Internet Working Group
Internet Draft
Intended status: Standards Track

Y. Jiang
Y. Luo
Huawei
E. Mallette
Bright House Networks
C. Shen
China Telecom
W. Cheng
China Mobile
Y. Shen
Juniper Networks
G. Zhou
China Unicom

Expires: April 2015 October 25, 2014

Multi-chassis PON Protection in MPLS
draft-jiang-pwe3-mc-pon-03.txt

Status of this Memo

Copyright Notice

Abstract

MPLS is being deployed deeper into operator networks, often to or past the access network node. Separately network access nodes such as PON OLTs have evolved to support first-mile access protection, where one or more physical OLTs provide first-mile diversity to the customer edge. Multi-homing support is needed on the MPLS-enabled PON OLT to provide resiliency for provided services. This document describes the multi-chassis PON protection architecture in MPLS and also proposes the ICCP extension to support it.

Table of Contents

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 [RFC2119].

2. Terminology

DSL Digital Subscriber Line

FTTx Fiber-to-the-x (FTTx, x = H for home, P for premises, C for curb)

ICCP Inter-Chassis Communication Protocol

OLT Optical Line Termination

ONU Optical Network Unit

MPLS Multi-Protocol Label Switching

PON Passive Optical Network

RG Redundancy Group

3. Introduction

MPLS is being extended to the edge of operator networks, as is described in the seamless MPLS use cases [SEAMLESS], and the MS-PW with PON access use case [RFC6456]. Combining MPLS with OLT access further facilitates a low cost multi-service convergence.

Tens of millions of FTTx lines have been deployed over the years, with many of those lines being some PON variant. PON provides operators a cost-effective solution for delivering high bandwidth (1Gbps or even 10Gbps) to a dozen or more subscribers simultaneously.

In the past, access technologies such as Passive Optical Network (PON) and Digital Subscriber Line (DSL) are usually used for subscribers, and no redundancy is provided in their deployment.

But with the rapid growth of mobile data traffic, more and more LTE small cells and Wi-Fi hotspots are deployed. PON is considered as a viable low cost backhaul solution for these mobile services. Besides its high bandwidth and scalability, PON further provides synchronization features, e.g., SyncE and IEEE1588 functionality, which can fulfill synchronization needs of mobile backhaul services.

The Broadband Forum specifies reference architecture for mobile backhaul network using MPLS transport in [TR-221] where PON can be the access technology, and is further working on PON-based mobile backhaul network architecture in [SD-331].

Unlike typical residential service where a single or handful of end-users hangs off of a single PON OLT port in a physical optical distribution network, a PON port that supports a dozen LTE small cells or Wi-Fi hotspots could be providing service to hundreds of simultaneous subscribers. Small cell backhaul often demands the economics of a PON first-mile and yet expects first-mile protection commonly available in point-to-point access portfolio.

Some optical layer of protection mechanisms, such as Trunk and Tree protection, are specified in [IEEE-1904.1] to avoid single point of failure in the access. They are called Type B and Type C protection respectively in [G983.1].

Trunk protection architecture is an economical PON resiliency mechanism, where the working OLT and the working link between the working splitter port and the working OLT (i.e., the working trunk fiber) is protected by a redundant protection OLT and a redundant trunk fiber between the protection splitter port and the protection OLT, however it only protects a portion of the optical path from OLT to ONUs. This is different from the more complex and costly Type C protection architecture where there is a working optical distribution network path from the working OLT and a complete protected optical distribution network path from the protection OLT to the ONUs. Figure 1 demonstrates a typical scenario of Trunk protection.

```
                        |                                 |
                        |<--Optical Distribution Network->|
                        |                                 |
                        |   branch               trunk    +-----+
                  +-----+   fibers               fibers   |     |
   Base     ------|     |     |                           . OLT |
   Stations ------| ONU |\    |                        ,''|  A  |
            ------|     |  \  V                     _-'   +-----+
                  +-----+    \                    .'
                            .  \  +----------+ ,-'
                  +-----+   .    \|         -'   Working
   Base     ------|     |   .     | Optical  |
   Stations ------| ONU |---------| Splitter |
            ------|     |   .    /|         -,   Protection
                  +-----+   .  /  +----------+ ''.,
                             /                     '-,    +-----+
                  +-----+  /                          ''.,|     |
   Base     ------|     |/                                | OLT |
   Stations ------| ONU |                                 |  B  |
            ------|     |                                 +-----+
                  +-----+
```

Figure 1 Trunk Protection Architecture in PON

Besides small cell backhaul, this protection architecture can also be applicable to other services, for example, DSL and Multi-System Operator (MSO) services. In that case, an ONU in Figure 1 can play the similar role as a Digital Subscriber Line Access Multiplexer (DSLAM) and dozens of Customer Premises Equipments (CPEs) or cable modems may be attached to it.

