Network Working Group
Internet-Draft
Intended status: Standards Track
Expires: September 7, 2011

D. Rao
A. Banerjee
H. Grover
Cisco Systems
March 6, 2011

IS-IS Extensions to support OTV
draft-drao-isis-otv-00

Abstract

This document specifies the IS-IS extensions necessary to support OTV [OTV]. The data formats and code points used for the extensions are described. Details regarding the usage of these extensions are described in the OTV specification document.

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 September 7, 2011.

Copyright Notice

Table of Contents

1. Overview

OTV [OTV] uses Layer-2 addresses carried in a routing protocol to provide a MAC-in-IP solution for extending Layer-2 domains transparently across a L2/L3 core network. To achieve the specified functions of OTV, a control plane is required to exchange reachability information among the different OTV Edge Devices. [OTV] refers to this control plane as the oURP and oMRP (Overlay Unicast Routing Protocol and Overlay Multicast Routing Protocol).

As one specific instance, IS-IS [IS-IS] [RFC1195] is used as a means to auto-discover overlay VPN members as well as to exchange Layer-2 unicast and multicast reachability information across the overlay. Thus, IS-IS serves as both the oURP and oMRP. This document specifies a set of TLVs and sub-TLVs to be added to [IS-IS] PDUs, and new PDU types, to support this proposed solution. Some of these TLVs are generic Layer-2 IS-IS extensions being leveraged, and some are specific to OTV. This draft does not propose any new forwarding mechanisms using this additional information carried within IS-IS.

1.1. Terminology

The terminology and acronyms defined in [OTV] are used herein with the same meaning.

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.

2. TLV, sub-TLV and PDU Extensions to IS-IS for OTV

This section specifies the extensions for PDUs, TLVs and sub-TLVs as needed by OTV.

2.1. Group Address TLV

OTV uses the Group Address (GADDR) TLV that is defined in [isis-trill]. However, the GADDR TLV as used by OTV is carried within a Multicast Group link state PDU instead of within an LSP.

The GADDR TLV carries sub-TLVs that in turn advertise multicast group listeners. The new sub-TLVs for this TLV defined for use by OTV are specified in the following subsections.

2.1.1. Group IPv4 Address sub-TLV

The Group IPv4 Address (GIP-ADDR) sub-TLV is IS-IS sub-TLV type 2 within the GADDR TLV and has the following format:

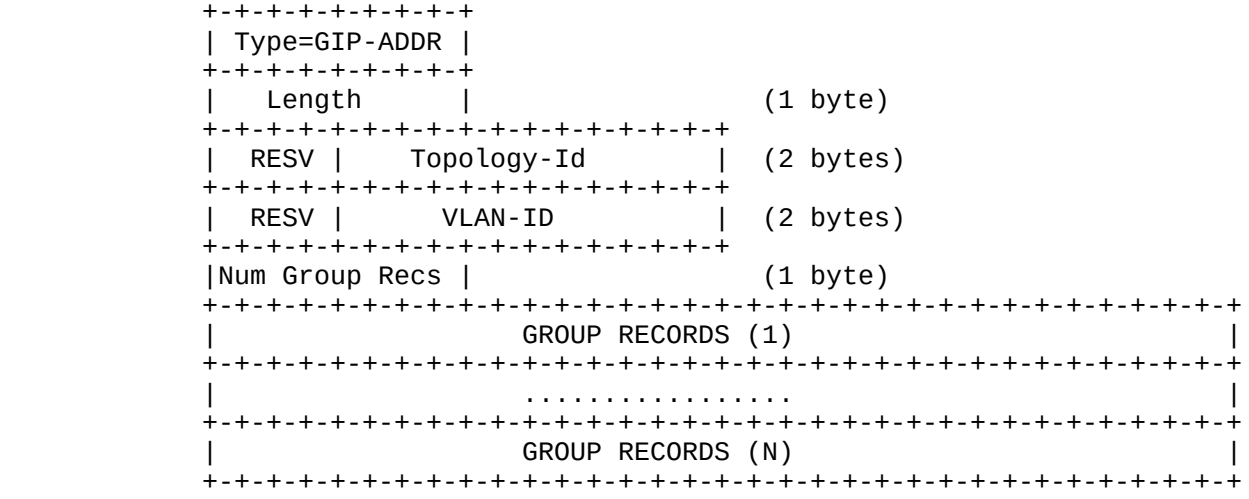

where each group record is of the form:

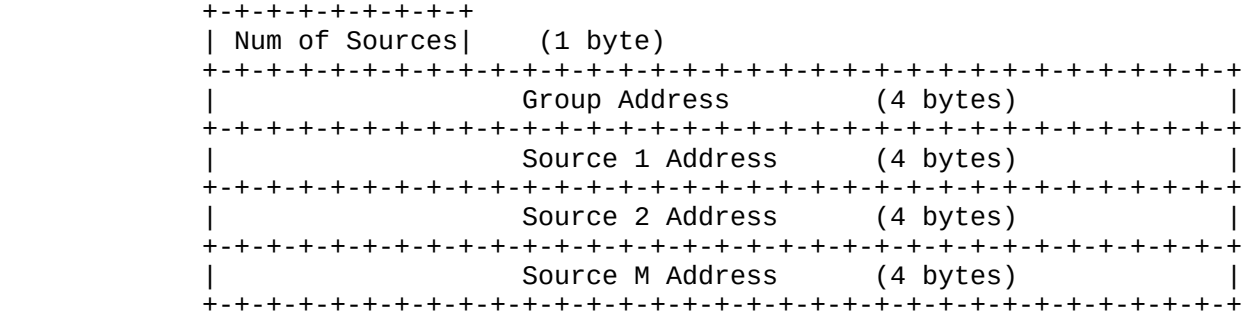

o Type: sub-TLV Type, set to 2 (GIP-ADDR).

o Length: Total number of bytes contained in the value field of the sub-TLV.

o Topology-Id: This carries the topology-id.

o RESV: Must be sent as zero on transmission and is ignored on receipt.

o VLAN-ID: This carries a 12-bit VLAN identifier that is valid for

all subsequent addresses in this sub-TLV, or the value zero if no VLAN is specified.

o Number of Group Records: This is of length 1 byte and lists the number of group records in this sub-TLV.

o Group Record: Each group record carries the number of sources. It then has a 32-bit IPv4 Group Address followed by 32-bit source IPv4 addresses. If the number of sources do not fit in a single sub-TLV, it is permitted to have the same group address repeated with different source addresses in another sub-TLV of another instance of the Group Address TLV.

The GIP-ADDR sub-TLV is carried only within a GADDR TLV and MUST be carried in a link state MGROUP PDU.

2.1.2. Group IPv6 Address sub-TLV

The Group IPv6 Address (GIPV6-ADDR) sub-TLV is IS-IS sub-TLV type 3 within the GADDR TLV and has the following format:

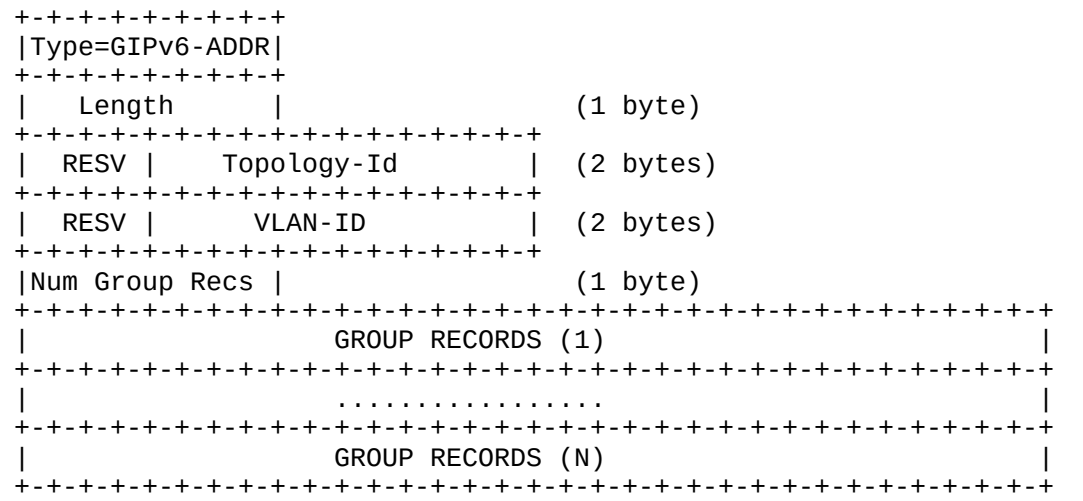

where each group record is of the form:

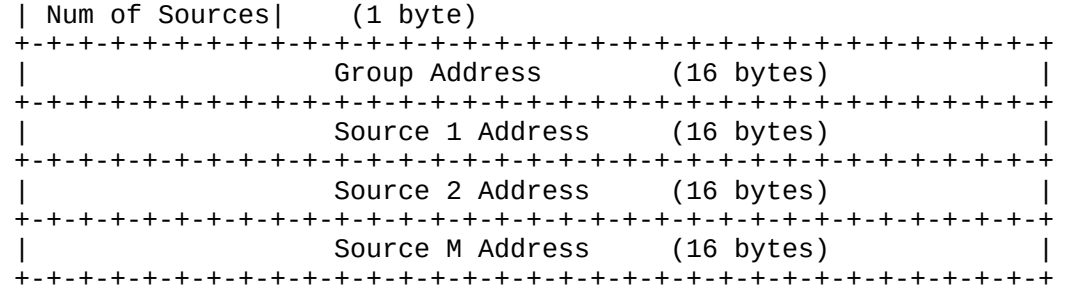

o Type: sub-TLV Type, set to 3 (GIPV6-ADDR).

o Length: Total number of bytes contained in the value field.

o Topology-Id: This carries the topology-id.

o RESV: Must be sent as zero on transmission and is ignored on receipt.

o VLAN-ID: This carries a 12-bit VLAN identifier that is valid for all subsequent addresses in this sub-TLV, or the value zero if no VLAN is specified.

o Number of Group Records: This of length 1 byte and lists the number of group records in this sub-TLV.

o Group Record: Each group record carries the number of sources. It then has a 128-bit IPv6 Group Address followed by 128-bit source IPv6 addresses. If the number of sources do not fit in a single sub-TLV, it is permitted to have the same group address repeated with different source addresses in another sub-TLV of another instance of the Group Address TLV.

The GIPV6-ADDR sub-TLV is carried only within a GADDR TLV and MUST be carried in a link state MGROUP PDU.

2.2. Multi-Topology aware Port Capability TLV

OTV uses the Multi-Topology aware Port Capability (MT-PORT-CAP) defined in [isis-layer2]. The sub-TLVs used by OTV are defined in the following sections.

