

TR-301

Architecture and Requirements for Fiber to the Distribution Point

Issue: 2 Amendment 3
Issue Date: September 2024

Notice

The Broadband Forum is a non-profit corporation organized to create guidelines for broadband network system development and deployment. This Technical Report has been approved by members of the Forum. This Technical Report is subject to change. This Technical Report is owned and copyrighted by the Broadband Forum, and all rights are reserved. Portions of this Technical Report may be owned and/or copyrighted by Broadband Forum members.

Intellectual Property

Recipients of this Technical Report are requested to submit, with their comments, notification of any relevant patent claims or other intellectual property rights of which they may be aware that might be infringed by any implementation of this Technical Report, or use of any software code normatively referenced in this Technical Report, and to provide supporting documentation.

Terms of Use

1. License

Broadband Forum hereby grants you the right, without charge, on a perpetual, non-exclusive and worldwide basis, to utilize the Technical Report for the purpose of developing, making, having made, using, marketing, importing, offering to sell or license, and selling or licensing, and to otherwise distribute, products complying with the Technical Report, in all cases subject to the conditions set forth in this notice and any relevant patent and other intellectual property rights of third parties (which may include members of Broadband Forum). This license grant does not include the right to sublicense, modify or create derivative works based upon the Technical Report except to the extent this Technical Report includes text implementable in computer code, in which case your right under this License to create and modify derivative works is limited to modifying and creating derivative works of such code. For the avoidance of doubt, except as qualified by the preceding sentence, products implementing this Technical Report are not deemed to be derivative works of the Technical Report.

2. NO WARRANTIES

THIS TECHNICAL REPORT IS BEING OFFERED WITHOUT ANY WARRANTY WHATSOEVER, AND IN PARTICULAR, ANY WARRANTY OF NONINFRINGEMENT AND ANY IMPLIED WARRANTIES ARE EXPRESSLY DISCLAIMED. ANY USE OF THIS TECHNICAL REPORT SHALL BE MADE ENTIRELY AT THE USER'S OR IMPLEMENTER'S OWN RISK, AND NEITHER THE BROADBAND FORUM, NOR ANY OF ITS MEMBERS OR SUBMITTERS, SHALL HAVE ANY LIABILITY WHATSOEVER TO ANY USER, IMPLEMENTER, OR THIRD PARTY FOR ANY DAMAGES OF ANY NATURE WHATSOEVER, DIRECTLY OR INDIRECTLY, ARISING FROM THE USE OF THIS TECHNICAL REPORT, INCLUDING BUT NOT LIMITED TO, ANY CONSEQUENTIAL, SPECIAL, PUNITIVE, INCIDENTAL, AND INDIRECT DAMAGES.

3. THIRD PARTY RIGHTS

Without limiting the generality of Section 2 above, BROADBAND FORUM ASSUMES NO RESPONSIBILITY TO COMPILE, CONFIRM, UPDATE OR MAKE PUBLIC ANY THIRD PARTY ASSERTIONS OF PATENT OR OTHER INTELLECTUAL PROPERTY RIGHTS THAT MIGHT NOW OR IN THE FUTURE BE INFRINGED BY AN IMPLEMENTATION OF THE TECHNICAL REPORT IN ITS CURRENT, OR IN ANY FUTURE FORM. IF ANY SUCH RIGHTS ARE DESCRIBED ON THE TECHNICAL REPORT, BROADBAND FORUM TAKES NO POSITION AS TO THE VALIDITY OR INVALIDITY OF SUCH ASSERTIONS, OR THAT ALL SUCH ASSERTIONS THAT HAVE OR MAY BE MADE ARE SO LISTED.

All copies of this Technical Report (or any portion hereof) must include the notices, legends, and other provisions set forth on this page.

Issue History

| Issue Number | Approval Date | Publication Date | Issue Editor | Changes |
|-----------------|-------------------|-------------------|--|---|
| 1 | 24 August 2015 | 17 September 2015 | Michael Shaffer, Alcatel-Lucent Dong Wei, Huawei Technologies | Original |
| 2 | 13 March 2017 | 5 May 2017 | Michael Shaffer, Alcatel-Lucent Dong Wei, Huawei Technologies | Reverse Powering, DPU management |
| 2 Corrigendum 1 | 6 March 2019 | 6 March 2019 | Ken Ko, ADTRAN | Corrections as listed in Executive Summary |
| 2 Amendment 1 | 30 April 2020 | 30 April 2020 | Aleksandra Kozarev, Intel | Corrections as listed in Executive Summary |
| 2 Amendment 2 | 15 December 2023 | 15 December 2023 | Herman Verbueken, Nokia | Include TR-301i2a1; Add XGS-PON, NG-PON2, HSP, and 25GS-PON support |
| 2 Amendment 3 | 03 September 2024 | 04 September 2024 | Herman Verbueken, Nokia | Include TR-301i2a2; Add requirements for IP Flow Information Export (IPFIX) Protocol. |

Comments or questions about this Broadband Forum Technical Report should be directed to info@broadband-forum.org.

Editor Herman Verbueken Nokia

**Physical Layer
Transmission WA
Director** Herman Verbueken Nokia

Table of Contents

| | |
|---|----|
| Executive Summary | 7 |
| 1 Purpose and Scope | 8 |
| 1.1 Purpose | 8 |
| 1.2 Scope | 8 |
| 2 References and Terminology | 10 |
| 2.1 Conventions | 10 |
| 2.2 References | 10 |
| 2.3 Definitions and Acronyms | 12 |
| 2.4 Abbreviations | 13 |
| 3 Technical Report Impact | 16 |
| 3.1 Energy Efficiency | 16 |
| 3.2 IPv6 | 16 |
| 3.3 Security | 16 |
| 3.4 Privacy | 16 |
| 4 Deployment Use Cases | 17 |
| 4.1 Overview of FTTdp | 17 |
| 4.1.1 <i>Deployment Scenarios</i> | 17 |
| 4.1.2 <i>DPU Size and Location</i> | 17 |
| 4.1.3 <i>Powering</i> | 18 |
| 4.1.4 <i>Voice Support</i> | 19 |
| 4.1.5 <i>Hybrid DPUs</i> | 19 |
| 4.1.6 <i>Remote Copper Reconfiguration</i> | 19 |
| 4.1.7 <i>Migration and Filtering</i> | 19 |
| 4.2 Main Highlights of Operators' Use Cases | 20 |
| 4.2.1 <i>DPU Size</i> | 20 |
| 4.2.2 <i>Location</i> | 20 |
| 4.2.3 <i>Backhaul Type</i> | 21 |
| 4.2.4 <i>Customer Premises Architecture</i> | 21 |
| 4.2.5 <i>Powering</i> | 21 |
| 4.2.6 <i>DSL Co-Existence/Support</i> | 21 |
| 4.2.7 <i>POTS</i> | 21 |
| 5 FTTdp Deployment Models | 22 |
| 5.1 Model 1: Point-to-Point Ethernet/TR-167 Backhaul | 23 |
| 5.2 Model 2: TR-156 Backhaul | 24 |
| 6 Fundamental Architectural and Topological Aspects | 26 |
| 7 DPU Environmental Aspects | 27 |
| 8 DPU Powering | 28 |
| 8.1 DPU Reverse Powering Requirements | 28 |
| 8.1.1 <i>DPU Dying Gasp Message</i> | 29 |
| 8.1.2 <i>Additional DPU Reverse Powering Requirements</i> | 34 |
| 8.1.3 <i>DPUs With RCR</i> | 34 |
| 8.2 Power Source Requirements | 35 |
| 8.3 Local and Forward powering requirements | 35 |
| 9 DPU Physical Interfaces | 36 |
| 9.1 DPU Copper Drop Physical Interface Requirements | 36 |
| 9.2 DPU Backhaul Physical Interface Requirements | 37 |
| 10 Traffic Management and Quality of Service (QoS) | 39 |
| 10.1 DPU QoS Management | 39 |

| | | |
|--------|--|----|
| 10.1.1 | DPU QoS Requirements | 41 |
| 11 | VLAN Handling | 43 |
| 11.1 | Deployment Model 1 DPU VLAN Requirements | 43 |
| 11.2 | Deployment Model 2 DPU VLAN Requirements | 44 |
| 12 | Multicast | 45 |
| 12.1 | Deployment Model 1 DPU Multicast Requirements | 45 |
| 13 | Ethernet OAM | 46 |
| 13.1 | DPU OAM Requirements | 46 |
| 13.2 | CPE OAM Requirements | 46 |
| 14 | Relay Agent and Intermediate Agent Operation | 47 |
| 14.1 | RA/IA Operation Requirements | 48 |
| 14.2 | G.fast specific Type-Length-Values (TLVs) | 48 |
| 15 | Diagnostics | 50 |
| 15.1 | Performance Monitoring | 50 |
| 15.1.1 | DPU Performance Monitoring Requirements | 50 |
| 15.2 | On Demand Diagnostics | 50 |
| 16 | Network Management | 52 |
| 16.1 | DPU Management Architecture | 52 |
| 16.2 | PMA Concepts | 55 |
| 16.3 | Management of Non-Reverse Powered DPUs | 56 |
| 16.4 | DPU Management Architecture Applied to Routable and Non-Routable Address Domains 56 | |
| 16.5 | DPU-PMA Discovery and NETCONF Session Establishment | 57 |
| 16.5.1 | FTTdp Discovery Architecture | 57 |
| 16.5.2 | Discovery Messages | 61 |
| 16.5.3 | Discovery Process | 64 |
| 16.6 | Management VLAN | 65 |
| 16.7 | Management Frame Handling | 65 |
| 16.7.1 | Management Frames In DPUs With Integrated Backhauls | 65 |
| 16.7.2 | Management Frames In DPUs With Non-Integrated Backhauls | 66 |
| 16.8 | Time Management | 66 |
| 16.9 | PMA Aggregator | 66 |
| 16.9.1 | PMAA Location | 67 |
| 16.9.2 | Northbound Interface | 67 |
| 16.9.3 | PMAA Security & Scalability | 67 |
| 16.9.4 | PMAA and PMA Addressing | 67 |
| 16.10 | Software image Management | 68 |
| 17 | Operations and Maintenance | 69 |
| 17.1 | DPU Installation | 69 |
| 17.1.1 | DPU Installation Requirements | 69 |
| 17.1.2 | DPU Startup With POTS From Exchange/Cabinet | 69 |
| 17.2 | CPE Installation | 70 |
| 17.2.1 | CPE Installation Requirements | 70 |
| | Appendix A – DHCPv6 and DHCP Vendor Specific Option formatting | 71 |
| A.1 | DHCPv6 option 17 formatting | 71 |
| A.2 | DHCP (IPv4) option formatting | 72 |
| | Appendix B Certificate Management | 75 |
| B.1 | Certificate Management | 75 |

List of Figures

| | |
|---|----|
| Figure 5-1 High level FTTdp architecture and TR-156/TR-167 co-existence | 23 |
| Figure 5-2 Deployment Model 1 (Pt-to-Pt Ethernet/TR-167 Backhaul) | 24 |
| Figure 5-3 Deployment Model 2 (TR-156 Backhaul) | 25 |
| Figure 9-1 ANCP With PMA In The HON | 36 |
| Figure 9-2 ANCP With The PMA External To The HON | 36 |
| Figure 10-1 Deployment Model 1 DPU Upstream Frame Handling | 40 |
| Figure 10-2 Deployment Model 1 Downstream Frame Handling | 40 |
| Figure 10-3 Deployment Model 2 DPU Upstream Frame Handling | 41 |
| Figure 10-4 Deployment Model 2 Downstream Frame Handling | 41 |
| Figure 14-1 RA/IA in a Model 1 DPU | 47 |
| Figure 14-2 RA/IA In The HON (Model 2 DPU) | 48 |
| Figure 16-1 Model 1 DPU Management Domains | 52 |
| Figure 16-2 Model 2 DPU Management Domains | 53 |
| Figure 16-3 DPU Management Architecture | 53 |
| Figure 16-4 Management Architecture Applied to Routable and Non-Routable IP Address Domains | 57 |
| Figure 16-5 PMAs Sharing a Common Address | 68 |
| Figure 16-6 PMAs With Unique Addresses | 68 |
| Figure 17-1 DPU Startup with Exchange/Cabinet POTS | 70 |

List of Tables

| | |
|--|----|
| Table 8-1 Dying Gasp Field | 30 |
| Table 8-2 Dying Gasp Message format, version 0 | 31 |
| Table 8-3 Secure Dying Gasp message format, version 1 | 32 |
| Table 8-4 Extended dying gasp codes in the Code Extensions field | 32 |
| Table 8-5 Example Source ID and SL fields | 32 |
| Table 14-1 G.fast Sub-TLVs | 49 |
| Table 16-1 Serial number formatting examples in the DUID unique identifier field | 59 |
| Table 16-2 SN Formatting examples in IDevID subject field | 60 |
| Table 16-3 DHCPv6 and DHCP options used for DPU Discover message | 61 |
| Table 16-4 TLVs for DPU Discover | 62 |
| Table 16-5 TLVs for PMA Information | 63 |
| Table A1-1 DHCPv6 Option 17 fields | 71 |
| Table A1-2 Format of BBF-specific TLV fields for DHCPv6 Option 17 | 72 |
| Table A2-1 DHCP Option 125 fields | 73 |
| Table A2-2 BBF-specific Option-data field | 73 |
| Table A2-3 Format of BBF-specific TLV fields for DHCP Option 125 | 74 |

Executive Summary

Through the use of G.fast [9] and VDSL2 [10] over short copper loops, it has become possible to provide broadband users with data rates approaching those of fiber access technologies. This capability allows service providers to provide ultra high-speed broadband service without the need to deploy fiber into the customer premises. Since the targeted copper loop lengths are typically less than 400 meters (250 meters with Reverse Power Feed), a new node type that supports very deep deployment in the access network is required. This Technical Report defines this new node type by detailing its position(s) in the network and functional requirements. In addition, the functional requirements for reverse power feeding of this node type and its management architecture are specified.

Issue 2 of this Technical Report expands on the requirements defined in Issue 1 through the addition of further powering, copper drop interface, and DPU management details.

Corrigendum 1 of this Technical Report addresses the following issues in TR-301 Issue 2 [32]:

- The formatting of the Dying Gasp message is clarified, and a note in Section 3.3 (Security) discusses security of the unencrypted Dying Gasp message.
- The language in R-50 is corrected.
- Duplicated rows in Table 14-1 are deleted.
- R-128 is deleted.
- R-134 is corrected.
- The requirements associated with keep-alive messages between the DPU and the PMA in Section 16.1 are updated.
- The requirements associated with formatting for the DUID Unique Identifier in Section 16.5.1.2 are clarified and updated.
- The requirements associated with formatting for the IDevID certificate's subject field in Section 16.5.1.3 are clarified and updated.

Amendment 1 to Issue 2 of this Technical Report specifies a new version of the DPU Dying Gasp message that adds protections against certain security vulnerabilities, includes the identity of the message source, and adds new codes for DPUs with local or forward powering.

- R-19 and R-22 are deleted.
- R-24 is updated.
- R-250...R-274 are added.
- The requirements associated with DPU Powering are clarified and updated.
- The requirements associated with the PMA Offer message are updated.

Amendment 2 to Issue 2 of this Technical Report broadens the applicability of TR-301 to include XGS-PON, NG-PON2, HSP, and 25GS-PON support.

[Amendment 3 to Issue 2 of this Technical Report adds requirements R-275 and R-276 for IP Flow Information Export \(IPFIX\) Protocol.](#)

1 Purpose and Scope

1.1 Purpose

This Technical Report provides the architectural basis and technical requirements that are needed to deploy FTTdp within a TR-101 and/or TR-178 architecture. To this end, a new node type, the DPU, is defined. This node, typically positioned at the Distribution Point (DP), supports one or more high-speed copper drops into the customer premises and uses a gigabit (or faster) fiber link to backhaul user data to a Higher Order Node (HON). A key aspect of the new node type is the ability for it to be reverse power fed from one or more copper drop pairs. To Reverse Power Feed (RPF), there needs to be power supply functionality at the customer premises, the requirements for which are also defined here.

1.2 Scope

This Technical Report defines the Distribution Point Unit (DPU), for use within the access network. All aspects of the introduction of the DPU into the network are considered and requirements are specified for the DPU and all affected nodes in the access network along with the RPF functionality.

The requirements in this Technical Report are defined within the framework of both the TR-101 [1] and TR-178 [6] architectures. The TR-101 architecture is required to support near term residential deployments while the TR-178 architecture is required for the evolution to a multi-service edge network that supports residential, business, and wholesale deployments.

The DPU supports a number of deployment scenarios ranging from complex multiport units to simplified multiport and single port DPUs. Complex DPUs are considered those that support the full set of functions for an Access Node (AN) described in TR-101 and/or TR-178. Simple DPUs are those that support a reduced set of functions relative to these. The focus of this Technical Report is the simple multiport and single port DPUs.

The DPU is differentiated from the Multi-Dwelling Unit (MDU) and Single-Family Unit (SFU) devices already deployed in today's service provider networks by the following characteristics:

- DPUs may be reverse powered over the copper drop interface.
- DPUs support the provisioning of services with zero manual intervention by the service provider.

Reverse power feeding of the DPU, and the loss of powering if all customers on a DPU turn off their Power Source Equipment (PSE), gives rise to management continuity issues. This Technical Report addresses these by specifying a management architecture that utilizes a management function that may reside at any location in the network that has reliable access to power; this is known as the Persistent Management Agent (PMA).

The following technologies on the DPU uplink side are covered in this Technical Report:

- Point-to-point fiber (e.g., IEEE802.3-2012 GbE, 10GbE [13])
- ITU-T PON Technologies (e.g., ITU-T Gigabit Passive Optical Network (GPON) [11], XG-PON [12], XGS-PON [33], NG-PON2 [34], HSP [35] and 25GS-PON [36])

Note: Wherever the term GPON is used within this Technical Report, it is representative of any ITU-T PON technology and 25GS-PON and is not intended to be restricted to only G.984 compliant systems.

