

Internet Engineering Task Force (IETF)
Request for Comments: 8024
Category: Standards Track
ISSN: 2070-1721

Y. Jiang, Ed.
Y. Luo
Huawei
E. Mallette, Ed.
Charter Communications
Y. Shen
Juniper Networks
W. Cheng
China Mobile
November 2016

Multi-Chassis Passive Optical Network (MC-PON) Protection in MPLS

Abstract

Multiprotocol Label Switching (MPLS) is being extended to the edge of operator networks including the network access nodes. Separately, network access nodes such as Passive Optical Network (PON) Optical Line Terminations (OLTs) have evolved to support first-mile access protection, where one or more physical OLTs provide first-mile diversity to the customer edge. Multihoming support is needed on the MPLS-enabled PON OLT to provide resiliency for provided services. This document describes the Multi-Chassis PON (MC-PON) protection architecture in MPLS and also specifies the Inter-Chassis Communication Protocol (ICCP) extension to support it.

Status of This Memo

This is an Internet Standards Track document.

This document is a product of the Internet Engineering Task Force (IETF). It represents the consensus of the IETF community. It has received public review and has been approved for publication by the Internet Engineering Steering Group (IESG). Further information on Internet Standards is available in Section 2 of RFC 7841.

Information about the current status of this document, any errata, and how to provide feedback on it may be obtained at <http://www.rfc-editor.org/info/rfc8024>.

Copyright Notice

Copyright (c) 2016 IETF Trust and the persons identified as the document authors. All rights reserved.

This document is subject to BCP 78 and the IETF Trust's Legal Provisions Relating to IETF Documents (http://trustee.ietf.org/license-info) in effect on the date of publication of this document. Please review these documents carefully, as they describe your rights and restrictions with respect to this document. Code Components extracted from this document must include Simplified BSD License text as described in Section 4.e of the Trust Legal Provisions and are provided without warranty as described in the Simplified BSD License.

Table of Contents

- 1. Introduction 3
 - 1.1. Conventions Used in This Document 5
 - 1.2. Terminology 5
- 2. ICCP Protocol Extensions 6
 - 2.1. Multi-Chassis PON Application TLVs 6
 - 2.1.1. PON Connect TLV 6
 - 2.1.2. PON Disconnect TLV 7
 - 2.1.3. PON Configuration TLV 8
 - 2.1.4. PON State TLV 9
- 3. Considerations on PON ONU Database Synchronization 9
- 4. Multi-Chassis PON Application Procedures 10
 - 4.1. Protection Procedure upon PON Link Failures 11
 - 4.2. Protection Procedure upon PW Failures 12
 - 4.3. Protection Procedure upon the Working OLT Failure 12
 - 4.4. Protection Procedure for a Dual-Homing PE 12
- 5. Security Considerations 13
- 6. IANA Considerations 13
- 7. References 14
 - 7.1. Normative References 14
 - 7.2. Informative References 14
- Acknowledgements 15
- Contributors 15
- Authors' Addresses 16

1. Introduction

Multiprotocol Label Switching (MPLS) is being extended to the edge of operator networks, as is described in the multi-segment pseudowires (PWs) with Passive Optical Network (PON) access use case [RFC6456]. Combining MPLS with Optical Line Termination (OLT) access further facilitates a low-cost, multi-service convergence.

Tens of millions of Fiber-to-the-x (FTTx) (x = H for home, P for premises, C for curb) 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 (1 Gbps or even 10 Gbps) to a dozen or more subscribers simultaneously.

In the past, access technologies such as 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 Long Term Evolution (LTE) small cells and Wi-Fi hotspots are deployed. PON is considered a viable low-cost backhaul solution for these mobile services. Besides its high bandwidth and scalability, PON further provides frequency and time-synchronization features, e.g., SyncE [G.8261] and IEEE 1588v2 [IEEE-1588] functionality, which can fulfill synchronization needs of mobile backhaul services.

The Broadband Forum specifies reference architecture for mobile backhaul networks using MPLS transport in [TR-221] where PON can be the access technology.

Unlike typical residential service where a single or handful of end-users hang off 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 a point-to-point access portfolio.

