INTERNET-DRAFT
Intended Status: Proposed Standard
Expires: January 13, 2014

Mingui Zhang
Peng Zhou
Huawei

July 12, 2013

ICCP Application TLVs for VPN Route Label Sharing
draft-zhang-pwe3-iccp-label-sharing-00.txt

Abstract

This document defines TLVs under Inter-Chassis Communication Protocol (ICCP) to include a new application: Label Sharing for Fast PE Protection. Egress PEs in the same Redundant Group utilize the ICCP connection to negotiate the "VPN route label" and the "BGP next hop" for each VPN.

Status of this Memo

This Internet-Draft is submitted to IETF 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/1id-abstracts.html

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

Copyright and License Notice

Table of Contents

1. Introduction

It’s common for Service Providers (SPs) to connect one CE to multiple PEs for the sake of reliability. In [LS], this feature is leveraged to realize a method for fast PE protection. There, egress PEs in the same Redundant Group (RG) share the same "VPN route label" for one VPN. These egress PEs use a virtual Next Hop (vNH) as their "BGP next hop". Primary and backup LDP LSP tunnels ended at the vNH are set up using IGP FRR [LFA] [MRT]. When the PLR redirects the failure affected packet to the backup egress PE, the VPN route label encapsulated in the packet can be recognized by the backup egress PE and the packet will be delivered naturally.

This document extends ICCP to include the "label sharing" method as a new application. The connection of ICCP is leveraged to synchronize the label and BGP next hop of each VPN for the PEs in one RG. TLVs are defined in the next section.

1.1. 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].

1.2. Terminology

vNH: virtual Next Hop
FRR: Fast ReRouting
PLR: Point of Local Repair

2. Label Sharing TLVs in ICCP

This section specifies the ICCP Connect, Disconnect and Application Data TLVs to be used by egress PEs for the label sharing application.

2.1. Label Sharing Connect TLV

This TLV is included in the RG Connect message to signal the establishment of Label Sharing application connection.

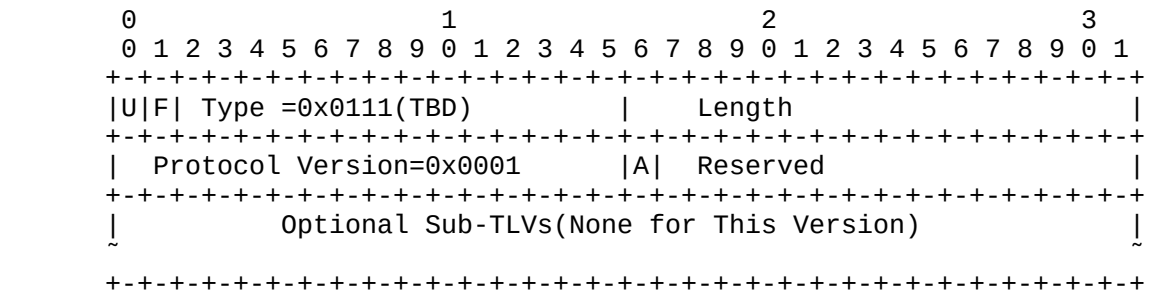

- U and F Bits

  Both are set to 0.

- Type

  set to 0x0111 (TBD) for "Label Sharing Connect TLV"

- Length

  Length of the TLV in octets excluding the U-bit, F-bit, Type, and Length fields.

- Protocol Version

  The version of this particular protocol for the purposes of ICCP. This is set to 0x0001.

- A bit

  Acknowledgement Bit. Set to 1 if the sender has received a Label Sharing Connect TLV from the recipient. Otherwise, set to 0.

- Reserved

  Reserved for future use.

- Optional Sub-TLVs

  There are no optional Sub-TLVs defined for this version of the protocol.

2.2. Label Sharing Disconnect TLV

This TLV is included in an RG Disconnect Message as the "Disconnect Code TLV" (See Section 6.3 of [ICCP]). It indicates that the connection for the Label Sharing application is to be terminated.