The following copper drop technologies are covered in this Technical Report:

- G.fast [9]
- VDSL2 [10]

A DPU may be deployed as part of a fiber rollout in an existing copper infrastructure area. While it is assumed that Plain Old Telephone Service (POTS) and Any Digital Subscriber Line Service (xDSL)

service from the central office or remote cabinet are not used by a user after activation of their FTTdp service, strategies for the coexistence with legacy services are considered. Specifically, the coexistence of CO/exchange or cabinet supplied POTS and xDSL on neighboring lines and the interaction with these services during FTTdp service establishment are addressed. In the case of a DPU that can be configured to support either G.fast or VDSL2, the transmission technology needs to be spectrally compatible with all services delivered via that DP.

2 References and Terminology

2.1 Conventions

In this Technical Report, several words are used to signify the requirements of the specification. These words are always capitalized. More information can be found be in RFC 2119 [2].

| | |
|-------------------|---|
| MUST | This word, or the term “REQUIRED”, means that the definition is an absolute requirement of the specification. |
| MUST NOT | This phrase means that the definition is an absolute prohibition of the specification. |
| SHOULD | This word, or the term “RECOMMENDED”, means that there could exist valid reasons in particular circumstances to ignore this item, but the full implications need to be understood and carefully weighed before choosing a different course. |
| SHOULD NOT | This phrase, or the phrase "NOT RECOMMENDED" means that there could exist valid reasons in particular circumstances when the particular behavior is acceptable or even useful, but the full implications need to be understood and the case carefully weighed before implementing any behavior described with this label. |
| MAY | This word, or the term “OPTIONAL”, means that this item is one of an allowed set of alternatives. An implementation that does not include this option MUST be prepared to inter-operate with another implementation that does include the option. |

2.2 References

The following references are of relevance to this Technical Report. At the time of publication, the editions indicated were valid. All references are subject to revision; users of this Technical Report are therefore encouraged to investigate the possibility of applying the most recent edition of the references listed below.

A list of currently valid Broadband Forum Technical Reports is published at www.broadband-forum.org.

| Document | Title | Source | Year |
|-------------------------|---|--------|------|
| [1] TR-101 Issue 2 | Migration to Ethernet-Based Broadband Aggregation | BBF | 2011 |
| [2] RFC 2119 | Key words for use in RFCs to Indicate Requirement Levels | IETF | 1997 |
| [3] TS 101 548-1 v2.3.1 | European Requirements for Reverse Powering of Remote Access Equipment | ETSI | 2019 |
| [4] TR-156 Issue 4 | Using GPON Access in the context of TR-101 | BBF | 2012 |
| [5] TR-167 Issue 3 | GPON-fed TR-101 Ethernet Access Node | BBF | 2010 |
| [6] TR-178 Issue 2 | Multi-service Broadband Network Architecture and Nodal Requirements | BBF | 2014 |
| [7] RFC 6241 | Network Configuration Protocol (NETCONF) | IETF | 2011 |

| | | | | |
|------|--------------|---|-------|------|
| [8] | RFC 7950 | The YANG 1.1 Data Modelling Language | IETF | 2016 |
| [9] | G.9701 | Fast access to user terminals (G.fast) - Physical layer specification | ITU-T | 2019 |
| [10] | G.993.2 | Very high speed digital user line transceivers 2 (VDSL2) | ITU-T | 2019 |
| [11] | G.984 Series | Gigabit-capable Passive Optical Networks | ITU-T | 2014 |
| [12] | G.987 Series | 10- Gigabit-capable Passive Optical Networks | ITU-T | 2014 |
| [13] | IEEE 802.3 | IEEE Standard for Ethernet | IEEE | 2012 |
| [14] | RFC 6320 | Protocol for Access Node Control Mechanism in Broadband Networks | IETF | 2011 |
| [15] | TR-147 | Layer 2 Control Mechanism For Broadband Multi-Service Architectures | BBF | 2008 |
| [16] | RFC 3376 | Internet Group Management Protocol V3 | IETF | 2002 |
| [17] | RFC 5280 | Internet X.509 Public Key Infrastructure Certificate and Certificate Revocation List (CRL) Profile | IETF | 2008 |
| [18] | RFC 7589 | Using the NETCONF Protocol over Transport Layer Security (TLS) with Mutual X.509 Authentication | IETF | 2015 |
| [19] | IEEE 802.1AR | Secure Device Identity | IEEE | 2009 |
| [20] | RFC 3315 | Dynamic Host Configuration Protocol for IPv6 (DHCPv6) | IETF | 2003 |
| [21] | RFC 6520 | Transport Layer Security (TLS) and Datagram Transport Layer Security (DTLS) Heartbeat Extension | IETF | 2012 |
| [22] | RFC 4604 | Using Internet Group Management Protocol Version 3 (IGMPv3) and Multicast Listener Discovery Protocol Version 2 (MLDv2) for Source-Specific Multicast | IETF | 2006 |
| [23] | RFC 8071 | NETCONF Call Home and RESTCONF Call Home | IETF | 2015 |
| [24] | RFC 1035 | Domain Names – Implementation and Specification | IETF | 1987 |
| [25] | RFC 1122 | Requirements for Internet Hosts – Communications Layers | IETF | 1989 |
| [26] | G.987.3 Amd2 | 10-Gigabit-capable passive optical networks (XG-PON): Transmission convergence (TC) layer specification | ITU-T | 2021 |
| [27] | ATIS-0300220 | Structure for the Representation of the Communications Industry Manufacturers, Suppliers, and Related Service Companies for Information Exchange | ATIS | 2016 |
| [28] | RFC 5277 | NETCONF Event Notifications | IETF | 2008 |
| [29] | RFC 2104 | HMAC: Keyed-Hashing for Message Authentication | IETF | 1997 |

| | | | | |
|------|--------------------------|---|----------------------|----------------------|
| [30] | PUB 180-4 | Secure Hash Standard (SHS) | FIPS | 2015 |
| [31] | RFC 5705 | Keying Material Exporters for Transport Layer Security (TLS) | IETF | 2010 |
| [32] | TR-301 Issue 2 | Architecture and Requirements for Fiber to the Distribution Point | BBF | 2017 |
| [33] | G.9807.1 | 10-Gigabit-capable symmetric passive optical network (XGS-PON) | ITU-T | 2023 |
| [34] | G.989.1 Amd1 | 40-Gigabit-capable passive optical networks (NG-PON2): General requirements | ITU-T | 2015 |
| [35] | G.9804.1 Amd1 | Higher speed passive optical networks – Requirements | ITU-T | 2021 |
| [36] | 25GS-PON | 25 Gigabit Symmetric Passive Optical Network (25GS-PON / 25G TDM PON) Specification | 25GS-PON MSA | 2022 |
| [37] | TR-384 | Cloud Central Office (CloudCO) Reference Architectural Framework | BBF | 2018 |
| [38] | TR-413 | SDN Management and Control Interfaces for CloudCO Network Functions | BBF | 2024 |
| [39] | RFC 7011 | Specification of the IP Flow Information Export (IPFIX) Protocol for the Exchange of Flow Information | IETF | 2013 |
| [40] | RFC 7012 | Information Model for IP Flow Information Export (IPFIX) | IETF | 2013 |

2.3 Definitions and Acronyms

The following terminology is used in this Technical Report.

| | |
|-----------------------------------|--|
| Certificate Path | A sequence of one or more certificates, with each certificate forming a link in a chain from the relying party's trust anchor to the subscriber's public key. |
| Certificate Validity Check | Certificate validity checks that each certificate is a valid certificate. It logically consists of a set of steps, some being optional depending on the application needs. The steps are Syntax Check, Integrity Check, Valid Period Check, Revocation Check, KeyUsage Check |
| Complex DPU | A DPU that complies with all the requirements for Ethernet Access Nodes found in TR-101 and/or TR-178. |
| Copper | Copper as used in TR-301 Issue 2 [32] is defined as any wire for telephone, Cable TV or Ethernet that may be used for point to point communications. It includes but is not limited to the following: Unshielded Twisted Pair (UTP) cables that includes legacy telephony wiring, Category 3, Category 5/5e and Category 6 wiring; Coaxial cable. |
| CPE | Customer Premises Equipment. |
| Device Certificate | In the context of TR-301 Issue 2 [32], the device certificate is an 802.1AR IDevID certificate. It is issued to uniquely and irrevocably identify the device manufactured. |
| DP | Distribution Point. The location in the access network where the multi-pair copper cables from the central office connect to the final copper drops into the customers' premises. |

| | |
|----------------------------|---|
| DPU | Distribution Point Unit. The node that typically resides at the DP in the Fiber To The Distribution Point architecture. |
| FTTdp | Fiber To The distribution point. An access network architecture that uses fiber to the DP to provide very high-speed digital subscriber line services. |
| HON | Higher Order Node. The first node upstream of the DPU. |
| Hybrid DPU | A DPU that supports both G.fast and VDSL2. |
| IDevID | The Secure Device Identifier installed on the device by the manufacturer. The IDevID is cryptographically bound to the device and is composed of the Secure Device Identifier Secret and the Secure Device Identifier Credential. See IEEE 802.1AR [19] |
| IDevID Certificate | An X.509 certificate forming the Secure Device Identifier Credential portion of the IDevID, as defined in IEEE 802.1AR [19]. |
| Path Validation | Path validation ensures that the certification path is properly constructed. |
| PFFF | Port Frame Forwarding Function. The functional component of a DPU that is responsible for the processing of user frames but is not part of the DPU backhaul. |
| Pinned Certificate | Pinning is the process of associating a host with its expected X509 certificate or public key. |
| PMA | Persistent Management Agent. A management proxy for the DPU that caches provisioning and last known status information for the DPU. |
| PSE | Power Source Equipment is the equipment at the customer premises that provides power to the DPU in a reverse power fed deployment. |
| RCR | Remote Copper Reconfiguration functionality allows the copper loop to be reconfigured such that it is physically disconnected from the incoming CO/cabinet copper lines and connected to the DPU, without a site visit. |
| Root CA Certificate | An unsigned or a self-signed public key certificate that identifies the root Certificate Authority (CA). It may also be used as a trust anchor. |
| RPF | Reverse Power Feed is the collective term used to describe the provision of power to the DPU from the customer premises. |
| Simple DPU | A DPU that complies with the requirements in this Technical Report. |
| Trust Anchor | An authoritative entity for which trust is assumed and not derived, such as an X.509 root CA certificate, intermediate certificate, or pinned certificate. |
| WiFi | A generic name used to refer to all versions of IEEE 802.11 |
| X.509 Certificate | A set of information that uniquely and securely identifies a key pair (public and private key) and an owner that is authorized to use the key pair. X.509 certificate format and usage is defined in RFC 5280. |

2.4 Abbreviations

This Technical Report uses the following abbreviations:

| | |
|-------|------------------------------|
| 10GbE | 10 Gigabit Ethernet |
| AC | Alternating Current |
| AFE | Analog Front End |
| AN | Access Node |
| ANCP | Access Node Control Protocol |
| ATA | Analog Terminal Adapter |
| BNG | Broadband Network Gateway |
| CA | Certificate Authority |
| CC | Continuity Check |

| | |
|--------|---|
| CO | Central Office |
| DC | Direct Current |
| DECT | Digital Enhanced Cordless Telecommunications |
| DSE | Disorderly Shutdown Event |
| DSLAM | Digital Subscriber Line Access Multiplexer |
| GDR | Gamma Data Rate |
| G.fast | Fast Access To Subscriber Terminals |
| FTTC | Fiber To The Cabinet |
| FTTP | Fiber To The Premises |
| FTU-O | G.fast Transceiver Unit - Office |
| FTU-R | G.fast Transceiver Unit - Remote |
| GbE | Gigabit Ethernet |
| GEM | GPON Encapsulation Method |
| GMP | Group Management Protocol |
| GPON | Gigabit Passive Optical Network |
| HSP | Higher Speed Passive Optical Networks – as defined in ITU-T Recommendation G.9804 series and their amendments |
| IA | Intermediate Agent |
| IGMP | Internet Group Management Protocol |
| IP | Internet Protocol |
| IPv4 | Internet Protocol version 4 |
| IPv6 | Internet Protocol version 6 |
| L2 | Layer 2 |
| LSB | Least Significant Bit |
| MAC | Media Access Control |
| MDSU | Metallic Detection Start Up |
| MDU | Multi-Dwelling Unit |
| MEP | Maintenance association End Point |
| MIP | Maintenance domain Intermediate Point |
| MSB | Most Significant Bit |
| NDR | Net Data Rate |
| NMS | Network Management System |
| NT | Network Termination |
| OAM | Operations Administration and Maintenance |
| ODN | Optical Distribution Network |
| OLT | Optical Line Termination |
| ONT | Optical Network Termination |
| OSS | Operations Support Systems |
| PMA | Persistent Management Agent |
| PON | Passive Optical Network |
| POTS | Plain Old Telephone Service |
| PRP | POTS Remote Copper Reconfiguration Protocol |
| QLN | Quiet Line Noise |
| QoS | Quality of Service |
| RA | Relay Agent |
| RCR | Remote Copper Reconfiguration |
| RFI | Radio Frequency Interference |

| | |
|--------|---|
| RG | Residential Gateway |
| SFU | Single Family Unit |
| SFP | Small Form-factor Pluggable Transceiver |
| TCP | Transmission Control Protocol |
| TLV | Type Length Value |
| TR | Technical Report |
| UC | Use Case |
| UDP | User Datagram Protocol |
| VDSL | Very High Bitrate Digital Subscriber Line |
| VLAN | Virtual Local Area Network |
| VoIP | Voice over IP |
| VSSN | Vendor-Specific Serial Number |
| WA | Work Area |
| xDSL | Any Digital Subscriber Line Service |
| XG-PON | 10 Gigabit Passive Optical Network |

3 Technical Report Impact

3.1 Energy Efficiency

This Technical Report has a minor impact on energy efficiency. A DPU can be regarded as a subtended node, which means that there will be a parent node, and so 2 active access elements per line rather than 1. However, the parent node will be sharing the backhaul interface between a number of DPU lines, which means the resulting increase per line will be small. Note that in the case of a PON-fed DPU, the parent node would be there anyway, but the DPU is in addition to the NT Module in the customer's premises, and so still counts as an additional active node. However this power increase will be offset by the emphasis on keeping the power consumption of the DPU itself as low as possible. This is done by design (keeping the DPU functionality to the absolute minimum), and the use of Low Power Modes and G.fast Discontinuous Operation. Limiting the power needed by the DPU facilitates reverse power feeding, and aids long-term reliability. Therefore any increase in the total power of the system should be small. Note however that there will also be some minor energy losses in the copper wire from the reverse or forward power feed.

3.2 IPv6

The DPU mainly operates at L2 and below and so is transparent to IPv4/v6 in the data plane. However, the DPU does have an IP address for management purposes. This can either be v4 or v6, but given the large number of DPUs and the need to avoid this being a public IP address (for security reasons), there is a case for it being a link-local IPv6 address.

3.3 Security

DPUs are typically located on pole-tops, underground chambers (footway boxes), building basements, rising mains, and pedestals. They are not usually in secure, locked enclosures. Unauthorized physical access to a pole-mounted device is unlikely, but an underground chamber and building basement are more vulnerable to intrusion. Where the DPU is a sealed for life unit, tapping into it would be difficult, but there is a general requirement that DPUs must not have any exposed, enabled craft ports.

It is recognized that reception and processing of a non-validated or non-authenticated Dying Gasp message creates a security vulnerability that may not be acceptable on some management networks, and that as a result it may not be permitted to reach its intended destination. This issue of the Technical Report defines a secure version of the message that includes authentication of the source and validation of the message content.

3.4 Privacy

This Technical Report has no impact on privacy.

4 Deployment Use Cases

4.1 Overview of FTTdp

This introduction summarizes the main features of FTTdp and the various deployment and migration options in order to put those key points in context.

The main objective of FTTdp is to provide much higher data rates than cabinet-based Very High Bitrate Digital Subscriber Line (VDSL) over the final part of the existing copper connection to the customer. Locating a new, high-speed access node at the DP and reusing the existing copper drops has several advantages over Fiber To The Premises (FTTP), namely:

- It avoids the need to install new infrastructure into and around the home, i.e., there is no need to install a new fiber cable between the DP and the home, or drill a hole in an external wall to take the fiber into the home, or install fiber between the entrance point and the Optical Network Termination (ONT).
- It allows customer self-install, which removes the need for a visit to the customer premises with its attendant cost, time, and logistical downsides.
- It reduces the time between receiving and being able to fulfill a customer order.

As part of the initial FTTdp architectural considerations, a large number of Use Cases were brought forward by different operators. These were mainly responsible for the detailed functional requirements in this Technical Report, but the Use Cases themselves have not been included in the published document. The key points arising from a Use Case analysis are covered in Section 4.2.

4.1.1 Deployment Scenarios

FTTdp can be used in one or more of the following deployment scenarios:

- A higher speed overlay in a region already served by VDSL or cable.
- High-speed services to users who are directly connected to the CO/exchange, i.e., with no intervening cabinet, and who are therefore unable to get Fiber To The Cabinet (FTTC)/VDSL.
- High speed services to customers who are in an FTTC deployment area, but their cabinet is too small to commercially justify a full Digital Subscriber Line Access Multiplexer (DSLAM).
- High-speed services to customers who are served from a VDSL enabled cabinet, but have a long drop side connection and so get a fairly low data rate.
- Making the final customer connection to a PON infrastructure over a copper tail. This is particularly relevant where the final drop is direct buried, i.e., not ducted, where the customer is unwilling to accept the disruption caused by installing fiber into and around their home, or where there is a wish to provide a self-install version of a 'fiber-rate' service.
- High-speed service distribution within an MDU.

4.1.2 DPU Size and Location

Since the DPU is typically located deeper in the network than the cabinet (in order to achieve higher speeds), then by definition it will have a smaller number of lines than a cabinet. One of the main locations is expected to be the copper distribution point (DP), as this is the closest point to the customer where there is an existing flexibility point. There are a range of DP sizes. Note that in many cases, not all the physical connections at a DP will be used.

Analysis of the Use Cases revealed the need for:

- a) Single line DPUs
- b) Small multiline DPUs, 4-16 lines
- c) Larger multiline DPUs, 17--48 lines

All the above are still 'simple' DPUs according to the definition in this document.

