup traffic. For example, service providers can create a backup MT that consists of links that are meant only for backup traffic. Service providers can then establish bypass LSPs, standby LSPs, using backup MT, thus keeping undeterministic backup traffic away from the primary

traffic.

3.1.3. Using MT for mLDP Protection

From the P2MP or MP2MP LSPs setup by using mLDP protocol, there is a need to setup a backup LSP to have an end to end protection for the primary LSP in the applications such as IPTV, where the end to end protection is a must. Since the mLDP LSP is setup following the IGP routes, the second LSP setup by following the IGP routes can not be guaranteed to have the link and node diversity from the primary LSP. By using MPLS-LDP-MT, two topologies can be configured with complete link and node diversity, where the primary and secondary LSP can be set up independently within each topology. The two LSPs setup by this mechanism can protect each other end-to-end.

3.1.4. Service Separation

MPLS-MT procedures allow establishing two distinct LSPs for the same FEC, by advertising separate label mapping for each configured topology. Service providers can implement CoS using MPLS-MT procedures without requiring to create separate FEC addresses for each class. MPLS-MT can also be used to separate multicast and unicast traffic.

3.1.5. An Alternative inter-AS VPN Solution

When the LSP is crossing multiple domains for the inter-AS VPN scenarios, the LSP setup process can be done by configuring a set of routers which are in different domains into a new single domain with a new topology ID using the LDP multiple topology. All the routers belong to this new topology will be used to carry the traffic across multiple domains and since they are in a single domain with the new topology ID, so the LDP LSP set up can be done without propagating VPN routes across AS boundaries.

3.2. Signaling Extensions

3.2.1. Topology-Scoped Prefix FEC

LDP assigns and binds a label to a 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 Multi-Topology LSPs for IP prefixes under a given topology scope, it is a natural requirement to extend LDP "Prefix" FEC element to include topology info. This infers that MT-ID becomes an attribute of Prefix FEC element, and all FEC-Label binding operations are performed under the context of given topology (MT-ID). Following

subsection proposes the extension to bind "Prefix FEC" to a topology.

3.2.1.1. New Address Families: MT IP

LDP base specification [[RFC5036](#)] ([section 3.4.1](#)) defines the "Prefix" FEC Element as follows:

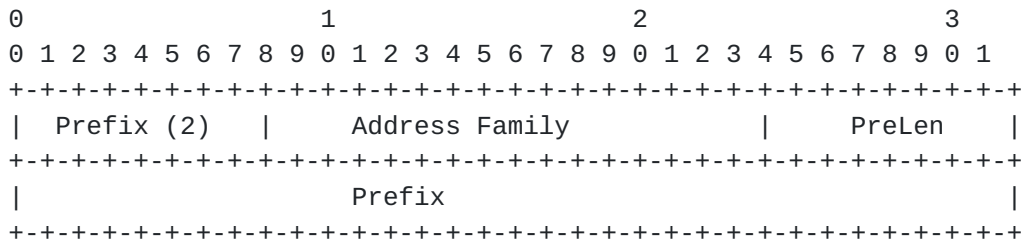


Figure 1: Prefix FEC Element Format [[RFC5036](#)]

Where "Prefix" encoding is as defined for given "Address Family", and whose length (in bits) is specified by the "PreLen" field.

To extend IP address families for MT, we propose two new Address Families named "MT IP" and "MT IPv6" that can be used to specify IP prefixes within a topology scope. The format of data associated with these new Address Family is:

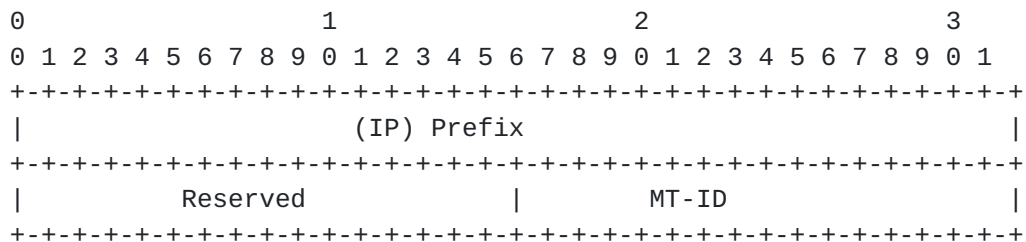


Figure 2: MT IP Address Family Format

Where "(IP) Prefix" is an IPv4 and IPv6 address/prefix for "MT IP" and "MT IPv6" AF respectively, and the field "MT-ID" corresponds to 16-bit Topology ID for given prefix.