Some optical layer protection mechanisms, such as Trunk and Tree protection, are specified in [IEEE-1904.1] to avoid a single point of failure in the access. They are called Type B and Type C protection, respectively, in [G.983.1].

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 Optical Network Units (ONUs). This is different from the more complex and costly Tree 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 depicts a typical scenario of Trunk protection.

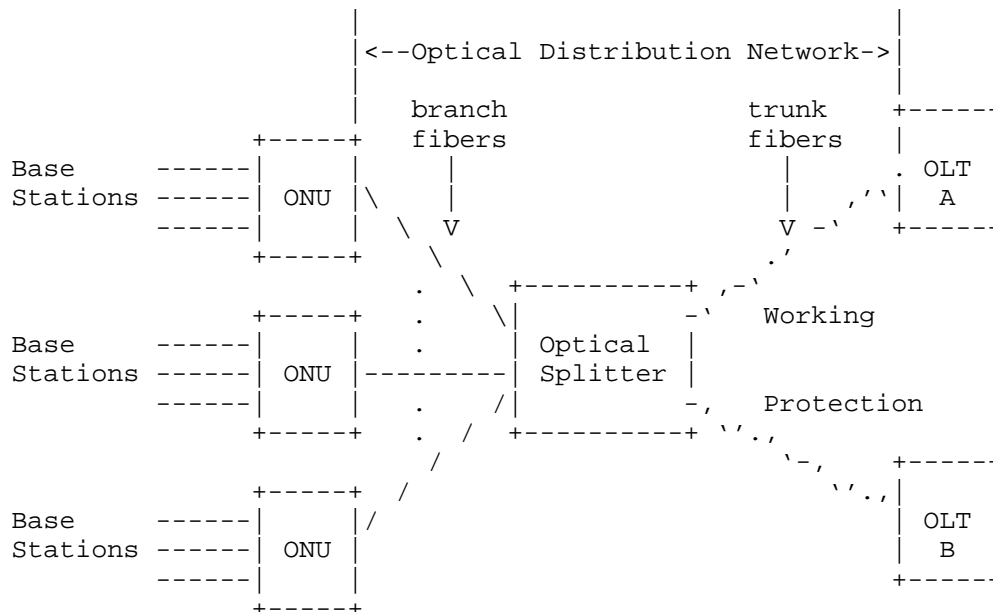


Figure 1: Trunk Protection Architecture in PON

Besides small-cell backhaul, this protection architecture can also be applicable to other services, for example, DSL and Multiple 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) or a Data Over Cable Service Interface Specification (DOCSIS) Remote Physical Layer (PHY) device [remote-phy], and it may further be attached with dozens of Customer Premises devices.

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

The PON architecture as depicted in Figure 1 can provide redundancy in its physical topology; however, all traffic, including link Operation Administration and Maintenance (OAM), is blocked on the

protection link, which frustrates end-to-end protection mechanisms such as those specified in ITU-T G.8031 [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 rapidly and reliably, especially when the OLTs also support MPLS.

ICCP [RFC7275] provides a framework for inter-chassis synchronization of state and configuration data between a set of two or more Provider Edges (PEs). Currently, ICCP only defines application-specific messages for Pseudowire Redundancy (PW-RED) and Multi-Chassis LACP (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.

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 [RFC2119].

1.2. 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: Multiprotocol Label Switching

