

TRILL Working Group
Internet Draft

Intended status: Standard Track

Tissa Senevirathne
Les Ginsberg
Satya Dillikar
CISCO

Ayan Banerjee
Consultant

Sam Aldrin
HuaaWei

Naveen Nimmu
Broadcom

September 11, 2012

Expires: March 2013

Multi Topology Encoding within TRILL data frames
draft-tissa-trill-mt-encode-02.txt

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), its areas, and its working groups. Note that other groups may also distribute working documents as Internet-Drafts.

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."

The list of current Internet-Drafts can be accessed at <http://www.ietf.org/ietf/lid-abstracts.txt>

The list of Internet-Draft Shadow Directories can be accessed at <http://www.ietf.org/shadow.html>

This Internet-Draft will expire on December 11, 2012.

Internet-Draft [draft-tissa-trill-mt-encode-02](#)

September 2012

Copyright Notice

Copyright (c) 2012 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.

Abstract

Two alternate methods of encoding Multi Topology Identifier within the TRILL data frames are presented. Methods proposed herein do not require overloading TRILL RBridge nickname to encode Multi Topology Identifier. A method that expands TRILL nickname space from 16bits to 24 bits is also presented in this draft.

Table of Contents

1.	Introduction.....	3
2.	Conventions used in this document.....	4
3.	Multi Topology Encoding.....	4
3.1.	Nickname construction.....	5
3.2.	Use of Next-Hop VLAN for MT Encoding.....	5
4.	Multi Topology Interoperability.....	6
4.1.	Interoperability during Migration.....	6
4.2.	Interoperability with RBridges with Non MT capable data plane.....	7
4.3.	Interoperability with MT unaware RBRdigs.....	8
5.	Backward compatibility.....	8
6.	IS-IS sub-TLV definition.....	8
6.1.	MT capability.....	8
7.	Security Considerations.....	9
8.	Assignment Considerations.....	9
8.1.	IANA Considerations.....	9
8.2.	IEEE Considerations.....	9
9.	References.....	9
9.1.	Normative References.....	9
9.2.	Informative References.....	9

10. Acknowledgments.....	9
Authors' Addresses.....	10

[1. Introduction](#)

Multi Topology is an attractive concept that allows creating different virtual topologies or overlays on top of a single physical topology. There are several important applications. Few such applications are listed below. The list below is by no means an exhaustive list and there are other applications that may utilize the MT framework.

Application specific topology (Storage topology vs. data topology)

Virtual topologies for different customers

Expansion of TRILL nickname space

Traffic Engineering

[RFC5120] presents Multi topology (MT) framework for IS-IS. MT has two important parts; IS-IS control plane extensions and Data plane encoding. [[RFC5120](#)] presents IS-IS control plane extensions. [TRILL-MT1] and [[TRILL-MT2](#)] presents required IS-IS sub-TLV definitions and the data plane encoding for TRILL.

In this document we propose methods that facilitate encoding of 16 bit Multi topology ID in to TRILL frames without reducing the effective TRILL nickname space. Additionally, the proposed scheme does not require definition of new Topology mapping sub-TLV. In other words, IS-IS control plane is similar to [[RFC5120](#)] and does not require additional complexity of mapping absolute Topology ID to abbreviated topology ID.

Methods proposed in this document encode the MT topology ID into TRILL data frames without modifying the TRILL header. Hence, MT capable, RBridges interfacing with non-MT capable RBridges can selectively not encode the proposed MT extensions on interfaces with non-MT capable RBridges. Non MT capable RBridges do not required to be in the base topology. They can be in any valid topology. Only restriction is non-MT capable RBridges can belong to a single

topology only. In its IS-IS HELLO messages, RrRidges exchange its MT capability and topology information. RBridges that are not capable of supporting proposed MT extension in data plane, MUST announce itself as non MT capable, but MAY advertise its association to a topology other than the base topology by including MT extensions proposed in [[RFC5120](#)]. MT encoding capability is announced by setting the proposed MT encoding capability bit in Port TRILL Version sub-TLV [[rfc6326bis](#)]. Presence of IS-IS Multi Topology TLV [[RFC5120](#)], indicates only the associated topology. MT encoding

capability indicates RBridges ability to support proposed data plane extensions. When MT capability is not set RBridge MUST not use the proposed data plane encoding methods, instead it must associate the announcing RBridge to the advertised topology or base topology in the absence of Multi-Topology TLV [[RFC5120](#)].