In some deployments, it is also possible that only some ONUs are needed to be protected.

The PON architecture depicted in Figure 1 can provide redundancy in its physical topology, however, all traffic including link OAM are blocked on the protection link which frustrates end to end protection mechanisms such as ITU-T G.8031. Therefore, some standard signaling mechanisms are needed between OLTs to exchange information, for example, PON link status, registered ONU information, and network status, so that protection and restoration can be done both rapidly and reliably, especially when the OLTs also support MPLS.

ICCP [ICCP] provides a framework for inter-chassis synchronization of state and configuration data between a set of two or more PEs. Currently ICCP only defines application specific messages for PW redundancy and mLACP, but it can be easily extended to support PON as an Attachment Circuit (AC) redundancy.

This document proposes the extension of ICCP to support Multi-chassis PON protection in MPLS.

4. ICCP Protocol Extensions

4.1. Multi-chassis PON Application TLVs

A set of multi-chassis PON application TLVs are defined in the following sub-sections.

4.1.1. PON Connect TLV

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

```
 0                   1                   2                   3
  0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
 +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
 |U|F|   Type=0x00XX             |    Length                     |
 +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
 |      Protocol Version         |A|        Reserved             |
 +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
 |                    Optional Sub-TLVs                          |
 ~                                                               ~
 |                                                               |
 +                                +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
 |             ...                |
 +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

- U and F Bits, both are set to 0.

- Type, set to 0x00XX for "PON 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 PON specific protocol for the purposes of inter-chassis communication. This is set to 0x0001.

- A Bit, Acknowledgement Bit. Set to 1 if the sender has received a PON 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.

4.1.2. PON Disconnect TLV

This TLV is included in the RG Disconnect message to indicate that the connection for the PON application is to be terminated.

```
 0                   1                   2                   3
  0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
 +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
 |U|F|   Type=0x00XX             |    Length                     |
 +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
 |                        Optional Sub-TLVs                      |
 +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

- U and F Bits, both are set to 0.

- Type, set to 0x00XX for "PON Disconnect TLV".

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

- Optional Sub-TLVs, there are no optional Sub-TLVs defined for this version of the protocol.

4.1.3. PON Configuration TLV

The "PON Configuration TLV" is included in the "RG Application Data" message, and announces an OLT’s system parameters to other members in the same RG.

```
 0                   1                   2                   3
  0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
 +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
 |U|F|   Type=0x00XX             |    Length                     |
 +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
 |                          System ID                            |
 |                                                              |
 |                                                              |
 +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
 |     System Priority           |             Port ID           |
 +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

- U and F Bits, both are set to 0.

- Type, set to 0x00XX for "PON Configuration TLV".

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

- System ID, 8 octets encoding the System ID used by the OLT, which is the Chassis MAC address. If a 6 octet System ID is used, the least significant 2 octets of the 8 octet field will be encoded as 0000.

- System Priority, 2 octets encoding the System Priority.

- Port ID, 2 octets PON Port ID.

Further configuration considerations such as multicast table and ARP table for static MAC addresses will be added in a next version.

4.1.4.PON State TLV

The "PON State TLV" is included in the "RG Application Data" message, and used by an OLT to report its PON states to other members in the same RG.

```
 0                   1                   2                   3
  0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
 +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
 |U|F|   Type=0x00XX             |    Length                     |
 +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
 |                              ROID                             |
 |                                                               |
 |                                                               |
 +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
 |                    Local PON Port state                       |
 +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
 |                    Remote PON Port state                      |
 +-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

- U and F Bits, both are set to 0.

- Type, set to 0x00XX for "PON State TLV"

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

- ROID, as defined in the ROID section of [ICCP].

- Local PON Port State, the status of the local PON port as determined by the sending OLT (PE). The last bit is defined as Fault indication of the PON Port associated with this PW (1 - in fault).

- Remote PON Port State, the status of the remote PON port as determined by the remote peer of the sending OLT (PE). The last bit is defined as Fault indication of the PON Port associated with this PW (1 - in fault).

4.1.5.PON ONU Database Sync TLV

This TLV is used to communicate the registered ONU database associated with a PON port between the active and standby OLT. This message is used to both transmit the PON ONU Database from working OLT to protect OLT and to communicate the PON ONU database status between protect OLT and working OLT.

```
0                   1                   2                   3
 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1 2 3 4 5 6 7 8 9 0 1
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|U|F|   Type=0x00XX             |    Length                     |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|                             ROID                              |
|                                                               |
|                                                               |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
|A|  Reserved   |                     OUI                       |
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-|
|                     ONU Database Entry1                       |
~                                                               ~
+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+-+
```

- U and F Bits, both are set to 0.