PON: Passive Optical Network

RG: Redundancy Group

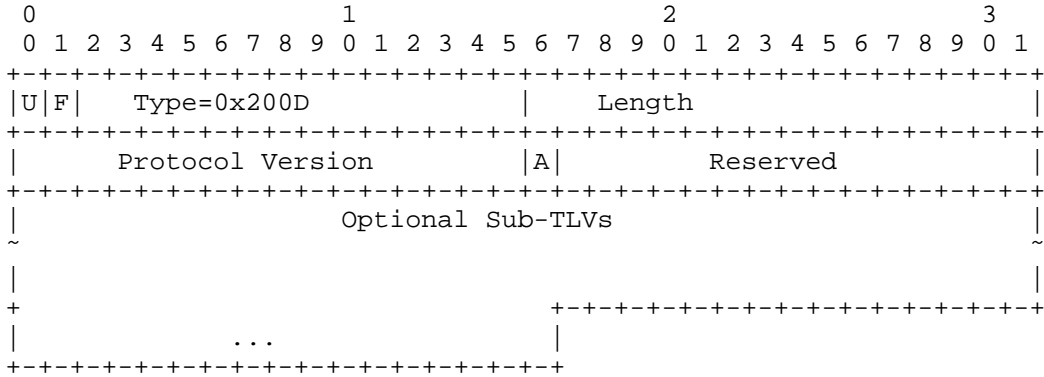
2. ICCP Protocol Extensions

2.1. Multi-Chassis PON Application TLVs

A set of MC-PON application Type-Length-Values (TLVs) are defined in the following subsections.

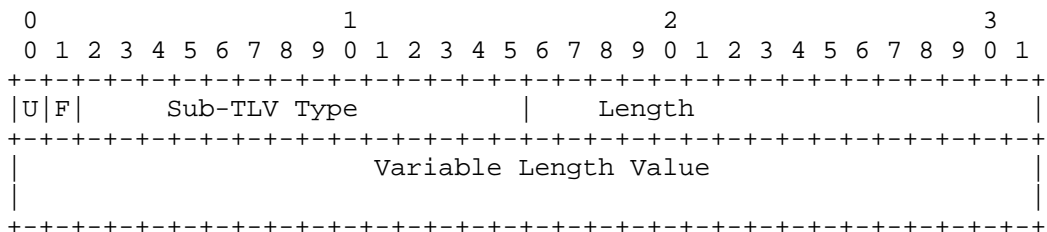
2.1.1. PON Connect TLV

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



- o U and F bits: both are set to 0.
- o Type: set to 0x200D for "PON Connect TLV".
- o Length: length of the TLV in octets excluding the U-bit, F-bit, Type, and Length fields.
- o Protocol Version: the version of this PON-specific protocol for the purposes of inter-chassis communication. This is set to 0x0001.
- o A bit: Acknowledgement bit. It MUST be set to 1 if the sender has received a PON Connect TLV from the recipient. Otherwise, set to 0.
- o Reserved: reserved for future use and MUST be set to zero.

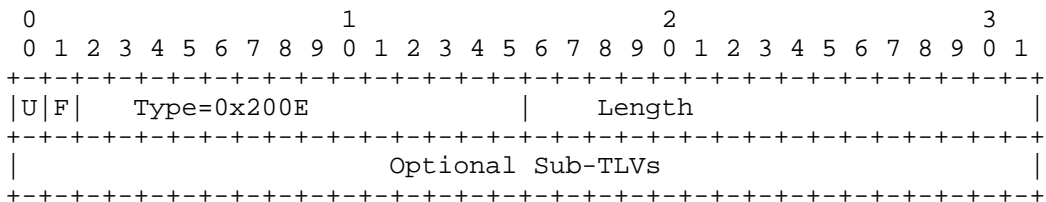
- o Optional Sub-TLVs: there are no optional Sub-TLVs defined for this version of the protocol. The structure of optional Sub-TLVs is defined as follows:



- o U bit: set to 1. The unknown Sub-TLV is silently ignored.
- o F bit: set to 0.
- o The optional Sub-TLV Type values will be allocated by IANA in a registry named "ICC RG Parameter Types" for Pseudowire Name Spaces (PWE3).
- o Length: length of the TLV in octets, excluding the U-bit, F-bit, Type, and Length fields.

2.1.2. PON Disconnect TLV

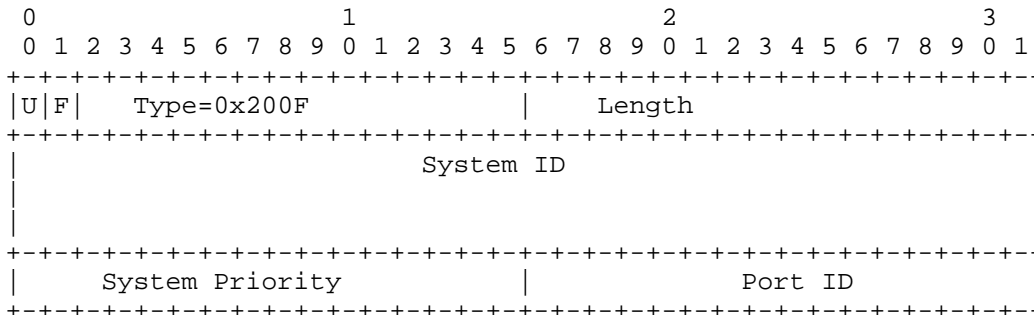
This TLV is included in the RG Disconnect message to indicate that the connection for the PON application is to be terminated.