TRILL protocol, as defined in [[RFC6325](#)], defines 16bit nickname space. 16bit nickname space allows up to 65536 unique nicknames. However, it has been discussed in the working group that, more the 65536 unique names are required in certain large deployments. Possible usage of Upperbits of Nickname is also being considered for encoding Multitoplogy, which further reduces the available nickname space. Presented in this document is a method that allows expanding the 16bit nickname space to a 24bit nickname space, without modifying the TRILL header defined in [[RFC6325](#)].

[2.](#) Conventions used in this document

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in [RFC-2119](#) [[RFC2119](#)].

[3.](#) Multi Topology Encoding

[RFC6325] TRILL Base Protocol, proposes encoding scheme of TRILL frames. TRILL frames contain outer MAC Header, 802.1QTAG, TRILL header and user Data. We propose to include multi topology ID after the 802.1Q TAG. Multi Toplogy ID is preceded by Ethernet Type MT-ETHTYPE.

Outer Ethernet Header:

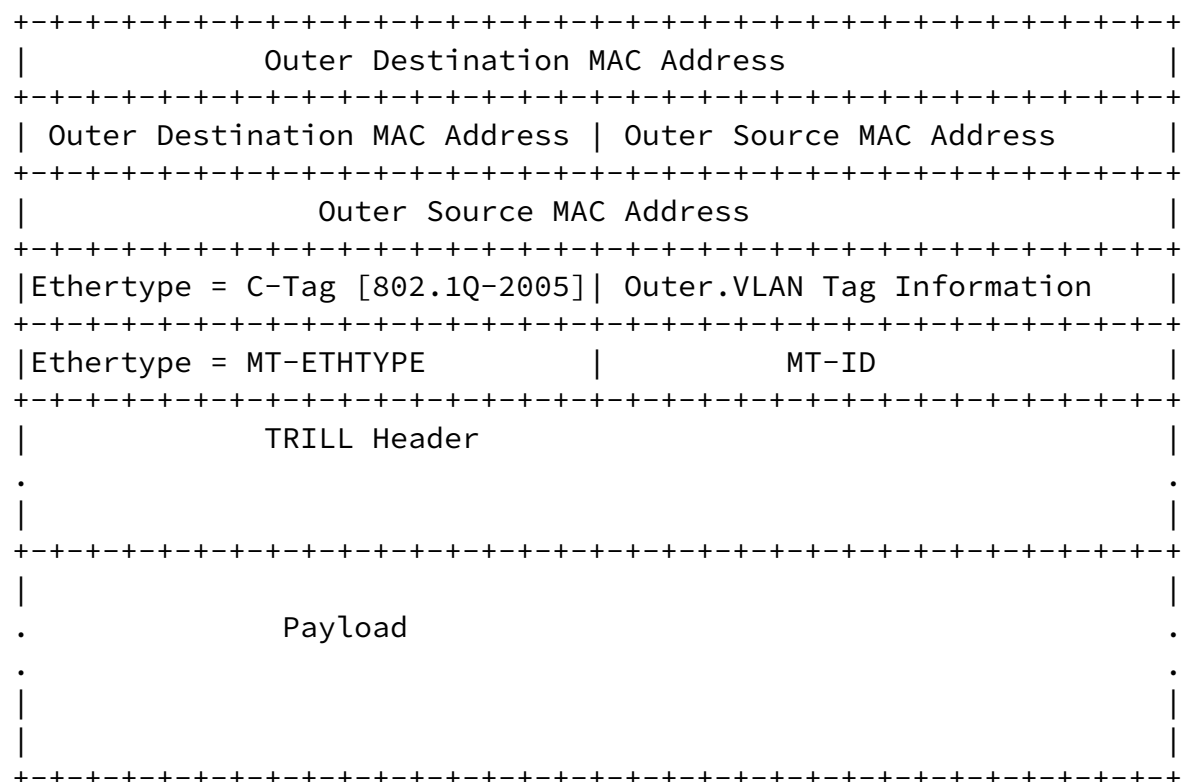


Figure 1 MT ID extensions in TRILL

Ethtype-MT : 16 bit Ethtype to encode Multi Topology ID.

MT-ID : 16bit Multi Topology ID
NICKN-Ethtype : 16 bit Ethtype to encode nickname expansion.

