

Internet Engineering Task Force (IETF)
Request for Comments: 8320
Category: Standards Track
ISSN: 2070-1721

A. Atlas
K. Tiruveedhula
C. Bowers
Juniper Networks
J. Tantsura
Individual
I.J. Wijnands
Cisco Systems, Inc.
February 2018

LDP Extensions to Support Maximally Redundant Trees

Abstract

This document specifies extensions to the Label Distribution Protocol (LDP) to support the creation of Label Switched Paths (LSPs) for Maximally Redundant Trees (MRTs). A prime use of MRTs is for unicast and multicast IP/LDP Fast Reroute, which we will refer to as "MRT-FRR".

The sole protocol extension to LDP is simply the ability to advertise an MRT Capability. This document describes that extension and the associated behavior expected for Label Switching Routers (LSRs) and Label Edge Routers (LERs) advertising the MRT Capability.

MRT-FRR uses LDP multi-topology extensions, so three multi-topology IDs have been allocated from the MPLS MT-ID space.

Status of This Memo

This is an Internet Standards Track document.

This document is a product of the Internet Engineering Task Force (IETF). It represents the consensus of the IETF community. It has received public review and has been approved for publication by the Internet Engineering Steering Group (IESG). Further information on Internet Standards is available in Section 2 of RFC 7841.

Information about the current status of this document, any errata, and how to provide feedback on it may be obtained at <https://www.rfc-editor.org/info/rfc8320>.

Copyright Notice

Copyright (c) 2018 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 (<https://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. Introduction	4
2. Requirements Language	5
3. Terminology	5
4. Overview of LDP Signaling Extensions for MRT	6
4.1. MRT Capability Advertisement	6
4.1.1. Interaction of MRT Capability and MT Capability	7
4.1.2. Interaction of LDP MRT Capability with IPv4 and IPv6	8
4.2. Use of the Rainbow MRT MT-ID	8
4.3. MRT-Blue and MRT-Red FECs	8
4.4. Interaction of MRT-Related LDP Advertisements with the MRT Topology and Computations	9
5. LDP MRT FEC Advertisements	10
5.1. MRT-Specific Behavior	10
5.1.1. ABR Behavior and Use of the Rainbow FEC	10
5.1.2. Proxy-Node Attachment Router Behavior	11
5.2. LDP Protocol Procedures in the Context of MRT Label Distribution	12
5.2.1. LDP Peer in RFC 5036	12
5.2.2. Next Hop in RFC 5036	13
5.2.3. Egress LSR in RFC 5036	13
5.2.4. Use of Rainbow FEC to Satisfy Label Mapping Existence Requirements in RFC 5036	15
5.2.5. Validating FECs in the Routing Table	15
5.2.6. Recognizing New FECs	15
5.2.7. Not Propagating Rainbow FEC Label Mappings	15
6. Security Considerations	16
7. Potential Restrictions on MRT-Related MT-ID Values Imposed by RFC 6420	16
8. IANA Considerations	17
9. References	18
9.1. Normative References	18
9.2. Informative References	19
Acknowledgements	21
Authors' Addresses	21

1. Introduction

This document describes the LDP signaling extensions and associated behavior necessary to support the architecture that defines how IP/LDP Fast Reroute can use MRTs [RFC7812]. The current document provides a brief description of the MRT-FRR architecture, focusing on the aspects most directly related to LDP signaling. The complete description and specification of the MRT-FRR architecture can be found in [RFC7812].

At least one common standardized algorithm (e.g., the MRT Lowpoint algorithm explained and fully documented in [RFC7811]) is required to be deployed so that the routers supporting MRT computation consistently compute the same MRTs. LDP depends on an IGP for computation of MRTs and alternates. Extensions to OSPF are defined in [OSPF-MRT]. Extensions to IS-IS are defined in [IS-IS-MRT].

MRT can also be used to protect multicast traffic (signaled via PIM or Multipoint LDP (mLDP)) using either global protection or local protection as described in [ARCH]. An MRT path can be used to provide node-protection for mLDP traffic via the mechanisms described in [RFC7715]; an MRT path can also be used to provide link protection for mLDP traffic.

For each destination, IP/LDP Fast Reroute with MRT (MRT-FRR) creates two alternate destination-based trees separate from the shortest-path forwarding used during stable operation. LDP uses the multi-topology extensions [RFC7307] to signal Forwarding Equivalency Classes (FECs) for these two sets of forwarding trees, MRT-Blue and MRT-Red.

In order to create MRT paths and support IP/LDP Fast Reroute, a new capability extension is needed for LDP. An LDP implementation supporting MRT MUST also follow the rules described here for originating and managing FECs related to MRT, as indicated by their multi-topology ID. Network reconvergence is described in [RFC7812] and the worst-case network convergence time can be flooded via the extension in [PARAM-SYNC].

