

MPLS Working Group
Internet-Draft
Intended status: Standards Track
Expires: June 18, 2015

A. Atlas
K. Tiruveedhula
C. Bowers
Juniper Networks
J. Tantsura
Ericsson
IJ. Wijnands
Cisco Systems, Inc.
December 15, 2014

LDP Extensions to Support Maximally Redundant Trees
draft-atlas-mpls-ldp-mrt-03

Abstract

This document specifies extensions to the Label Distribution Protocol(LDP) to support the creation of label-switched paths for Maximally Redundant Trees (MRT). 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 LSRs and LERs advertising the MRT Capability.

MRT-FRR uses LDP multi-topology extensions and requires three different multi-topology IDs to be allocated from the LDP MT-ID space.

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 June 18, 2015.

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.

Table of Contents

1.	Introduction	3
2.	Requirements Language	4
3.	Terminology	4
4.	Overview of LDP Signaling Extensions for MRT	5
4.1.	MRT Capability Advertisement	5
4.1.1.	Interaction of MRT Capability and MT Capability	6
4.1.2.	Interaction of LDP MRT Capability with IPv4 and IPv6	6
4.2.	Use of the Rainbow MRT MT-ID	7
4.3.	MRT-Blue and MRT-Red FECs	7
5.	LDP MRT FEC Advertisements	7
5.1.	MRT-specific behavior	8
5.1.1.	ABR behavior and use of the Rainbow FEC	8
5.1.2.	Proxy-node attachment router behavior	9
5.2.	LDP protocol procedures in the context of MRT label distribution	10
5.2.1.	LDP peer in RFC5036	10
5.2.2.	Next hop in RFC5036	10
5.2.3.	Egress LSR in RFC5036	11
5.2.4.	Use of Rainbow FEC to satisfy label mapping existence requirements in RFC5036	12
5.2.5.	Validating FECs in routing table	13
5.2.6.	Recognizing new FECs	13
5.2.7.	Not propagating Rainbow FEC label mappings	13
6.	Security Considerations	13
7.	IANA Considerations	14
8.	Acknowledgements	14
9.	References	14
9.1.	Normative References	14
9.2.	Informative References	15
	Authors' Addresses	16

1. Introduction

This document describes the LDP signaling extension and associated behavior necessary to support the architecture that defines how IP/LDP Fast-Reroute can use MRTs [[I-D.ietf-rtgwg-mrt-frr-architecture](#)]. It is necessary to be familiar with the architecture in [[I-D.ietf-rtgwg-mrt-frr-architecture](#)] to understand how and why the LDP extensions for behavior are needed.

At least one common standardized algorithm (e.g. the MRT Lowpoint algorithm explained and fully documented in [[I-D.ietf-rtgwg-mrt-frr-algorithm](#)]) is required 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 [[I-D.atlas-ospf-mrt](#)]. Extension to IS-IS are defined in [[I-D.li-isis-mrt](#)].

MRT can also be used to protect multicast traffic (signalled via PIM or mLDP) using either global protection or local protection [[I-D.atlas-rtgwg-mrt-mc-arch](#)]. An MRT path can be used to provide node-protection for mLDP traffic via the mechanisms described in [[I-D.wijnands-mpls-mldp-node-protection](#)]; 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 [[I-D.ietf-rtgwg-mrt-frr-architecture](#)] and the worst-case network convergence time can be flooded via the extension in Section 7 of [[I-D.atlas-ospf-mrt](#)].

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", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [[RFC2119](#)]

3. Terminology

For ease of reading, some of the terminology defined in [[I-D.ietf-rtgwg-mrt-frr-architecture](#)] is repeated here.

Redundant Trees (RT): 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. These can be computed in 2-connected graphs.

Maximally Redundant Trees (MRT): 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. Any RT is an MRT but many MRTs are not RTs. The two MRTs are referred to as MRT-Blue and MRT-Red.