2.2.1. Site Capability sub-TLV

The Site Capability sub-TLV (SITE-CAP) is type 250 within the MT-PORT-CAP TLV and carries information about or relevant to the site this edge device belongs to. This is used in OTV to aid in Authoritative Edge Device election. It has the following format:

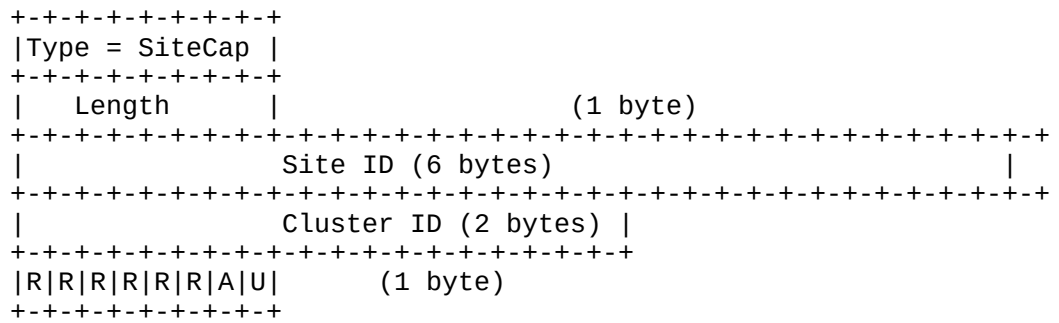

o Type: sub-TLV Type, set to Site Capability sub-TLV 250.

o Length: The size of the value.

o Site Id: The site-id of the site.

o Cluster Id: The cluster-id within the site.

o R: Must be sent as zero on transmission and is ignored on receipt.

o U bit: Denotes if the site is a unicast only site.

o A bit: Denotes if the edge device is capable of being an OTV

Authoritative Edge Device.

The Site Capability sub-TLV is carried only within the MT-PORT-CAP TLV and this is carried in an IIH PDU. There MUST be only one occurrence of this sub-TLV in an IIH PDU.

2.2.2. Site Group IPv4 sub-TLV

The Site Group IPv4 sub-TLV (SITE-GRP-IPV4) is type 251 within the MT-PORT-CAP TLV and carries information about the overlays active on this device. This is used in OTV to aid in Authoritative Edge Device election. It has the following format:

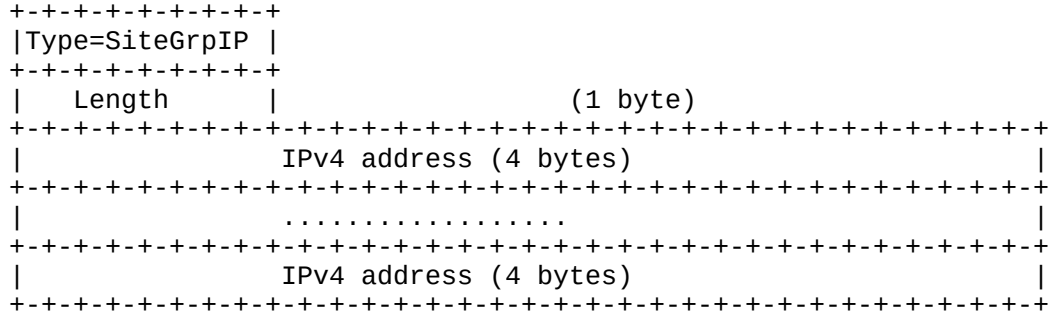

o Type: sub-TLV Type, set to Site Group IPv4 sub-TLV 251.

o Length: The size of the value.

o Value: The list of IPv4 addresses used by the site.

The Site Group IPv4 sub-TLV is carried within an IIH PDU. There may be more than one occurrence of this sub-TLV in an IIH PDU.

2.2.3. Site Group IPv6 sub-TLV

The Site Group IPv6 sub-TLV (SITE-GRP-IPV6) is type 252 and carries information about the overlays active on this device. This is used in OTV to aid in Authoritative Edge Device election. It has the following format:

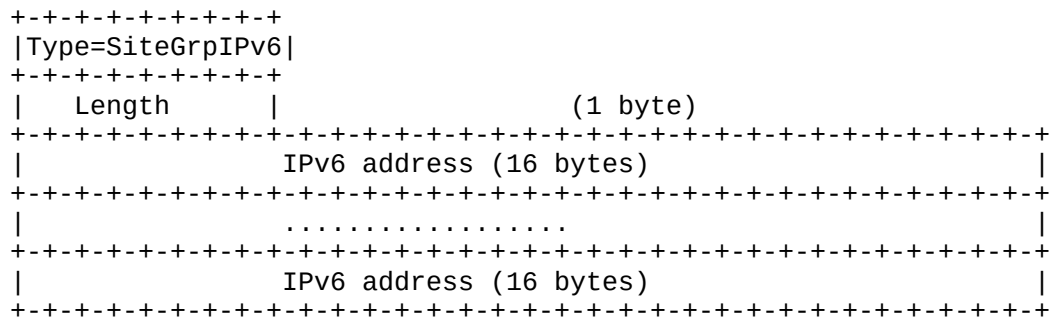

o Type: sub-TLV Type, set to Site Group IPv6 sub-TLV 252.

o Length: The size of the value.

o Value: The list of IPv6 addresses used by the site.

The Site Group IPv6 sub-TLV is carried within an IIH PDU. There may be more than one occurrence of this sub-TLV in an IIH PDU.

2.2.4. Adjacency Server IPv4 sub-TLV

The Adjacency Server IPv4 sub-TLV (ADJ-SVR-IPV4) is type 253 within the MT-PORT-CAP TLV and carries information about the capability of the sites in OTV. It has the following format:

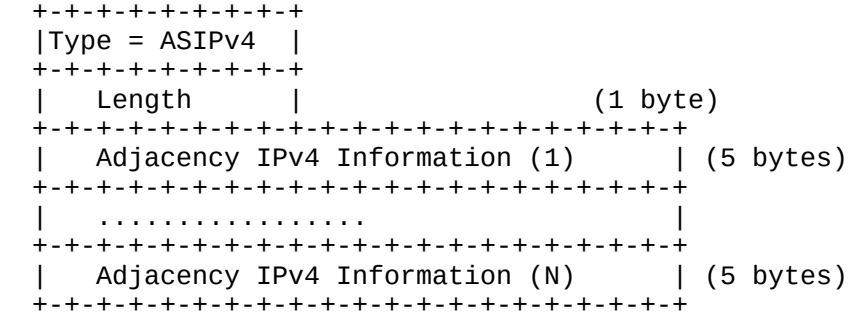

where each adjacency IPv4 information is of the form:

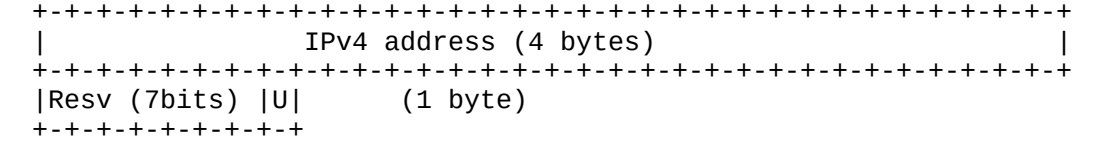

o Type: sub-TLV Type, set to Adjacency Server IP sub-TLV 253.

o Length: The size of the value, 5*n, where there are n adjacency server information blocks.

o IPv4 Address: The IPv4 addresses used by the sites.

o Reserved: Must be sent as zero on transmission and is ignored on receipt.

o U bit: Denotes if the site is a unicast only site.

The Adjacency Server IPv4 sub-TLV is carried within an IIH PDU. There may be more than one occurrence of this sub-TLV in an IIH PDU.

2.2.5. Adjacency Server IPv6 sub-TLV

The Adjacency Server IPv6 sub-TLV (ADJ-SVR-IPV6) is type 254 within the MT-PORT-CAP TLV and carries information about the capability of the sites in OTV. It has the following format:

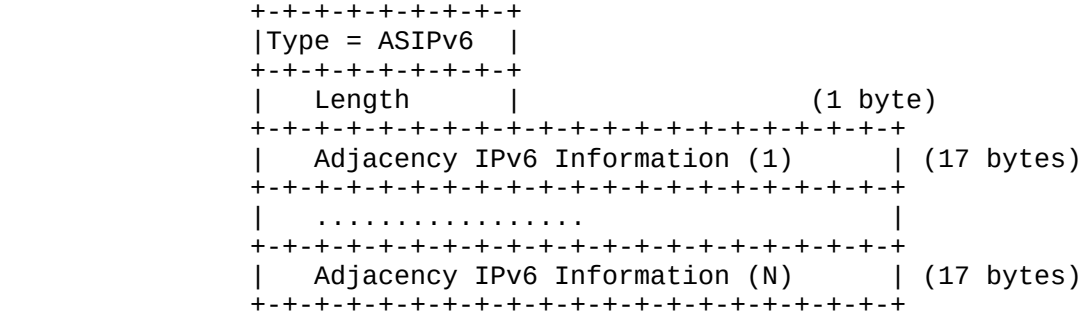

where each adjacency IPv6 information is of the form:

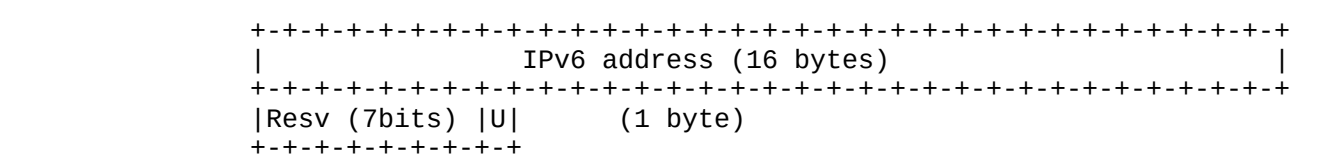

o Type: sub-TLV Type, set to Adjacency Server IPv6 sub-TLV 254.

o Length: The size of the value.

o Value: The IPv6 addresses used by the sites.

o Reserved: Must be sent as zero on transmission and is ignored on receipt.

o U bit: Denotes if the site is a unicast only site.

The Adjacency Server IPv6 sub-TLV is carried within an IIH PDU. There may be more than one occurrence of this sub-TLV in an IIH PDU.

2.3. Group Membership Active Source TLV

The Group Active Source (GMAS) TLV is IS-IS TLV type 146 and has the u following format:

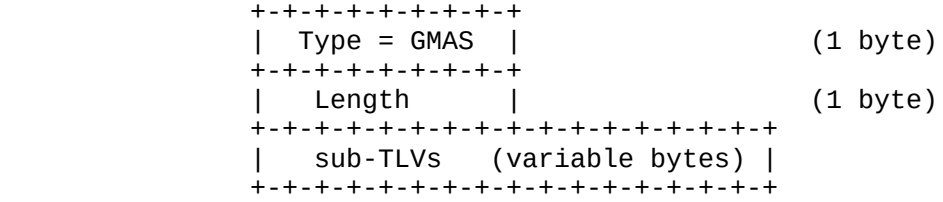

o Type: TLV Type, set to GMAS-TLV 146.

o Length: Total number of bytes contained in the value field, which includes the length of the sub-TLVs carried in this TLV.

o sub-TLVs: The Group Active Source TLV value contains sub-TLVs formatted as described in [RFC5305]. The sub-TLVs for this TLV are specified in the following subsections.

The GMAS TLV MUST be carried within a Multicast Group link state PDU.