IP/LDP Fast Reroute using MRTs can provide 100% coverage for link and node failures in an arbitrary network topology where the failure doesn't partition the network. It can also be deployed incrementally; an MRT Island is formed of connected supporting routers and the MRTs are computed inside that island.

2. Requirements Language

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "NOT RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in BCP 14 [RFC2119] [RFC8174] when, and only when, they appear in all capitals, as shown here.

3. Terminology

For ease of reading, some of the terminology defined in [RFC7812] is repeated here. Please refer to Section 3 of [RFC7812] for a more complete list.

Redundant Trees (RTs): A pair of trees where the path from any node X to the root R along the first tree is node-disjoint with the path from the same node X to the root along the second tree. Redundant trees can always be computed in 2-connected graphs.

Maximally Redundant Trees (MRTs): A pair of trees where the path from any node X to the root R along the first tree and the path from the same node X to the root along the second tree share the minimum number of nodes and the minimum number of links. Each such shared node is a cut-vertex. Any shared links are cut-links. In graphs that are not 2-connected, it is not possible to compute RTs. However, it is possible to compute MRTs. MRTs are maximally redundant in the sense that they are as redundant as possible given the constraints of the network graph.

MRT-Red: MRT-Red is used to describe one of the two MRTs; it is used to describe the associated forwarding topology and MPLS Multi-Topology Identifier (MT-ID).

MRT-Blue: MRT-Blue is used to describe one of the two MRTs; it is used to describe the associated forwarding topology and MPLS MT-ID.

Rainbow MRT: It is useful to have an MPLS MT-ID that refers to the multiple MRT forwarding topologies and to the default forwarding topology. This is referred to as the "Rainbow MRT MPLS MT-ID" and is used by LDP to reduce signaling and permit the same label to always be advertised to all peers for the same (MT-ID, Prefix).

MRT Island: The set of routers that support a particular MRT Profile and the links connecting them that support MRT.

Island Border Router (IBR): A router in the MRT Island that is connected to a router not in the MRT Island, both of which are in a common area or level.

Island Neighbor (IN): A router that is not in the MRT Island but is adjacent to an IBR and in the same area/level as the IBR.

There are several places in this document where the construction "red(blue) FEC" is used to cover the case of the red FEC and the case of the blue FEC, independently. As an example, consider the sentence "When the ABR requires best-area behavior for a red(blue) FEC, it MUST withdraw any existing label mappings advertisements for the corresponding Rainbow FEC and advertise label mappings for the red(blue) FEC." This sentence should be read as applying to red FECs. Then it should be read as applying to blue FECs.

4. Overview of LDP Signaling Extensions for MRT

Routers need to know which of their LDP neighbors support MRT. This is communicated using the MRT Capability Advertisement. Supporting MRT indicates several different aspects of behavior, as listed below.

1. Sending and receiving multi-topology FEC elements, as defined in [RFC7307].
2. Understanding the Rainbow MRT MT-ID and applying the associated labels to all relevant MT-IDs.
3. Advertising the Rainbow MRT FEC to the appropriate neighbors for the appropriate prefix.
4. If acting as LDP egress for a prefix in the default topology, also acting as egress for the same prefix in MRT-Red and MRT-Blue.
5. For a FEC learned from a neighbor that does not support MRT, originating FECs for MRT-Red and MRT-Blue with the same prefix. This MRT Island egress behavior is to support an MRT Island that does not include all routers in the area/level.

4.1. MRT Capability Advertisement

A new MRT Capability Parameter TLV is defined in accordance with the LDP Capability definition guidelines [RFC5561].

The LDP MRT Capability can be advertised during LDP session initialization or after the LDP session is established. Advertisement of the MRT Capability indicates support of the

procedures for establishing the MRT-Blue and MRT-Red Label Switched Paths (LSPs) detailed in this document. If the peer has not advertised the MRT Capability, then it indicates that LSR does not support MRT procedures.

If a router advertises the LDP MRT Capability to its peer, but the peer has not advertised the MRT Capability, then the router MUST NOT advertise MRT-related FEC-label bindings to that peer.

The following is the format of the MRT Capability Parameter.

```

      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
+-----+-----+-----+-----+-----+-----+-----+-----+
|U|F| MRT Capability (0x050E) | Length (= 1) |
+-----+-----+-----+-----+-----+-----+-----+-----+
|S| Reserved |
+-----+-----+-----+-----+-----+-----+-----+-----+

```

MRT Capability TLV Format

Where:

U-bit: The unknown TLV bit MUST be 1. A router that does not recognize the MRT Capability TLV will silently ignore the TLV and process the rest of the message as if the unknown TLV did not exist.