- Type, set to 0x00XX for "PON ONU Database Sync TLV"

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

- ROID, defined in the ROID section of [ICCP].

- A bit, Acknowledgement bit. Set to 1 if the receiver has received a PON ONU Database Sync. Otherwise, set to 0.

- Reserved, reserved for future use.

- OUI, the 3-byte [IEEE-802.3] organization unique identifier that uniquely identifies the format for describing the registered ONU database information. There are multiple PON standards and are varying implementations within a given PON standard which likely have different required information, format, etc., related to the ONU Database Entry.

- ONU Database Entry, there may be one or more ONU Database Entries transmitted in the PON ONU Database Sync TLV, each of which would describe a registered ONU. The format of the ONU Database Entry is outside the scope of this document and will be defined by the relevant PON standard organization.

5. PON ONU Database Synchronization

Without an effective mechanism to communicate the registered ONUs between the working and protection OLT, all registered ONUs would be de-registered and go through re-registration during a switchover, which would significantly increase protection time. To enable faster switchover capability, the work OLT must be able to communicate the registered ONUs associated with an ROID to the protection OLT.

The PON ONU Database Synchronization would begin once the ICCP PON Application enters OPERATIONAL state. The working OLT, the one with the working link member for the ROID, would begin transmitting the database of actively registered ONUs to the protection OLT for the same ROID. Each instance of the PON ONU Database Sync TLV describes a set of ONU Database Entries. Each ONU Database Entry would describe a registered ONU.

The transmission of PON ONU Database Descriptors for a given ROID is only unidirectional - from the working OLT to the protection OLT. The protection OLT would only be responsible for acknowledging the received message to provide a reliable database synchronization mechanism. As ONUs register and deregister from the working OLT, the working OLT would transmit PON ONU Database Synchronization TLV including only the updated ONU Database Entries.

If protected ONUs and unprotected ONUs are miscellaneously attached to the same splitter, only the protected ONUs needs to be synchronized. The specific ONUs which needs to be synchronized can be policy driven and provisioned in the management plane, or by some other signaling options.

6. Multi-chassis PON application procedures

Two typical MPLS protection network architectures for PON access are depicted in Fig.2 and Fig.3 (their PON access segments are the same as in Fig.1 and thus omitted for simplification). OLTs with MPLS functionality are connected to a single PE (Fig.2) or dual home PEs (Fig.3) respectively, i.e., the working OLT to PE1 by a working PW and the protection OLT to PE1 or PE2 by a protection PW, thus these devices constitute an MPLS network which provides PW transport services between ONUs and a CE.

```
                  +-----+
                  |     |
                  |OLT  -,
                  | A   | '.,
                  +-----+    ', PW1
                               '',
                                  '.,   +-----+          +-----+
                                     ', |     |          |     |
                                       '. PE1 ------------  CE |
                                     .''|     |          |     |
                                  ,-'   +-----+          +-----+
                                 ,'
                                .
                  +-----+    .'' PW2
                  |     | ,-'
                  |OLT  -'
                  | B   |
                  +-----+
                 Figure 2 An MPLS Network with a Single PE

                  +-----+               +-----+
                  |     |     PW1       |     |
                  |OLT  ----------------- PE1 -,
                  | A   |               |     | ',
                  +-----+               +--/--+   ',
                                           |        '.
                                           |          '. +-----+
                                           |            ''     |
                                           |             |  CE |
                                           |             .     |
                                           |           ,'+-----+
                                           |        ,-'
                  +-----+               +--\--+   ,'
                  |     |     PW2       |     | .'
                  |OLT  ----------------- PE2 -'
                  | B   |               |     |
                  +-----+               +-----+
               Figure 3 An MPLS Network with Dual-homing PEs
```

Faults may be encountered in PON access links, or in the MPLS network (including the working OLT). Procedures for these cases are described in this section (it is assumed that both OLTs and PEs are working in independent mode of PW redundancy [RFC6870]).

6.1. Protection procedure upon PON link failures

When a fault is detected on a working PON link, a working OLT MUST turn off its associated PON interface so that the protection trunk link to the protection OLT can be activated, then it MUST send an LDP fault notification message (i.e., with the status bit "Local AC (ingress) Receive Fault " being set) to its peer PE on the remote end of the PW. At the same time, the working OLT MUST send an ICCP message with PON State TLV with local PON Port State being set to notify the protection OLT of the PON fault.

Upon receiving a PON state TLV where Local PON Port state is set, a protection OLT MUST activate the protection PON link in the protection group, and advertise a notification message for the protection PW with the Preferential Forwarding status bit of active to the remote PE.