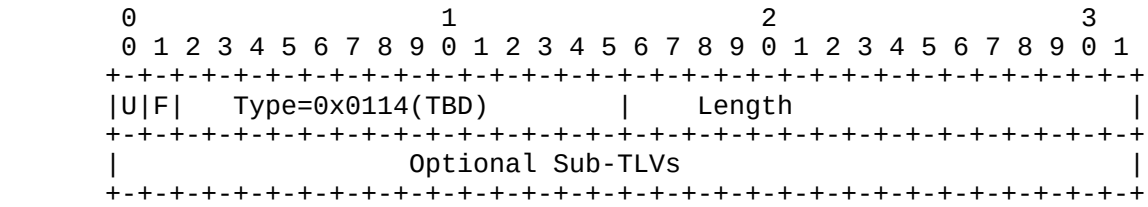

- U and F Bits

  Both are set to 0.

- Type

  set to 0x0114 (TBD) for "Label Sharing Disconnect TLV"

- Length

  Length of the TLV in octets excluding the U-bit, F-bit, Type, and Length fields.

- Optional Sub-TLVs

  The only optional Sub-TLV defined for this version of the protocol is the "Label Sharing Disconnect Cause" TLV defined next:

2.2.1. Label Sharing Disconnect Cause TLV

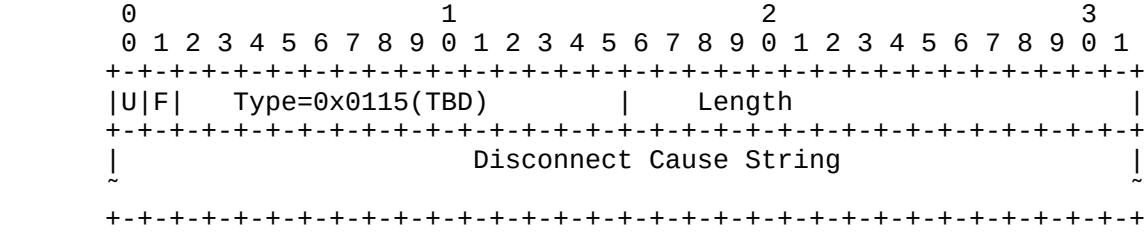

- U and F Bits

  Both are set to 0.

- Type

  set to 0x0115 (TBD) for "Label Sharing Disconnect Cause TLV"

- Length

Length of the TLV in octets excluding the U-bit, F-bit, Type, and Length fields.

- Disconnect Cause String

Variable length string specifying the reason for the disconnect. Used for network management.

2.3. Label Sharing Application Data TLVs

The following TLVs are included in the RG Application Data message to deliver the information that need be synchronized among RG members.

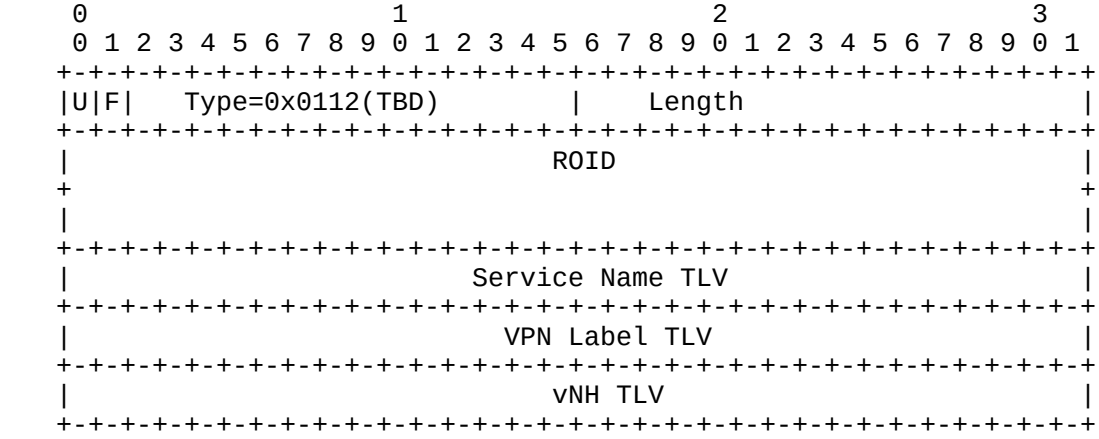

- U and F Bits

  Both are set to 0.