- o U and F bits: both are set to 0.
- o Type: set to 0x200E for "PON Disconnect TLV".
- o Length: length of the TLV in octets excluding the U-bit, F-bit, Type, and Length fields.
- o Optional Sub-TLVs: there are no optional Sub-TLVs defined for this version of the protocol.

2.1.3. PON Configuration TLV

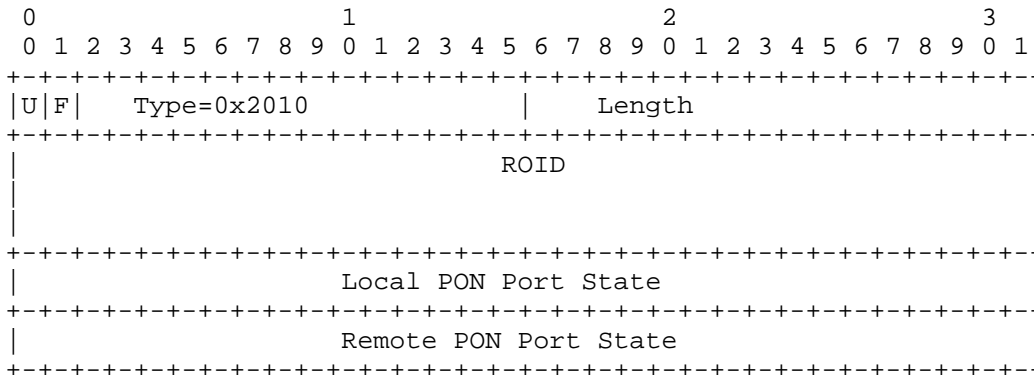
The "PON Configuration TLV" is included in the "RG Application Data" message and announces an OLT's system parameters to other members in the same RG.



- o U and F bits: both are set to 0.
- o Type: set to 0x200F for "PON Configuration TLV".
- o Length: length of the TLV in octets excluding the U-bit, F-bit, Type, and Length fields.
- o System ID: 8 octets encoding the System ID used by the OLT, which is the chassis Media Access Control (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.
- o System Priority: a 2-octet value assigned by management or administration policy; the OLT with the numerically lower value of System Priority has the higher priority.
- o Port ID: 2-octet PON Port ID.

2.1.4. 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.



- o U and F bits: both are set to 0.
- o Type: set to 0x2010 for "PON State TLV".
- o Length: length of the TLV in octets excluding the U-bit, F-bit, Type, and Length fields.
- o ROID: Redundant Object ID (ROID) as defined in Section 4.3 of [RFC7275].
- o 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; 0 - in normal).
- o Remote PON Port State: the status of the remote PON port as determined by the remote peer of the sending OLT (i.e., the sending PE). The last bit is defined as Fault indication of the PON Port associated with this PW (1 - in fault; 0 - in normal).

3. Considerations on 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 working and protection OLTs need to know about the protected ONUs. To enable service continuity, a mechanism

needs to be employed such that the operational state and significant configuration data of both the protected ONU and the services provisioned to it can be distributed to the working and protection OLT.

The specific ONU's configuration and operational data can be synchronized by some policy mechanism or provisioned in the management plane. Alternatively, said synchronization could occur by some other signaling options. Describing how to synchronize the configuration objects associated with both protected ONU as well as the services constructed to the ONU (e.g., ONU MAC address, IPv4 addresses, IPv6 addresses, VLAN identifiers, etc.) is outside of the scope of this document.

4. Multi-Chassis PON Application Procedures

Two typical MPLS protection network architectures for PON access are depicted in Figures 2 and 3 (their PON access segments are the same as in Figure 1 and thus omitted for simplification). OLTs with MPLS functionality are connected to a single PE (Figure 2) or dual-homing PEs (Figure 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 that provides PW transport services between ONUs and a Customer Edge (CE), and the PWs can provide protection for each other.