There are 2 locations where larger (subtype 'c' above) DPUs may be appropriate. The first is the basement of MDUs. The second is a new location. Although the initial focus was on G.fast deployment at the existing copper DP, further analysis and the performance of initial G.fast equipment has led to a view that, in some geographies, DPUs could be located somewhat further back in the network, namely between the DP and the cabinet. There would obviously be some performance reduction as a result, but data rates much greater than those of VDSL2 could still be provided. One of these larger DPUs would serve several (copper) DPs. This means that fewer DPUs would be needed and it would only be necessary to install one backhaul fiber for this 'cluster', rather than requiring one to each DP. These 2 factors should reduce the installation cost, and allow deployment of an initial critical mass of DPUs to be achieved more quickly. Note that this does not preclude, and in many cases would also involve, subtype 'a' and/or 'b' DPUs in the same network.

In some scenarios, subtype 'b' DPUs are located on the top of poles or in small underground chambers (sometimes known as footway boxes). There is no protective enclosure (such as a cabinet) at these locations and so these DPUs need to be environmentally sealed with the appropriate thermal and weather resistant properties. This is also likely to be the case for pedestal deployments.

Deployment indoors may be less environmentally challenging, but the requirements for passive cooling and security still apply.

Finally, it was also recognized that there are scenarios, for example very large MDUs, where significantly more than 48 lines might be needed. However in this case it was agreed that it would be more appropriate to use an access node as defined in TR-101/TR-178, but with G.fast line cards; these are known as 'complex' DPUs, and are out of scope of this Technical Report.

4.1.3 Powering

DPUs may be powered in one of three ways:

- 1) Reverse power, where the DPU draws its power from the customer premises via the copper lines between those premises and the DPU. The reverse power feed capacity and DPU power consumption need to be such that the DPU can be fully operational when only a single customer is connected. When there is more than one active customer on a given DPU, the DPU draws roughly equal power from each line. Any back-up battery would be located in the customer premises.
- 2) Forward power, where the DPU draws its power from a network power node which typically powers multiple DPUs via one or more copper lines between the power node and the DPUs. This may be a newly installed power cable (put in with the fiber feed), or might re-use existing spare copper pairs. In this case, any back-up battery would be located at the network power node.
- 3) Local power, where the DPU draws its power from a local Alternating Current (AC) mains source. In this case, any back-up battery would be located near the DPU.

The best method to power a DPU depends on several factors:

- 1) For smaller DPUs, reverse powering might be appropriate.
- 2) When copper backhaul to a nearby network power node is available, forward powering might be viable.
- 3) When local AC mains power is already available, local powering might be viable.

4.1.4 Voice Support

There is no requirement to support baseband voice from the CO/exchange on the same pair that is providing service from a reverse powered DPU. To do so would be very difficult because of the conflict between the RPF, and the forward Direct Current (DC) voltage feed and DC signaling. Instead, the voice service could be delivered as a derived service (Voice over IP [VoIP]), terminating in an Analog Terminal Adapter (ATA), Digital Enhanced Cordless Telecommunications (DECT) base-station, or transported to a Smartphone over WiFi. In the case of an ATA, there may be a requirement to re-inject the voice onto the in-premises wiring so that existing analogue phones can still be used. If the reverse power feed runs over the same wiring, this would require a signaling conversion dongle to be attached to each phone, and the RPF power source to have certain safety features. Being able to detect a (off-hook) phone with a missing dongle would be a particular need, as this is a potential safety risk.

If there is a requirement to provide a lifeline capability for the derived voice service, then the RPF needs to have battery backup so that the DPU can continue to be powered for the required time during a mains power failure. Some of the G.fast/VDSL2 low-power modes are specified so as to reduce the power consumed to a minimum during battery backup, as this extends the time for which battery backup operation can last.

In the absence of reverse powering, then baseband analogue voice from the CO/exchange or cabinet can continue to be offered, if the operator so chooses, but this has an impact on Remote Copper Reconfiguration (RCR) and the band filtering needed in the DPU.

4.1.5 Hybrid DPUs

The initial focus of FTTdp was on a pure G.fast-based DPU. However, some use cases included the DPU being able to offer both G.fast and VDSL2. One reason for these G.fast/VDSL2 hybrid DPUs is to use FTTdp to offload VDSL2 customers from a cabinet that has run out of VDSL2 ports; this could provide the same service from the DP without needing to change out the CPE. G.fast could then offer an upsell opportunity. Another application would be to operate individual, very long lines from the DPU with VDSL2 instead of G.fast.

The architecture supports hybrid DPUs, but does not specify detailed requirements as to how these should be implemented. In particular, a vendor could either use a dual mode chip, or have completely separate G.fast and VDSL2 modules that just happened to share the same box. There are, however, implications such as Analog Front End (AFE), and filtering associated with this choice.

4.1.6 Remote Copper Reconfiguration

After initial installation of the DPU, connecting a customer to a DPU-based service should not need a visit to, or applying jumpers at, the DPU. Disconnection from CO-based or cabinet-based services and connection to DPU-based services is done remotely under management control and/or the detection of a reverse power feed. Note also that the DPU may need to continue to transparently support legacy voice, CO/exchange-based ADSL, and cabinet-based VDSL on lines which pass through the DPU but are not taking a DPU-based service. Finally, there is the need to be able to remotely reconnect any line to a non-DPU-based service.

4.1.7 Migration and Filtering

Nearly all G.fast deployment scenarios will need to take into account the installed base of VDSL, from the point of view of both co-existence and migration. There are a number of possible migration paths, which may or may not involve serving VDSL2, in addition to G.fast, from a hybrid DPU.

The main migration scenarios are:

- A. Deploying G.fast only DPUs, and then upselling VDSL2 cabinet customers to a G.fast, DPU-based service.
- B. Offloading some VDSL2 cabinet customers to a similar VDSL2-based service from a hybrid DPU, and then upselling them to a G.fast, DPU-based service.
- C. Moving all VDSL2 cabinet customers to the same VDSL2-based service from a hybrid DPU, and then upselling them to a G.fast, DPU-based service.
- D. Moving all VDSL2 customers to a G.fast only DPU, and then offering a somewhat better, G.fast delivered, VDSL2-like service with upsell to higher rate G.fast services.

The spectral co-existence requirements for these different scenarios are as follows. In all cases the G.fast is only located at the DP.

1. In case A, G.fast must be spectrally compatible with VDSL2 from the cabinet.
2. In case B, G.fast must be spectrally compatible with VDSL2 from the same DPU. If there was also still VDSL2 present on the same cable from the cabinet, then there would be very significant disruption of the cabinet VDSL, if they used the same spectrum. One way to avoid this is by ensuring there are no cabinet VDSL customers left on the Copper to that DPU. Alternatively, the cabinet VDSL could continue to use the spectrum up to 17 MHz, with the DPU VDSL using 17-30 MHz, and G.fast starting at 30 MHz. In practice these frequency ranges would not be contiguous of course, requiring significant guard bands (depending on the quality of filtering). The latter approach would significantly decrease the G.fast capacity, making upsell more difficult.
3. In case C, G.fast must be spectrally compatible with VDSL2 from the same DPU. There would be no point in increasing the VDSL2 spectrum up to 30 MHz. The VDSL2 performance with the 17MHz profile would be better anyway because of the much shorter reach; increasing the VDSL2 spectrum would improve this performance still further, but at the expense of reducing the G.fast capacity significantly, making upsell much harder.
4. Case D would allow the entire VDSL spectrum (above 2.2 MHz) to be used for G.fast.

Given the above considerations, the use of any VDSL profile above 17 MHz in an FTTdp deployment area is strongly discouraged.

4.2 Main Highlights of Operators' Use Cases

This section summarizes the main common features of the Use Cases (UCs) submitted by Operators to guide the development of this Technical Report; the Use Cases themselves have not been included in the published document. This is not a comprehensive summary but highlights points of convergence in the FTTdp deployment needs of the Operators.

4.2.1 DPU Size

DPU sizes of 8 and 16 ports are those most required. Smaller sizes are also represented as well as larger units of ~48 ports, but to a lesser extent.

4.2.2 Location

Both outdoor (pole, underground) and indoor (basement, floor) locations are needed. This has clear implications on the environmental class and dissipation constraints of the solutions suitable for these different types of environment.

4.2.3 Backhaul Type

Operator Use Cases indicated a fairly even split between point-to-point (GbE/10GbE) and point-to-multipoint (e.g., GPON) types of backhaul.

4.2.4 Customer Premises Architecture

The most required option for the customer premises is to have the G.fast NT Module integrated with the Residential Gateway (RG).

The RPF PSE may be embedded in the RG, or it may be external.

Customer self-install is required in almost all the UCs.

4.2.5 Powering

Reverse Power Feed (RPF) is required in almost all the UCs, but there are also Use Cases that need forward and local powering.

4.2.6 DSL Co-Existence/Support

Coexistence with ADSL systems deployed at the CO is required in all UCs.

Coexistence with VDSL2 systems deployed at the Cabinet or at the DP is required in a fair number of UCs.

4.2.7 POTS

Coexistence with VoIP services re-injected as baseband analogue signals on the in-home wiring is required for some Use Cases, but this Technical Report does not require the support of POTS from the CO/exchange or cabinet in conjunction with RPF.

5 FTTdp Deployment Models

FTTdp is typically deployed nearer to the end-user than FTTC. This can result in a very large number of active network nodes that need to be installed, provisioned, powered, and managed. Therefore, the goal is to keep the DPU as simple as possible to minimize power consumption (in particular to facilitate reverse powering) and ease the problem of management scale. It is also recognized that DPUs should be able to be incorporated into whatever fiber backhaul infrastructure network providers have already installed, in particular both point-to-point fiber and PON backhaul need to be supported.

These high-level business needs led to the following architectural principles:

- DPUs are connected to a HON that provides aggregation and those access functions that are not supported in the DPU itself.
- The architecture and DPU functionality need to support both PON and point-to-point fiber backhaul.
- DPUs are managed by a PMA, which acts as a management proxy for the DPU when it is unpowered.
- Where the backhaul is PON-based, it must be possible to operate TR-156 ONTs and TR-167 ONUs on the same Optical Distribution Network (ODN) as the new TR-301 Issue 2 [32] DPUs. Ideally this would not require any changes to the Optical Line Termination (OLT) service provisioning and functionalities. The OLT may however require additional management capabilities, depending on the architecture adopted, to manage the DPUs (e.g., in the case of a PMA located on the OLT itself).

There is a difference between the TR-156-compliant backhaul and the point-to-point Ethernet/TR-167 backhaul case. In the TR-156 case, the OLT has visibility of the user ports, and can perform various functions on behalf of the DPU. For the point-to-point Ethernet and TR-167 backhaul cases, these functions (such as Virtual Local Area Network [VLAN] tag manipulation and user port identification) have to be done in the DPU itself. The point-to-point Ethernet/TR-167 backhaul case is known as “Model 1”, and the TR-156 backhaul case is known as “Model 2”. These two models are described in more detail later in this section.

The high level FTTdp architecture is illustrated in Figure 5-1.

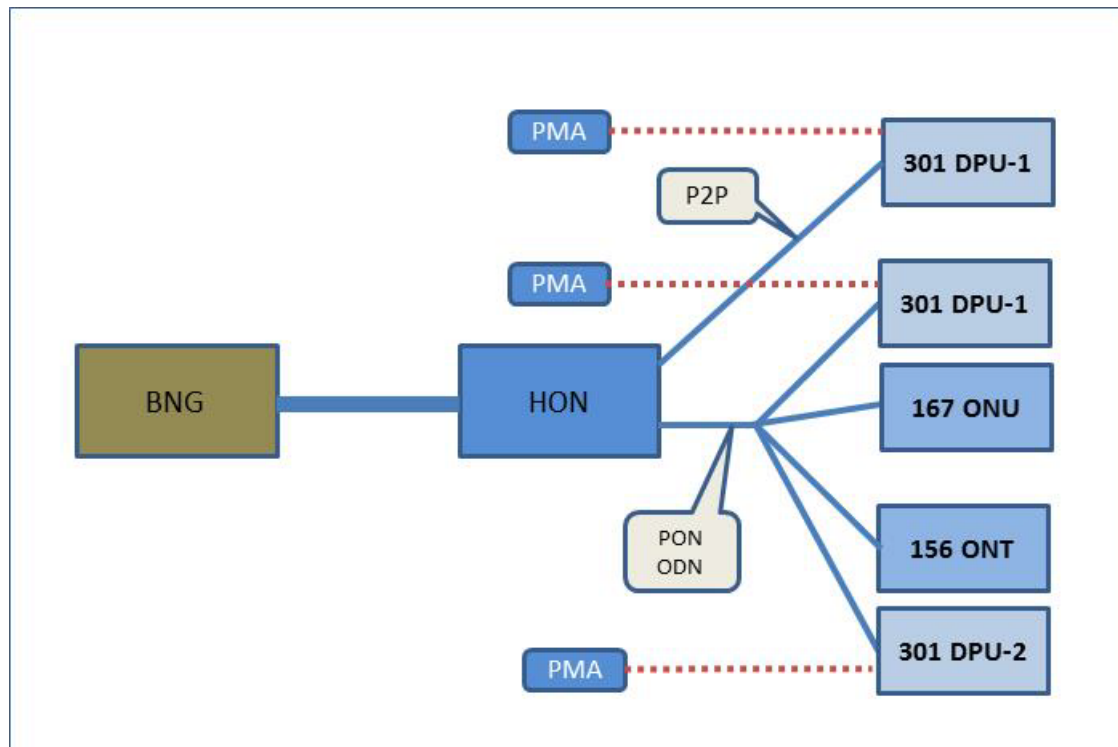


Figure 5-1 High level FTTdp architecture and TR-156/TR-167 co-existence

Although there are some differences in DPU functionality that depend on the backhaul interface type, it is of course open to vendors to make a single DPU product type, which is then configured according to the deployment model. It is likely that a DPU with an integrated GPON backhaul could be realized as a single device that may be configured to operate in both Model 1 and Model 2 (ref Figure 5-2 and Figure 5-3) deployments. It is also likely that a DPU for Model 1 deployments could use a Small Form-factor Pluggable Transceiver (SFP) to support both TR-167-compliant GPON and point-to-point Ethernet backhauled. This would then allow a single device to be used for both backhaul types, and for the migration between backhaul types after initial installation. It should be noted, however, that sealed units may limit the ability to perform field replacement of backhaul interfaces.

5.1 Model 1: Point-to-Point Ethernet/TR-167 Backhaul

When point-to-point Ethernet or TR-167 [5]-compliant GPON backhaul are used, the HON performs the functions of an aggregation node as defined in TR-101 and/or TR-178. Traffic for all user ports in a DPU share a common interface to the backhaul and any per user port tagging functions are performed by the DPU. Refer to Figure 5-2 for a depiction of a Model 1 deployment. At the V reference point, unicast traffic is either single-tagged (S-tag) or double-tagged (S-tag + C-tag). Since only 1:1 VLANs are supported for unicast traffic in this model, there is a unique tag or tag stack for each user port for unicast traffic between the DPU and the Broadband Network Gateway (BNG). Additionally, N:1 VLANs are supported for the purpose of multicast delivery.

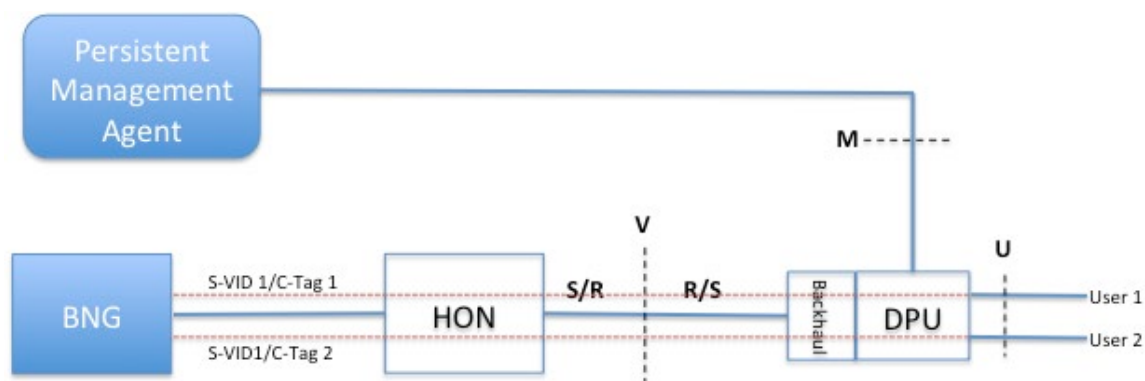


Figure 5-2 Deployment Model 1 (Pt-to-Pt Ethernet/TR-167 Backhaul)

For downstream multicast and broadcast frames, S-VLANs may be used to achieve efficiency in the use of GPON backhaul; the DPU provides Internet Group Management Protocol (IGMP) transparent snooping functionality to ensure multicast traffic is only sent to the appropriate ports. Additionally, relay and intermediate agent functions that require user port location information are performed in the DPU (see Section 0).

In this model, the PMA is responsible for the management of the following functions:

- Tag addition, translation, and removal at the user port.
- Upstream and downstream priority queue configuration.
- Upstream and downstream frame to priority queue mapping.
- All copper drop transceiver provisioning and monitoring (G.fast, VDSL2).
- User port state including RPF.
- Multicast whitelist.
- Equipment command and control including software image download and restart.
- Circuit ID.
- Provisioning of Intermediate and Relay Agents.

5.2 Model 2: TR-156 Backhaul

In the second deployment model, TR-156 [4]-compliant GPON is used as the DPU backhaul. Rather than having direct access to the physical user ports, the backhaul uses a virtual Ethernet interface per DPU user port. Frames are forwarded unchanged between the physical user port and the virtual Ethernet interface and carried to the HON by GPON Encapsulation Method (GEM) ports that are unique to each virtual Ethernet Interface. The HON is able to perform Media Access Control (MAC) learning on a per-DPU user port basis. Both 1:1 and N:1 VLAN models are supported along with multicast as described in TR-156.




Figure 5-3 Deployment Model 2 (TR-156 Backhaul)

In this model, the PMA is responsible for the management of the following functions:

- All copper drop transceiver provisioning and monitoring (G.fast, VDSL2).
- User port state including RPF.

6 Fundamental Architectural and Topological Aspects

This section provides a description of the FTTdp architecture and methods of deployment.

In all cases, G.fast or VDSL2 transceiver technologies are used at the U-reference point. For DPU backhaul, PON or point-to-point Ethernet may be used.