According to [RFC6870], the remote PE(s) can match the local and remote Preferential Forwarding status and select PW2 as the new active PW to which to send traffic.

6.2. Protection procedure upon PW failures

Usually MPLS networks have its own protection mechanism such as LSP protection or Fast Reroute (FRR). But in a link sparse access or aggregation network where protection for a PW is impossible in its LSP layer, the following PW layer protection procedures can be enabled.

When a fault is detected on its working PW (e.g., by VCCV BFD), a working OLT SHOULD turn off its associated PON interface and then send an ICCP message with PON State TLV with local PON Port State being set to notify the protection OLT of the PON fault.

Upon receiving a PON state TLV where Local PON Port state is set, the protection OLT MUST activate its PON interface to the protection trunk fiber. At the same time, the protection OLT MUST send a notification message for the protection PW with the Preferential Forwarding status bit of active to the remote PE, so that traffic can be switched to the protection PW.

6.3. Protection procedure upon the working OLT failure

As depicted in Fig. 2, a service is provisioned with a working PW and a protection PW, both PW terminated on PE1. If PE1 lost its connection to the working OLT, it SHOULD send a LDP notification message on the protection PW with the Request Switchover bit set.

Upon receiving a LDP notification message from its remote PE with the Request Switchover bit set, a protection OLT MUST activate its optical interface to the protection trunk fiber and activate the associated protection PW, so that traffic can be reliably switched to the protection trunk PON link and the protection PW.

In the case of Fig.3, PW-RED State TLV [ICCP] can be used by PE1 to notify PE2 the faults in all the scenarios, and PE2 operates the same as described in Section 5.1 to 5.3.

7. Security Considerations

Security considerations as described in [ICCP] apply.

8. IANA Considerations

These values are requested from the registry of "ICC RG parameter type":

```
0x00X0         PON Connect TLV
0x00X1         PON Disconnect TLV
0x00X2         PON Configuration TLV
0x00X3         PON State TLV
0x00X4         PON ONU Database Sync TLV
```

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

[RFC6870] Muley, P., Aissaoui, M., "Pseudowire Preferential Forwarding Status Bit", RFC 6870, February 2013

[ICCP] Martini, L. and et al, "Inter-Chassis Communication Protocol for L2VPN PE Redundancy", RFC 7275, June 2014

9.2. Informative References

[RFC6456] Li, H., Zheng, R., and Farrel, A., "Multi-Segment Pseudowires in Passive Optical Networks", RFC 6456, November 2011

[SEAMLESS] Leymann, N., and et al, "Seamless MPLS Architecture", draft-ietf-mpls-seamless-mpls-04, Work in progress

[G983.1] ITU-T, "Broadband optical access systems based on Passive Optical Networks (PON)", ITU-T G.983.1, January, 2005

[IEEE-1904.1] IEEE Std. 1904.1, "Standard for Service Interoperability in Ethernet Passive Optical Networks (SIEPON)", IEEE Computer Society, June, 2013

[IEEE-802] IEEE Std. 802, "IEEE Standard for Local and Metropolitan Area Networks: Overview and Architecture", IEEE Computer Society, December, 2001 with amendments

[TR-221] BBF TR-221, "Technical Specifications for MPLS in Mobile Backhaul Networks", the Broadband Forum, October, 2011

[SD-331] BBF SD-331, "Architecture and Technical Requirements for PON-Based Mobile Backhaul Networks", the Broadband Forum, Work in progress

10. Acknowledgments

The authors would like to thank Min Ye, Hongyu Li, Wei Lin, Xifeng Wan, Yannick Legoff and Shrinivas Joshi for their valuable discussions and comments.

Authors' Addresses

Yuanlong Jiang
Huawei Technologies Co., Ltd.
Bantian, Longgang district
Shenzhen 518129, China
Email: jiangyuanlong@huawei.com

Yong Luo
Huawei Technologies Co., Ltd.
Bantian, Longgang district
Shenzhen 518129, China
Email: dennis.luoyong@huawei.com

Edwin Mallette
Bright House Networks
4145 S. Falkenburg Road
Tampa, FL 33578 USA
Email: edwin.mallette@gmail.com

Chengbin Shen
China Telecom
Email: shencb@sttri.com.cn

Yimin Shen
Juniper Networks
10 Technology Park Drive
Westford, MA 01886, USA
Email: yshen@juniper.net

Weiqiang Cheng
China Mobile
No.32 Xuanwumen West Street
Beijing 100053, China
Email: chengweiqiang@chinamobile.com

Guangtao Zhou
China Unicom
No.9 Shouti South Road
Beijing 100048, China
Email: zhouguangtao@chinaunicom.cn