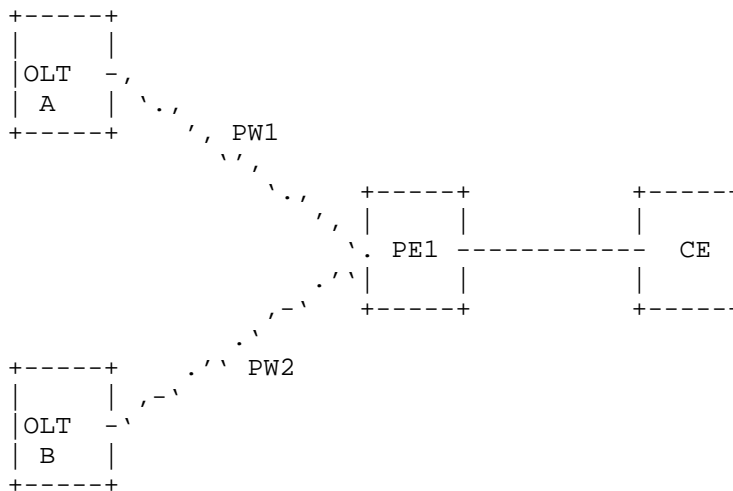


Figure 2: An MPLS Network with a Single PE

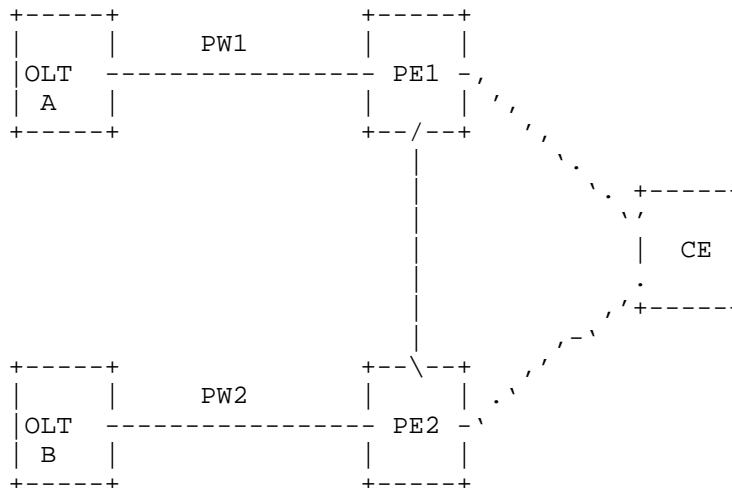


Figure 3: An MPLS Network with Dual-Homing PEs

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 the independent mode of PW redundancy [RFC6870]).

4.1. Protection Procedure upon PON Link Failures

When a fault is detected on a working PON link, a working OLT switches to the corresponding protection PON link attached with its protection OLT, i.e., the working OLT turns off its faulty PON interface so that the protection trunk link to its protection OLT can be activated. Then, the working OLT 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 over which data traffic is sent.

4.2. Protection Procedure upon PW Failures

Usually, MPLS networks have their own protection mechanism such as Label Switched Path (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 Virtual Circuit Connectivity Verification (VCCV) Bidirectional Forwarding Detection (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.

4.3. Protection Procedure upon the Working OLT Failure

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

Upon receiving an 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.

4.4. Protection Procedure for a Dual-Homing PE

In the case of Figure 3, the PW-RED State TLV as described in Section 7.1 of [RFC7275] can be used by PE1 to notify PE2 of the faults in all the scenarios, and PE2 operates the same as described in Sections 4.1 to 4.3 of this document.

5. Security Considerations

Similar to ICCP itself, this ICCP application SHOULD only be used in well-managed and highly monitored service provider PON access networks in a single administrative domain, including the implementation of rogue ONU attachment detection and mitigation via device authentication. Thus, many of the security considerations as described in [RFC7275] apply here as well.