F-bit: The forward unknown TLV bit MUST be 0 as required by Section 3 of [RFC5561].

MRT Capability: 0x050E

Length: The length (in octets) of the TLV. Its value is 1.

S-bit: The State bit MUST be 1 if used in the LDP Initialization message. MAY be set to 0 or 1 in the dynamic Capability message to advertise or withdraw the capability, respectively, as described in [RFC5561].

4.1.1. Interaction of MRT Capability and MT Capability

An LSR advertising the LDP MRT Capability MUST also advertise the LDP Multi-Topology (MT) Capability. If an LSR negotiates the LDP MRT Capability with an LDP neighbor without also negotiating the LDP MT Capability, the LSR MUST behave as if the LDP MRT Capability was not negotiated and respond with the "MRT Capability negotiated without MT Capability" status code in the LDP Notification message (defined in

the document). The E-bit of this Notification should be set to 0 to indicate that this is an Advisory Notification. The LDP session SHOULD NOT be terminated.

4.1.2. Interaction of LDP MRT Capability with IPv4 and IPv6

The MRT LDP Capability Advertisement does not distinguish between IPv4 and IPv6 address families. An LSR that advertises the MRT LDP Capability is expected to advertise MRT-related FEC-label bindings for the same address families for which it advertises shortest-path FEC-label bindings. Therefore, an LSR advertising MRT LDP Capability and shortest-path FEC-label bindings for IPv4 only (or IPv6 only) would be expected to advertise MRT-related FEC-label binding for IPv4 only (or IPv6 only). An LSR advertising the MRT LDP Capability and shortest-path FEC-label bindings for BOTH IPv4 and IPv6 is expected to advertise MRT-related FEC-label bindings for BOTH IPv4 and IPv6. In this scenario, advertising MRT-related FEC-label bindings only for IPv4 only (or only for IPv6) is not supported.

4.2. Use of the Rainbow MRT MT-ID

Section 10.1 of [RFC7812] describes the need for an Area Border Router (ABR) to have different neighbors use different MPLS labels when sending traffic to the ABR for the same FEC. More detailed discussion of the Rainbow MRT MT-ID is provided in Section 5.1.1.

Another use for the Rainbow MRT MT-ID is for an LSR to send the Rainbow MRT MT-ID with an IMPLICIT_NULL label to indicate penultimate-hop-popping for all three types of FECs (shortest path, red, and blue). The EXPLICIT_NULL label advertised using the Rainbow MRT MT-ID similarly applies to all the types of FECs. Note that the only scenario in which it is generally useful to advertise the implicit or explicit null label for all three FEC types is when the FEC refers to the LSR itself. See Section 5.2.3 for more details.

The value of the Rainbow MRT MPLS MT-ID (3945) has been assigned by IANA from the MPLS MT-ID space.

4.3. MRT-Blue and MRT-Red FECs

To provide MRT support in LDP, the MT Prefix FEC is used. [RFC7812] defines the Default MRT Profile. Section 8 specifies the values in the "MPLS Multi-Topology Identifiers" registry for the MRT-Red and MRT-Blue MPLS MT-IDs associated with the Default MRT Profile (3946 and 3947).

As described in Section 8.1 of [RFC7812], when a new MRT Profile is defined, new and unique values should be allocated from the "MPLS Multi-Topology Identifiers" registry, corresponding to the MRT-Red and MRT-Blue MT-ID values for the new MRT Profile.

The MT Prefix FEC encoding is defined in [RFC7307] and is used without alteration for advertising label mappings for MRT-Blue, MRT-Red, and Rainbow MRT FECs.

4.4. Interaction of MRT-Related LDP Advertisements with the MRT Topology and Computations

[RFC7811] and [RFC7812] describe how the MRT topology is created based on information in IGP advertisements. The MRT topology and computations rely on IGP advertisements. The presence or absence of MRT-related LDP advertisements does not affect the MRT topology or the MRT-Red and MRT-Blue next hops computed for that topology.

As an example, consider a network where all nodes are running MRT IGP extensions to determine the MRT topology, which is then used to compute MRT-Red and MRT-Blue next hops. The network operator also configures the nodes in this network to exchange MRT-related LDP advertisements in order to distribute MPLS labels corresponding to those MRT next hops. Suppose that, due to a misconfiguration on one particular link, the MRT-related LDP advertisements are not being properly exchanged for that link. Since the MRT-related IGP advertisements for the link are still being distributed, the link is still included in the MRT topology and computations. In this scenario, there will be missing MPLS forwarding entries corresponding to paths that use the misconfigured link.