For MT LDP, the "Prefix" FEC element's "Address Family" will be set to "MT IP" or "MT IPv6", and the FEC element will be encoded as follows:

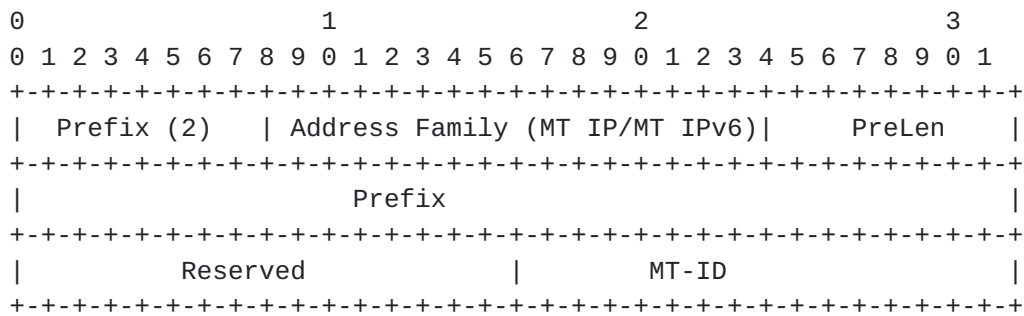


Figure 3: MT Prefix FEC Element Format

Where 16-bit "MT-ID" field defines the Topology ID, and the definition and usage of "Prefix" and "PreLen" field is same as defined for IP/IPv6 AF. The value of MT-ID 0 corresponds to default topology and MUST be ignored on receipt so as to not cause any conflict/confusion with existing non-MT procedures.

The proposed "Prefix FEC Element" with "MT IP" Address Family can be used in any LDP message and procedures that currently specify and allow the use of "Prefix FEC" element with IP/IPv6 Address Family.

This document does not limit the use of these new AF only to LDP "Prefix FEC element", and these can be used in other FECs and signaling as required. For example, mLDP MP FECs, as specified in [mLDP], can be extended to use these new address families to make MP FECs to become MT aware.

[Editors Note - RFC[5036] doesn't specify the handling of unknown Address Family. After we have introduced the two new address family here, RFC[5036] need to be updated to add the handling procedure for the unknow address families.

3.2.1.2. IGP MT-ID Mapping and Translation

The non-reserved non-special IGP MT-ID values can be used/carried in LDP as-is and need no translation. However, there is a need for translating reserved/special IGP MT-ID values to corresponding LDP MT-IDs. The corresponding special/reserved LDP MT-ID values are defined in later [section 9](#).

3.2.2. LDP MT Capability Advertisement

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

MTR for given IP address family.

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 Prefix FEC Element" as defined in [RFC5918].

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

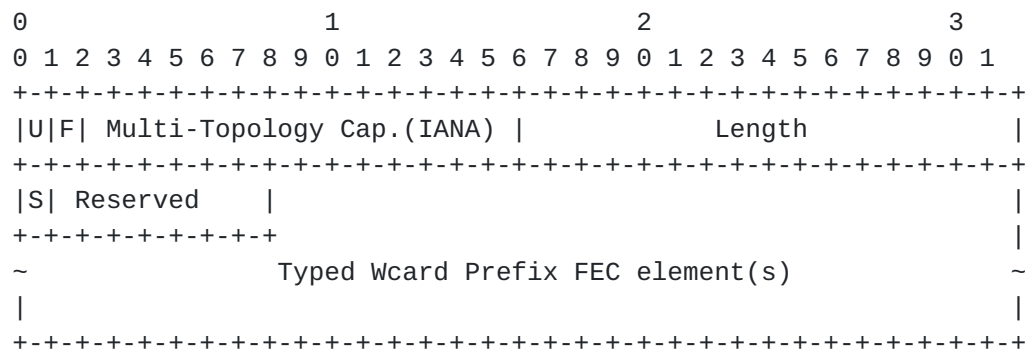


Figure 4: Multi-Topology Capability TLV Format

Where:

U- and F-bits: MUST be 1 and 0, respectively, as per [Section 3](#) of LDP Capabilities [RFC5561].