ITU-T G.9701 [9] and ETSI TS 101 548-1 [3] are relevant to this BBF architecture document.

Section 5.3 of ETSI TS 101 548-1 [3] provides a reference model for the end-to-end FTTdp architecture.

7 DPU Environmental Aspects

Due to their close proximity to the customer premises, DPUs are often deployed in environmentally challenging locations such as pole mounts, outside building wall mounts, and underground chambers. Therefore, temperature and humidity extremes as well as physical security need to be considered in DPU implementations. Additionally, with the potential for large geographically dispersed numbers of DPUs deployed in networks, visits by service provider technicians should be kept to a minimum.

- R-1** The DPU **MUST** be passively cooled.
- R-2** The DPU **MUST** contain an internal temperature sensor.
- R-3** The DPU internal temperature **MUST** be able to be read on demand from the PMA.
- R-4** The DPU **MUST** generate a temperature alarm when the measured internal temperature (T0) is greater than a configured threshold (T1).
- R-5** The DPU **MUST** generate a temperature alarm when the measured internal temperature (T0) is less than a configured threshold (T2).
- R-6** The DPU **MUST** undergo a thermal shutdown when the measured internal temperature (T0) is greater than a configured threshold (T3).
- R-7** The DPU **SHOULD** reduce power consumption when the measured internal temperature (T0) is greater than a configured threshold (T4).
- R-8** The DPU **SHOULD** contain a sensor that detects the opening of the enclosure, the status of which may be read on demand by the PMA and **SHOULD** be able to raise an alarm.

8 DPU Powering

DPUs require access to a power source. However, many DPUs will be deployed in close proximity to the customer premises to achieve very high copper drop line rates. This can result in the deployment of large numbers of DPUs within a service provider access network. Providing local power to each DPU under such circumstances becomes problematic due to the need to gain power utility access, and provide power functionalities such as surge suppression and, in some cases, battery backup. A solution to this problem is to power these DPUs over the copper drops from the customer premises. Since DPUs often serve more than one customer premises, they should be able to equitably distribute the power drawn over multiple copper drops. Additionally, DPUs must have the ability to operate the uplink, common circuitry, and the appropriate copper drop transceiver when only one line is providing power.

Powering DPUs over copper drops results in additional requirements, e.g., the DPU needs to know the state of the remote power source. Changes in power source state such as the transition to and from battery backup, and impending total loss of power, must be communicated to the DPU. Upon receipt of these state change notifications, the DPU may need to respond with internal actions, such as the reduction of power consumption while operating on power from battery backup. Further, the power source must have the ability to detect the presence of a DPU, the absence of fault conditions (e.g., short or open circuit), and that there are no other copper connected devices on the pair before applying power, thereby avoiding damage to other equipment.

In order to support reverse power feeding, ETSI TS 101 548-1 [3] defines the Metallic Detection Start Up (MDSU) protocol based on the detection of resistive signatures located in the DPU. Some service provider deployments require the removal of exchange-based POTS on the copper drop between the DPU and the customer premise before the RPF can start up. ETSI TS 101 548-1 [3] defines the POTS Remote Copper Reconfiguration Protocol (PRP), an optional extension of MDSU in case this functionality is required. Detailed specifications are defined in ETSI TS 101 548-1 [3]. Different service provider implementation options are described in section 5 of ETSI TS 101 548-1 [3].

If the DPU operates only on RPF, then requirements are defined in section 8.1. If the DPU operates only on local or forward powering, then requirements are defined in section 8.3. If the DPU operates with a combination of local or forward powering and RPF (e.g., RPF as a backup powering in case local or forward powering fails), then the dying gasp message relates to the failure of the backup powering, as defined in either section 8.1 (RPF backup) or section 8.3 (local or forward powering backup).

8.1 DPU Reverse Powering Requirements

R-9 A DPU MUST be able to operate all of its transceivers concurrently.

R-10 A DPU MUST prevent any CO/exchange or cabinet power feed entering the DPU from being connected to a reverse powered customer drop.

R-11 A DPU MUST be able to be powered in at least one of three ways:

- Reverse Power from the customer premises.
- Forward Power from a Network Node.
- Local Power from AC mains source.

R-12 A reverse powered DPU MUST be able to operate when there is only 1 power source providing power. This includes powering the central DPU and backhaul functions necessary to support that user's service, in addition to powering the transceiver of the user supplying the power. However, there is no need for vectoring for a single connected user.

Note: This requirement applies both in the case where the reverse power is provided via mains electrical power in the customer premises location, and in the case where it is operating on CPE battery power, albeit with reduced power mode on the copper link and reduced backhaul capacity.

R-13 An increase or orderly decrease in the number of lines providing RPF MUST NOT degrade the operation of the DPU by interrupting service on the active copper lines (i.e., causing them to lose sync) and on its active uplink interfaces. It is acceptable that these copper lines experience a small number of errors and/or a small Gamma Data Rate (GDR) loss for a limited period.

Note: An orderly decrease in the number of lines providing RPF is the result of a shutdown with RPF-Dying Gasp from the PSE. In case of a Disorderly Shutdown Event (DSE) of a copper line concurrent with the decrease of feeding lines, only the error behavior typical of DSEs is acceptable. Examples for DSE are disconnecting the line or NT/CPE loss of power.

R-14 A DPU MUST take a roughly equal share of power (as measured at the DPU) from all connected, powered-on PSEs with lines operating in L0 full power mode.

R-15 DPU power consumption SHOULD scale appropriately with traffic demand, including but not limited to, support of G.fast discontinuous operation, and low power modes for both G.fast and VDSL2 access technologies.

R-16 Under all combinations of changes in the RPF powering state of the lines on a DPU, the share of power demand taken from any line MUST NOT cause the current on that line to exceed the limits for the power supply class according to the time and current specified in Note 2 a) and b) of clause 7.5.1.1 of ETSI TS 101 548-1 [3].

R-17 The DPU and PMA MUST maintain the following current PSE status for each of its user lines:

- PSE-UKN: Shutdown with unknown reason.
- PSE-DGL: Shutdown with dying gasp (normal shutdown).
- PSE-OHP: Shutdown with dying gasp with off-hook phone.
- PSE-PWR: Powered with unknown PSE powering method.
- PSE-BAT: Battery powered.
- PSE-ACM: Mains powered.

R-18 The DPU MUST support the receipt of the following Dying Gasp indications from the PSE and the NT Module on each one of its user lines:

- PSE-DGL: Shutdown with dying gasp (normal shutdown).
- PSE-OHP: Shutdown with dying gasp with off-hook phone.

8.1.1 DPU Dying Gasp Message

In addition to receiving Dying Gasp messages from PSEs, the DPU sends Dying Gasp to the PMA when it loses its own power source. Previous issues of this Technical Report specify the original version of the DPU Dying Gasp message. This issue specifies a new version of the message that adds protections against certain security vulnerabilities, includes the identity of the message source, and adds new codes for DPUs with local or forward powering.

R-250 The DPU MUST support sending the Dying Gasp message, with the version sub-field set to 1 as defined in Table 8-1 and using the format shown in Table 8-3, to its PMA immediately prior to shutting down.

R-251 The PMA MUST support receiving and correctly interpreting the Dying Gasp message with the version sub-field set to either 0 or 1.

| Bits | Sub-Field | Value | Meaning |
|------|-----------|-------|---|
| 15 | Version | 0 | When this bit is zero, the original Dying Gasp format in Table 8-2 is used. |
| | | 1 | When this bit is one, the secure Dying Gasp format in Table 8-3 is used. |

| | | | |
|-------|-----------------|--------|--|
| 14-12 | Dying Gasp Code | 000 | DPU-LLPS: DPU shutdown because last line stopped providing power due to normal shutdown (reason: power lost at PSE). The Line ID field identifies the last line. |
| | | 001 | Unexpected shutdown, cause unknown |
| | | 010 | DPU-MLPF: DPU shutdown because insufficient power provided. The RPF Line State fields identify the lines not providing power and their PSE RPF states. |
| | | 011 | Safe temperature of DPU exceeded |
| | | 100 | DPU-LPPS: DPU shutdown because local power was lost. |
| | | 101 | DPU-FPPS: DPU shutdown because forward power was lost. |
| | | 110 | Dying gasp code specified in the Code Extensions field. If the Version bit is zero, this code is invalid. |
| | | 111 | DPU-LLPF: DPU shutdown because last line stopped providing power due to fault condition (reason: unknown or off-hook phone). The Line ID and RPF Line State fields identify the last line and its PSE RPF state. |
| 11-0 | Line ID | 1-4095 | When the Dying Gasp code is DPU-LLPS or DPU-LLPF, this field identifies the last line. |

Table 8-1 Dying Gasp Field

Note: Some of these codes depend on the DPU having received a Dying Gasp, or other state change information, from the PSE.

DPU-LLPS is a non-fault condition. This corresponds to the last power feed having disappeared without any accompanying PSE fault Code. This is normal, expected behavior and so should not raise any alarm (from the Network Management System [NMS]).

All the other dying gasp messages in Table 8-1 are, or may be, faults, and so would be expected to trigger some follow-up action by the NMS.

R-252 The DPU-LLPS and DPU-LLPF codes include a Line ID identifying the last line active before the shutdown condition occurred. The valid range of Line ID values is from 1 to 4095, with the values corresponding to the RPF Line State fields in Table 8-2 and Table 8-3.

R-253 With other Dying Gasp Codes, the LineID sub-field is not used and MUST be set to zeros by the transmitter and ignored by the receiver.

The Dying Gasp message is sent to a configurable port on the PMA using User Datagram Protocol (UDP) and always contains the state of all the RPF lines on the DPU.

R-19 REQUIREMENT DELETED

R-20 The DPU Dying Gasp MUST be sent to a UDP port on the PMA IP address.

R-21 The Dying Gasp UDP port must be configurable in the DPU by the PMA.

R-22 REQUIREMENT DELETED

R-23 The Dying Gasp message MUST contain the RPF line state for all the RPF lines on the DPU regardless of the Dying Gasp Code.

8.1.1.1 Dying Gasp Message Format Version 0

The Dying Gasp message uses one of two formats depending on the value of the Version sub-field of the Dying Gasp field. When the Version sub-field is set to 0¹, the message format is shown in Table

¹ Note that a DPU conforming to this Working Text cannot use the Version 0 format of the Dying Gasp
September 2024

8-2. The Dying Gasp field is provided using big endian ordering in the two octets following the UDP checksum. The octets following the Dying Gasp field are populated with the RPF Line States for each line on the DPU beginning with Line 1, with each octet containing 2 RPF Line States. Within each octet, the first RPF Line State occupies the 4 MSBs and the second RPF Line State occupies the 4 LSBs.

The Dying Gasp message is truncated at the octet boundary following the last RPF Line State field required to report the state of each line on the DPU. If the DPU has an odd number of lines, the 4 LSB's of the last octet contain zeros.

| Offset Octets | 0 | 1 | 2 | 3 |
|---------------|--|-------------------|-------------------|-------------------|
| 0 | Dying Gasp field (with Version sub-field set to 0) | | RPF-1/RPF-2 | RPF-3/RPF-4 |
| 4 | RPF-5/RPF-6 | RPF-7/RPF-8 | RPF-9/RPF-10 | RPF-11/RPF-12 |
| ... | ... | | | |
| 2044 | RPF-4085/RPF-4086 | RPF-4087/RPF-4088 | RPF-4089/RPF-4090 | RPF-4091/RPF-4092 |
| 2048 | RPF-4093/RPF-4094 | RPF-4095/0000 | Not Applicable | |

Table 8-2 Dying Gasp Message format, version 0

The length of the Dying Gasp message for a large DPU (e.g., 3000 lines) can exceed the size of a normal Ethernet MTU. As a result, the Dying Gasp message for large DPUs will be fragmented on networks that do not support jumbo Ethernet frames.

R-254 The DPU MUST use the format in Table 8-2 for the Dying Gasp message version 0.

The following requirements apply to all versions of the Dying Gasp message.

R-254 The DPU MUST format the Dying Gasp field within the Dying Gasp message in big endian order.

R-256 The DPU MUST format the RPF Line State fields within the Dying Gasp message in big endian order.

8.1.1.2 Dying Gasp Message Format Version 1

When the Version sub-field of the Dying Gasp field is set to 1, the format of the Dying Gasp message is shown in Table 8-3. The Dying Gasp field and RPF Line State fields have the same format and requirements as specified above for the Version 0 message format. This message format also includes the following fields which are described below:

- Code Extensions
- Source ID Length
- Source ID
- Hashed Message Authentication Code (HMAC)

R-257 The DPU MUST use the format in Table 8-3 for the Dying Gasp message version 1.

As in Version 0, the Dying Gasp message is truncated at the octet boundary following the last RPF Line State field required to report the state of each line on the DPU. In Version 1, a new RPF Line State code is defined for the 4 LSBs of the last octet if the DPU has an odd number of lines. If this is the case, the 4 LSBs of the last octet contain the value “PNP: Port Not Present”.

The following requirement applies to Version 1 of the Dying Gasp message.

R-24 The following RPF Line State Codes **MUST** be supported in the RPF Line State field of the Dying Gasp Message.

- 0 - PSE-UKN: Shutdown with unknown reason.
- 1 - PSE-DGL: Shutdown with dying gasp (normal shutdown).
- 2 - PSE-OHP: Shutdown with dying gasp with off-hook phone.
- 3 - PSE-PWR: Powered with unknown PSE powering method.
- 4 - PSE-BAT: Battery powered.
- 5 - PSE-ACM: Mains powered.
- 6-13 - Reserved for future use.
- 14 - PNP: Port Not Present (Version 1 only).
- 15 - No RPF Line Present.

| Offset Octets | 0 | 1 | 2 | 3 |
|---------------|--|-------------------|-------------------|-----------------------|
| 0 | Dying Gasp field (with Version sub-field set to 1) | | Code Extensions | Source ID Length (SL) |
| 4 | Hashed Message Authentication Code (HMAC) (32 octets) | | | |
| 36 | Source ID (variable length SL octets) | | | |
| 36+SL | RPF-1/RPF-2 | RPF-3/RPF-4 | RPF-5/RPF-6 | RPF-7/RPF-8 |
| ... | ... | | | |
| 2080+SL | RPF-4089/RPF-4090 | RPF-4091/RPF-4092 | RPF-4093/RPF-4094 | RPF-4095/PNP |

Table 8-3 Secure Dying Gasp message format, version 1

The Code Extensions field shown in Table 8-4 is added to specify additional Dying Gasp codes.

R-258 When the Version sub-field is set to 1 and the Dying Gasp Code sub-field is set to 110, the extended Dying Gasp code is specified in the Code Extensions field.

R-259 If the Version sub-field is set to 1 and the Dying Gasp Code sub-field is different from '110' then the Code Extensions field **MUST** be set to '0' by the transmitter and ignored by the receiver.

| Bits | Value | Meaning |
|------|-------|---------------------------------------|
| 7-0 | 0 | No dying gasp code extension provided |
| 7-0 | 1-255 | Reserved |

Table 8-4 Extended dying gasp codes in the Code Extensions field

The SourceID field contains either the serialNumber attribute or the CN attribute of the IDevID certificate's subject field, depending on which attribute contains the vendor's enterprise identifier (see Section 16.5.1.3). The use of the serialNumber attribute is recommended. The use of the CN attribute is provided for backwards compatibility with TR-301 Issue 2 [32].

An example showing how the Source ID and SL fields are created from the Enterprise ID and DPU serial number is shown in Table 8-5 .

| Enterprise ID | SN format | DPU serial number | SL field | Source ID field |
|---------------|---------------------------------|-------------------|----------|-------------------|
| 3561 | String | BBFX87654321 | 17 | 3561-BBFX87654321 |
| 3561 | VSSN converted to decimal value | BBFX05397FB1 | 13 | 3561-87654321 |

Table 8-5 Example Source ID and SL fields

The HMAC field contains a Message Authentication Code that is generated by the DPU and verified by the PMA to check the integrity of the fields in the Dying Gasp message. The mechanism used for this purpose is the keyed-Hash Message Authentication Code algorithm defined in RFC-2104 [29], using SHA-256 as defined in FIPS PUB 180-4 [30]. The key used for the HMAC algorithm is derived in accordance with the keying material exporter function defined in RFC 5705 [31]. The input values for the keying material exporter function are:

- Label = "EXPORTER-BBF- Dying-Gasp".
- Application context = "Dying-Gasp"
- Length = 32

The keying material exporter function uses these inputs together with parameter values from the current TLS connection to create a key whose value can be duplicated in the DPU and PMA, but which cannot be recreated without the TLS connection parameters that are unique to each DPU/PMA pair. The DPU uses the key and the message fields to generate the HMAC and the PMA uses the same key and message fields to recreate the HMAC. If the HMAC created by the PMA matches the received value, the PMA accepts the Dying Gasp message as authenticated and unmodified. If the two values do not match, the PMA rejects the message.

The following requirements apply when a DPU sends a Dying Gasp message with the Version bit set to 1.

R-260 The DPU MUST include the HMAC to protect the integrity of the payload fields in the Dying Gasp message.

R-261 When the PMA receives the Dying Gasp message, it MUST use the received HMAC to validate the integrity of the payload fields.

R-262 If the HMAC generated by the PMA does not match the HMAC received in the Dying Gasp message, the PMA MUST reject the Dying Gasp message.

R-263 The DPU and PMA MUST use the keyed-Hash Message Authentication Code algorithm defined in RFC-2104 and SHA-256 as defined in FIPS 180-4 to generate the HMAC.

R-264 When generating the HMAC, the DPU and PMA MUST concatenate the payload fields in the following order and provide the result to the HMAC generation function as a byte array:

1. Dying Gasp field in big endian ordering
2. Code Extensions field
3. Source ID Length field
4. Source ID field as a UTF-8 encoded byte array
5. RPF Line State fields in byte order, starting with RPF-1/RPF-2 and ending with the byte containing the last line supported in the DPU.

R-265 The DPU and PMA MUST export the key used to generate the HMAC from the TLS connection supporting the current DPU/PMA NETCONF connection, using the keying material export function as defined in RFC 5705 with the following input parameters:

- Label = "EXPORTER-BBF-Dying-Gasp"
- Application context = "Dying-Gasp"
- Length = 32

R-266 The DPU and PMA MUST use an output length of 32 bytes when generating the HMAC.