Note that the situation is analogous to the interaction of normal LDP advertisements and IGP advertisements for shortest-path forwarding. Deactivating the distribution of labels for normal shortest-path FECs on a link does not change the topology on which the Shortest Path First (SPF) algorithm is run by the IGP.

"LDP IGP Synchronization" [RFC5443] addresses the issue of the LDP topology not matching the IGP topology by advertising the maximum IGP cost on links where LDP is not fully operational. This makes the IGP topology match the LDP topology. As described in Section 7.3.1 of [RFC7812], MRT is designed to be compatible with the LDP IGP synchronization mechanism. When the IGP advertises the maximum cost on a link where LDP is not fully operational, the link is excluded from MRT Island formation, which prevents the MRT algorithm from creating any paths using that link.

5. LDP MRT FEC Advertisements

This section describes how and when labels for MRT-Red and MRT-Blue FECs are advertised. In order to provide protection paths that are immediately usable by the point of local repair in the event of a failure, the associated LSPs need to be created before a failure occurs.

In this section, we will use the term "shortest-path FEC" to refer to the usual FEC associated with the shortest-path destination-based forwarding tree for a given prefix as determined by the IGP. We will use the terms "red FEC" and "blue FEC" to refer to FECs associated with the MRT-Red and MRT-Blue destination-based forwarding trees for a given prefix as determined by a particular MRT algorithm.

We first describe label distribution behavior specific to MRT. Then, we provide the correct interpretation of several important concepts in [RFC5036] in the context of MRT FEC label distribution.

[RFC5036] specifies two different Label Distribution Control Modes (Independent and Ordered), two different Label Retention Modes (Conservative and Liberal), and two different Label Advertisement Modes (Downstream Unsolicited and Downstream on Demand). The current specification for LDP MRT requires that the same Label Distribution Control, Label Retention, and Label Advertisement modes be used for the shortest-path FECs and the MRT FECs.

5.1. MRT-Specific Behavior

5.1.1. ABR Behavior and Use of the Rainbow FEC

Section 10.1 of [RFC7812] describes the need for an ABR to have different neighbors use different MPLS labels when sending traffic to the ABR for the same FEC. The method to accomplish this using the Rainbow MRT MT-ID is described in detail in [RFC7812]. Here we provide a brief summary. To those LDP peers in the same area as the best route to the destination, the ABR advertises two different labels corresponding to the MRT-Red and MRT-Blue forwarding trees for the destination. An LDP peer receiving these advertisements forwards MRT traffic to the ABR using these two different labels, depending on the FEC of the traffic. We refer to this as "best-area advertising and forwarding behavior", which is identical to normal MRT behavior.

For all other LDP peers supporting MRT, the ABR advertises a FEC-label binding for the FEC, which is in the Rainbow MRT MT-ID, with the label that corresponds to that FEC in the default forwarding tree for the destination. An LDP peer receiving this advertisement forwards MRT traffic to the ABR using this label, for both MRT-Red

and MRT-Blue traffic. We refer to this as "non-best-area advertising and forwarding behavior".

The use of the Rainbow-FEC by the ABR for non-best-area advertisements is RECOMMENDED. An ABR MAY advertise the label for the default topology in separate MRT-Blue and MRT-Red advertisements. An LSR advertising the MRT Capability MUST recognize the Rainbow MRT MT-ID and associate the advertised label with the specific prefix with the MRT-Red and MRT-Blue MT-IDs associated with all MRT Profiles that advertise LDP as the forwarding mechanism.

Due to changes in topology or configuration, an ABR and a given LDP peer may need to transition from best-area advertising and forwarding behavior to non-best-area behavior for a given destination, and vice versa. When the ABR requires best-area behavior for a red(blue) FEC, it MUST withdraw any existing label mappings advertisements for the corresponding Rainbow FEC and advertise label mappings for the red(blue) FEC. When the ABR requires non-best-area behavior for a red(blue) FEC, it MUST withdraw any existing label mappings for both red and blue FECs and advertise label mappings for the corresponding Rainbow FEC label-binding.

In this transition, an ABR should never advertise a red(blue) FEC before withdrawing the corresponding Rainbow FEC (or vice versa). However, should this situation occur, the expected behavior of an LSR receiving these conflicting advertisements is defined as follows:

- If an LSR receives a label mapping advertisement for a Rainbow FEC from an MRT LDP peer while it still retains a label mapping for the corresponding red or blue FEC, the LSR MUST continue to use the label mapping for the red or blue FEC, and it MUST send a Label Release message corresponding to the Rainbow FEC label advertisement.
- If an LSR receives a label mapping advertisement for a red or blue FEC while it still retains a label mapping for the corresponding Rainbow FEC, the LSR MUST continue to use the label mapping for the Rainbow FEC, and it MUST send a Label Release message corresponding to the red or blue FEC label advertisement.