- Type

  set to 0x0112 (TBD) for "Label Sharing Information TLV"

- Length

  Length of the MAC address, which is 6 octets.

- ROID

  As defined in the ROID section of [ICCP].

- Sub-TLVs

  i Service Name TLV

ii VPN Label TLV

iii vNH TLV

2.3.1. Service Name TLV

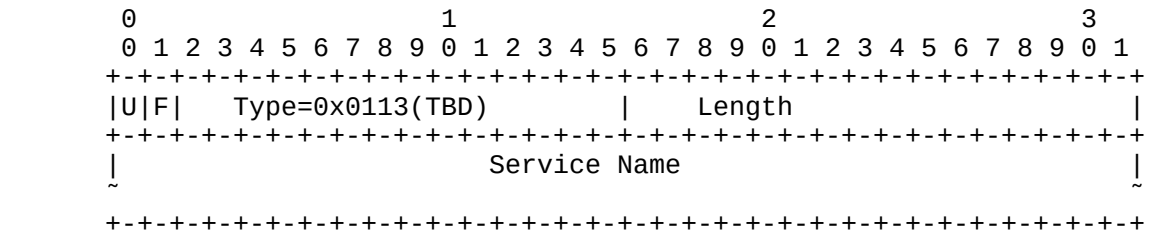

- U and F Bits

  Both are set to 0.

- Type

  set to 0x0113 (TBD) for "Service Name TLV"

- Length

  Length of the TLV in octets excluding the U-bit, F-bit, Type, and Length fields.

- Service Name

  The name of the VPN service instance encoded in UTF-8 format and up to 80 character in length.

2.3.2. VPN Label TLV

The PE with the highest priority (with its MAC address as the tiebreaker) assigns the shared VPN label for a VPN. In a well configured network, PEs in the same RG will be configured to have the same range of VPN labels for sharing. When the ranges of the VPN labels are different, the VPN label is chosen from the intersection of the ranges.

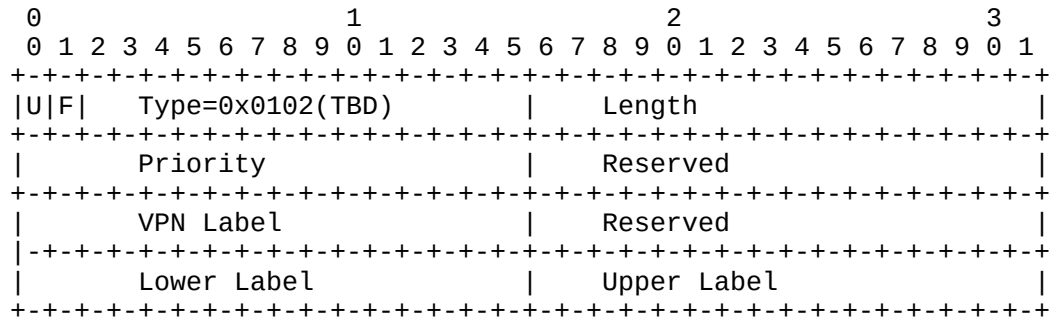

- U and F Bits

  Both are set to 0.

- Type

  set to 0x0112 (TBD) for "VPN Label TLV"

- Length

  Length of the TLV in octets excluding the U-bit, F-bit, Type, and Length fields.

- Priority

  The priority that the sender has for the VPN label in this TLV. When there are more than one sender who has the highest priority, the MAC address of the sender used as the tiebreaker.

- Reserved

  Reserved for future use.