R-267 The DPU MUST format the HMAC within the Dying Gasp message in big endian order.

In order to defend against replay attacks, duplicate Dying Gasp messages are disallowed. Since it is possible that the payload fields in consecutive Dying Gasp messages may be identical, the HMAC value, and the key used to generate it, needs to be different. If the DPU loses power after sending the Dying Gasp this will be the case, but even if it remains powered it needs to terminate the TLS connection and create a new one to generate a new key.

R-268 If after sending a Dying Gasp message the DPU remains powered, it MUST terminate the current NETCONF, TLS, and TCP connections. It MUST then establish a new NETCONF connection to the PMA using Step 4 of the discovery process specified in Section 16.5.

R-269 The PMA MUST reject any Dying Gasp message in which the HMAC and payload field contents duplicate those of a previous message.

8.1.2 Additional DPU Reverse Powering Requirements

R-25 The DPU SHOULD take a roughly equal share of power (as measured at the DPU) from all connected, powered-on PSEs with lines operating in G.fast L2.1 Normal low power mode.

R-26 The DPU SHOULD take a roughly equal share of power (as measured at the DPU) from all connected, powered-on PSEs with lines operating in G.fast L2.1 Battery low power mode.

R-27 The DPU SHOULD take a roughly equal share of power (as measured at the DPU) from all connected, powered-on PSE with lines operating in G.fast L2.2 Battery low power mode.

R-28 The DPU MUST be able to be configured on a per port basis to allow only minimal power extraction, compliant with Section 7.5.2.2 of ETSI TS 101 548-1 [3], when the PSE reverse power feed is operating on battery power and sufficient power is available from one or more remaining reverse power feeding lines for normal operation of the DPU.

R-29 The DPU MUST notify the PMA of the receipt of a Dying Gasp from a user line.

R-30 The DPU start up protocol MUST operate irrespective of the presence of MELT-P signatures encountered in FTTdp deployments.

R-31 The DPU reverse powering functionality MUST comply with ETSI TS 101 548-1 [3].

R-32 Upon detecting there is no power on the user line, the DPU MUST be configurable on a per line basis to take one of the following actions:

- FL3: Force the line to the L3 link state (no service).
- NL0: Force the line to an enabled low power link state (limited service).
- ESO: Force the line to support emergency services only (all link states allowed).
- SRV: No service limitations forced on the line (all link states allowed).

R-33 The DPU MUST be able to automatically put a provisioned port in-service when reverse powering is detected on that line.

R-34 The DPU MUST be able to continue to operate in the presence of micro-interruptions of the reverse power feeding, up to a maximum duration of 10 ms per feeding line at a repetition rate of 20 seconds for the short range class in the following two cases:

One active line with reverse power and one micro-interruption.

Two active lines with reverse power applying a synchronized micro-interruption on each line.

R-35 A DPU SHOULD prevent electrical noise and Radio Frequency Interference (RFI) on the CO/cabinet side of the DPU in the frequency bands occupied by DPU hosted data services from reaching the customer drop. This applies to all drops whether or not the drop is connected to a DPU hosted service.

R-36 A DPU MUST be able to accept Reverse Power Feeding from user lines regardless of the tip to ring polarity of the received DC voltage.

8.1.3 DPUs With RCR

The following requirements apply to DPUs that support RCR.

R-37 In the absence of reverse power feeding, the DPU MUST maintain copper continuity without significant impairments on the associated non active FTTdp user's interface, so that xDSL access from CO or cabinet can be provided.

R-38 In the presence of reverse power feeding, the DPU MUST automatically become active on FTTdp user's interface that provides powering and it MUST disconnect the corresponding copper line from CO or cabinet.

Note: the way this function is implemented may have an impact on the long-term reliability of the DPU.

8.2 Power Source Requirements

R-39 The PSE MUST comply with one of the power classes defined in ETSI TS 101 548-1 [3].

R-40 The PSE MUST send a Dying Gasp indication to the DPU immediately before it removes power from the line.

R-41 The PSE MUST promptly remove power from a line upon the detection of a fault.

R-42 At a minimum, the PSE MUST support the detection of the following fault conditions:

- Presence of an unprotected off hook telephone.
- Presence of a short circuit.
- Presence of an open circuit.
- Presence of a foreign voltage.

Note: the promptness in the power removal is dictated by a trade-off between safety considerations and the time needed to send Dying Gasp indications; the most critical cases being the presence of an unprotected off hook telephone and of a short circuit.

R-43 Before providing the DPU with sufficient power to reach an operational state, the PSE MUST verify that all the following conditions are met:

- Absence of voltage on the line.
- Absence of unprotected off hook telephone.
- Absence of short circuit.
- Absence of open circuit.
- The detection of a DPU that supports reverse powering.

R-44 At the CPE side reverse powering functionality MUST comply with ETSI TS 101 548-1 [3].

8.3 Local and Forward powering requirements

A DPU may be operating only on local or forward powering, because the DPU does not support RPF or the RPF is disabled on all the lines.

R-270 In this case, the DPU MUST send a dying gasp message (as defined in section 8.1.1) upon the following events:

- Unexpected shutdown, cause unknown (dying gasp code 001)
- Safe temperature of DPU exceeded (dying gasp code 011)
- DPU-LPPS: DPU shutdown because local power was lost (dying gasp code 100).
- DPU-FPPS: DPU shutdown because forward power was lost (dying gasp code 101).
- Extended dying gasp code specified in the Code Extensions field (dying gasp code 110). This code may be sent only if the Dying Gasp code specified in the Code Extensions field is applicable to Local or Forward powering requirements.

R-271 The DPU MUST NOT send a dying gasp message with other dying gasp codes defined in Section 8.1.1.

R-272 The Dying gasp message MUST NOT contain the RPF field codes.

9 DPU Physical Interfaces

A DPU contains two types of physical interfaces used for data transport by the DPU:

1. Copper drop interfaces that provide the user broadband service (U-O).
2. A backhaul interface that connects the DPU to the HON (R/S).

This section provides requirements for both types of interface.

9.1 DPU Copper Drop Physical Interface Requirements

FTTdp will often be deployed so as to provide an upgraded service in existing broadband networks. This means that the DPU copper drop interface needs to be spectrally compatible with other DSL technologies that may already exist in the same wiring bundle. Additionally, DPUs require a method for isolating a user's local loop from CO/exchange or cabinet supplied battery and service when that user is connected to DPU provided service. This may be performed by Remote Copper Re-configuration or Auto Configuration at power up.

To enable the support of traffic shaping at the BNG, DPUs need to provide timely information on the current attainable data rate on a per line basis. Historically, Access Node Control Protocol (ANCP) [14] has been used to provide this information to the BNG and DPUs support this capability through the PMA using one of 2 options. In the first option, the PMA is integrated into the HON as depicted in Figure 9-1.




Figure 9-1 ANCP With PMA In The HON

In the second option, the PMA is external to the HON as depicted in Figure 9-2.




Figure 9-2 ANCP With The PMA External To The HON

R-45 A DPU MUST support at least one of the following customer facing copper drop technologies:

- G.fast
- VDSL2

R-46 A DPU that supports G.fast MUST be able to notch out specific frequencies in the FM bands, amateur radio bands and safety of life frequencies.

- R-47** A multi-line DPU MUST support crosstalk cancellation (vectoring) between all the pairs of a given technology on that DP. There is no requirement for crosstalk cancellation between multiple DPUs.
- R-48** A DPU MUST be able to be spectrally compatible with the following:
- ADSL/ADSL2/ADSL2plus from the CO/exchange or cabinet
 - VDSL2 from the cabinet or CO/exchange or cabinet for all profiles up to and including 17 MHz
 - VDSL2 from the same DP for all profiles up to and including 17 MHz
- R-49** The DPU MUST be able to report the attainable Net Data Rate (NDR) and the port state of each line to the PMA at a configurable time interval and whenever the attainable NDR on any port changes by more than a configurable threshold.
- R-50** The DPU MUST be able to report the NDR and the port state of each line to the PMA at a configurable time interval and whenever the NDR on any port changes by more than a configurable threshold.
- R-51** The DPU MUST be able to losslessly correct noise induced erasures of up to at least 10ms duration on all its lines simultaneously by means of PHY layer retransmission.
- Note:** the memory required may be constrained by the downstream data rate of the DPU uplink.
- R-52** The DPU MUST be able to losslessly correct periodic 1 ms noise induced erasures, due to REIN with 120 Hz repetition frequency, on all lines simultaneously by means of retransmission
- R-53** Erasure handling MUST be handled within the DPU itself (e.g., it must not rely on retransmission memory outside the DPU, or flow control extending beyond the DPU as part of retransmission).
- R-54** If the line is forced to the L0 state (AdminStatus=UP), then performance monitoring counters MUST be active, irrespective of the actual power management state of the line.
- R-55** If the line is forced to the L3 state (AdminStatus=DOWN), then all performance monitoring counters MUST be frozen, including the UAS counter.

A G.fast copper drop should support synchronous timing to support:

1. Carrier grade voice services where derived voice is converted to analogue voice via an ATA in the G.fast RG. Note that 'carrier grade voice services' in this context are defined as those services that support modem and fax transmission in addition to voice.
 2. The use of a G.fast copper drop for mobile backhaul.
- R-56** A DPU that supports carrier grade voice or mobile backhaul MUST support frequency sync at the PHY layer via SyncE on any point to point Ethernet backhaul.
- R-57** A DPU that supports carrier grade voice or mobile backhaul MUST support frequency sync at the PON PHY layer on any PON backhaul.
- R-58** A DPU that supports carrier grade voice or mobile backhaul MUST transfer the sync information to the G.fast Network Timing Reference.
- R-59** An NT module/CPE that supports carrier grade voice or mobile backhaul MUST interwork with the ITU-T G.827x time profiles i.e., those appropriate to end-systems.
- R-60** A DPU that supports carrier grade voice or mobile backhaul MUST support the ITU-T G.827x time profile i.e., appropriate to intermediate systems.

9.2 DPU Backhaul Physical Interface Requirements

- R-61** A DPU MUST support a fiber backhaul with a minimum bandwidth of 1 Gbps in both the upstream and downstream directions.
- R-62** DPUs having more than 16 ports SHOULD support a second uplink interface.
- R-63** Both uplink interfaces MUST be of the same type and speed.
- R-64** DPUs with a second, point-to-point, uplink interface MUST support Link Aggregation.

10 Traffic Management and Quality of Service (QoS)

10.1 DPU QoS Management

An analysis of the line rates, both on the customer and network side, and service mixes led to the conclusion that having no differential packet treatment with regard to queuing and forwarding in the DPU could lead to non-trivial amounts of jitter. However a very simple scheme, with only 4 levels of traffic priority per direction addresses this particular problem. This can be implemented with 4 shared strict priority queues in the upstream direction and 4 strict priority queues per user port in the downstream direction. No need for more complicated queuing disciplines, e.g., weighted round robin, has been identified. Further, although there may need to be shaping and policing at some point in the network, there is no business need for this functionality to be in the DPU itself; having it there would significantly increase the amount of configuration needed, and add non-trivial functionality.

The DPU is mainly a Layer 2 device and so packet classification is done on the basis of VID and/or .1p bit value along with the ability to apply VID and .1p bit values at the user port based on a limited set of criteria.

As discussed in Section 5, FTTdp supports two deployment models. Each of these models results in different frame handling within the DPU. The DPU contains 2 frame-handling subsystems:

1. The Port Frame Forwarding Function (PFFF)
2. The DPU Backhaul

A virtual Ethernet interface is used to represent the interface between the two subsystems. Within this interface are 4 virtual priority queues. Neither the virtual Ethernet interface nor the virtual queues always exist in a physical implementation. They simply provide a convenient paradigm for the discussion of frame forwarding that may be carried over into the data model for each subsystem. In this way, the subsystems share a common view of frame forwarding provisioning even though they may use different management interfaces.

As depicted in Figure 10-1, a Model 1 DPU PFFF receives upstream frames from the U reference point interface, adds or translates VLAN tags, and forwards them to the appropriate virtual queue. The DPU backhaul then performs any required uplink specific adaptation and forwards the frames over the uplink. As depicted in Figure 10-2, downstream frames are placed in the appropriate virtual priority queue by the DPU Backhaul and the PFFF places the frame into the correct physical priority queue at the U reference point interface. During the forwarding of the downstream frames, the PFFF removes or translates tags. Frame forwarding actions of the PFFF are provisioned by the PMA. Frame forwarding actions of the DPU backhaul are provisioned by the HON as required.

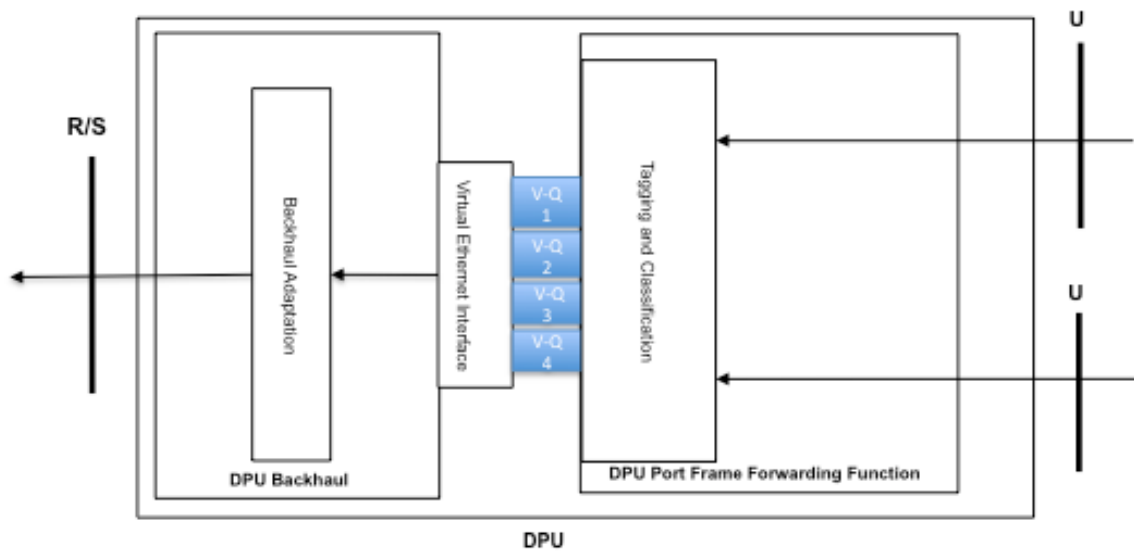


Figure 10-1 Deployment Model 1 DPU Upstream Frame Handling




Figure 10-2 Deployment Model 1 Downstream Frame Handling

Figure 10-3 depicts the internal division of upstream frame handling functionality for a DPU in a Model 2 deployment. The PFFF receives upstream frames and forwards them unchanged to the virtual Ethernet interface. The DPU Backhaul then adds or translates tags and associates the frames with a GEM port as provisioned by the HON. Next, the DPU backhaul places frames into the correct upstream queue based on GEM port.

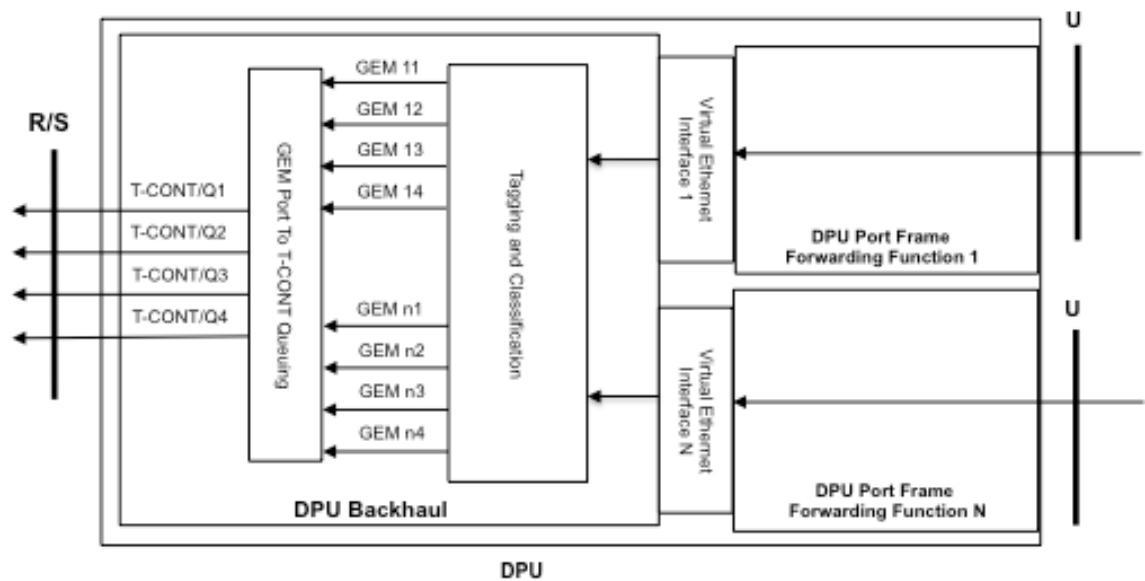


Figure 10-3 Deployment Model 2 DPU Upstream Frame Handling

For downstream frames (Figure 10-4), the DPU Backhaul receives the frames, removes or translates tags, and places them into the correct queue based upon the GEM port on which they were received. The PFFF then forwards the frames to the user unchanged.