2.3.1. The Group MAC Active Source sub-TLV

The Group MAC Source (GMAS-MAC) sub-TLV is IS-IS sub-TLV type 4 within the GMAS TLV. It is used in OTV to create multicast distribution trees and has the following format:

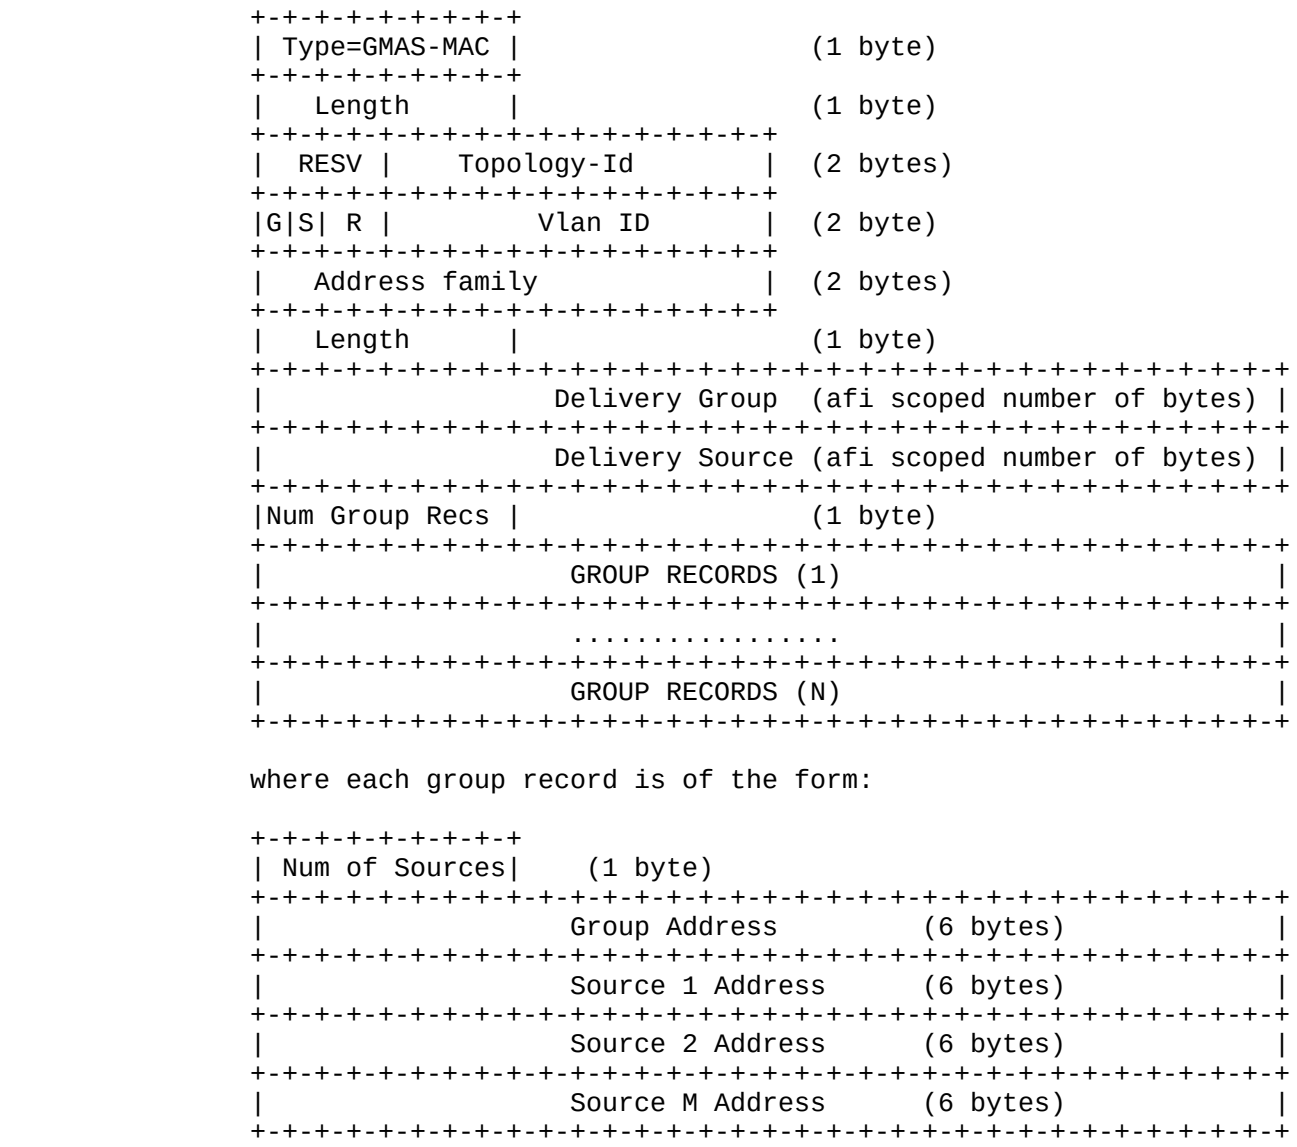

where each group record is of the form:

o Type: sub-TLV Type, set to 4 (GMAS-MAC) of length 1 byte.

o Length: Total number of bytes contained in the value field.

o Topology-Id: This carries the topology-id.

o RESV: Must be sent as zero on transmission and is ignored on receipt.

o G (1 bit): Delivery Group is set

o S (1 bit): Delivery Source is set

o R (2 bits) : Must be sent as zero on transmission and is ignored on receipt.

o VLAN-ID: This carries a 12-bit VLAN identifier that is valid for all subsequent MAC addresses in this sub-TLV, or the value zero if no VLAN is specified.

o Address Family: Describes the Address family of the Delivery Source/Group information. It is set to 1 for IPv4 and 2 for IPv6.

o Length: Gives the length of the Delivery Source and Delivery Group field.

o Delivery Group: Describes the group used to deliver packets.

o Delivery Source: Describes the source address used to deliver packets.

o Number of Group Records: This is of length 1 byte and lists the number of group records in this sub-TLV.

o Group Record: Each group record has a one byte which carries the number of sources. It then has a 48-bit multicast Group Address followed by 48-bit source MAC addresses. An address being a group multicast address or unicast source address can be checked using the multicast bit in the address. If the number of sources do not fit in a single sub-TLV, it is permitted to have the same group address repeated with different source addresses in another sub-TLV of another instance of the Group Active Source TLV.

The GMAS-MAC sub-TLV is carried within the GMAS TLV and MUST be carried in a link state MGROUP PDU.

2.3.2. Group IPv4 Active Source sub-TLV

The Group IPv4 Address (GMAS-IP) sub-TLV is IS-IS sub-TLV type 5 within the GMAS TLV. It is used in OTV to create multicast distribution trees and has the following format:

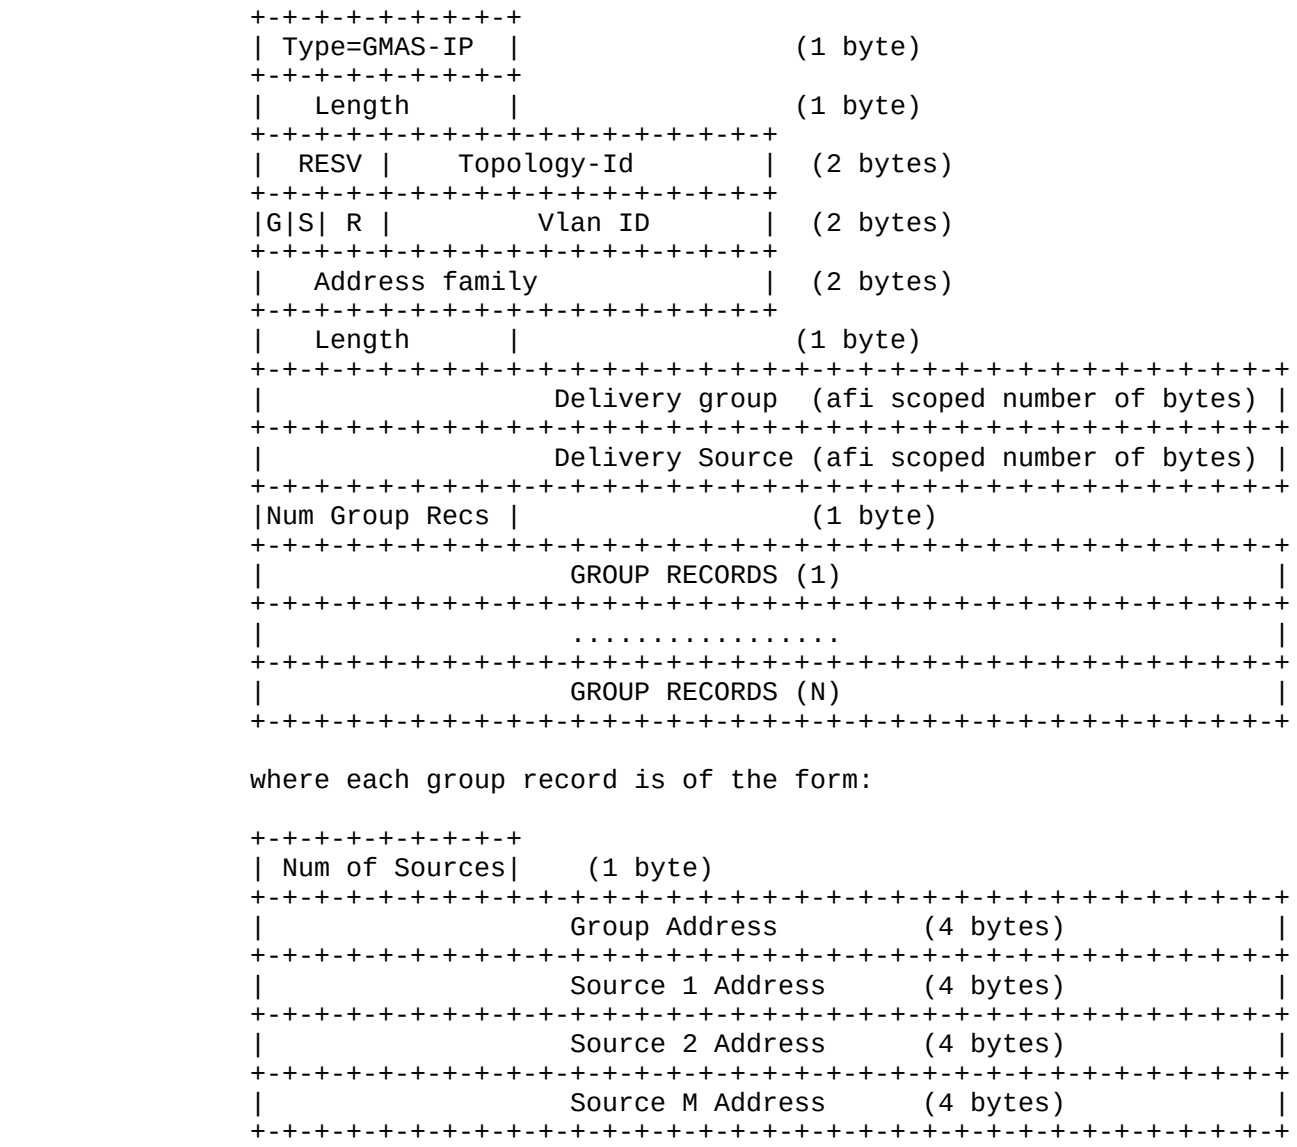

where each group record is of the form:

o Type: sub-TLV Type, set to 5 (GMAS-IP).

o Length: Total number of bytes contained in the value field of the sub-TLV.

o Topology-Id: This carries the topology-id.

o RESV: Must be sent as zero on transmission and is ignored on

receipt.

o G (1 bit): Delivery Group is set

o S (1 bit): Delivery Source is set

o R (2 bits) : Must be sent as zero on transmission and is ignored on receipt.

o VLAN-ID: This carries a 12-bit VLAN identifier that is valid for all subsequent MAC addresses in this sub-TLV, or the value zero if no VLAN is specified.

o Address Family: Describes the Address family of the Delivery Source/Group information.

o Length: Gives the length of the Delivery Source and Delivery Group field.

o Delivery Group: Describes the group used to deliver packets.

o Delivery Source: Describes the source address used to deliver packets.

o Number of Group Records: This is of length 1 byte and lists the number of group records in this sub-TLV.

o Group Record: Each group record has a one byte which carries the number of sources. It is followed by a 32-bit IPv4 Group Address followed by 32-bit source IPv4 addresses. If the number of sources do not fit in a single sub-TLV, it is permitted to have the same group address repeated with different source addresses repeated in another sub-TLV of another instance of the Group Active Source TLV.

