Network Working Group
Internet-Draft
Intended status: Standards Track
Expires: April 21, 2014

Z. Li
N. Wu
Q. Zhao
Huawei Technologies
A. Atlas
C. Bowers
Juniper Networks
J. Tantsura
Ericsson

October 18, 2013

# Intermediate System to Intermediate System (IS-IS) Extensions for Maximally Redundant Trees (MRT) draft-li-isis-mrt-00

## Abstract

This document describes necessary extensions to IS-IS to support the distributed computation of Maximally Redundant Trees (MRT). Some example uses of the MRTs include IP/LDP Fast-Reroute and global protection or live-live for multicast traffic. The extensions indicate what MRT profile(s) each router supports. Different MRT profiles can be defined to support different uses and to allow transition of capabilities. An extension is introduced to flood MRT-Ineligible links, due to administrative policy.

## 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 April 21, 2014.

Table of Contents



## 1. Introduction

The IS-IS protocol is specified in [ISO10589], with extensions for supporting IPv4 and IPv6 specified in [RFC1195] and [RFC5308]. Each Intermediate System (IS) (router) advertises one or more IS-IS Link State Protocol Data Units (LSPs) with routing information. Each LSP is composed of a fixed header and a number of tuples, each consisting

of a Type, a Length, and a Value. Such tuples are commonly known as TLVs, and are a good way of encoding information in a flexible and extensible format.





[I-D.ietf-rtgwg-mrt-frr-architecture] gives a complete solution for IP/LDP fast-reroute using Maximally Redundant Trees (MRT) to provide alternates. This document describes the necessary signaling extensions for supporting MRT-FRR used in IS-IS routing domain.

## 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

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 R 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 R 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.

MRT Island: From the computing router, the set of routers that support a particular MRT profile and are connected via MRT- eligible links.

GADAG: Generalized Almost Directed Acyclic Graph - a graph which is the combination of the ADAGs of all blocks. Transforming a network graph into a GADAG is part of the MRT algorithm.

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.

# 4. Overview of IS-IS Signaling Extensions for MRT





As stated in [I-D.enyedi-rtgwg-mrt-frr-algorithm], it is necessary for each MRT-Capable router to compute MRT next hops in a consistent fashion. This is achieved by using same MRT profile and selecting the unique root in an MRT Island which is connected by MRT-Eligible links. Each of these issues will be discussed in following sections separately.

## 4.1. Supporting MRT Profiles

The contents and requirements of an MRT profile has been defined in [I-D.ietf-rtgwg-mrt-frr-architecture]. The parameters and behavioral rules contained in an MRT profile define one router's MRT capabilities. Based on common capabilities, one unified MRT Island is built.

The MRT-Capable router MUST advertise its corresponding MRT profiles by IS-IS protocol extension within IS-IS routing domain. The capabilities of advertiser MUST conform to the profile it claimed completely, especially the MT-IDs, the algorithm and the corresponding forwarding mechanism. This advertisement MUST have level scope. One router MAY support multiple MRT profiles and it MUST advertise these profiles in corresponding IS-IS level. The MT-IDs used in one supported MRT Profile MUST NOT overlap with those MT-IDs used in a different supported MRT Profile.

The default MRT Profile is defined in [I-D.ietf-rtgwg-mrt-frr-architecture]. Its behavior is intended to support IP/LDP unicast and multicast Fast-Reroute. MRT-Capable routers SHOULD support the default MRT profile.

## 4.2. Electing GADAG Root

As per [I-D.enyedi-rtgwg-mrt-frr-algorithm], a GADAG root MUST be selected for one MRT Island. An unique GADAG root in common-sense among MRT Island routers is a necessity to do MRT computation. Since the selection of the GADAG root can affect the alternates and the traffic through it, the selection rules give network operator a knob to control the alternates and the traffic inside the MRT Island. Relevant discussion for the relationship between GADAG root role and MRT Island alternates is out of the scope of this document.





Each MRT-Capable router MUST advertise its priority for GADAG root selection. One router can only have one priority in the same MRT Island. It can have multiple priorities for different MRT Islands it supports. Routers that are marked as overloaded([RFC3787]) are not qualified as candidate for root selection. A GADAG root is selected by comparing their priorities. The router with the highest priority among those available candidates is the GADAG root with higher system ID as the tie-breaker if priorities are same.