- VPN Label

  The VPN label to be shared among the RG.

- Lower/Upper Label

  The lower/upper bound of a valid VPN label.

2.3.3. vNH TLV

When a VPN route is distributed to ingress PEs by BGP, the IP address of the vNH will be used as the BGP next hop. Thus, tunnels terminated at the vNH will be set up. The PE with the highest priority (with its

MAC address as the tiebreaker) determines the IP address of the vNH.

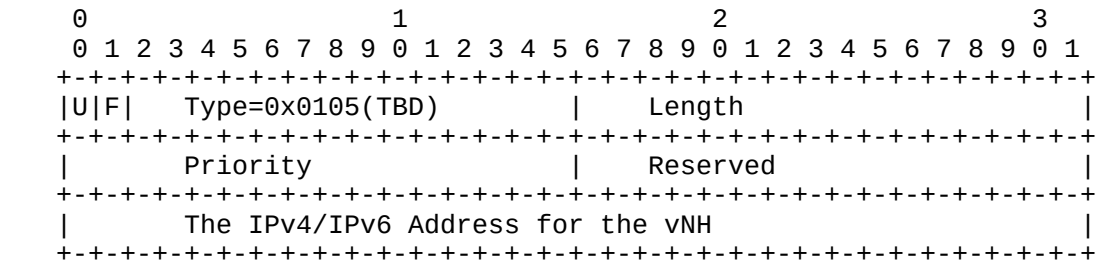

- U and F Bits

  Both are set to 0.

- Type

  set to 0x0105 (TBD) for "Service Name TLV"

- Length

  Length of the TLV in octets excluding the U-bit, F-bit, Type, and Length fields. Lengths for the IPv4 and IPv6 Addresses TLVs are different.

- Priority

  The priority that the sender has for the IPv4/IPv6 address for the vNH in this TLV. When there are more than one sender who has the highest priority, the MAC address of these senders will be used as the tiebreaker.

- Reserved

  Reserved for future use.

- IPv4/IPv6 Address for the vNH

  The IPv4/IPv6 address that the sender wants the vNH to use. The IPv4/IPv6 address of vNH TLV sent out by sender with the highest priority will be used as the IPv4/IPv6 address of the vNH by all the PEs in the same RG.

3. Security Considerations

This document raises no new security issues.

4. IANA Considerations

The types used by the application TLVs defined in Section 3 should be assigned.

5. References

5.1. Normative References

[ICCP] L. Martini, S. Salam, et al, "Inter-Chassis Communication Protocol for L2VPN PE Redundancy", draft-ietf-pwe3-iccp-11.txt, work in progress.

[LS] M. Zhang, P. Zhou, "Label Sharing for Fast PE Protection", draft-zhang-l3vpn-label-sharing-00.txt, work in progress.

5.2. Informative References

[LFA] Filsfils, C., Ed., Francois, P., Ed., Shand, M., Decraene, B., Uttaro, J., Leymann, N., and M. Horneffer, "Loop-Free Alternate (LFA) Applicability in Service Provider (SP) Networks", RFC 6571, June 2012.

[MRT] A. Atlas, Ed., R. Kebler, et al, "An Architecture for IP/LDP Fast-Reroute Using Maximally Redundant Trees", draft-ietf-rtgwg-mrt-frr-architecture-02.txt, work in progress.

Author's Addresses

Mingui Zhang
Huawei Technologies Co., Ltd
Huawei Building, No.156 Beiqing Rd.
Z-park, Shi-Chuang-Ke-Ji-Shi-Fan-Yuan, Hai-Dian District,
Beijing 100095 P.R. China

Email: zhangmingui@huawei.com

Peng Zhou
Huawei Technologies Co., Ltd
Huawei Building, No.156 Beiqing Rd.
Z-park, Shi-Chuang-Ke-Ji-Shi-Fan-Yuan, Hai-Dian District,
Beijing 100095 P.R. China

Email: Jewpon.zhou@huawei.com