The GMAS-IP TLV is carried within the GMAS TLV and MUST be carried in a link state MGROUP PDU.

2.3.3. Group IPv6 Active Source sub-TLV

The Group IPv6 Active Source (GMAS-IPV6) sub-TLV is IS-IS sub-TLV type 6. It is used in OTV to create multicast distribution trees and has the following format:

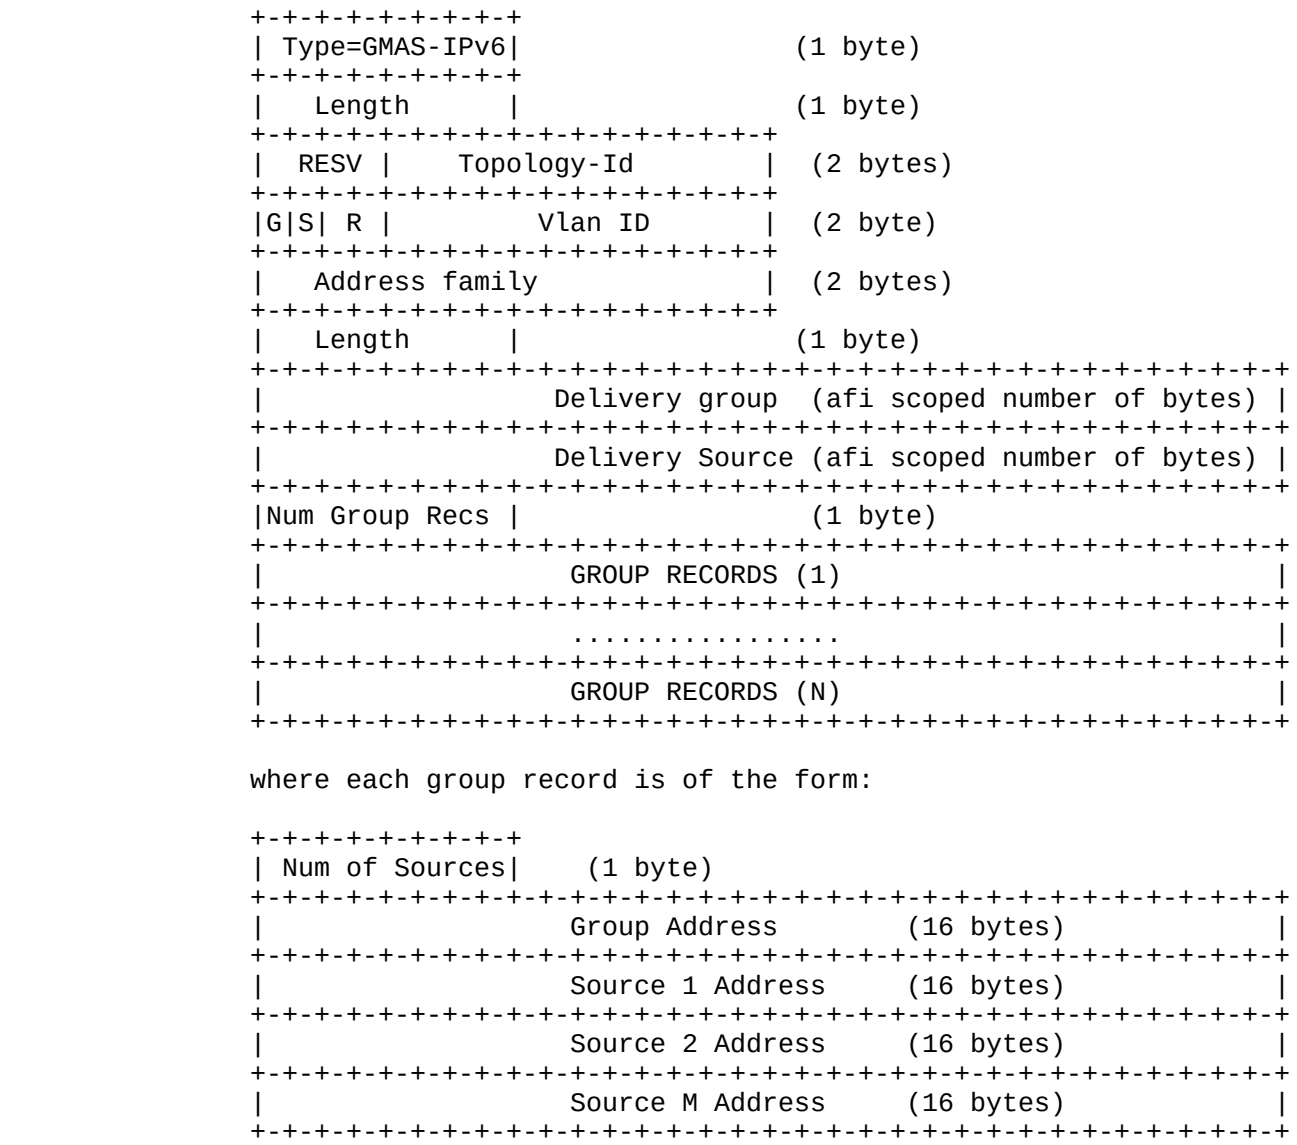

where each group record is of the form:

o Type: sub-TLV Type, set to 6 (GMAS-IPV6).

o Length: Total number of bytes contained in the value field.

o Topology-Id: This carries the topology-id.

o RESV: Must be sent as zero on transmission and is ignored on receipt.

o G (1 bit): Delivery Group is set

o S (1 bit): Delivery Source is set

o R (2 bits) : Must be sent as zero on transmission and is ignored on receipt.

o VLAN-ID: This carries a 12-bit VLAN identifier that is valid for all subsequent MAC addresses in this sub-TLV, or the value zero if no VLAN is specified.

o Address Family: Describes the Address family of the Delivery Source/Group information.

o Length: Gives the length of the Delivery Source and Delivery Group field.

o Delivery Group: Describes the group used to deliver packets.

o Delivery Source: Describes the source address used to deliver packets.

o Number of Group Records: This of length 1 byte and lists the number of group records in this sub-TLV.

o Group Record: Each group record has one byte which carries the number of sources. It is followed by a 128-bit multicast IPv6 Group Address followed by 128-bit source IPv6 addresses. If the number of sources do not fit in a single sub-TLV, it is permitted to have the same group address repeated with different source addresses repeated in another sub-TLV in another instance of the Group Address TLV.

The GMAS-IPV6 sub-TLV is carried within the GMAS TLV and MUST be carried in a link state MGROUP PDU.

2.4. PDU Extensions to IS-IS

2.4.1. Multicast Group PDU

This section specifies three new IS-IS PDUs, the Multicast Group (MGROUP) PDU, for carrying a list of attached or joined multicast groups. The Multicast Group Complete Sequence Number (MGROUP-CSNP) PDU and the Multicast Group Partial Sequence Number (MGROUP-PSNP) PDU packets are also defined to be used with the new MGROUP-PDU to perform database exchange on the MGROUP PDU packets. Only Level-1 PDUs are defined.

The Multicast Group (MGROUP) PDU can be used to advertise a set of

attached, or joined, multicast groups. The MGROUP PDU is formatted identical to a Level-1 Link State PDU, as described in Section 9.3 of [IS-IS]. One field, PDU Type, is changed to 19, to signify this PDU is carrying multicast group information, rather than unicast reachability information.

The Multicast Group PDU carries TLVs indicating multicast membership information. There are three sub-TLVs of the GADDR TLV defined in this document, that MAY be present in this PDU, namely, GMAC-ADDR, GIP-ADDR, and GIPV6-ADDR sub-TLVs. Furthermore, it MAY carry the Authentication TLV (TLV 10) and the Interested VLANs sub-TLV of the Capability TLV.

One or more TLVs MAY be carried in a single MGROUP PDU. Future multicast address TLVs MAY be defined using other type codes, and be carried in an MGROUP PDU.

2.4.2. Multicast Group Partial Sequence Number PDU

The Multicast Group Partial Sequence Number (MGROUP-PSNP) PDU is type 29. The MGROUP-PSNP performs a function analogous to the PSNP but applies to MGROUP-PDUs.

2.4.3. Multicast Group Complete Sequence Number PDU

The Multicast Group Complete Sequence Number PDU (MGROUP-CSNP) PDU is type 22. The MGROUP-CSNP performs a function analogous to the CSNP but applies to MGROUP-PDUs.

2.4.4. MGROUP PDU related changes to Base protocol

In this section, we describe the changes to the base protocol due to the introduction of the MGROUP, MGROUP-PSNP, MGROUP-CNSP PDUs.

2.4.4.1. Enhancements to the flooding process

OTV introduces a second instance of the Update Process which applies to MGROUP-PDUs. Operation of the MGROUP update process is identical to that defined in [IS-IS] but MGROUP-PDUs, MGROUP-PSNPs, and MGROUP-CSNPs are used in place of LSPs, PSNPs, and CSNPs respectively.

For example, on P2P links CSNP is exchanged on the formation of an adjacency. In a similar fashion a MGROUP-CSNP MUST also be exchanged between the neighbors at the same time. This gets the initial MGROUP-database synchronization going. After this similar actions of the base protocol specifications for the regular database synchronization will be maintained to keep the MGROUP-database synchronized. There need not be any more correlation between the

updates of the LSP and the MGROUP-PDU.

Similarly, on LAN links the DIS is responsible for sending periodic CSNP transmissions. The DIS in this case will also be responsible for sending periodic MGROUP-CSNP transmissions. The update and flooding process will work in parallel for the two databases and there is no further synchronization between them.

In general, the database synchronization is performed in parallel with no interactions between the messages. However, the initial triggers that start a CSNP exchange are correlated, in the sense it also triggers a MGROUP-CSNP exchange.