Multi-Topology Capability: Capability TLV type (IANA assigned)

S-bit: MUST be 1 if used in LDP "Initialization" message. MAY be set to 0 or 1 in dynamic "Capability" message to advertise or withdraw the capability respectively.

Typed Wcard Prefix FEC element(s): One or two elements specified as the "Capability data".

Length: The length (in octets) of TLV. The value of this field MUST be 6 with one FEC element specification, and 11 for two FEC element specifications.

The encoding of Typed Wcard Prefix FEC element, as defined in [RFC5561], 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) | Type=Prefix (2) | Len = 2 | Address... |
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
| ...Family      |
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+

```

Figure 5: Typed Wildcard Prefix FEC Element [[RFC5561](#)]

Where:

Address Family: MUST be set to "MT IP" or "MT IPv6"

3.2.3. Procedures

To announce its MT capability for given IP address family, an LDP speaker MAY send "MT Capability" with exactly one Typed Wildcard Prefix FEC element with corresponding "Address Family" field (i.e. set to "MT IP" for IPv4 and set to "MT IPv6" for IPv6 address family). To announce its MT capability for both IPv4 and IPv6 address family, an LDP speaker MAY send "MT Capability" with two Typed Prefix FEC elements in it, where Address Family is set to "MT IP" in one element and set to "MT IPv6" in other element.

- 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 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 messages that include MT identifier to this LSR.
- o If an LSR receives a Label Mapping message with MT parameter from downstream LSR-D and its upstream LSR-U has not advertised MT capability, an LSP for the MT will not be established.
- o We propose to add a new notification event to signal the upstream that the downstream is not capable.
- 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. The existing LSP is treated as LSP for default MT (ID 0).

- o If an LSR is changed from MT capable to non-MT capable, it may initiate withdraw of all label mapping for existing LSPs of all non-default MTs. Alternatively, it may wait until the routing update to withdraw FEC and release the label mapping for existing LSPs of specific MT.
- o There will be case where IGP is MT capable but MPLS is not and the handling procedure for this case is TBD.

3.2.4. LDP Sessions

Depending on the number of label spaces supported, if a single global label space is supported, there will be one session supported for each pair of peer, even there are multiple topologies supported between these two peers. If there are different label spaces supported for different topologies, which means that label spaces overlap with each other for different MTs, then it is suggested to establish multiple sessions for multiple topologies between these two peers. In this case, multiple LSR-IDs need to be allocated beforehand so that each multiple topology can have its own label space ID.

[Editors Note - This section requires further discussion]

3.2.5. Reserved MT ID Values

Certain MT topologies are assigned to serve pre-determined purposes:

Default-MT: Default topology. This corresponds to OSPF default IPv4 and IPv6, as well as ISIS default IPv4. A value of 0 is proposed.

ISIS IPv6 MT: ISIS default MT-ID for IPv6.

Wildcard-MT: This corresponds to All-Topologies. A value of 65535 (0xffff) is proposed.

We propose a new IANA registry "LDP Multi-Topology ID Name Space" under IANA "LDP Parameter" namespace to keep 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 MT-ID" to the sender.

[Editors Note - This section requires further discussion].

3.3. MT Applicability on FEC-based features

3.3.1. Typed Wildcard Prefix FEC Element

[RFC-5918](#) extends base LDP and defines Typed Wildcard FEC Element framework [[RFC5918](#)]. 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 proposed in document do not require any extension in procedures for Typed Wildcard Prefix FEC element, and these procedures apply as-is to MT Prefix wildcarding. The MT extensions, though, allow use of "MT IP" or "MT IPv6" in the Address Family field of the Typed Wildcard Prefix 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.

This document extends Typed Wildcard Prefix FEC element encoding for MT 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)|Type=Prefix (2)|  Len = 4      |  AF = MT IP ..|
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+
|... or MT IPv6 |      MT-ID                      |
+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+--+

```

Figure 6: Typed Wildcard Prefix FEC Element for MT

The proposed format 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 "Wildcard Topology" MT-ID as defined in [section 5.4](#). For instance, upon local un-configuration of topology "x", an LSR may send wildcard label withdraw with MT-ID "x" to withdraw all its labels from peer that were advertised under the scope of topology "x". On the other hand, upon some global configuration change, an LSR may send wildcard label withdraw with MT-ID set to "Wildcard Topology" to withdraw all its labels under all topologies from the peer.