5.1.2. Proxy-Node Attachment Router Behavior

Section 11.2 of [RFC7812] describes how MRT provides FRR protection for multi-homed prefixes using calculations involving a named proxy-node. This covers the scenario where a prefix is originated by a router in the same area as the MRT Island, but outside of the MRT Island. It also covers the scenario of a prefix being advertised by multiple routers in the MRT Island.

In the named proxy-node calculation, each multi-homed prefix is represented by a conceptual proxy-node that is attached to two real proxy-node attachment routers. (A single proxy-node attachment router is allowed in the case of a prefix advertised by a same area router outside of the MRT Island, which is singly connected to the MRT Island.) All routers in the MRT Island perform the same calculations to determine the same two proxy-node attachment routers for each multi-homed prefix. Section 5.9 of [RFC7811] describes the procedure for identifying one proxy-node attachment router as "red" and one as "blue" with respect to the multi-homed prefix, and computing the MRT red and blue next hops to reach those red and blue proxy-node attachment routers.

In terms of LDP behavior, a red proxy-node attachment router for a given prefix MUST originate a label mapping for the red FEC for that prefix, while the blue proxy-node attachment router for a given prefix MUST originate a label mapping for the blue FEC for that prefix. If the red(blue) proxy-node attachment router is an Island Border Router (IBR), then when it receives a packet with the label corresponding to the red(blue) FEC for a prefix, it MUST forward the packet to the Island Neighbor (IN) whose cost was used in the selection of the IBR as a proxy-node attachment router. The IBR MUST swap the incoming label for the outgoing label corresponding to the shortest-path FEC for the prefix advertised by the IN. In the case where the IN does not support LDP, the IBR MUST pop the incoming label and forward the packet to the IN.

If the proxy-node attachment router is not an IBR, then the packet MUST be removed from the MRT forwarding topology and sent along the interface(s) that caused the router to advertise the prefix. This interface might be out of the area/level/AS.

5.2. LDP Protocol Procedures in the Context of MRT Label Distribution

[RFC5036] specifies the LDP label distribution procedures for shortest-path FECs. In general, the same procedures can be applied to the distribution of label mappings for red and blue FECs, provided that the procedures are interpreted in the context of MRT FEC label distribution. The correct interpretation of several important concepts in [RFC5036] in the context of MRT FEC label distribution is provided below.

5.2.1. LDP Peer in RFC 5036

In the context of distributing label mappings for red and blue FECs, we restrict the LDP peer in [RFC5036] to mean LDP peers for which the LDP MRT Capability has been negotiated. In order to make this distinction clear, in this document we will use the term "MRT LDP

peer" to refer to an LDP peer for which the LDP MRT Capability has been negotiated.

5.2.2. Next Hop in RFC 5036

Several procedures in [RFC5036] use the next hop of a (shortest-path) FEC to determine behavior. The next hop of the shortest-path FEC is based on the shortest-path forwarding tree to the prefix associated with the FEC. When the procedures of [RFC5036] are used to distribute label mapping for red and blue FECs, the next hop for the red(blue) FEC is based on the MRT-Red(Blue) forwarding tree to the prefix associated with the FEC.

For example, Appendix A.1.7 of [RFC5036] specifies the response by an LSR to a change in the next hop for a FEC. For a shortest-path FEC, the next hop may change as the result of the LSR running a shortest-path computation on a modified IGP topology database. For the red and blue FECs, the red and blue next hops may change as the result of the LSR running a particular MRT algorithm on a modified IGP topology database.

As another example, Section 2.6.1.2 of [RFC5036] specifies that when an LSR is using LSP Ordered Control, it may initiate the transmission of a label mapping only for a (shortest-path) FEC for which it has a label mapping for the FEC next hop, or for which the LSR is the egress. The FEC next hop for a shortest-path FEC is based on the shortest-path forwarding tree to the prefix associated with the FEC. In the context of distributing MRT LDP labels, this procedure is understood to mean the following. When an LSR is using LSP Ordered Control, it may initiate the transmission of a label mapping only for a red(blue) FEC for which it has a label mapping for the red(blue) FEC next hop, or for which the LSR is the egress. The red or blue FEC next hop is based on the MRT-Red or Blue forwarding tree to the prefix associated with the FEC.