2.4.4.2. Enhancements to Graceful Restart

During graceful restart, the normal hello operations as described in RFC 5306 will be followed. The enhancements will take place such that CSNP and PSNP triggers will necessitate a parallel MGROUP-CSNP and MGROUP-PSNP exchange and update process will be triggered in parallel for the MGROUP-PDUs. After the databases containing information from both LSPs and MGROUP-PDUs have been obtained, the restart process is deemed complete.

2.4.4.3. Enhancements to the maximum sequence number reached

In the event, LSPs reach the maximum sequence number, ISO/IEC 10589 states the rules for the process to shut down and its duration. With the introduction of the MGROUP-PDU, the same process now applies when LSPs from either database reach the maximum sequence number.

2.4.4.4. Enhancements to SPF

The MGROUP-PDU advertises a set of attached, or joined, multicast groups. These groups act as leaves of the advertising nodes. As a result, there are no new requirements of running a SPF if only information within the MGROUP-PDU changes.

3. Acknowledgements

The authors would like to thank Dino Farinacci and Les Ginsberg for their input and useful comments on various aspects of the extensions.

4. Security Considerations

This document adds no additional security risks to IS-IS, nor does it provide any additional security for IS-IS.

5. IANA Considerations

This document specifies a set of new IS-IS TLVs and PDU types, which need to be reflected in the IS-IS TLV code-point registry. IANA is requested to allocate the necessary registry code points listed below.

5.1. Codepoints

TLV Codepoints

Description	Type	IIH	LSP	SNP
GMAS	146	-	X	-

Sub-TLV Codepoints

MT Port Capability TLV

Description	Sub-TLV#
SITE-CAP	250
SITE-GRP-IPV4	251
SITE-GRP-IPV6	252
ADJ-SVR-IPV4	253
ADJ-SVR-IPV6	254

Group Address TLV

Description	Sub-TLV#
GIP-ADDR	2
GIPV6-ADDR	3

IS-IS PDU Codepoints

IANA is requested to allocate three new IS-IS PDUs from the IS-IS PDUs registry, namely the MGROUP PDU, the MGROUP-CSNP PDU and the MGROUP-PSNP PDU [suggested PDU values below].

IS-IS PDUs Registry:

Mnemonic	PDU#	Reference
MGROUP PDU	19	This document
MGROUP-CSNP PDU	22	This document
MGROUP-PSNP PDU	29	This document

5.2. New Sub-Registry

IANA is requested to create a new sub-TLV IS-IS sub-registry for sub-TLVs within the Group Membership Active Source (GMAS) TLV. The codepoints are requested to be allocated as listed below.

Registry Name: IS-IS Group Membership Active Source Type Codes

Reference: This document

Registration Procedures: Expert Review [RFC5226]

Registry:

Value	GMAS Type Code	Reference
0-3	Reserved	This document
4	GMAS-MAC	This document
5	GMAS-IP	This document
6	GMAS-IPV6	This document
4-255	Unassigned	This document

6. Normative References

[IS-IS] ISO/IEC 10589, "Intermediate System to Intermediate System Intra-Domain Routing Exchange Protocol for use in Conjunction with the Protocol for Providing the Connectionless-mode Network Service (ISO 8473)", 2005.

[OTV] Grover, H., "OTV: Overlay Transport Virtualization, draft-hasmit-otv-01.txt, work in progress", 2010.

[isis-layer2] Banerjee, A. and D. Ward, "Extensions to IS-IS for Layer-2 Systems, draft-ietf-isis-layer2-11.txt, work in progress", 2011.

[isis-trill] Eastlake, D., "TRILL Use of IS-IS, draft-ietf-isis-trill-05.txt, work in progress", 2011.

Authors' Addresses

Dhananjaya Rao
Cisco Systems
170 W Tasman Drive
San Jose, CA 95138
US

Email: dhrao@cisco.com

Ayan Banerjee
Cisco Systems
170 W Tasman Drive
San Jose, CA 95138
US

Email: ayabaner@cisco.com

Hasmit Grover
Cisco Systems
170 W Tasman Drive
San Jose, CA 95138
US

Email: hasmit@cisco.com

TRILL Working Group
INTERNET-DRAFT
Intended status: Proposed Standard

Donald Eastlake
Huawei
Vishwas Manral
HP Networking
Li Yizhou
Sam Aldrin
Huawei
Dave Ward
Cisco

Expires: January 12, 2013
July 13, 2012

TRILL: RBridge Channel Support
<draft-ietf-trill-rbridge-channel-08.txt>

Abstract

This document specifies a general channel mechanism for sending messages, such as BFD (Bidirectional Forwarding Detection) messages, between RBridges (Routing Bridges) and between RBridges and end stations in an RBridge campus through extensions to the TRILL (TRansparent Interconnection of Lots of Links) protocol.

Status of This Memo

This Internet-Draft is submitted to IETF in full conformance with the provisions of BCP 78 and BCP 79.

Distribution of this document is unlimited. Comments should be sent to the TRILL working group mailing list.

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/1id-abstracts.html. The list of Internet-Draft Shadow Directories can be accessed at http://www.ietf.org/shadow.html.

Table of Contents

1. Introduction

RBridge campuses provide transparent least-cost forwarding using the TRILL (TRansparent Interconnection of Lots of Links) protocol that builds on IS-IS (Intermediate System to Intermediate System) routing [IS-IS] [RFC1195] [RFC6326bis]. Devices that implement TRILL are called RBridges (Routing Bridges) or TRILL Switches. However, the TRILL base protocol standard [RFC6325] provides only for TRILL Data messages and TRILL IS-IS messages.

This document specifies a general channel mechanism for the transmission of other messages within an RBridge campus, such as BFD (Bidirectional Forwarding Detection, [RFC5880]) messages, (1) between RBridges and end stations that are directly connected on the same link and (2) between RBridges. This mechanism supports a requirement to be able to operate with minimal configuration.

1.1 RBridge Channel Requirements

It is anticipated that various protocols operating at the TRILL layer will be desired in RBridge campuses. For example, there is a need for rapid response continuity checking with a protocol such as BFD [RFC5880] [RFC5882] and for a variety of optional reporting.

To avoid the requirement to design and specify a way to carry each such protocol, this document specifies a general channel for sending messages between RBridges in a campus at the TRILL level by extending the TRILL protocol. To accommodate a wide variety of protocols, this RBridge Channel facility accommodates all the regular modes of TRILL Data transmission including single and multiple hop unicast as well as VLAN scoped multi-destination distribution.

To minimize any unnecessary burden on transit RBridges and to provide a more realistic test of network continuity and the like, RBridge Channel messages are designed to look like TRILL Data frames and, in the case of multi-hop messages, can normally be handled by transit RBridges as if they were TRILL Data frames; however, to enable processing at transit RBridges when required by particular messages, they may optionally use the RBridge Channel Alert TRILL extended header flags [RFCext] that causes a transit RBridge implementing the flag to more closely examine a flagged frame.

This document also specifies a format for sending RBridge Channel messages between RBridges and end stations that are directly connected over a link, in either direction, when provided for by the protocol involved. For the most part, this format is the same as the format that is TRILL Data encapsulated for inter-RBridge channel messages.

Each particular protocol using the RBridge Channel facility will likely use only a subset of the facilities specified herein.

1.2 Relation to the MPLS Generic Channel

The RBridge Channel is similar to the MPLS Generic Channel specified in [RFC5586]. Instead of using a special MPLS label to indicate a special channel message, an RBridge Channel message is indicated by a special multicast Inner.MacDA and inner Ethertype (see Section 2.1).

1.3 Terminology

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].

The terminology and acronyms of [RFC6325] are used in this document with the additions listed below.

BFD - Bidirectional Forwarding Detection

CHV - Channel Header Version

MH - Multi-Hop

NA - Native

SL - Silent

2. Inter-RBridge Channel Messages

Channel messages between RBridges are transmitted as TRILL Data frames. (For information on channel messages that can be transmitted between RBridges and end stations that are directly connected by a link, see Section 4.) Inter-RBridge Channel messages are identified as such by their Inner.MacDA, which is the All-Egress-RBridges multicast address, together with their Inner Ethertype, which is the RBridge-Channel Ethertype. This Ethertype is part of and starts the RBridge Channel Header.

The diagram below shows the overall structure of a RBridge Channel Message frame on a link between two RBridges:

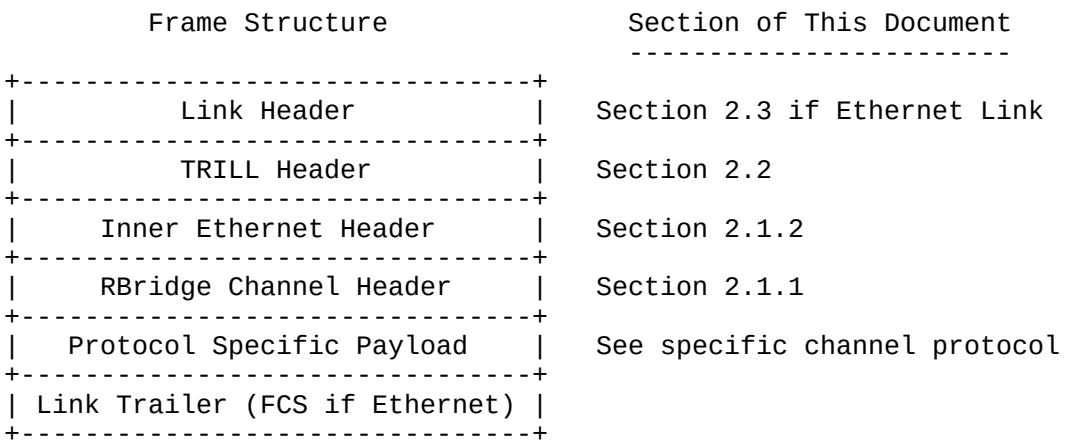

Figure 1. RBridge Channel Frame Structure

Optionally, some channel messages may require examination of the frame by transit RBridges that support the RBridge Channel feature, to determine if they need to take any action. To indicate this, such messages use a RBridge Channel Alert extended TRILL header flag as further described in Section 3 below.

The Sections 2.1 and 2.2 below describe the Inner frame and the TRILL Header for frames sent in an RBridge Channel. As always, the Outer link header and trailer are whatever is needed to get a TRILL Data frame to the next hop RBridge, depending on the technology of the link, and can change with each hop for multi-hop messages. Section 2.3 describes the outer Link Header for Ethernet links. And Section 2.4 discusses some special considerations for the first hop transmission of RBridge Channel messages.

Section 3 describes some details of RBridge Channel message processing. Section 4 provides the specifications for native RBridge Channel frames between RBridges and end stations that are directly

connected over a link. Section 5 describes how support for RBridge Channel protocols is indicated. And Sections 6, 7, and 8 give congestion, allocation (IANA and IEEE), and security considerations respectively.

2.1 The RBridge Channel Message Inner Frame

The encapsulated inner frame within an RBridge Channel message frame is as shown below.