3.3.2. End-of-LIB

[RFC5919] specifies extensions and procedures for an LDP speaker to signal its convergence for given FEC type towards a peer. The procedures defined in [[RFC5919](#)] apply as-is to MT Prefix FEC element. This means that an LDP speaker MAY signal its IP convergence using

Typed Wildcard Prefix FEC element, and its MT IP convergence per topology using MT Typed Wildcard Prefix FEC element (as defined in earlier section).

3.4. 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 is the feature that allows the existing data plane techniques to be used; and the specified signaling mechanisms do not provide any way for the data plane to associate a given packet with a context-specific label space.

3.5. Security Consideration

No specific security issues with the proposed solutions are known. The proposed extension in this document does not introduce any new security considerations beyond that already apply to the base LDP specification [[RFC5036](#)] and [[RFC5920](#)].

3.6. 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: 0x510 from LDP registry "TLV Type Name Space").
- o New Status Code: "Invalid Topology ID" (requested code point: 0x50 from LDP registry "Status Code Name Space") as follows:

Registry:		
Range/Value	E	Description
-----	---	-----
0x00000050	0	Invalid MT-ID

Figure 7: Invalid Topology ID

- 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 8: Address Family Numbers

- o New registry "LDP Multi-Topology (MT) ID Name Space" under "LDP Parameter" namespace. The registry is defined as:

Range/Value	Name
-----	-----
0	Default Topology (ISIS and OSPF)
1-4095	Unassigned
4096	ISIS IPv6 routing topology (i.e. ISIS MT ID
#2)	
4097-65534	Reserved (for future allocation)
65535	Wildcard Topology (ISIS or OSPF)

Figure 9: LDP Multi-Topology (MT) ID Name Space

3.7. Acknowledgement

The authors would like to thank Dan Tappan, Nabil Bitar, Huang Xin, Daniel King and Eric Rosen for their valuable comments on this draft.

4. References

4.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), March 1997.
- [RFC3692] Narten, T., "Assigning Experimental and Testing Numbers Considered Useful", [BCP 82](#), [RFC 3692](#), January 2004.
- [RFC5036] Andersson, L., Minei, I., and B. Thomas, "LDP Specification", [RFC 5036](#), October 2007.
- [RFC5919] Asati, R., Mohapatra, P., Chen, E., and B. Thomas, "Signaling LDP Label Advertisement Completion", [RFC 5919](#), August 2010.
- [RFC5918] Asati, R., Minei, I., and B. Thomas, "Label Distribution Protocol (LDP) 'Typed Wildcard' Forward Equivalence Class (FEC)", [RFC 5918](#), August 2010.

4.2. 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

Huaimo Chen
Huawei Technology
125 Nagog Technology Park
Acton, MA 01719
US

Email: huaimochen@huawei.com

Emily Chen
Huawei Technology
No. 5 Street, Shangdi Information, Haidian
Beijing
China

Email: chenying220@huawei.com

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

Email: lilianyuan@chinamobile.com

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

Email: lichenyj@chinamobile.com

Lu Huang
China Mobile
53A, Xibianmennei Ave.
Xunwu District, Beijing 01719
China

Email: huanglu@chinamobile.com

Luyuang Fang
Cisco Systems
300 Beaver Brook Road
Boxborough, MA 01719
US

Email: lufang@cisco.com

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

Email: czhou@cisco.com

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

Email: E-mail: skraza@cisco.com

Ning So
Verizon Business
2400 North Glenville Drive
Richardson, TX 75082
USA

Email: Ning.So@verizonbusiness.com

Raveendra Torvi
Juniper Networks
10, Technoogy Park Drive
Westford, MA 01886-3140
US

Email: rtorvi@juniper.net