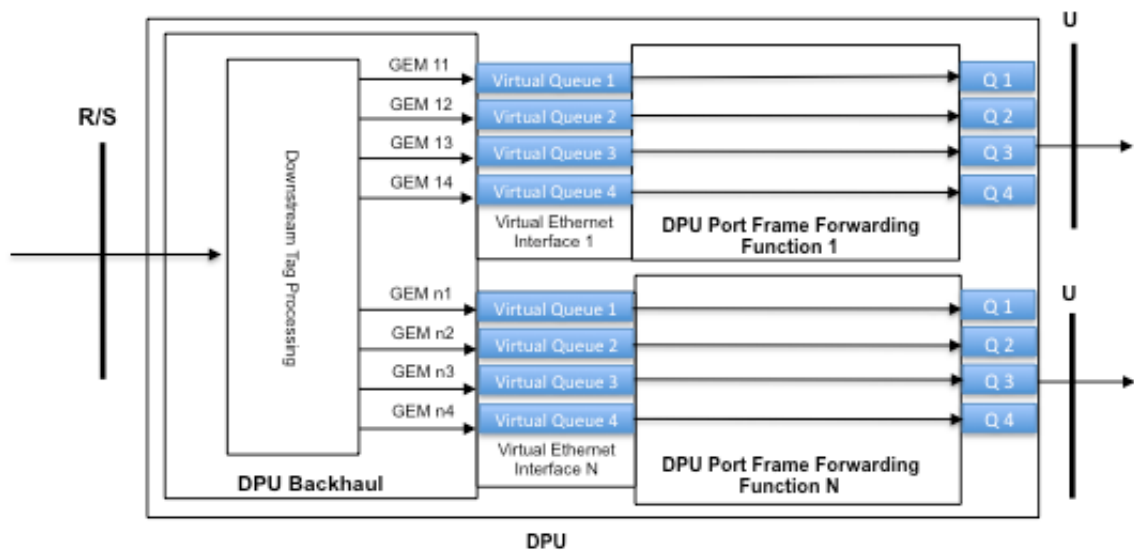


Figure 10-4 Deployment Model 2 Downstream Frame Handling

10.1.1 DPU QoS Requirements

The requirements in this section are expressed in terms of user ports rather than lines so that the same requirements can apply to both G.fast and VDSL-based DPUs.

- R-65** In the downstream direction, the DPU MUST support 4 strict priority queues per user port.
- R-66** In the upstream direction, the DPU MUST support 4 strict priority queues shared between all users on that DPU.

Note: Since there is no policing requirement, there is a need to specify a minimum buffer size. It is suggested to be at least 1 Mbytes per customer-facing port.

- R-67** The DPU MUST support forwarding frames in the upstream direction to the appropriate queue on the basis of VID, .1p bits and combinations thereof.
- R-68** The DPU MUST support forwarding frames to the appropriate queue in the downstream direction on the basis of VID, .1p bits and combinations thereof.
- R-69** The DPU MUST support configuration of the actions required by R-67 and R-68 by the PMA.

11 VLAN Handling

Deployment Model 1 (see Section 5.1) requires the DPU PFFF to support 1:1 VLANs in the upstream direction, the DPU PFFF always adds one or two tags to untagged frames, or translates an incoming tag, or translates an incoming tag and adds a tag.

- For single-tagged VLANs at the V reference point, the DPU is provisioned to either add an S-tag, or match and translate an incoming tag into an S-tag.
- For double-tagged VLANs at V reference point, the DPU is provisioned to either add a C-tag to untagged traffic, or match and translate a single tag into a C-tag, then add the S-tag.
- For the case where the VLANs are double-tagged at the U reference point, the DPU is provisioned to match and translate the outer tag into an S-tag.

These tagging operations are provisioned on a per user port basis. The mapping of upstream frames to upstream priority queues is based upon the VID and/or p-bit value in the frame tag *after* the tag manipulation has occurred.

In the downstream direction the DPU PFFF needs the ability to remove or translate tags on a per user port basis. In this case, the mapping of downstream frames to downstream queues is based on frame tag VID and/or p-bit values *prior* to tag manipulation.

In all cases, the addition, removal, or translation of frames must be performed based on a limited but arbitrary combination of criteria provided by the PMA.

Deployment model 2 (see Section 5.2) requires the DPU PFFF to transparently forward frames between a physical user port and the backhaul virtual Ethernet interface associated with that user port.

11.1 Deployment Model 1 DPU VLAN Requirements

- R-70** The DPU PFFF MUST support the addition, or translation of up to two Ethernet VLAN tags in the upstream direction on a per user port basis.
- R-71** The DPU PFFF MUST support the removal, or translation of up to two Ethernet VLAN tags in the downstream direction on a per user port basis.
- R-72** The DPU PFFF MUST support Ethernet VLAN tag addition, removal, or translation based on an arbitrary combination of: user port, VID, and received P-bit markings.
- R-73** The DPU PFFF SHOULD support deriving the P-bit markings in the upstream direction based on an arbitrary combination of user port, VID, and received DSCP value.
- R-74** The DPU PFFF MUST support Ethernet tag addition, removal, or translation based on EtherType.
- R-75** The DPU PFFF MUST support the addition, removal, or translation of two Ethernet tags in the downstream direction on a per user port basis.
- R-76** The DPU PFFF MUST support both 0x8100 (C-tag) and 0x88A8 (S-tag) Ethertype values.
- R-77** The DPU MUST NOT alter any VLAN tags beyond the outer two, and MUST treat any additional tags as part of the payload.
- R-78** The DPU PFFF MUST perform any necessary VID and P-bit manipulations before performing the mapping into upstream queues.
- R-79** The DPU PFFF MUST perform the mapping of frames into downstream queues prior to any necessary VID and P-bit manipulations.
- R-80** The DPU PFFF MUST support multiple P-bit values being used in the same VLAN.
- R-81** The DPU PFFF MUST NOT prevent multiple VLANs from using the same P-bits.
- R-82** The DPU PFFF MUST support at least 4 simultaneously active VLANs per user port.

R-83 The DPU PFFF MUST support at least 16 simultaneous tagging operation rules per user port.

11.2 Deployment Model 2 DPU VLAN Requirements

R-84 The DPU PFFF MUST transparently forward upstream frames from the user port to the corresponding virtual Ethernet interface.

R-85 The DPU PFFF MUST transparently forward downstream frames from the virtual Ethernet interface to the user port.

R-86 The DPU Backhaul MUST support at least 4 simultaneously active VLANs per user port.

R-87 The DPU Backhaul MUST support at least 16 simultaneous tagging operation rules per user port.

12 Multicast

12.1 Deployment Model 1 DPU Multicast Requirements

Model 1 DPUs support both IGMP and MLD functionality based on RFC 4604 [22]. As RFC 4604, the following requirements use the term Group Management Protocol (GMP) to indicate both IGMP and MLD.

- R-88** The DPU PFFF MUST support enabling and disabling GMP snooping on a per user port basis.
- R-89** The DPU PFFF MUST support processing GMP packets on a per VLAN basis.
- R-90** The DPU PFFF MUST support the identification and processing of user-initiated GMP messages. When this function is disabled on a port and/or VLAN, these messages MUST be transparently forwarded.
- R-91** The DPU PFFF MUST support an IGMP v3 (as per RFC 3376) transparent snooping function. This feature MUST be configurable on a per VLAN basis.
- R-92** The DPU PFFF MUST support an MLD v2 (as per RFC 4604) transparent snooping function. This feature MUST be configurable on a per VLAN basis.
- R-93** The transparent snooping function MUST be able to snoop the multicast source IP address and destination IP group address in GMP packets, and set the corresponding MAC group address filters as specified in R-94.
- R-94** The transparent snooping function MUST be able to dynamically create and delete MAC-level Group Filter entries to enable/disable selective multicast forwarding from network-facing VLANs to user-facing ports.
- R-95** The transparent snooping function MUST be able to translate the upstream VLAN of GMP packets to the configured VLAN pertaining to that 1:n multicast VLAN.
- R-96** The DPU PFFF MUST allow the configuration of IP multicast groups and/or ranges of multicast groups per multicast VLAN based on:
 - Source address matching
 - Group address matching
- R-97** The DPU PFFF MUST support matching groups conveyed by GMP messages to a provisioned list of multicast groups corresponding to a multicast VLAN associated with the receiving user port.
- R-98** When no match is found by R-97, the GMP message MUST be either forwarded or dropped, based on configuration.
- R-99** When there is a match found by R-97, the GMP message MUST be forwarded within a multicast VLAN, and frames from the matching group forwarded to the requesting port and VLAN.
- R-100** The DPU PFFF MUST be able to configure, on a per user port basis, the maximum number of simultaneous multicast groups allowed.

Note: Transparent forwarding of GMP messages in N:1 VLANs might result in network flooding and is therefore discouraged.

Note: IGMP V3 report messages may carry membership information for multiple multicast groups. Therefore, a single IGMP report message may carry membership information on groups “matching” a multicast VLAN as well as on groups “not matching” a multicast VLAN

13 Ethernet OAM

DPUs are often sealed units, and may be located in fairly inaccessible locations. Therefore, fairly comprehensive Operations Administration and Maintenance (OAM) and diagnostics functions are needed to allow remote management to determine the location and possible nature of a problem.

Ethernet OAM can be used to determine connectivity/reachability and some aspects of performance. The below requirements support one-off, on-demand reachability tests, periodic connectivity monitoring via continuity checks (CCs), and the loopback of user traffic for performance and data integrity testing.

13.1 DPU OAM Requirements

The following requirements apply to DPUs in Model 1 deployments. DPUs in Model 2 deployments provide Ethernet OAM support in the DPU backhaul and must comply with the OAM requirements defined in TR-156.

- R-101** The DPU MUST support configuring a Maintenance association End Point (MEP) on its backhaul interface.
- R-102** The DPU MUST support configuring a MEP on each of its end-user facing ports.
- R-103** The DPU SHOULD support configuring a Maintenance domain Intermediate Point (MIP) on each of its end-user facing ports.
- R-104** The Maintenance Domain level of each MEP and MIP MUST be configurable.
- R-105** The DPU MUST support IEEE802.1ag and Y.1731 Loopback and Link Trace.
- R-106** The DPU MUST support IEEE802.1ag and Y.1731 CC.
- R-107** The DPU MUST be able to establish an IEEE802.3ah EFM OAM session with the G.fast NT Module/CPE (i.e., OAM discovery and exchange of state and configuration information).
- R-108** The DPU MUST support IEEE802.3ah OAM clause 57.2.9 active mode.
- R-109** The DPU MUST support configuring IEEE802.3ah loopback (enable and disable) on the CPE via each end-user facing interface.
- R-110** The DPU MUST forward looped traffic from the CPE to its WAN interface on the normal data path.
- R-111** The DPU MUST support Y.1731 Frame Loss Measurement.
- R-112** The DPU MUST support Y.1731 Frame Delay and Frame Delay Variation Measurements.

13.2 CPE OAM Requirements

- R-113** The G.fast NT Module/CPE MUST support configuring IEEE802.3ah Passive mode.
- R-114** The G.fast NT Module/CPE MUST be able to configure Loopback (enable/disable) when requested by an Active Mode DPU.
- R-115** The G.fast NT Module/CPE SHOULD support 802.1ag/Y.1731.
- R-116** The G.fast NT Module/CPE SHOULD support US and DS byte counters on its WAN interface.

14 Relay Agent and Intermediate Agent Operation

The Relay Agent (RA) and Intermediate Agent (IA) functions provide the ability to insert port information in upstream session initiation requests arriving at the user port of a DPU. Examples of these requests are PADI for PPPoE and Discovery messages for DHCPv4 and v6. This information is derived from preconfigured information elements associated with each user port. These information elements contain strings that may assign port/customer identifiers that are dependent on an operator's conventions.

DPUs in a Model 2 deployment use a TR-156 compliant backhaul that permits the HON to have visibility of user port information through their one to one association with GEM ports. Therefore, the RA/IA functions are performed in the HON as depicted in Figure 14-2.

In accordance with TR-156, the access loop logical ports are identified using a syntax that relates to the parameters of the HON (e.g., access node logical name, chassis, rack, slot, port) and ONUID for each DPU on the PON interface. For example (taken from R-127/TR-156):

HON-Access-Node-Identifier eth Slot/Port/ONU-ID/Slot/Port[:VLAN-ID]

The resulting syntax is different from the syntax used by DPUs in a Model 1 / TR-167 deployment (ref Figure 14-1). In accordance with TR-167, the access loop logical port syntax relates to the parameters of the DPU. For example:

DPU-Access-Node-Identifier Slot/Port[:VLAN-ID]

In some cases it may be desirable for DPUs in a Model 2 / TR-156 deployment to use an access loop syntax that relates to the parameters of the DPU instead of the HON. To do this, a one-to-one mapping is needed between the ONUID used by the HON, the logical association towards the DPU, and the DPU port identifier. Such association can be done by the management layer. It could also be done by the network elements (e.g., HON), however the detailed specification is beyond the scope of this document.

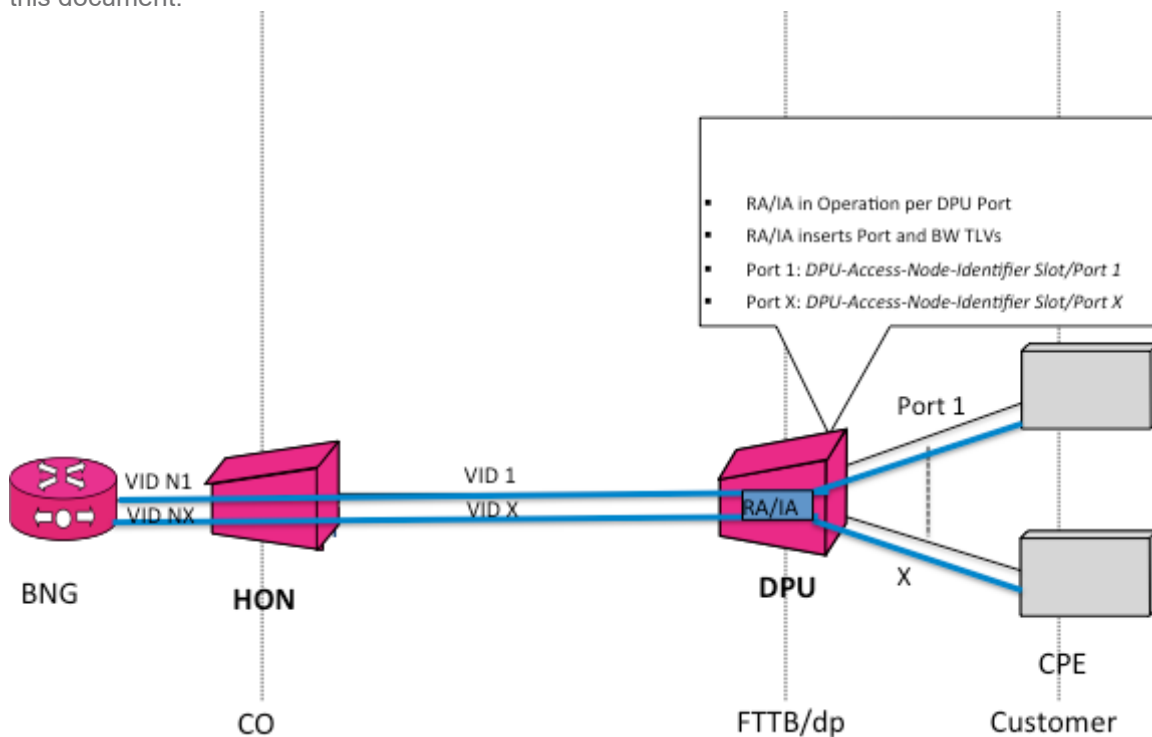


Figure 14-1 RA/IA in a Model 1 DPU

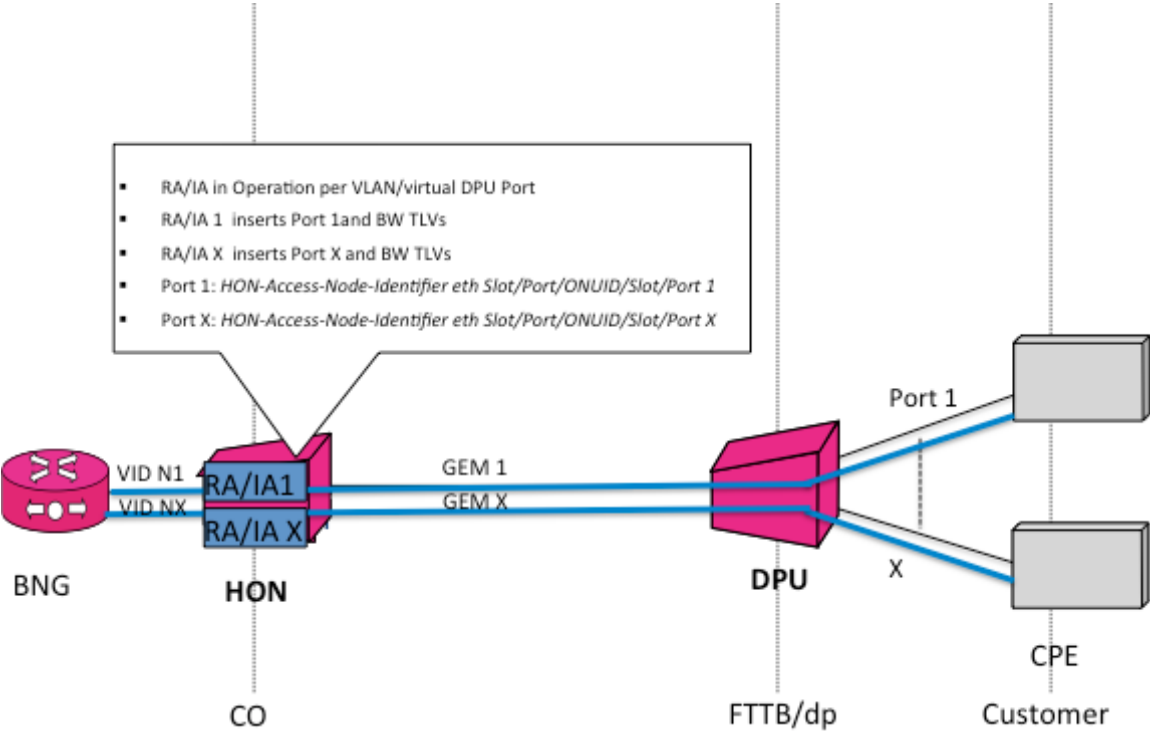


Figure 14-2 RA/IA In The HON (Model 2 DPU)

14.1 RA/IA Operation Requirements

R-117DPUs in a Model 1 deployment MUST comply with section 5.4.7 of TR-178.

R-118HONs in a Model 2 deployment MUST comply with section 5.4.7 of TR-178.

Note: This requirement may imply an interface between the PMA and HON that is not within the scope of this specification.

14.2 G.fast specific Type-Length-Values (TLVs)

It is appropriate to use a small subset of the Line Parameter already defined for other DSL technologies to report Net Data Rate and Attainable Net Data rate. VDSL2 specific parameters are already part of TR-101 and so are not repeated in this document.