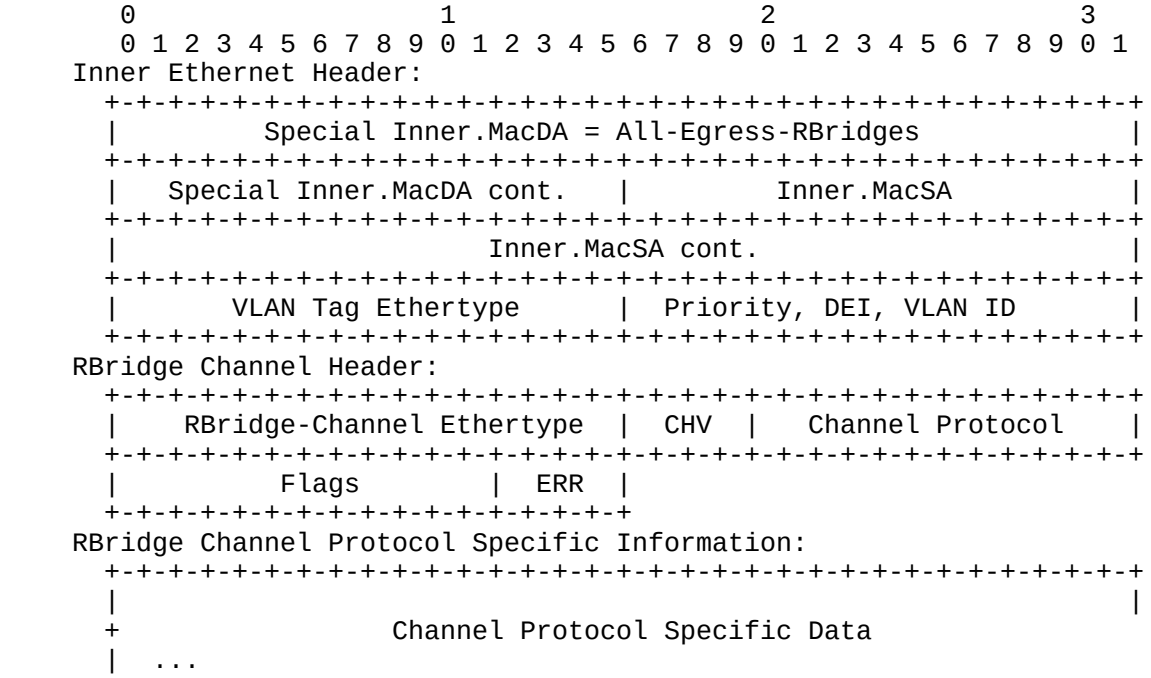

Figure 2. RBridge Channel Inner Frame Header Fields

The Channel Protocol Specific Data contains the information related to the specific channel protocol used in the channel message. Details of that data are outside the scope of this document, except in the case of the RBridge Channel Error protocol specified below.

2.1.1 RBridge Channel Header

As shown in Figure 2, the RBridge Channel header starts with the RBridge-Channel Ethertype (see Section 7.2). Following that is a four-byte quantity with four sub-fields as follows:

CHV: A 4-bit field that gives the RBridge Channel Header Version. This document specifies version zero.

Channel Protocol: A 12-bit unsigned integer that specifies the particular RBridge Channel protocol to which the message applies.

Flags: Provides 12 bits of flags described below.

ERR: A 4-bit unsigned integer used in connection with error reporting at the RBridge Channel level as described in Section 3.

The flag bits are numbered from 0 to 11 as shown below.

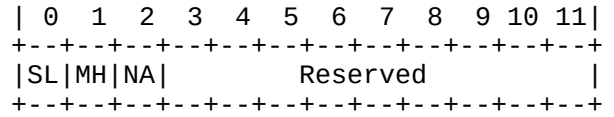

Figure 3. Channel Header Flag Bits

Bit 0, which is the high order bit in network order, is defined as the SL or Silent bit. If it is a one, it suppresses RBridge Channel Error messages (see Section 3).

Bit 1 is the MH or Multi-Hop bit. It is used to inform the destination RBridge protocol that the message may be multi-hop (MH=1) or was intended to be one-hop only (MH=0).

Bit 2 is the NA or Native bit. It is used as described in Section 4 below.

Reserved: Bits reserved for future specification that MUST be sent as zero and ignored on receipt.

The RBridge Channel Protocol field specifies the protocol that the channel message relates to. The initial defined value is listed below.

Protocol	Name - Section of this Document
0x001	RBridge Channel Error - Section 3

IANA Considerations for RBridge Channel protocol numbers are provided in Section 7. These include provisions for Private Use protocol numbers. Because different uses of Private Use RBridge Channel protocol numbers may conflict, such use MUST be within a private network. It is the responsibility of the private network manager to

avoid conflicting use of these code points and unacceptable burdens within the private network from their use.

2.1.1 Inner Ethernet Header

The special Inner.MacDA is the All-Egress-RBridges multicast MAC address to signal that the frame is intended for the egress (decapsulating) RBridge itself (or the egress RBridges themselves if the frame is multi-destination). (This address is called the All-ESADI-RBridges address in [RFC6325].) The RBridge-Channel Ethertype indicates that the frame is an RBridge Channel message. The only other Ethertype currently specified for use with the All-Egress-RBridges Inner.MacDA is L2-IS-IS to indicate an ESADI frame [RFC6325]. In the future additional Ethertypes may be specified for use with the All-Egress-RBridges multicast address.

The RBridge originating the channel message selects the Inner.MacSA. The Inner.MacSA MUST be set by the originating RBridge to a MAC address unique within the campus owned by the originating RBridge. This MAC address can be considered, in effect, the MAC address of a virtual internal end station that handles the RBridge Channel frames originated by or destined for that RBridge. It MAY be the same as the Inner.MacSA used by the RBridge when it originates ESADI frames [RFC6325].

2.1.3 Inner.VLAN Tag

As with all frames formatted to be processed as a TRILL Data frame, an Inner.VLAN tag is present. Use of a VLAN tag Ethertype other than 0x8100 or stacked tags is beyond the scope of this document but is an obvious extension.

Multi-destination RBridge Channel messages are, like all multi-destination TRILL Data messages, VLAN scoped so the Inner.VLAN ID MUST be set to the VLAN of interest. To the extent that distribution tree pruning is in effect in the campus, such channel messages may only reach RBridges advertising that they have connectivity to that VLAN.

For channel messages sent as known unicast TRILL Data frames the default value for the Inner.VLAN ID is VLAN 1 but particular RBridge Channel protocols MAY specify other values.

The Inner.VLAN also specifies a three-bit frame priority for which the following recommendations apply:

1. For one-hop channel messages critical to network connectivity, such as one-hop BFD for rapid link failure detection in support of TRILL IS-IS, the RECOMMENDED priority is 7.

2. For single and multi-hop unicast channel messages important to network operation but not critical for connectivity, the RECOMMENDED priority is 6.

3. For other unicast channel messages and all multi-destination channel messages, it is RECOMMENDED that the default priority zero be used. In any case, priorities higher than 5 SHOULD NOT be used for such frames.

There is one additional bit in a VLAN tag value between the 12-bit VLAN ID and 3-bit priority, the Drop Eligibility Indicator (DEI, [ClearCorrect]). It is RECOMMENDED that this bit be zero for the first two categories of channel messages listed immediately above. The setting of this bit for channel messages in the third category may be dependent on the channel protocol and no general recommendation is made for that case.

2.2 The TRILL Header for RBridge Channel Messages

After the outer Link Header (that, for an Ethernet link, ends with the TRILL Ethertype) and before the encapsulated frame, the channel message's TRILL Header initially appears as follows:

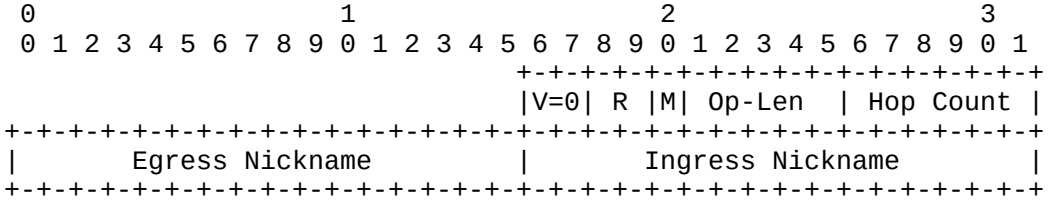

Figure 4. RBridge Channel TRILL Header Fields

The TRILL Header version V MUST be zero, the R bits are reserved, the M bit is set appropriately as the channel message is to be forwarded as known destination unicast (M=0) or multi-destination (M=1) regardless of the fact that the Inner.MacDA is always the All-Egress-RBridges multicast address, and Op-Len is set appropriately for the length of the TRILL Header extensions area, if any, all as specified in [RFC6325].

When an RBridge Channel message is originated, the Hop Count field defaults to the maximum value, 0x3F, but particular RBridge Channel protocols MAY specify other values. For messages sent a known number

of hops, such as one-hop messages or a two-hop self-addressed message intended to loop back through an immediate neighbor RBridge, setting the Hops field to the maximum value and checking the Hop Count field on receipt provides an additional validity check as discussed in [RFC5082].

The RBridge originating a channel message places a nickname that it holds into the ingress nickname field.

There are several cases for the egress nickname field. If the channel message is multi-destination, then the egress nickname designates the distribution tree to use. If the channel message is a multi-hop unicast message, then the egress nickname is a nickname of the target RBridge; this includes the special case of a message intended to loop back from an immediate neighbor where the originator places one of its own nicknames in both the ingress and egress nickname fields. If the channel message is a one-hop unicast message, there are two possibilities for the egress nickname.

- The egress nickname can be set to a nickname of the target neighbor RBridge.
- The special nickname Any-RBridge may be used. RBridges supporting the RBridge Channel facility MUST recognize the Any-RBridge special nickname and accept TRILL Data frames having that value in the egress nickname field as being sent to them as the egress. Thus, for such RBridges, using this egress nickname guarantees processing by an immediate neighbor regardless of the state of nicknames.

2.3 Ethernet Link Header and Trailer

An RBridge Channel frame has the usual link header and trailer for a TRILL Data frame depending on the type of link on which it is sent.

For an Ethernet link [RFC6325] the Outer.MacSA is the MAC address of the port from which the frame is sent. The Outer.MacDA is the MAC address of the next-hop RBridge port for unicast RBridge Channel messages or the All-RBridges multicast address for multi-destination RBridge Channel messages. The Outer.VLAN tag specifies the Designated VLAN for that hop and the priority should be the same as in the Inner.VLAN tag; however, the output port may have been configured to strip VLAN tags, in which case no Outer.VLAN tag appears on the wire. And the link trailer is the Ethernet FCS.

2.4 Special Transmission and Rate Considerations

If a multi-hop RBridge Channel message is received by an RBridge, the criteria and method of forwarding it are the same as for any TRILL Data frame. If it is so forwarded, it will be on a link that was included in the routing topology because it was in the Report state as specified in [RFC6327].

However, special considerations apply to single hop messages because, for some RBridge Channel protocols, it may be desirable to send RBridge Channel messages over a link that is not yet fully up. In particular, it is permissible, if specified by the particular channel protocol, for the source RBridge that has created an RBridge Channel message to attempt to transmit it to a next hop RBridge when the link is in the Detect or Two-Way states, as specified in [RFC6327], as well as when it is in the Report state. Such messages can also be sent on point-to-point links that are not in the Up state.