When the current root is out of service or new router with higher priority joined into the MRT Island, the GADAG root MUST be re-selected. A new MRT computation will be triggered because of such a topology change.

## 4.3. Advertising MRT-Ineligible Links for MRT

For certain administrative or management reason, some links may not be involved into MRT computation. In this scenario, MRT-Capable router MUST claim those MRT-Ineligible links are out of MRT Island scope. If such claim splits current MRT Island then MRT computation has to be done inside the modified MRT Island which the computing router belongs to.

## 4.4. Triggering an MRT Computation

An MRT Computation can be triggered through topology changes or MRT capability changes of any router in the MRT Island. It is always triggered for a given MRT Profile in the corresponding level. First, the associated MRT Island is determined. Then, the GADAG Root is selected. Finally, the actual MRT algorithm is run to compute the transit MRT-Red and MRT-Blue topologies. Additionally, the router MAY choose to compute MRT-FRR alternates or make other use of the MRT computation results.

Prefixes can be attached and detached and have their associated MRT-Red and MRT-Blue next-hops computed without requiring a new MRT computation.

# 5. MRT Capability Advertisement

MRT-Capable router MUST identify its MRT capabilities through IS-IS Link State Packet(LSP) in level scope.





## 5.1. Advertising MRT Capability in IS-IS LSP

One new M-bit is introduced into TLV 229 to identify router is MRT-Capable. Structure of TLV 229 is stated in [RFC5120] as pictured below:

TYPE: 229

LENGTH: total length of the value field, it SHOULD be 2 times the number of MT components.

VALUE: one or more 2-byte MT components, structured as follows:

No. of Octets

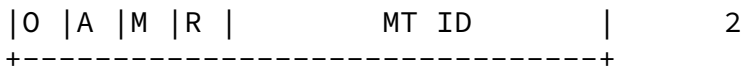


Bit M identifies the originator is of MRT-Capable. The MRT-Blue and the MRT-Red alternates will be calculated for the MT identified by MT-ID.

This M-bit MUST be set and checked in LSP fragment 0. An MRT-Capable router MUST advertise this TLV with M-bit set for corresponding MT. For instance, if M-bit is set for MT-ID #0, MRT alternates will be calculated for standard topology.

If only M-bit is advertised for MRT-Capabilities without any other MRT information then the router is regarded as supporting default MRT profile with default GADAG root priority.

## 5.2. MRT Profile sub-TLV in IS-IS Router CAPABILITY TLV

One new MRT Profile sub-TLV is introduced into IS-IS Router CAPABILITY TLV[RFC4971] to advertise MRT capabilities. Since MRT is per level scope, the S-bit and D-bit of IS-IS Router CAPABILITY TLV MUST be set zero. Structure of MRT Profile sub-TLV is pictured as below:

TYPE: TBD

LENGTH: 8 octets

VALUE:

MT ID (2 octet with 4 bits reserved)





Profile ID (1 octet)

GADAG priority (1 octet)

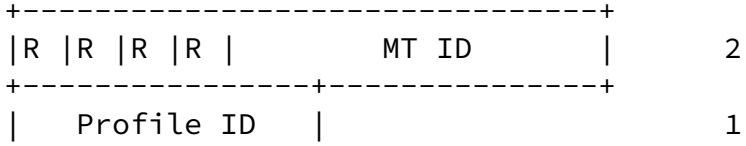

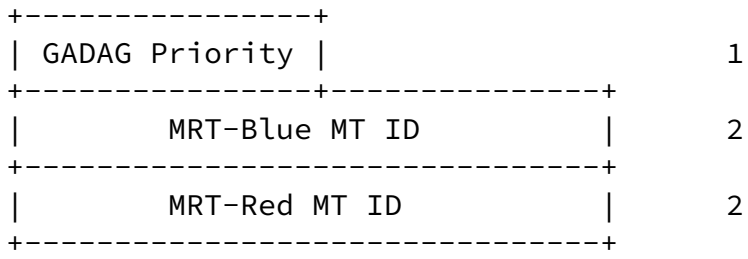


12-bit MT ID represents the base MT topology which MRT computation is based on. Profile ID represents the MRT profile this router supports and GADAG Priority is the priority for root selection. The range of this priority is [0, 255] with 128 as the default value. Higher numerical value means higher priority.