R-119The RA/IA function MUST support the Sub TLVs in Table 14-1.

| Sub-TLV | Message Type | Information | Reference |
|---------|--|---|--------------------------------|
| 129 | Net data rate (NDR) upstream | Reports the net data rate upstream | ITU-T G.997.2 Section 7.11.1.1 |
| 130 | Net data rate (NDR) downstream | Reports the net data rate downstream | ITU-T G.997.2 Section 7.11.1.1 |
| 133 | Attainable net data rate (ATTNDR) upstream | Reports the attainable net data rate upstream | ITU-T G.997.2 Section 7.11.2.1 |
| 134 | Attainable net data rate (ATTNDR) downstream | Reports the attainable net data rate downstream | ITU-T G.997.2 Section 7.11.2.1 |
| 155 | Expected throughput (ETR) upstream | Reports the expected throughput upstream | ITU-T G.997.2 Section 7.11.1.2 |
| 156 | Expected throughput (ETR) downstream | Reports the expected throughput downstream | ITU-T G.997.2 Section 7.11.1.2 |

| Sub-TLV | Message Type | Information | Reference |
|---------|--|---|--------------------------------|
| 157 | Attainable expected throughput (ATTETR) upstream | Reports the attainable expected throughput upstream | ITU-T G.997.2 Section 7.11.2.2 |
| 158 | Attainable expected throughput (ATTETR) downstream | Reports the attainable expected throughput downstream | ITU-T G.997.2 Section 7.11.2.2 |
| 159 | Gamma data rate (GDR) upstream | Reports the gamma data rate upstream | ITU-T G.997.2 Section 7.11.1.3 |
| 160 | Gamma data rate (GDR) downstream | Reports the gamma data rate downstream | ITU-T G.997.2 Section 7.11.1.3 |
| 161 | Attainable gamma data rate (ATTGDR) upstream | Reports the attainable gamma data rate upstream | ITU-T G.997.2 section 7.11.2.3 |
| 162 | Attainable gamma data rate (ATTGDR) downstream | Reports the attainable gamma data rate downstream | ITU-T G.997.2 section 7.11.2.3 |

Table 14-1 G.fast Sub-TLVs

15 Diagnostics

15.1 Performance Monitoring

The ability to maintain a quality customer experience over time requires the ongoing collection of data on various aspects of DPU performance. These include, but are not limited to: user Ethernet frame counts, uplink PHY/MAC data and copper drop data. The DPU is responsible for keeping the current and the previous 15 minute interval for each PM counter. All other history counters beyond the 1 previous interval are maintained in the PMA.

15.1.1 DPU Performance Monitoring Requirements

The following requirements apply to the DPU performance monitoring capabilities:

R-120 The DPU MUST support the full set of counters that are defined in G.997.2 for G.fast and G.997.1 for VDSL2 on a per user port basis.

R-121 The DPU MUST support counters for the total number of bytes received and bytes transmitted on its backhaul interface.

R-122 The DPU MUST support counters for the total number of bytes received and bytes transmitted on each activated user port.

R-123 The DPU MUST support collecting data for the current 15 minute interval while simultaneously storing the previous 15 minute interval totals for all counts.

R-124 The DPU MUST support the retrieval of the previous 15 minute interval counts upon request from the PMA.

R-125 The DPU MUST support the retrieval of the current 15 minute interval counts upon request from the PMA at any time during that interval.

R-126 The DPU MUST use the same 15 minute interval for all its counters.

R-127 The DPU MUST support the synchronization of its 15 minute interval with the PMA.

R-128 REQUIREMENT DELETED

R-129 The DPU MUST be able to include periodically reported but not binned data in the 15 minute count report.

R-130 The DPU MUST accumulate the number of seconds spent in L0, L2.1, L2.1Bat, and L2.2 power states per activated drop line in each 15 minute period.

R-131 The PMA MUST upload each set of accumulated 15-minute counts from its powered up DPU before that set is overwritten.

R-132 The PMA MUST be able to store 96 sets of 15 minute counts.

R-133 The DPU MUST accumulate a 15 minute count of SRA events.

R-134 The DPU SHOULD log the rate change of each SRA event.

R-135 The DPU MUST measure the average power consumption on a per line basis over a 15 minute interval.

R-275 For deployment in a CloudCO architecture (TR-384[37]), the DPU MUST support IP Flow Information Export (IPFIX) Protocol as described in RFC 7011[39]/RFC 7012[40].

15.2 On Demand Diagnostics

- R-136** The DPU MUST report, on demand from the PMA, the DPU average power consumption since the beginning of the current 15 minute interval.
- R-137** The DPU MUST be able to measure Quiet Line Noise (QLN) and Hlog on a per provisioned line basis.
- R-138** The DPU MUST be able to determine downstream NDR, ATT NDR, and margin for each provisioned line.
- R-139** The DPU MUST be able to determine upstream NDR, ATT NDR, and margin for each provisioned line.
- R-140** The DPU MUST be able to detect a short circuit between the 2 legs of any drop pair.
- R-141** The DPU MUST be able to detect an open circuit on either leg of any drop pair.
- R-142** The DPU SHOULD be able to determine the distance (from the DPU) of short and open circuits with an accuracy of 10m or better.
- R-143** The DPU SHOULD be able to make a coarse measurement of DC voltage on any drop pair.
- R-144** When a ground connection is available, the DPU MUST be able to detect the presence of AC mains on any drop pair and report it to the PMA.
- R-145** When a ground connection is available, the DPU MUST be able to report 'safe to touch' on demand in the absence of AC mains on any drop pair.
- R-146** When a ground connection is available, the DPU MUST support SELT as defined in [G996.2] to do the following:
- Detect bridged taps.
 - Detect High-Resistance joints.
 - Measure Loop length.
- R-147** When a ground connection is available, the DPU MUST support MELT as defined in [G996.2] to do the following:
- Measure loop capacitance.
 - Measure loop resistance.
 - Pair identification via a configurable, voice-band tone.
- R-148** Sealed unit DPUs SHOULD contain a humidity sensor.
- R-149** If the DPU has an internal humidity sensor, it MUST raise an alarm if the humidity raises above a configurable threshold level.

16 Network Management

16.1 DPU Management Architecture

A DPU may have one of several different uplink technologies. These include but are not limited to GPON and point-to-point Ethernet. While these uplink technologies all include their own embedded management interfaces, expanding those interfaces to include common DPU management functions would result in a unique management interface per uplink technology. This increases complexity in large deployments that use a mix of uplink technologies. It can also increase the complexity of DPUs that use small form factor pluggable uplink transceivers by causing the management interface to change with the transceiver. This mix of possible DPU management interfaces would also make interoperability difficult to achieve since it would require a unique interoperability-testing ecosystem for each uplink technology. For these reasons, a DPU management architecture that uses a single, uplink technology agnostic management interface has been defined. While the management of uplink specific functions remains uplink technology specific, the management of common DPU functions is accomplished using a single management protocol and data model. Figure 16-1 depicts the functional split from a Model 1 DPU management domain perspective. Figure 16-2 depicts the functional split from a Model 2 DPU management domain perspective.

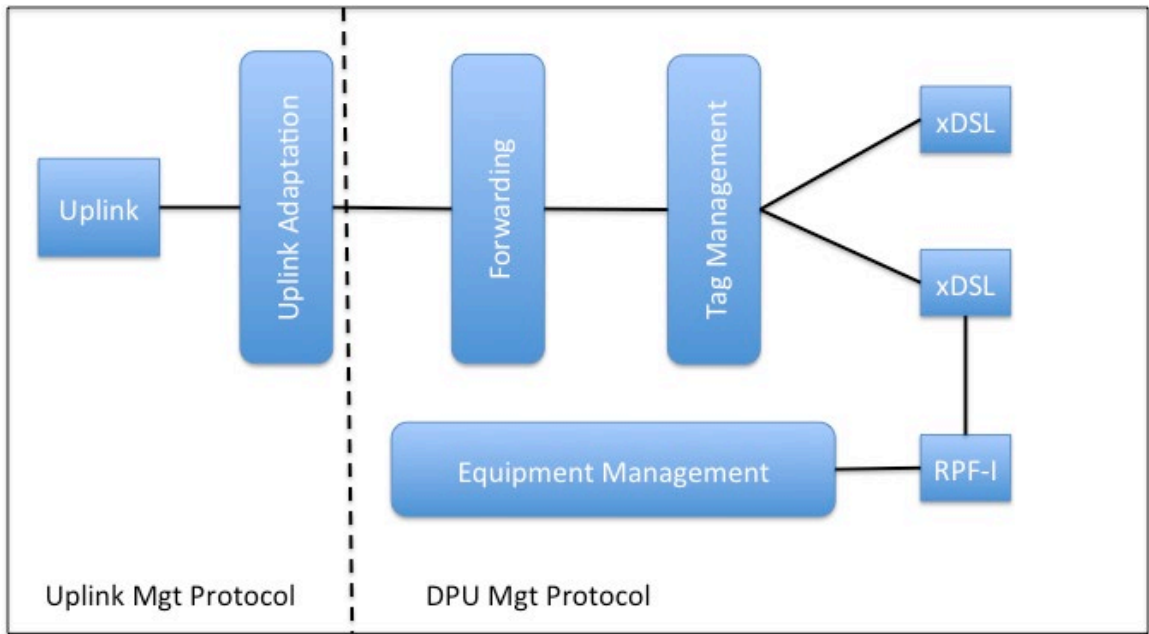


Figure 16-1 Model 1 DPU Management Domains

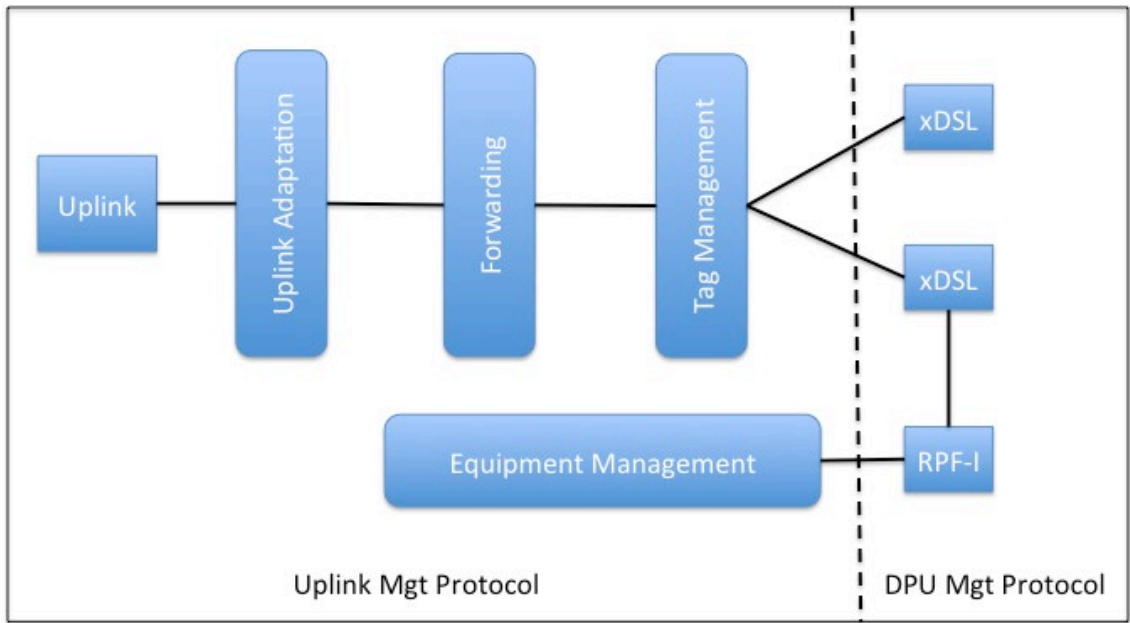


Figure 16-2 Model 2 DPU Management Domains

One aspect of FTTdp that must be accounted for in the management architecture is reverse powering. Reverse powering of the DPU means that it can be powered down at any time without the Network Operator’s advance knowledge or control. Most network management systems would treat such a spontaneous loss of power as a fault condition and raise an alarm, which is clearly not appropriate for FTTdp. This gave rise to the concept of a PMA.




Figure 16-3 DPU Management Architecture

In this architecture, there is a one to one correspondence between a PMA instance and the DPU it manages. Instances are aggregated up to higher order management systems. A PMA instance may be distributed across multiple compute platforms, but this distribution is not visible to either the DPU it manages or the higher order management systems. The PMA and aggregation functions may be realized as part of a PMA server that manages multiple DPUs, an intermediate system that aggregates

PMAs residing on lower order systems, or as part of the NMS function. They may also be implemented as virtualized functions on one or more network virtualization infrastructures.

R-150 There MUST be only 1 PMA associated with a given DPU at any one time.

R-151 There MUST be only 1 DPU associated with a given PMA at any one time.

This section describes the overall DPU management architecture (ref Figure 16-3) and specifies a set of high-level functional requirements. The PMA uses NETCONF [7] plus a set of YANG [8] data models to manage the DPU. The detailed YANG models are documented in separate Technical Reports.

R-152 The PMA and DPU MUST support NETCONF.

R-153 The PMA and DPU MUST support YANG.

NETCONF requires a persistent, reliable connection that supports authentication, data integrity, confidentiality, and replay protection. For the PMA-DPU connection, TLS is used to provide the functionality to meet the connection requirements of NETCONF.

R-276 The YANG modules to be used for managing the DPU are listed in TR-413[38].

R-154 The PMA and DPU MUST support TLS.

Permanent connectivity between the DPU and PMA is essential to the management of the DPU when it is powered up. There is a need for a regular keep-alive between these 2 entities to know when there is management plane connectivity; it is not sufficient to rely on the Dying Gasp to deduce that management connectivity has been lost.

R-155 REQUIREMENT DELETED

R-156 REQUIREMENT DELETED

R-229 The PMA MUST periodically transmit Transmission Control Protocol (TCP) keep-alive messages, as defined in RFC 1122 [25], to the DPU on the TCP connection established for call home when the connection is otherwise idle.

R-230 The DPU MUST respond to keep-alive messages from the PMA on the TCP connection established for call home.

R-157 The PMA MUST raise an alarm to the NMS upon detection of a keep-alive failure unless this was immediately preceded by a dying gasp from the DPU without a fault code.

R-158 If a DPU loses contact with its PMA it MUST continue to operate using its last known configuration until PMA connectivity is re-established.

R-231 The PMA SHOULD support configuration of the timing of TCP keep-alive messages and the threshold for detecting keep-alive failures.

R-232 The DPU MUST periodically transmit TCP keep-alive messages, as defined in RFC 1122 [25], to the PMA on the TCP connection established for call home when the connection is otherwise idle.

R-233 The PMA MUST respond to keep-alive messages from the DPU on the TCP connection established for call home.

R-234 In the event of a keep-alive failure, the DPU MUST re-initiate the NETCONF Call Home procedure as specified in the discovery process in Section 16.5.3, beginning with step 6b.

R-235 If the DPU fails to reconnect to the PMA, it MUST restart the discovery process beginning with the DPU Discovery message.

R-236 The DPU SHOULD support configuration of the timing of TCP keep-alive messages and the threshold for detecting keep-alive failures.

16.2 PMA Concepts

The fundamental purpose of the PMA is to allow the Operations Support Systems (OSS)/NMS to perform operations on a given DPU whether or not that DPU is currently accessible. This includes the following:

- Firmware download and management.
- Initial provisioning.
- Configuration.
- Test and diagnostics.
- Statistics gathering.
- Event reporting.

Of course, some of these operations have limited capabilities when the DPU is without power. For example, statistics gathering is limited to the history stored in the PMA since the current information is not available from the DPU.

The OSS/NMS still needs to be able to ascertain the true power state of the DPU as a whole, and of each given line, for example for diagnostics purposes. It can choose to take into account the power state of a DPU for various processes, e.g., a new firmware download, but this is not required.

The PMA is purely a functional entity and is not tied to any single platform or location. It may be hosted anywhere within a service provider's network that is always powered. The PMA may be deployed on the HON, an EMS, NMS, OSS or as a cloud based service. For example, a given service provider may choose to deploy the PMA within the OLT that is acting as the HON for a DPU. As shown in Figure 16-4, PMAs may be deployed in different locations, even within the same network.

The PMA assumes that there is a secure transport layer available between the PMA and DPU. This means that tasks such as PDU fragmentation and retransmission are assumed to occur in protocol layers below the actual management protocol.

The PMA introduces a new concept of a pending action. When the OSS/NMS attempts to carry out an action that needs the DPU to be powered up to be completed, the PMA acknowledges receipt and understanding of the action, but does not indicate it as complete until it has actually happened, i.e., when the DPU is next powered up. When a failure (e.g., power failure) occurs in the course of a DPU firmware upgrade, the PMA should gracefully recover.

R-159 The PMA MUST accept all valid OSS/NMS management commands on a given DPU irrespective of whether that DPU (or any of its lines) is powered up.

R-160 The PMA MUST store configuration changes and firmware downloads.

R-161 When the DPU is reachable, the PMA MUST keep the actual DPU configuration continually aligned with PMA DPU configuration representation.

R-162 When the DPU is unreachable, the PMA MUST:

- Keep the last known configuration of the DPU.
- Store a representation of the changes to configurations of the DPU until the DPU becomes reachable.

R-163 When the DPU becomes reachable again, the PMA MUST synchronize its provisioned configuration with the DPU configuration in a timely fashion.

R-164 The PMA MUST report to the NMS the status of DPU synchronization upon request.

R-165 The PMA MUST support the provisioning of a DPU that has not yet been installed.

R-166 The PMA MUST support the provisioning of services on a DPU that has not yet been installed.

R-167 The PMA MUST support the provisioning of ports and interfaces that have not yet been installed on an existing DPU.

R-168 The PMA MUST support the provisioning of services on ports and interfaces that have not yet been installed on an existing DPU.

R-169 The DPU MUST support the reporting of the following events via NETCONF Event Notifications (RFC 5277 [28]):

- Interface State Changes.
- Threshold Crossing Alerts.
- Equipment Alerts/Dying Gasps.

R-170 The PMA MUST support the receipt of events via NETCONF Event Notifications (RFC 5277):

R-171 The PMA MUST be able to pass on any event reported by the DPU to the NMS.

R-172 The PMA MUST be able to log the decrypted NETCONF messages exchanged between the PMA and DPU.