RBridge Channel messages represent a burden on the RBridges and links in a campus and should be rate limited, especially if they are sent as high priority, multi-destination, or multi-hop frames or have an RBridge Channel Alert extended header flag set. See Section 6, Congestion Considerations.

3. Processing RBridge Channel TRILL Data Messages

RBridge Channel TRILL Data messages are designed to look like and, to the extent practical, be forwarded as regular TRILL Data frames. On receiving a channel message, an RBridge performs the usual initial tests on the frame and makes the same forwarding and/or decapsulation decisions as for a regular TRILL Data frame [RFC6325] with following exceptions for RBridges implementing the RBridge Channel facility:

1. An RBridge implementing the RBridge Channel facility MUST recognize the Any-RBridge egress nickname in TRILL Data frames, decapsulating such frames if they meet other checks. (Such a frame cannot be a valid multi-destination frame because the Any-RBridge nickname is not a valid distribution tree root.)

2. If an RBridge Channel Alert extended header flag is set, then the RBridge MUST process the RBridge Channel message as described below even if it is not egressing the frame. If it is egressing the frame, then no additional processing beyond egress processing is needed even if an RBridge Channel Alert flag is set.

3. On decapsulation, the special Inner.MacDA value of All-Egress-RBridges MUST be recognized to trigger checking the Inner.Ethertype and processing as an RBridge Channel message if that Ethertype is RBridge-Channel.

RBridge Channel messages SHOULD only be sent to RBridges that advertise support for the channel protocol involved as described in Section 5.

All RBridges supporting the RBridge Channel facility MUST recognize the RBridge-Channel inner Ethertype.

3.1 Processing the RBridge Channel Header

Knowing that it has an RBridge Channel message, the egress RBridge, and any transit RBridge if an RBridge Channel Alert bit is set in the TRILL Header, looks at the CHV (RBridge Channel Header Version) and Channel Protocol fields.

If any of the following conditions occur at an egress RBridge, the frame is not processed, an error may be generated as specified in Section 3.2, and the frame is discarded. The behavior is the same if the frame is being processed at a transit RBridge because the critical RBridge Channel Alert flag is set [RFCext]. However, if these conditions are detected at a transit RBridge examining the message because the non-critical RBridge Channel Alert flag is set

[RFCext] but the critical RBridge Channel Alert flag is not set, no error is generated and the frame is still forwarded normally.

Error Conditions:

1. The Ethertype is not RBridge-Channel and not any other Ethertype known to the RBridge as usable with the All-Egress-RBridges Inner.MacDA, or the frame is so short that the Ethertype is truncated.

2. The CHV field is non-zero or the frame is so short that the version zero Channel Header is truncated.

3. The Channel Protocol field is a reserved value or a value unknown to the processing RBridge.

4. The ERR field is non-zero and Channel Protocol is a value other than 0x001.

5. The RBridge Channel Header NA flag is set to one indicating that the frame should have been received as a native frame rather than a TRILL Data frame.

If the CHV field and NA flag are zero and the processing RBridge recognizes the Channel Protocol value, it processes the message in accordance with that channel protocol. The processing model is as if the received frame starting with and including the TRILL Header is delivered to the Channel protocol along with a flag indicating whether this is (a) transit RBridge processing due to an RBridge Channel Alert flag being set or (b) egress processing.

Errors within a recognized Channel Protocol are handled by that channel protocol itself and do not produce RBridge Channel Error frames.

3.2 RBridge Channel Errors

A variety of problems at the RBridge Channel level cause the return of an RBridge Channel Error frame unless one of the following apply: (a) the "SL" (Silent) flag is a one in the channel message for which the problem was detected, (b) the processing is due to the non-critical RBridge Channel Alert bit being set, (c) the frame in error appears, itself, to be an RBridge Channel error frame (has a non-zero ERR field or a Channel Protocol of 0x001), or (d) the error is suppressed due to rate limiting.

An RBridge Channel Error frame is a multi-hop unicast RBridge Channel message with the ingress nickname set to a nickname of the RBridge

detecting the error, and the egress nickname set to the value of the ingress nickname in the channel message for which the error was detected. No per-hop transit processing is specified for such error frames, so the RBridge Channel Alert extended header flags SHOULD, if an extension is present, be set to zero. The SL and MH flags SHOULD be set to one, the NA flag MUST be zero, and the ERR field MUST be non-zero as described below. For the protocol specific data area, an RBridge Channel Message Error frame has at least the first 256 bytes (or less if less are available) of the erroneous decapsulated channel message starting with the TRILL Header. (Note: The TRILL Header does not include the TRILL Ethertype that is part of the Link Header on Ethernet Links.)

The following values for ERR are specified:

ERR	RBridge Channel Error Code Meaning
0	- No error
1	Frame too short (truncated Ethertype or Channel Header)
2	Unrecognized Ethertype
3	Unimplemented value of CHV
4	Wrong value of NA flag
5	Channel Protocol is reserved or unimplemented
6-14	- Available for allocation, see Section 7.
15	Reserved (see Note)

Note: Intended to be allocated by Standards Action for an error code expansion feature when it appears likely that all other available error codes are being allocated.

All RBridges implementing the RBridge Channel feature MUST recognize the RBridge Channel Error protocol value (0x001). They MUST NOT generate an RBridge Channel Error message in response to a RBridge Channel Error message, that is, a channel message with a protocol value of 0x001 or with a non-zero ERR field.

4. Native RBridge Channel Frames

Other sections of this document specify non-native RBridge Channel messages and their processing, that is, RBridge Channel messages formatted as TRILL Data frames and sent between RBridges. This section specifies the differences for native RBridge Channel messages.

If provided for by the channel protocol involved, native RBridge channel messages may be sent between end-stations and RBridges that are directly connected over a link, in either direction. On an Ethernet link, such native frames have the RBridge-Channel Ethertype and are like the encapsulated frame inside an RBridge Channel message except as follows:

1. TRILL does not require the presence of a VLAN tag on such native RBridge channel frames. However, port configuration, link characteristics, or the channel protocol involved may require such tagging.

2. If the frame is unicast, the destination MAC address is the unicast MAC address of the RBridge or end-station port that is its intended destination. If the frame is multicast by an end station to all the RBridges on a link that support an RBridge Channel protocol that uses this transport, the destination MAC address is the All-Edge-RBridges multicast address (see Section 7). A native RBridge Channel frame received at an ingress RBridge with a destination MAC address that is a unicast address different from that of the port or multicast address different from All-Edge-RBridges, is discarded. If the frame is multicast by an RBridge to all the devices that TRILL considers to be end stations on a link that support an RBridge Channel protocol that uses this transport, the destination MAC address is the TRILL-End-Stations multicast address (see Section 7). A native RBridge Channel frame received at an end station with a destination MAC address that is a unicast address different from that of the port or multicast address different from TRILL-End-Stations, is discarded.

3. The RBridge-Channel outer Ethertype must be present. In the future there may be other protocols using the All-Edge-RBridges and/or TRILL-End-Stations multicast addresses on native frames distinguished by different Ethertypes.

4. The NA or native bit in the RBridge Channel Header flags MUST be a one.

5. There might be additional tags present between the Outer.MacDA, Outer.MacSA pair and the RBridge-Channel Ethertype.

The RBridge Channel protocol number space for native RBridge Channel messages and TRILL Data formatted RBridge Channel messages is the same. If provided for by the channel protocol involved, the receipt of a native RBridge Channel frame MAY lead to the generation and transmission of one or more Inter-RBridge Channel frames. The decapsulation and processing of a TRILL Data RBridge Channel frame MAY, if provided for by the channel protocol involved, result in the sending of one or more native RBridge channel frames to one or more end stations. Thus, there could be an RBridge Channel protocol that involved an RBridge Channel message sent from an origin RBridge where the message is created, through one or more transit RBridges and from the last as a native RBridge channel message to and end station or the reverse of such a path; however, to do this the RBridge channel protocol involved must be implemented at the RBridge where the channel message is changed between a native frame and a TRILL Data format frame and that RBridge must change the channel message itself, at a minimum complementing the NA flag in the Channel Header and making appropriate MAC address changes.

An erroneous native channel message results in a native RBridge channel error message under the same conditions for which an TRILL Data RBridge Channel message would result in a TRILL Data RBridge channel error message. However, in a native RBridge Channel error message, the NA flag MUST be one. Also, since there is no TRILL Header in native RBridge Channel protocol frames, the beginning part of the frame in which the error was detected that is included in native RBridge Channel error frames starts with the RBridge Channel Header (including the RBridge-Channel Ethertype). The destination MAC address of such error messages is set to the source MAC address of the native RBridge Channel message that was in error.

There is no mechanism to stop end stations from directly exchanging native RBridge Channel messages but such usage is beyond the scope of this document.

5. Indicating Support for RBridge Channel Protocols

Support for RBridge Channel protocols is indicated by the presence of one or more TLVs and/or sub-TLVs in an RBridge's LSP as documented in [RFC6326bis].

RBridge Channel protocols 0 and 0xFFF are reserved and protocol 1, the RBridge Channel error protocol, MUST be implemented as part of the RBridge Channel feature. Thus, if an RBridge supports the RBridge Channel feature, it should be advertising support for protocol 1 and not advertising support for protocols 0 or 0xFFF in its LSP. However, indication of support or non-support for RBridge Channel protocol 1 is ignored on receipt and support for it is always assumed, if support for any RBridge Channel is indicated in the RBridge's LSP.

6. Congestion Considerations

The bandwidth resources used by RBridge Channel protocols are recommended to be small compared to the total bandwidth of the links they traverse. When doing network planning, the bandwidth requirements for TRILL data, TRILL IS-IS, the TRILL ESADI protocol, RBridge Channel traffic, and any other link local traffic need to be taken into account.

Specifications for particular RBridge Channel protocols MUST consider congestion and bandwidth usage implications and provide guidance on bandwidth or packet frequency management. RBridge Channel protocols can have built-in bandwidth management in their protocols. Outgoing channel messages SHOULD be rate-limited, by configuring the underlying protocols or otherwise, to prevent aggressive connectivity verification or other applications consuming excessive bandwidth, causing congestion, or becoming denial-of-service attacks.

If these conditions cannot be followed, an adaptive loss-based scheme SHOULD be applied to congestion-control outgoing RBridge Channel traffic, so that it competes fairly, taking into account packet priorities and drop eligibility as indicated in the Inner.VLAN, with TCP or similar traffic within an order of magnitude. One method of determining an acceptable bandwidth for RBridge Channel traffic is described in [RFC5348]; other methods exist. For example, bandwidth or packet frequency management can include any of the following: a negotiation of transmission interval/rate such as that provided in BFD [RFC5880], a throttled transmission rate on "congestion detected" situations, a gradual ramp-up after shutdown due to congestion and until basic connectivity is verified, and other mechanisms.

Connectivity chcking applications such as BFD [RFC5880] SHOULD be rate-limited to below 5% of the bit-rate of the associated link or links. For this purpose, the mean or sustained bit-rate of the link or links is used.