MRT-Red: MRT-Red is used to describe one of the two MRTs; it is used to describe the associated forwarding topology and MT-ID. Specifically, MRT-Red is the decreasing MRT where links in the GADAG are taken in the direction from a higher topologically ordered node to a lower one.

MRT-Blue: MRT-Blue is used to describe one of the two MRTs; it is used to describe the associated forwarding topology and MT-ID. Specifically, MRT-Blue is the increasing MRT where links in the GADAG are taken in the direction from a lower topologically ordered node to a higher one.

Rainbow MRT MT-ID: It is useful to have an MT-ID that refers to the multiple MRT topologies and to the default topology. This is referred to as the Rainbow MRT 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: From the computing router, the set of routers that support a particular MRT profile and are connected via MRT-eligible links.

Island Border Router (IBR): A router in the MRT Island that is connected to a router not in the MRT Island and both routers 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..

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 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 LSP paths 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 (IANA)      |      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: TBA-MRT-LDP-1 (To Be Allocated by IANA)

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

S-bit: The State 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, 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 LDP MRT Capability with an LDP neighbor without also negotiating the LDP MT Capability, the LSR MUST behave as if LDP MRT Capability has not been 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 which 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 [[I-D.ietf-rtgwg-mrt-frr-architecture](#)] 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 MT-ID (TBA-MRT-LDP-2) will be assigned by IANA from the LDP MT-ID space. Prototype experiments have used the value 3999.

4.3. MRT-Blue and MRT-Red FECs

To provide MRT support in LDP, the MT Prefix FEC is used. [[I-D.ietf-rtgwg-mrt-frr-architecture](#)] contains the IANA request for the MRT-Red and MRT-Blue MT-IDs associated with the Default 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.

5. LDP MRT FEC Advertisements

This sections describes how and when labels for MRT-Red and MRT-Blue FECs are advertised. The associated LSPs must be created before a failure occurs, in order to provide protection paths which are

immediately usable by the point of local repair in the event of a failure.

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.

5.1. MRT-specific behavior

5.1.1. ABR behavior and use of the Rainbow FEC

Section 10.1 of [I-D.ietf-rtgwg-mrt-frr-architecture] 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. The method to accomplish this using the Rainbow MRT MT-ID is described in detail in [I-D.ietf-rtgwg-mrt-frr-architecture]. 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 Rainbow MRT MT-ID scoped FEC with the label corresponding to 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.

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 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 [[I-D.ietf-rtgwg-mrt-frr-architecture](#)] 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 a multiple routers in the MRT Island.

In the named proxy-node calculation, each multi-homed prefix is represented by a conceptual proxy-node which 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. The resulting graph in the computation consists of the MRT Island with the proxy-node representing the multi-homed prefix directly attached to the two proxy-node attachment routers. Conceptually, one then runs the MRT algorithm on this simplified graph to determine the MRT-red and blue next-hops to reach the proxy-node, which gives the next-hops to reach the prefix. In this manner, one can see that one of the two proxy-node attachment routers will always have a MRT-red next-hop to the proxy-node while the other will always have the MRT-blue next-hop to the proxy-node. We will refer to these as the red and blue proxy-node attachment routers respectively. (In practice, the MRT-red and blue next-hops to reach the proxy-node can then be determined in a more computationally

efficient manner based on the MRT-red and blue next-hops to reach the proxy-node attachment routers, as described in [[I-D.ietf-rtgwg-mrt-frr-algorithm](#)].)

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 a 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 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 [RFC5036](#)

In the context of distributing label mappings for red and blue FECs, we restrict 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 [RFC5036](#)

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 how 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 [RFC5036](#)

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 loop detection procedures 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 area border router (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 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 [RFC5036](#)

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 which 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 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

Section 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, when MRT LDP capability has been enabled, when an LSR learns a new shortest path FEC, it should generate "Recognize New FEC" events for the corresponding red and blue FECs, in addition to the "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.