[3.1.](#) Nickname construction

Each RBridge that is compatible with the proposed scheme, First check for presence of Nick-Ethtype, if present extract the EG-NICKN-MSB and IG-NICKN-MSB. EG-NICKN-MSB and IG-NICKN-MSB are then concatenated with the Egress TRILL nickname and the Ingress TRILL nickname to form 24bit nickname space. The derived nickname MUST be utilized for forwarding.

[3.2.](#) Use of Next-Hop VLAN for MT Encoding

Alternatively, Next-Hop VLAN may be utilized to encode MT ID in point to point only networks. Next Hop VLAN on TRILL outer header is

independent of the inner VLAN. On Point-Point links Next Hop VLAN is only required to be of local significance. Hence, we propose to map topologies to Next-Hop VLAN per link basis.

For sake of simplicity, we propose to map topology-id to Next-VLAN based on local policies such as configuration.

[4.](#) Multi Topology Interoperability

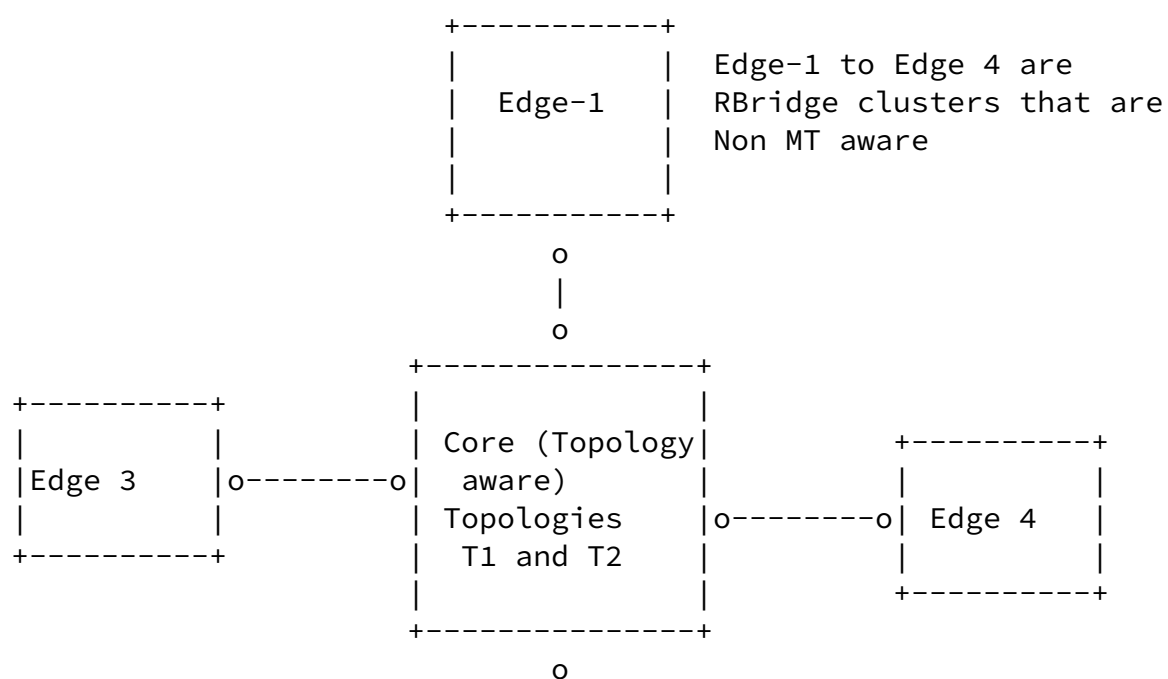
There are three possible scenarios of Interoperability with RBridges that are non-MT capable.

1. Interoperability During migration.
2. Interoperability with RBridges that are non-MT capable in the data plane. (i.e. Software is MT aware and supports the extensions specified here-in but, data plane is not capable of supporting the proposed encoding methods).
3. Interoperability with Rbridges that are MT unaware in both Control and data planes.

[4.1.](#) Interoperability during Migration

We recommend upgrading from the core to the edge, as depicted in

the figure below. With this approach, different clusters of RBridges may belong to different topologies or to the same topology. RBridges in the core provide connectivity to RBridge clusters at the edge in a topology aware manner.






Figure 1 MT aware interconnect

In the above diagram Core may be configured to connect Edge 1 and Edge 2 to a different Topology than the topology of Edge 3 and Edge 4.