Incoming RBridge Channel messages MAY be rate-limited as a protection against denial-of-service attacks. This throttling of incoming messages SHOULD honor packet priorities and drop eligibility indications as indicate in the Inner.VLAN, preferentially discarding drop eligible and lower priority packets.

7. Allocation Considerations

The following subsections give IANA and IEEE allocation considerations. In this document, the allocation procedure specifications are as defined in [RFC5226].

7.1 IANA Considerations

IANA is requested to allocate a previously unassigned TRILL Nickname as follows:

Any-RBridge	TBD (0xFFCO suggested)

IANA is requested to add "All-Egress-RBridges" to the TRILL Parameter Registry as an alternative name for the "All-ESADI-RBridges" multicast address.

IANA is requested to allocate two previously unassigned TRILL Multicast address as follows:

TRILL-End-Stations	TBD (01-80-C2-00-00-45 suggested)
All-Edge-RBridges	TBD (01-80-C2-00-00-46 suggested)

IANA is requested to create an additional sub-registry in the TRILL Parameter Registry for RBridge Channel Protocols, with initial contents as follows:

Protocol	Description	Reference
0x000	Reserved, not to be allocated	(This document)
0x001	RBridge Channel Error	(This document)
0x002-0x0FF	Available (1)	
0x100-0xFF7	Available (2)	
0xFF8-0xFFE	Private Use	
0xFFF	Reserved, not to be allocated	(This document)

(1) RBridge Channel protocol code points from 0x002 to 0x0FF require a Standards Action, as modified by [RFC4020], for allocation.

(2) RBridge Channel protocol code points from 0x100 to 0xFF7 are RFC Required to allocate a single value or IESG Approval to allocate multiple values.

IANA is requested to create an additional sub-registry in the TRILL Parameter Registry for RBridge Channel Header Flags with initial contents as follows:

Flag Bit	Mnemonic	Allocation
0	SL	Silent
1	MH	Multi-hop
2	NA	Native
3-11	-	Available for allocation

Allocation of an RBridge Channel Header Flag is based on IETF Review.

IANA is requested to create an additional sub-registry in the TRILL Parameter Registry for RBridge Channel Error codes with initial contents as listed in Section 3.2 above and with available values allocated by Standards Action as modified by [RFC4020].

7.2 IEEE Registration Authority Considerations

The IEEE Registration Authority has assigned the Ethertype <TBD> for RBridge-Channel.

8. Security Considerations

No general integrity, authentication, or encryption mechanisms are provided herein for RBridge Channel messages. If these services are required for a particular RBridge Channel protocol, they MUST be supplied by that channel protocol. See, for example, the BFD Authentication mechanism [RFC5880].

See [RFC6325] for general TRILL Security Considerations. As stated therein, no protection is provided by TRILL against forging of the ingress nickname in a TRILL Data formatted channel message or the Outer.MacSA in a native RBridge Channel frame on an Ethernet link. This may result in misdirected return responses or error messages. However, link level security protocols may be used to authenticate the origin station on a link and protect against attacks on links. See also Section 6 above concerning congestion.

If indication of RBridge Channel Protocol support are improperly absent from an RBridge's LSP, it could deny all RBridge Channel services, for example some BFD services, for the RBridge in question. If a particular RBridge channel protocol is incorrectly not advertised as supported, it could deny the service of that channel protocol to the RBridge in question.

Incorrect indication of RBridge Channel Protocol support or incorrect assertion of support for a channel protocol could encourage RBridge channel messages to be sent to an RBridge that does not support the channel feature or the particular channel protocol used. The inner frame of such messages could be decapsulated and that inner frame could be sent out all ports that are appointed forwarders for the frame's Inner.VLAN. However, this is unlikely to cause much harm; in particular, there are two possibilities as follows: (a) If end stations do not recognize the RBridge-Channel Ethertype of the frame, they will drop it. (b) If end stations do recognize the RBridge-Channel Ethertype and the channel protocol indicated in the frame, they should refuse to process the frame due to an incorrect value of the RBridge Channel Header NA flag.

9. References

The following sections list normative and informative references for this document.

9.1 Normative References

[IS-IS] - ISO/IEC 10589:2002, Second Edition, "Intermediate System to Intermediate System Intra-Domain Routing Exchange Protocol for use in Conjunction with the Protocol for Providing the Connectionless-mode Network Service (ISO 8473)", 2002.

[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.

[RFC4020] - Kompella, K. and A. Zinin, "Early IANA Allocation of Standards Track Code Points", BCP 100, RFC 4020, February 2005.

[RFC5226] - Narten, T. and H. Alvestrand, "Guidelines for Writing an IANA Considerations Section in RFCs", BCP 26, RFC 5226, May 2008.

[RFC5348] - Floyd, S., Handley, M., Padhye, J., and J. Widmer, "TCP Friendly Rate Control (TFRC): Protocol Specification", RFC 5348, September 2008.

[RFC6325] - Perlman, R., Eastlake 3rd, D., Dutt, D., Gai, S., and A. Ghanwani, "Routing Bridges (RBridges): Base Protocol Specification", RFC 6325, July 2011.

[RFC6327] - Eastlake 3rd, D., Perlman, R., Ghanwani, A., Dutt, D., and V. Manral, "Routing Bridges (RBridges): Adjacency", RFC 6327, July 2011.

[RFCext] - D. Eastlake, A. Ghanwani, V. Manral, Y. Li, C. Bestler, "TRILL: TRILL Header Extension", draft-ietf-trill-rbridge-extension, in RFC Editor’s queue.

[RFC6326bis] - Eastlake, D., A. Banerjee, D. Dutt, A. Ghanwani, R. Perlman, "TRILL Use of IS-IS", draft-eastlake-isis-rfc6326bis, work in progress.

9.2 Informative References

[RFC5082] - Gill, V., Heasley, J., Meyer, D., Savola, P., Ed., and C. Pignataro, "The Generalized TTL Security Mechanism (GTSM)", RFC 5082, October 2007

[RFC5586] - Bocci, M., Ed., Vigoureux, M., Ed., and S. Bryant, Ed., "MPLS Generic Associated Channel", RFC 5586, June 2009.

[RFC5880] - D. Katz, D. Ward, "Bidirectional Forwarding Detection (BFD)", June 2010.

[RFC5882] - D. Katz, D. Ward, "Generic Application of Bidirectional Forwarding Detection (BFD)", June 2010.

[ClearCorrect] - D. Eastlake, M. Zhang, A. Ghanwani, A. Banerjee, V. Manral, "TRILL: Clarifications, Corrections, and Updates", draft-ietf-trill-clear-correct, work in progress.

Appendix: Change History

RFC Editor: please delete this appendix before publication.

Changes from -00 to -01

1. Spell out more acronyms.
2. Add reference to "Guidelines for the Use of OAM" draft.
3. Move definition of Alert flag to draft-ietf-trill-rbridge-options and refer to it as an extended header flag.
4. Change name of "Egress-RBridges" multicast address to "All-Egress-RBridges". Merge with All-ESADI-RBridges (i.e., they are two names for the same MAC address).
5. Add detailed bit vector description for indicating support of RBridge channel protocols. Add GENAPP and an APPsub-TLV to hold one or more bit vectors.
6. Assorted editorial changes.

Changes from -01 to -02

1. Update for drafts that have been issued as RFCs.
2. Change to specification of Inner.VLAN in RBridge channel messages.
3. Remove GENAPP and move RBridge Channels supported information to another document.
4. Clarify native RBridge Channel error messages.
5. Assorted editorial changes.

Changes from -02 to -03

1. Liberalize restrictions on RBridge acceptance of native RBridge Channel messages. These are typically messages and should generally be accepted unless in a VLAN not enabled at the port or the like.
2. Change multi-cast address used by end stations in sending a native

RBridge Channel message to all RBridges on the link from All-Egress-RBridges to All-Edge-RBridges to avoid possible confusion if such a frame were encapsulated resulting in an All-Egress-RBridges Inner.MacDA.

3. Reword references to "two-hop echo" and the like for clarity. (This meant an echo frame that went to an immediate neighbor and back.)

4. Add reference to and move some material to the RFC 6326bis draft.

5. Assorted editorial changes.

Changes from -03 to -04

1. Update for the replacement of the CFI bit by the DEI bit (see [ClearCorrect]).

2. Update for the existence of both critical and non-critical RBridge Channel alert flags.

3. Update author information.

4. Assorted editorial changes.

Changes from -04 to -05

1. Clarify the distinction between native and non-native RBridge Channel messages and that native channel messages are only intended to be transmitted between RBridge and end stations on the same link.

2. Add a paragraph to the Security Considerations section about forged ingress nicknames / source MAC addresses in channel messages.

3. Add acknowledgements section.

4. Replace "OAM" references with "BFD" references in Abstract and Introduction.

5. Very minor editorial changes.

Changes from -05 to -06

1. Improve wording in 2.1.1 re CHV values.

2. Revert "Ext-Len" to "Op-Len".

3. Fix typos and make minor editorial changes.

Changes from -06 to -07

1. Add bit numbers at top of figures where they were missing.

2. Add figure numbers and captions.

3. Add text to Section 2.1.1 concerning Private Use RBridge Channel protocol numbers.

4. Change IANA Considerations for the allocation of multiple RBridge Channel protocol numbers in the 0x100 to 0xFF7 range from IETF Review to IESG Approval.

5. Add text that the intended use for ERR code 15 is for some future error code expansion feature should more error codes be required and indicate that protocol numbers 0x000 and 0xFFF are not to be allocated.

6. Captialize the first occurrence of "must" in Section 7.

7. Add statement that directly connected end-stations are not blocked from communicating with each other using channel messages but such messages are beyond the scope of this document.

8. Re-order and add some references to the Securty Considertions section.

9. Typo fixes and various editorial changes.

Changes from -07 to -08

1. Add congestion considerations section.

2. Minor editorial changes.

Acknowledgments

The authors gratefully acknowledge the comments and contributions of the follows, listed is alphabetic order: Stewart Bryant, Somnath Chatterjee, Adrian Farrel, Stephen Farrell, Miguel A. Garcia, Anoop Ghanwani, Brian Haberman, Rakesh Kumar, Barry Leiba, and Tissa Senevirathne.

This document was prepared with raw nroff. All macros used were defined in the document source files.

Authors' Addresses

Donald Eastlake 3rd
Huawei R&D USA
155 Beaver Street
Milford, MA 01757 USA

Tel: +1-508-333-2270
EMail: d3e3e3@gmail.com

Vishwas Manral
HP Networking
19111 Pruneridge Avenue
Cupertino, CA 95014 USA

Tel: +1-408-477-0000
EMail: vishwas.manral@hp.com

Yizhou Li
Huawei Technologies
101 Software Avenue,
Nanjing 210012, China

Phone: +86-25-56622310
Email: liyizhou@huawei.com

Sam Aldrin
Huawei Technologies
2330 Central Expressway
Santa Clara, CA 95050 USA

Phone: +1-408-330-5000
Email: sam.aldrin@huawei.com

Dave Ward
Cisco Systems
170 W. Tasman Drive
San Jose, CA 95134 USA

EMail: dward@cisco.com

Copyright, Disclaimer, and Additional IPR Provisions