7. IANA Considerations

IANA is requested to allocate a value for the new LDP Capability TLV (the first free value in the range 0x0500 to 0x05FF) from the LDP registry "TLV Type Name Space": MRT Capability TLV (TBA-MRT-LDP-1).

Value	Description	Reference	Notes / Reg. Date
-----	-----	-----	-----
TBA-MRT-LDP-1	MRT Capability TLV	[This draft]	

IANA is requested to allocate a value for the new LDP Status Code (the first free value in the range 0x00000032-0x00000036) from the LDP registry "Status Code Name Space": "MRT Capability negotiated without MT Capability" (TBA-MRT-LDP-3).

Value	E	Description	Reference	Notes / Reg. Date
-----	-	-----	-----	-----
TBA-MRT-LDP-3	0	MRT Capability negotiated without MT Capability	[This draft]	

IANA is requested to allocate a value from the MPLS Multi-Topology Identifiers Name Space [[RFC7307](#)]: Rainbow MRT MT-ID (TBA-MRT-LDP-2).

Value	Purpose	Reference
-----	-----	-----
TBA-MRT-LDP-2	Rainbow MRT MT-ID	[This draft]

8. Acknowledgements

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

9. References

9.1. Normative References

[I-D.ietf-rtgwg-mrt-frr-algorithm]
 Enyedi, G., Csaszar, A., Atlas, A., Bowers, C., and A. Gopalan, "Algorithms for computing Maximally Redundant Trees for IP/LDP Fast-Reroute", [draft-rtgwg-mrt-frr-algorithm-01](#) (work in progress), July 2014.

[I-D.ietf-rtgwg-mrt-frr-architecture]

Atlas, A., Kebler, R., Bowers, C., Enyedi, G., Csaszar, A., Tantsura, J., Konstantynowicz, M., and R. White, "An Architecture for IP/LDP Fast-Reroute Using Maximally Redundant Trees", [draft-rtgwg-mrt-frr-architecture-04](#) (work in progress), July 2014.

[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.

[RFC7307] Zhao, Q., Raza, K., Zhou, C., Fang, L., Li, L., and D. King, "LDP Extensions for Multi-Topology", [RFC 7307](#), July 2014.

9.2. Informative References

[I-D.atlas-ospf-mrt]

Atlas, A., Hegde, S., Bowers, C., and J. Tantsura, "OSPF Extensions to Support Maximally Redundant Trees", [draft-atlas-ospf-mrt-02](#) (work in progress), July 2014.

[I-D.atlas-rtgwg-mrt-mc-arch]

Atlas, A., Kebler, R., Wijnands, I., Csaszar, A., and G. Enyedi, "An Architecture for Multicast Protection Using Maximally Redundant Trees", [draft-atlas-rtgwg-mrt-mc-arch-02](#) (work in progress), July 2013.

[I-D.li-isis-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(MRT)", [draft-li-isis-mrt-01](#) (work in progress), July 2014.

[I-D.wijnands-mpls-mldp-node-protection]

Wijnands, I., Rosen, E., Raza, K., Tantsura, J., Atlas, A., and Q. Zhao, "mLDP Node Protection", [draft-wijnands-mpls-mldp-node-protection-04](#) (work in progress), June 2013.

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

Authors' Addresses

Alia Atlas
Juniper Networks
10 Technology Park Drive
Westford, MA 01886
USA

Email: akatlas@juniper.net

Kishore Tiruveedhula
Juniper Networks
10 Technology Park Drive
Westford, MA 01886
USA

Email: kishoret@juniper.net

Chris Bowers
Juniper Networks
1194 N. Mathilda Ave.
Sunnyvale, CA 94089
USA

Email: cbowers@juniper.net

Jeff Tantsura
Ericsson
300 Holger Way
San Jose, CA 95134
USA

Email: jeff.tantsura@ericsson.com

IJsbrand Wijnands
Cisco Systems, Inc.

Email: ice@cisco.com