Again, similar to ICCP, activity on the attachment circuits may cause security threats or be exploited to create denial-of-service attacks. In many passive optical networks, the optical paths between OLT and ONUs traverse publicly accessible facilities including public attachments (e.g., telephone poles), which opens up the risk of excessive link bouncing by optical layer impairment. While ICCP for MC-PON interconnects in the MPLS domain and does not traverse the PON network, risks do include introduction of a malicious ONU that could cause, for example, excessive link bouncing. This link bouncing could result in increased ICCP exchanges similar to the malicious CE case described in [RFC7275]. Operators of such networks should take additional care to restrict unauthorized ONUs and to limit the impact of link bouncing at the OLT, as these could result in service impairment.

6. IANA Considerations

IANA maintains a top-level registry called "Pseudowire Name Spaces (PWE3)". It has a subregistry called "ICC RG Parameter Types". The following values have been allocated from this subregistry:

0x200D	PON Connect TLV
0x200E	PON Disconnect TLV
0x200F	PON Configuration TLV
0x2010	PON State TLV

7. References

7.1. Normative References

- [RFC2119] Bradner, S., "Key words for use in RFCs to Indicate Requirement Levels", BCP 14, RFC 2119, DOI 10.17487/RFC2119, March 1997, <<http://www.rfc-editor.org/info/rfc2119>>.
- [RFC6870] Muley, P., Ed. and M. Aissaoui, Ed., "Pseudowire Preferential Forwarding Status Bit", RFC 6870, DOI 10.17487/RFC6870, February 2013, <<http://www.rfc-editor.org/info/rfc6870>>.
- [RFC7275] Martini, L., Salam, S., Sajassi, A., Bocci, M., Matsushima, S., and T. Nadeau, "Inter-Chassis Communication Protocol for Layer 2 Virtual Private Network (L2VPN) Provider Edge (PE) Redundancy", RFC 7275, DOI 10.17487/RFC7275, June 2014, <<http://www.rfc-editor.org/info/rfc7275>>.

7.2. Informative References

- [G.8031] International Telecommunications Union, "Ethernet Linear Protection Switching", ITU-T Recommendation G.8031, January 2015.
- [G.8261] International Telecommunications Union, "Timing and synchronization aspects in packet networks", ITU-T Recommendation G.8261, August 2013.
- [G.983.1] International Telecommunications Union, "Broadband optical access systems based on Passive Optical Networks (PON)", ITU-T Recommendation G.983.1, January 2005.
- [IEEE-1588] IEEE, "IEEE Standard for a Precision Clock Synchronization Protocol for Networked Measurement and Control Systems", IEEE Std 1588-2008, DOI 10.1109/IEEESTD.2008.4579760, July 2008.
- [IEEE-1904.1] IEEE, "Standard for Service Interoperability in Ethernet Passive Optical Networks (SIEPON)", IEEE Std 1904.1-2013, DOI 10.1109/IEEESTD.2013.6605490, June 2013.

[remote-phy]

CableLabs, "Remote PHY Specification", DCN: CM-SP-R-PHY-I05-160923, September 2016.

[RFC6456] Li, H., Zheng, R., and A. Farrel, "Multi-Segment Pseudowires in Passive Optical Networks", RFC 6456, DOI 10.17487/RFC6456, November 2011, <<http://www.rfc-editor.org/info/rfc6456>>.

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

Acknowledgements

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

Contributors

The following people made significant contributions to this document:

Chengbin Shen
China Telecom
1835 South Pudong Road
Shanghai 200122, China
Email: shencb@sttri.com.cn

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

Authors' Addresses

Yuanlong Jiang (editor)
Huawei
Bantian, Longgang district
Shenzhen 518129
China

Email: jiangyuanlong@huawei.com

Yong Luo
Huawei
Bantian, Longgang district
Shenzhen 518129
China

Email: dennis.luoyong@huawei.com

Edwin Mallette (editor)
Charter Communications
4145 S. Falkenburg Road
Tampa, FL 33578
United States of America

Email: edwin.mallette@gmail.com

Yimin Shen
Juniper Networks
10 Technology Park Drive
Westford, MA 01886
United States of America

Email: yshen@juniper.net

Weiqiang Cheng
China Mobile
No.32 Xuanwumen West Street
Beijing 100053
China

Email: chengweiqiang@chinamobile.com