5.2.3. Egress LSR in RFC 5036

Procedures in [RFC5036] related to Ordered Control label distribution mode rely on whether or not an LSR may act as an egress LSR for a particular FEC in order to determine whether or not the LSR may originate a label mapping for that FEC. The status of being an egress LSR for a particular FEC is also used in the loop detection procedures described in [RFC5036]. Section 2.6.1.2 of [RFC5036] specifies the conditions under which an LSR may act as an egress LSR with respect to a particular (shortest-path) FEC:

1. The (shortest-path) FEC refers to the LSR itself (including one of its directly attached interfaces).

2. The next hop router for the (shortest-path) FEC is outside of the Label Switching Network.
3. (Shortest-path) FEC elements are reachable by crossing a routing domain boundary.

The conditions for determining an egress LSR with respect to a red or blue FEC need to be modified. An LSR may act as an egress LSR with respect to a particular red(blue) FEC under any of the following conditions:

1. The prefix associated with the red(blue) FEC refers to the LSR itself (including one of its directly attached interfaces).
2. The LSR is the red(blue) proxy-node attachment router with respect to the multi-homed prefix associated with the red(blue) FEC. This includes the degenerate case of a single red and blue proxy-node attachment router for a single-homed prefix.
3. The LSR is an ABR AND the MRT LDP peer requires non-best-area advertising and forwarding behavior for the prefix associated with the FEC.

Note that condition 3 scopes an LSR's status as an egress LSR with respect to a particular FEC to a particular MRT LDP peer. Therefore, the condition "Is LSR egress for FEC?" that occurs in several procedures in [RFC5036] needs to be interpreted as "Is LSR egress for FEC with respect to Peer?"

Also note that there is no explicit condition that allows an LSR to be classified as an egress LSR with respect to a red or blue FEC based only on the primary next hop for the shortest-path FEC not supporting LDP or not supporting LDP MRT Capability. These situations are covered by the proxy-node attachment router and ABR conditions (conditions 2 and 3). In particular, an Island Border Router is not the egress LSR for a red(blue) FEC unless it is also the red(blue) proxy-node attachment router for that FEC.

Also note that, in general, a proxy-node attachment router for a given prefix should not advertise an implicit or explicit null label for the corresponding red or blue FEC, even though it may be an egress LSR for the shortest-path FEC. In general, the proxy-node attachment router needs to forward red or blue traffic for that prefix to a particular loop-free island neighbor, which may be different from the shortest-path next hop. The proxy-node attachment router needs to receive the red or blue traffic with a non-null label to correctly forward it.

5.2.4. Use of Rainbow FEC to Satisfy Label Mapping Existence Requirements in RFC 5036

Several procedures in [RFC5036] require the LSR to determine if it has previously received and retained a label mapping for a FEC from the next hop. In the case of an LSR that has received and retained a label mapping for a Rainbow FEC from an ABR, the label mapping for the Rainbow FEC satisfies the label mapping existence requirement for the corresponding red and blue FECs. Label mapping existence requirements in the context of MRT LDP label distribution are modified as: "Has LSR previously received and retained a label mapping for the red(blue) FEC (or the corresponding Rainbow FEC) from the red(blue) next hop?"

As an example, this behavior allows an LSR that has received and retained a label mapping for the Rainbow FEC to advertise label mappings for the corresponding red and blue FECs when operating in Ordered Control label distribution mode.

5.2.5. Validating FECs in the Routing Table

In [RFC5036], an LSR uses its routing table to validate prefixes associated with shortest-path FECs. For example, Section 3.5.7.1 of [RFC5036] specifies that "an LSR receiving a Label Mapping message from a downstream LSR for a Prefix SHOULD NOT use the label for forwarding unless its routing table contains an entry that exactly matches the FEC Element." In the context of MRT FECs, a red or blue FEC element matches a routing table entry if the corresponding shortest-path FEC element matches a routing table entry.

5.2.6. Recognizing New FECs

Appendix A.1.6 of [RFC5036] describes the response of an LSR to the "Recognize New FEC" event, which occurs when an LSR learns a new (shortest-path) FEC via the routing table. In the context of MRT FECs, if the MRT LDP Capability has been enabled, then when an LSR learns a new shortest-path FEC, the LSR should generate "Recognize New FEC" events for the corresponding Red and Blue FECs in addition to the normally generated "Recognize New FEC" event for the shortest-path FEC

5.2.7. Not Propagating Rainbow FEC Label Mappings

A label mapping for the Rainbow FEC should only be originated by an ABR under the conditions described in Section 5.1.1. A neighbor of the ABR that receives a label mapping for the Rainbow FEC MUST NOT propagate a label mapping for that Rainbow FEC.

6. Security Considerations

The labels distributed by the extensions in this document create additional forwarding paths that do not follow shortest-path routes. The transit label swapping operations defining these alternative forwarding paths are created during normal operations (before a failure occurs). Therefore, a malicious packet with an appropriate label injected into the network from a compromised location would be forwarded to a destination along a non-shortest path. When this technology is deployed, a network security design should not rely on assumptions about potentially malicious traffic only following shortest paths.