Those routers which do not want to be involved into GADAG root selection can have priority 0. If all routers in MRT Island carry the same priority then the one with the highest system ID has to be chosen as GADAG root.

If the MRT-Blue MT-ID is 0, then the value specified in the associated MRT Profile is assumed. If the MRT-Red MT-ID is 0, then the value specified in the associated MRT profile is assumed. The MRT-Blue MT-ID and MRT-Red MT-ID MUST NOT be the reserved values for MT-ID([RFC5120]). The value for MRT-Blue MT-ID and MRT-Red MT-ID MUST be different except for 0. As stated above, the MRT-Blue MT-ID and MRT-Red MT-ID MUST NOT overlap among profiles if multiple MRT-Profile sub-TLVs are advertised.

This sub-TLV can occur multiple times if this router support multiple MRT profiles. This can happen during transition or to support multiple uses of MRT which prefer different profiles.

### 5.3. MRT-Ineligible Links sub-TLV in IS-IS Router CAPABILITY TLV

As a matter of policy, some links may not be available for the MRT computation, which can prevent alternates or traffic using these links. For instance, policy can be made to prevent fast-rerouted traffic from taking those links.





For a link to be excluded from the MRT computation, it MUST be

advertised as sub-TLV in IS-IS Router CAPABILITY TLV which is in level scope with B-bit and D-bit unset. The MRT-Ineligible Link sub-TLV is structured as below:

TYPE: TBD

LENGTH: from 9 to 255 octets

VALUE:

MT ID (2 octet with 4 bits reserved)

System ID and pseudo-node number (7 octet for each MRT-Ineligible Link)

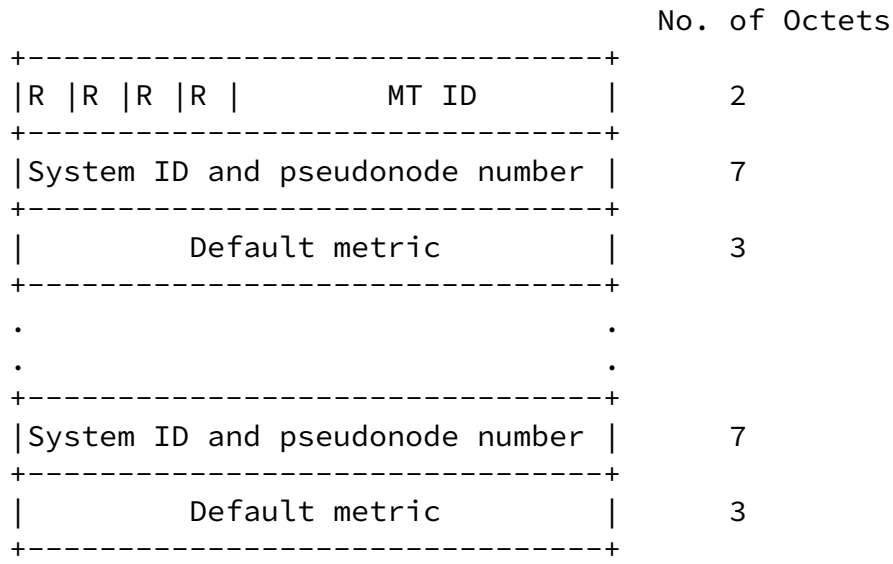


Each MRT-Ineligible Link is identified by neighbor's System ID and pseudo-node number and Default metric, same as IS Reachability TLV. This sub-TLV MAY occur multiple times if multiple links are ineligible.

## 6. Handling MRT Capability Sending and Receiving

The M-bit which identifies router's MRT capability MUST be advertised in LSP fragment 0. Those MRT related sub-TLVs SHOULD be ignored when MRT Capability bit is unset. When changes in MRT capabilities are received, an MRT computation SHOULD be triggered but MAY be delayed for a while to allow reception of all MRT-related information.

### 6.1. Advertising MRT extension

MRT sub-TLVs are encapsulated in the Router Capability TLV and advertised through LSP PDU for the level-wide. MRT sub-TLVs are optional. If one router does not support MRT, it MUST NOT advertise those sub-TLVs.

Since the advertisement scope of the MRT sub-TLV is level-wide, the D-Bit and S-Bit of the Router Capability TLV MUST be set as 0 when it is advertised. If other sub-TLVs in the Router Capability TLV need different values for those two bits, there MUST be an independent Router Capability TLV for MRT sub-TLVs.