R Bridges in Edge 1 - 4 are not required to be MT capable or aware. R Bridges in the core associate the corresponding links to the appropriate topology.

[4.2.](#) Interoperability with R Bridges with Non MT capable data plane

R Bridges with Non MT capable data plane may implement MT support by dedicating a separate link for each topology.

Alternatively, R Bridges, on point-point links may assign a different next hop VLAN for different topologies and derive topology ID based on VLAN. Use Next-Hop VLAN reduces the need for multiple physical

links. This method may be utilized as a permanent method for MT encoding in Point-Point only networks.

[4.3.](#) Interoperability with MT unaware R Bridges

MT aware R Bridges identify MT unaware R Bridges with either not presence of capability flags (pre RFC6326bis) or MT capability flags not being set ([Section 6.1.](#)). In such an event MT aware R Bridges MUST only forward traffic related to the base topology to MT unaware R Bridges. Additionally, proposed encoding MUST be removed prior to forwarding to MT unaware R Bridges.

[5.](#) Backward compatibility

The proposed methods are encoded as part of the outer header of the TRILL frame. An RBridge that is aware of the proposed extensions when interfacing with an RBridge that is not capable of the proposed extensions MUST remove the proposed encoding from the outer header, prior to transmission of TRILL frames on those links that has R Bridges that are not capable of the proposed extensions.

[6.](#) IS-IS sub-TLV definition

[6.1.](#) MT capability

We propose to define two MT capability flags within Port TRILL Version sub-TLV.

1. MT Encoding capability
2. MT to NH-VLAN Encoding capability

MT Encoding capability flag indicates the RBridge is capable encoding MT ID using ETHTYPE-MT as defined in [section 3](#).

MT to NH-VLAN Encoding capability flag indicates the announcing RBridge is capable of using NH-VLAN to MT ID mapping as presented in [section 3.2](#).

When both of the flags are set RBRridge SHOULD select MT Encoding capability.

[7.](#) Security Considerations

TBD

[8.](#) Assignment Considerations

[8.1.](#) IANA Considerations

IANA is requested to allocate MT Encoding capability Flag, MT to NH-VLAN Encoding capability Flags and Nickname MSB capability flag under Port TRILL version sub-TLV.

[8.2.](#) IEEE Considerations

IEEE is requested to assign new Ether Type to represent MT-ETHTYPE defined in [section 3](#).

[9.](#) References

[9.1.](#) Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", [BCP 14](#), [RFC 2119](#), March 1997.
- [RF5120] Przygienda, T. et.al , "M-ISIS: Multi Topology (MT) Routing in Intermediate System to Intermediate System (IS-ISs)", [RFC 5120](#), February, 2008.

[9.2.](#) Informative References

- [TRILL-MT1] Manral,V. et.al., "Multiple Topology Routing Extensions for Transparent Interconnection of Lots of Links (TRILL)", [draft-manral-isis-trill-multi-topo-03](#), Work in Progress, December, 2011.
- [TRILL-MT2] Eastlake, D. et.al., "Multi Topology TRILL", [draft-eastlake-trill-rbridge-multi-topo-02](#), Work in Progress, January, 2012.
- [rfc6326bis] Eastlake, D. et.al., "Transparent Interconnection of Lots of Links (TRILL) Use of IS-IS", [draft-eastlake-isis-rfc6326bis-02](#), Work in Progress, December, 2011.

[10.](#) Acknowledgments

Special acknowledgment to Dinesh Dutt, for encouragements, support and coming up with the initial idea.

Authors' Addresses

Tissa Senevirathne
CISCO Systems
375 East Tasman Drive,
San Jose, CA 95134

Phone: +1-408-853-2291
Email: tsenevir@cisco.com

Les Ginsberg
CISCO Systems
510 McCarthy Blvd,
Milpitas, CA 95035.

Email: ginsberg@cisco.com

Satya Dillikar
CISCO Systems
375 East Tasman Drive,
San Jose, CA 95134

Email: dsatya@cisco.com

Ayan Banerjee
Consultant

Email: ayabaner@gmail.com

Sam Aldrin
HuaWei Technologies
2330 Central Expressway
Santa Clara, CA 95051, USA

Email: aldrin.ietf@gmail.com

Naveen Nimmu
Broadcom th 9 Floor, Building no 9, Raheja Mind space
Hi-Tec City, Madhapur,
Hyderabad - 500 081 India.

Email: naveen@broadcom.com