It should be noted that the creation of non-shortest forwarding paths is not unique to MRT. For example, RSVP-TE [RFC3209] can be used to construct forwarding paths that do not follow the shortest path.

7. Potential Restrictions on MRT-Related MT-ID Values Imposed by RFC 6420

As discussed in the introduction, in addition to unicast-forwarding applications, MRT can be used to provide disjoint trees for multicast traffic distribution. In the case of PIM, this is accomplished by using the MRT red and blue next hops as the PIM Reverse Path Forwarding (RPF) topology, the collection of routes used by PIM to perform the RPF operation when building source trees. The PIM Multi-Topology ID (MT-ID) Join Attribute defined in Section 5.2 of [RFC6420] can be used to establish MRT-based multicast distribution trees. [RFC6420] limits the values of the PIM MT-ID from 1 through 4095.

For the purpose of reducing management overhead and simplifying troubleshooting, it is desirable to be able to use the same numerical value for the PIM MT-ID as for the MPLS MT-ID for multicast and unicast applications using MRT routes constructed using the same MRT Profile. In order to enable this simplification, the MPLS MT-ID values assigned in this document fall in the range 1 through 4095. The "MPLS Multi-Topology Identifiers" registry reflects this by listing the values from 3948 through 3995 as for MRT-related MPLS MT-ID values. This allows for 51 MRT-related MPLS MT-ID values that can be directly mapped to PIM MT-ID values, which accommodates 25 MRT Profiles with red and blue MT-ID pairs, with one extra for the Rainbow MPLS MT-ID value. [RFC7307] designates the MT-ID range 6-3995 as "Unassigned for future IGP topologies". As shown in the IANA Considerations, the guidance for the range 3948-3995 has been changed to "Unassigned (for future MRT-related values)".

8. IANA Considerations

IANA has allocated a value for the new LDP Capability TLV from the "Label Distribution Protocol (LDP) Parameters" registry under "TLV Type Name Space": MRT Capability TLV (0x050E).

Value	Description	Reference	Notes / Reg. Date
-----	-----	-----	-----
0x050E	MRT Capability TLV	RFC 8320	

IANA has allocated a value for the new LDP Status Code from the "Label Distribution Protocol (LDP) Parameters" registry under "Status Code Name Space": MRT Capability negotiated without MT Capability (0x00000034). The Status Code E-bit is set to 0.

Value	E	Description	Reference	Notes / Reg. Date
-----	-	-----	-----	-----
0x00000034	0	MRT Capability negotiated without MT Capability	RFC 8320	

IANA has allocated three values from the "MPLS Multi-Topology Identifiers" registry [RFC7307]:

- 3945 Rainbow MRT MPLS MT-ID
- 3946 Default Profile MRT-Red MPLS MT-ID
- 3947 Default Profile MRT-Blue MPLS MT-ID

Also, IANA has changed the Purpose field of the "MPLS Multi-Topology Identifiers" registry for MT-ID range 3948-3995 to "Unassigned (for future MRT-related values)". The registration procedure for the entire registry remains Standards Action [RFC8126]. The current registry is shown below:

Value	Purpose	Reference
-----	-----	-----
0	Default/standard topology	[RFC7307]
1	IPv4 in-band management	[RFC7307]
2	IPv6 routing topology	[RFC7307]
3	IPv4 multicast topology	[RFC7307]
4	IPv6 multicast topology	[RFC7307]
5	IPv6 in-band management	[RFC7307]
6-3944	Unassigned (for future IGP topologies)	
3945	Rainbow MRT MPLS MT-ID	RFC 8320
3946	Default Profile MRT-Red MPLS MT-ID	RFC 8320
3947	Default Profile MRT-Blue MPLS MT-ID	RFC 8320
3948-3995	Unassigned (for future MRT-related values)	RFC 8320
3996-4095	Reserved for Experimental Use	[RFC7307]
4096-65534	Unassigned (for MPLS topologies)	
65535	Wildcard Topology	[RFC7307]

9. References

9.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <<https://www.rfc-editor.org/info/rfc2119>>.
- [RFC5036] Andersson, L., Ed., Minei, I., Ed., and B. Thomas, Ed., "LDP Specification", RFC 5036, DOI 10.17487/RFC5036, October 2007, <<https://www.rfc-editor.org/info/rfc5036>>.
- [RFC5561] Thomas, B., Raza, K., Aggarwal, S., Aggarwal, R., and JL. Le Roux, "LDP Capabilities", RFC 5561, DOI 10.17487/RFC5561, July 2009, <<https://www.rfc-editor.org/info/rfc5561>>.
- [RFC6420] Cai, Y. and H. Ou, "PIM Multi-Topology ID (MT-ID) Join Attribute", RFC 6420, DOI 10.17487/RFC6420, November 2011, <<https://www.rfc-editor.org/info/rfc6420>>.