When MRT related information is changed for the router or existing IS-IS LSP mechanisms are triggered for refreshing or updating, MRT sub-TLVs MUST be advertised if the router is MRT-Capable.

For administrative policies or reasons, certain links can not be involved into MRT Computation. MRT-Ineligible sub-TLV is used to advertise these links among MRT Island.

### 6.2. Parsing MRT extension

MRT extension MUST NOT affect the peer setup and the routing calculation of the standard topology.

MRT sub-TLVs SHOULD be validated like other sub-TLVs when received. MRT sub-TLVs SHOULD also be taken for the checksum calculation and authentication.

If MT-ID conflict is found for MRT-Red or MRT-blue from multiple sub-TLVs then those associated sub-TLVs MUST be ignored.

Links advertised in MRT-Ineligible sub-TLV MUST be precluded from MRT Computation. The removal of those links may change the computing router's MRT Island significantly.

## 7. Backwards Compatibility

The M-bit for MRT capability, the MRT Profile sub-TLV and the MRT-Ineligible Link sub-TLV defined in this document SHOULD NOT introduce any interoperability issues. Routers that do not support these MRT extensions SHOULD silently ignore them. Alternates or traffic MUST NOT be affected in current IS-IS routing domain.

## 8. Security Considerations

This IS-IS extension is not believed to introduce new security concerns.

## 9. IANA Considerations

Please allocate a value from the IS-IS Router CAPABILITY TLV[RFC4971] for the MRT Profile sub-TLV, and for the MRT-Ineligible Link sub-TLV.

## 10. References

### 10.1. Normative References

[ISO10589] ISO. Intermediate System to Intermediate System Routing Exchange Protocol for Use in Conjunction with the Protocol for Providing the Connectionless-Mode Network Service. ISO 10589, 1992.

[I-D.enyedi-rtgwg-mrt-frr-algorithm] Envedi, G., Csaszar, A., Atlas, A., cbowers@juniper.net, c., and A. Gopalan, "Algorithms for computing Maximally Redundant Trees for IP/LDP Fast- Reroute", draft-enyedi-rtgwg-mrt-frr-algorithm-03 (work in progress), July 2013.

[I-D.ietf-rtgwg-mrt-frr-architecture] Atlas, A., Kebler, R., Envedi, G., Csaszar, A., Tantsura, J., Konstantynowicz, M., and R. White, "An Architecture for IP/LDP Fast-Reroute Using Maximally Redundant Trees", draft-ietf-rtgwg-mrt-frr-architecture-03 (work in progress), July 2013.

[RFC3137] Retana, A., Nguyen, L., White, R., Zinin, A., and D. McPherson, "OSPF Stub Router Advertisement", RFC 3137, June 2001.

[RFC3787] Parker, J., "Recommendations for Interoperable IP Networks using Intermediate System to Intermediate System (IS-IS)", RFC 3787, May 2004.

## 10.2. Infomative References


[RFC1195] Callon, R., "Use of OSI IS-IS for routing in TCP/IP and dual environments", RFC 1195, December 1990.

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

[RFC3787] Parker, J., "Recommendations for Interoperable IP Networks using Intermediate System to Intermediate System (IS-IS)", RFC 3787, May 2004.







[RFC4971] Vasseur, JP., Shen, N., and R. Aggarwal, "Intermediate System to Intermediate System (IS-IS) Extensions for Advertising Router Information", RFC 4971, July 2007.

[RFC5120] Przygienda, T., Shen, N., and N. Sheth, "M-ISIS: Multi Topology (MT) Routing in Intermediate System to Intermediate Systems (IS-ISs)", RFC 5120, February 2008.

[RFC5308] Hopps, C., "Routing IPv6 with IS-IS", RFC 5308, October 2008.


## Authors' Addresses


Zhenbin Li
Huawei Technologies
Huawei Bld., No.156 Beiqing Rd.
Beijing 100095
China

Email: lizhenbin@huawei.com

Nan Wu
Huawei Technologies
Huawei Bld., No.156 Beiqing Rd.
Beijing 100095
China

Email: eric.wu@huawei.com

Quintin Zhao
Huawei Technologies
125 Nagog Technology Park
Acton, MA  01719
USA

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

Email: akatlas@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