- [RFC7307] Zhao, Q., Raza, K., Zhou, C., Fang, L., Li, L., and D. King, "LDP Extensions for Multi-Topology", RFC 7307, DOI 10.17487/RFC7307, July 2014, <<https://www.rfc-editor.org/info/rfc7307>>.
- [RFC7811] Enyedi, G., Csaszar, A., Atlas, A., Bowers, C., and A. Gopalan, "An Algorithm for Computing IP/LDP Fast Reroute Using Maximally Redundant Trees (MRT-FRR)", RFC 7811, DOI 10.17487/RFC7811, June 2016, <<https://www.rfc-editor.org/info/rfc7811>>.
- [RFC7812] Atlas, A., Bowers, C., and G. Enyedi, "An Architecture for IP/LDP Fast Reroute Using Maximally Redundant Trees (MRT-FRR)", RFC 7812, DOI 10.17487/RFC7812, June 2016, <<https://www.rfc-editor.org/info/rfc7812>>.
- [RFC8126] Cotton, M., Leiba, B., and T. Narten, "Guidelines for Writing an IANA Considerations Section in RFCs", BCP 26, RFC 8126, DOI 10.17487/RFC8126, June 2017, <<https://www.rfc-editor.org/info/rfc8126>>.
- [RFC8174] Leiba, B., "Ambiguity of Uppercase vs Lowercase in RFC 2119 Key Words", BCP 14, RFC 8174, DOI 10.17487/RFC8174, May 2017, <<https://www.rfc-editor.org/info/rfc8174>>.

9.2. Informative References

- [ARCH] Atlas, A., Kebler, R., Wijnands, IJ., Csaszar, A., and G. Enyedi, "An Architecture for Multicast Protection Using Maximally Redundant Trees", Work in Progress, draft-atlas-rtgwg-mrt-mc-arch-02, July 2013.
- [IS-IS-MRT] Li, Z., Wu, N., Zhao, Q., Atlas, A., Bowers, C., and J. Tantsura, "Intermediate System to Intermediate System (IS-IS) Extensions for Maximally Redundant Trees (MRTs)", Work in Progress, draft-ietf-isis-mrt-03, June 2017.
- [OSPF-MRT] Atlas, A., Hegde, S., Bowers, C., Tantsura, J., and Z. Li, "OSPF Extensions to Support Maximally Redundant Trees", Work in Progress, draft-ietf-ospf-mrt-03, June 2017.
- [PARAM-SYNC] Bryant, S., Atlas, A., and C. Bowers, "Routing Timer Parameter Synchronization", Work in Progress, draft-ietf-rtgwg-routing-timer-param-sync-00, October 2017.

- [RFC3209] Awduche, D., Berger, L., Gan, D., Li, T., Srinivasan, V., and G. Swallow, "RSVP-TE: Extensions to RSVP for LSP Tunnels", RFC 3209, DOI 10.17487/RFC3209, December 2001, <<https://www.rfc-editor.org/info/rfc3209>>.
- [RFC5443] Jork, M., Atlas, A., and L. Fang, "LDP IGP Synchronization", RFC 5443, DOI 10.17487/RFC5443, March 2009, <<https://www.rfc-editor.org/info/rfc5443>>.
- [RFC7715] Wijnands, IJ., Ed., Raza, K., Atlas, A., Tantsura, J., and Q. Zhao, "Multipoint LDP (mLDP) Node Protection", RFC 7715, DOI 10.17487/RFC7715, January 2016, <<https://www.rfc-editor.org/info/rfc7715>>.

Acknowledgements

The authors would like to thank Ross Callon, Loa Andersson, Stewart Bryant, Mach Chen, Greg Mirsky, Uma Chunduri, and Tony Przygienda for their comments and suggestions.

Authors' Addresses

Alia Atlas
Juniper Networks
10 Technology Park Drive
Westford, MA 01886
United States of America

Email: akatlas@juniper.net

Kishore Tiruveedhula
Juniper Networks
10 Technology Park Drive
Westford, MA 01886
United States of America

Email: kishoret@juniper.net

Chris Bowers
Juniper Networks
1194 N. Mathilda Ave.
Sunnyvale, CA 94089
United States of America

Email: cbowers@juniper.net

Jeff Tantsura
Individual
United States of America

Email: jefftant.ietf@gmail.com

IJsbrand Wijnands
Cisco Systems, Inc.

Email: ice@cisco.com