16.3 Management of Non-Reverse Powered DPUs

The main reason for the PMA concept is to allow management of DPUs to be undertaken regardless of their powering state. Although many DPUs are expected to be reverse powered, there are several Use Cases, especially for larger DPUs, which do not involve reverse powering.

However the PMA architecture is also relevant to non-reverse powered DPUs for the following reasons:

- The use of NETCONF/YANG provides the DPU data model and interoperability between the PMA and DPU.
- While the number of large DPUs will be less than if they were all 8/16 lines, in absolute terms there may be many more large DPUs deployed than existing access nodes such as DSLAMs. PMA aggregation therefore still helps with the scalability.
- There are deployments which may use a mixture of forward or local powered DPUs and Reverse Power Fed DPUs. Using PMAs with both provides a unified approach to management.
- One of the main functions of the PMA is to allow upgrades to be delayed until the DPU is powered up. Even with non-reverse powered DPUs, there may be some benefit in phasing upgrades for scaling reasons. The PMA can relieve the OSS of the responsibility of managing this phasing.

The same DPU management architecture, in particular use of the PMA, is specified regardless of whether the DPU is reverse powered.

16.4 DPU Management Architecture Applied to Routable and Non-Routable Address Domains

Although the DPU and PMA communicate using NETCONF over TLS over TCP/IP, in many cases the DPU does not require assignment of a routable IP address; avoiding a (publically) routable address provides better security. Where the PMA and the DPU can communicate with each other in the same subnet without a router between them, the DPU can use a non-routable (e.g., link-local) address. This is shown in Figure 16-4. In the figure, most of the DPUs are located in the same management subnet as the PMAs serving them, so these DPUs may use self-assigned link-local addresses.

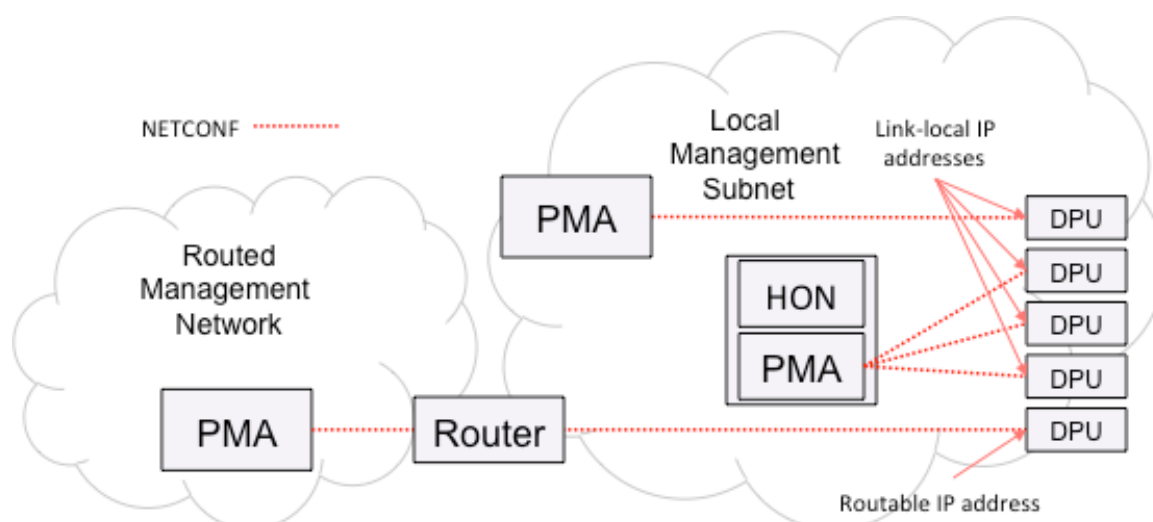


Figure 16-4 Management Architecture Applied to Routable and Non-Routable IP Address Domains

In other cases, the PMA may be in a separate subnet from the DPU within the management network, necessitating a routed connection between it and the DPU. In this instance the DPU has a routable IP address. This case is also shown in Figure 16-4. Note that the cases shown in Figure 16-4 are compatible with each other and may be deployed in the same network.

16.5 DPU-PMA Discovery and NETCONF Session Establishment

When a DPU is powered up, it needs to discover the correct PMA with which to associate. Both the installation process and PMA discovery must be supported without requiring manual configuration of IP addresses or other parameters in the DPU. PMA discovery requires the following steps:

- Upon powering up, a DPU creates a link-local IPv6 address on its management interface. The DPU sends a DPU Discover formatted as a set of options within a DHCPv6 Information-request message to discover the address of the PMA to which it should establish a connection. The DPU may also send other IPv4 DHCP and/or DHCPv6 messages to discover PMA Information and to establish a routable IPv4 or IPv6 address for the DPU.
- A DHCP/DHCPv6 server responds to a DPU Discover with a PMA Offer, which embeds the required PMA Information as sub-options within a DHCP or DHCPv6 message.
- 6.
- The DPU initiates a TCP connection to the PMA as defined in NETCONF Call Home and RESTCONF Call Home RFC 8071 [23]. The PMA responds by initiating TLS to the DPU. As part of the TLS session initiation, the DPU and the PMA authenticate each other.
- 7.
- The PMA initiates a NETCONF session over the established TLS session.
- 8.

These steps are defined in detail in the following subsections.

16.5.1 FTTdp Discovery Architecture

16.5.1.1 Use of DHCP

DHCP plays a central role in DPU-PMA discovery, as the DHCP (or DHCPv6 or stateless DHCPv6) server provides the PMA reachability information to the DPU. Since the DPU to PMA mapping is client-specific, the DHCP server needs access to the mapping information. The mapping mechanism is out of scope, however several example mechanisms are shown below:

- DHCP Relay may be used to relay DPU Discover to a centralized server associated with a network management database where the mapping is performed.
- The DHCP/DHCPv6 server may use a different out of scope mechanism to communicate with a centralized network management database where the mapping is performed.
- The DHCP/DHCPv6 server may be pre-configured with the subset of DPU-PMA mappings required for a localized network scope.

R-173A DHCP or DHCPv6 or stateless DHCPv6 server used for DPU-PMA discovery **MUST** support client-specific mapping of each DPU to its PMA.

If the DPU and PMA are within the same Layer 2 broadcast domain, they can communicate using IPv6 link-local IP addresses and there is no need for the DPU to establish a routable IP address. Otherwise, the DPU establishes a routable IPv4 or IPv6 address (using SLAAC, DHCPv6 or DHCP) to communicate with the PMA.

The PMA Information sub-option (section 16.5.2.2) used to convey PMA reachability information includes fields for Domain name, IPv4 and IPv6 addresses. While the address version used will typically match the DHCP version, there are valid exceptions. For instance, IPv6 link-local addressing may be used for DPU-PMA communication in a management network that uses IPv4 for other applications including DHCP.

In some deployment scenarios, more than one PMA may be offered to the DPU by the DHCP server. For example, a deployment that centralizes PMAs in a network data center may have a pool of available PMAs, any of which may be assigned to a DPU. If multiple PMAs are offered, the DPU attempts to establish connections to them in the order in which the PMAs are listed in the PMA Offer message. In other scenarios, the PMA’s IP address may be resolved from the domain name using DNS in order to support PMA clustering, load sharing, or other features.

16.5.1.2 DPU Identification

The DPU must identify itself to the DHCP server so it can be mapped to the correct PMA. The DPU Discover message provides two mechanisms to do this. The first mechanism is the DHCP Unique Identifier (DUID) as defined in RFC 3315 [20], which provides an identification code that is both globally unique and stable – no two DPUs should have the same DUID, and a DPU’s DUID should not change over time. RFC 3315 defines three ways to construct the DUID. In the DPU, DUID construction is limited to the DUID-EN option, with the vendor’s IANA registered private enterprise number followed by a unique serial number assigned by the vendor.

The DUID’s unique identifier field (i.e., unique serial number assigned by the vendor) is formatted in one of two ways.

- If the serial number formatted as a string can be interpreted as defined in G.987.3 [26], with 4 characters matching the vendor’s ID code followed by 8 hexadecimal integers, then the unique identifier is formatted as the 4-byte unsigned integer corresponding to the hexadecimal value in the last 8 characters. This formatting is used whether or not the serial number is actually constructed from a Vendor-ID code concatenated with a Vendor-Specific Serial Number (VSSN) per G.987.3.
- If the serial number cannot be interpreted as defined above, it is formatted as an ASCII character string.

ASCII string-based serial numbers are differentiated from 4-byte unsigned integers by adding null characters if necessary to extend them beyond 4 bytes in length. Examples of both types of formatting are shown in Table 16-1.

| Serial number | Format | DUID Unique Identifier (hex bytes) |
|---------------|-------------|------------------------------------|
| BBFX31324158 | 4-byte UINT | 31:32:41:58 |

| Serial number | Format | DUID Unique Identifier (hex bytes) |
|---------------|--------|------------------------------------|
| 12AX | String | 31:32:41:58:00 |
| BBFX12AX | String | 42:42:46:58:31:32:41:58 |

Table 16-1 Serial number formatting examples in the DUID unique identifier field

The second mechanism is the registration ID as defined in G.984.3 [11] and TR-156 [4]. The registration ID is used in some PON deployments to define the network location in which the ONU is installed, thus allowing ONUs to be swapped or selected from a pool at the time of installation. Registration ID can be used in the same way for DPU identification in DPU-PMA discovery. If the registration ID is configured in an integrated ONU/DPU device, the same configured parameter can be used by the ONU and by the DPU discovery functions.

R-174 The DPU MUST contain a DHCP Unique Identifier (DUID).

R-175 REQUIREMENT DELETED

R-237 The DUID MUST be in DUID-EN format as defined in RFC 3315.

R-238 If the DPU’s serial number matches the following three rules:

- 1. It is exactly 12 characters long.
- 2. The first 4 characters match the DPU vendor’s ID code as specified in ATIS-0300220 [27].
- 3. The last 8 characters are each in the range [0-9,A-F].

Then the unique identifier in the DUID-EN MUST be a 4-byte unsigned integer with the value of the hexadecimal number formed by the last 8 characters.

If the DPU’s serial number does not match all three of the above formatting rules in R-238, then requirements R-239 through R-241 apply:

R-239 The unique identifier in the DUID-EN MUST be formatted as an ASCII string.

R-240 The unique identifier MUST be equal to the DPU’s serial number, including vendor ID if any.

R-241 The unique identifier MUST be at least 5 characters long. Shorter serial numbers MUST be padded to a total length of 5 characters using null (0x00) characters.

R-176 The DPU SHOULD support configuration of a registration ID as defined in TR-156.

R-177 If the DPU includes an ONU in which a registration ID is configured, the DPU SHOULD use the same registration ID for DPU-PMA discovery.

R-178 For a DPU with a non-pluggable GPON backhaul, the serial number in the DUID SHOULD be identical to the GPON ONU VSSN.

16.5.1.3 Initiation of NETCONF connection and DPU-PMA Authentication

Before the DPU and the PMA can establish a NETCONF connection, they need to authenticate each other’s identities. The mutual authentication takes place using X.509 certificates, with each node tracing the other node’s certificates to a trust anchor. To validate against a trust anchor, the DPU contains one or more trusted certificates (e.g., pinned certificates, CA root certificates, intermediate certificates, or issuing certificates) that are either pre-loaded in the DPU or provided to the DPU by a secure means outside the scope of this specification.

The PMA and the DPU each contain unique certificates. The PMA’s certificate is unique to each PMAA, and the DPU’s certificate is unique to each DPU device. Each node (DPU and PMA) provides the certificates necessary to allow the other node to trace its identity back to the applicable trust anchor.

The vendor's enterprise identifier and the DPU's serial number together form a globally unique identifier in the IDevID certificate's subject field. These two values are concatenated in either the CN or the serialNumber attribute and separated by a hyphen. The serial number is formatted in one of two ways:

- The recommended format is a string containing all characters of the serial number, including the vendor ID if any, as it would appear on a label. When using this format, the CN attribute is not present and the concatenated string appears in the serialNumber attribute.
- Serial numbers containing a Vendor ID and VSSN as defined in G.987.3 [26] can be formatted as defined in TR-301 Issue 2 [32], in which the serial number is the decimal value equal to the 4-byte unsigned integer in the VSSN field. When using this format, the CN attribute contains the concatenated enterprise ID and VSSN, and the serialNumber attribute contains only the VSSN. The Vendor ID is not included. This format is provided for backwards compatibility with TR-301 Issue 2 [32].

Additional subject field attributes are out of scope. Examples of the subject field DN using both formats are shown in Table 16-2.

| SN format | Serial number | Subject field DN |
|---------------------------------|---------------|--|
| String | BBFX87654321 | /O=Broadband Forum/C=US/ST=CA/L=Fremont /serialNumber =3561-BBFX87654321 |
| VSSN converted to decimal value | BBFX05397FB1 | /O=Broadband Forum/C=US/ST=CA/L=Fremont /CN=3561-87654321/serialNumber=87654321 |

Table 16-2 SN Formatting examples in IDevID subject field

R-179 The DPU MUST provide an IDevID certificate to the PMA.

R-180 The DPU's certificate MUST be unique to the individual DPU unit.

R-242 The DPU's serial number SHOULD be formatted as a character string in the IDevID subject field.

If the serial number is formatted as a character string then requirements R-243 through R-245 apply:

R-243 The subject field MUST include the serialNumber attribute.

R-244 The serialNumber attribute MUST include the following elements, concatenated in order:

- The vendor's IANA registered enterprise number, formatted as a decimal value string;
- Hyphen (ASCII 0x2D);
- The DPU's serial number, formatted as a character string and including the vendor ID if present.

R-245 The subject field MUST NOT include the CN attribute.

R-246 For backwards compatibility with TR-301 Issue 2 [32], if the DPU's serial number is a concatenation of a 4-byte Vendor-ID and a 4-byte VSSN as specified in G.987.3 for ONUs, the serial number in the IDevID subject field MAY be formatted as a decimal value converted from the 4-byte unsigned integer in the VSSN field.

If the serial number in the IDevID subject field is formatted as a decimal value converted from the VSSN field then requirements R-247 to R-249 apply:

R-247 The subject field MUST include both the CN attribute and the serialNumber attribute.

R-248 The CN attribute MUST include the following elements, concatenated in order:

- The vendor's IANA registered enterprise number, formatted as a decimal value string;
- Hyphen (ASCII 0x2D);
- The device's serial number, formatted as a decimal value equal to the 4-byte unsigned integer in the VSSN field.

R-249 The serial number MUST have the same format in both the CN and serialNumber attributes.

R-181 REQUIREMENT DELETED

- R-182** The PMA MUST ensure that the presented DPU certificate has a valid chain of trust to a preconfigured trust anchor, and that the reference identifier from the certificate matches a preconfigured value before establishing a NETCONF connection.
- R-183** The PMA MUST provide an X.509 certificate which is unique to the individual PMAA associated with the PMA.
- R-184** The DPU MUST verify the integrity of the certificate presented by the PMA, either by performing path certificate validation per RFC 7589 and RFC 5280, OR by another trusted mechanism such as matching the presented certificate fingerprint against the configured certificate fingerprint.
- R-185** The PMA MUST verify the integrity of the certificate presented by the DPU, either by performing path certificate validation per RFC 7589 and RFC 5280, OR by another trusted mechanism such as matching the presented certificate fingerprint against the configured certificate fingerprint.

RFC 7589 [18] specifies the use of NETCONF over TLS with mutual X.509 authentication. In the FTTdp architecture, the DPU is the server for both NETCONF and TLS and the PMA is the client for both protocols. However, since the DPU is the node on which power may be cycled, it initiates the TCP connection to the PMA using NETCONF Call Home [23] rather than waiting for the PMA to initiate the connection.

16.5.2 Discovery Messages

The discovery process uses two message types, both of which are carried within DHCP or DHCPv6 messages. DPU Discover is sent from the DPU to initiate the process. PMA Offer is sent from the DHCP/DHCPv6 server to provide information on one or more PMAs to the DPU. While the discovery process may be based on either v4 or v6 IP message types, the support of IPv6 is mandatory and IPv4 is optional.

16.5.2.1 DPU Discover message

The DPU Discover message is sent by the DPU to identify itself and to request the IP address of the PMA with which it should connect. It can be formatted as any of several DHCP or DHCPv6 messages.

- R-186** A DPU MUST send DPU Discover formatted as a DHCPv6 Information-request message as part of the discovery process.
- R-187** A DPU MAY send DPU Discover formatted as a DHCPv6 Solicit, DHCPINFORM, and/or DHCPDISCOVER message as part of the discovery process.

The DPU Discover message includes the options shown in Table 16-3 and described below.

| Description | DHCPv6 option | | | DHCP option (IPv4) | | |
|-----------------------------|---------------|-----------------------------|---------|--------------------|--|---------|
| | # | Name | Ref | # | Name | Ref |
| DHCP Unit Identifier (DUID) | 1 | Client Identifier | RFC3315 | 61 | Client Identifier | RFC4361 |
| Requested parameters | 6 | Option Request | RFC3315 | 55 | Parameter Request List | RFC2132 |
| BBF-specific information | 17 | Vendor-Specific Information | RFC3315 | 125 | Vendor-Identifying Vendor-Specific Information | RFC3925 |

Table 16-3 DHCPv6 and DHCP options used for DPU Discover message

The DHCP Unique Identifier (DUID) is mandatory in all DPU Discover messages. Note that in addition to the DUID, the DPU Discover message can include a Registration ID as described below

R-188 A DPU MUST provide a DUID as defined in RFC3315 as option 1 in all DHCPv6 DPU Discover messages.

R-189 A DPU MUST provide a DUID as defined in RFC4361 as option 61 in all DHCP DPU Discover messages.

The requested parameters option is used in DHCPv6 and DHCP client messages to identify to the DHCP server the network information sought by the client. For DPU Discover, the requested information includes the vendor-specific information option. The DPU may request other information in addition to this option.

R-190 A DPU MUST request option 17 (Vendor-specific Information) in option 6 of all DHCPv6 DPU Discover messages.

R-191 A DPU MUST request option 125 (Vendor-Identified Vendor-specific information) in option 55 of all DHCP DPU Discover messages.

The Vendor-specific Information option is used to provide BBF-specific information as identified by the BBF enterprise-number (3561) assigned by IANA. The formatting used for the Vendor-specific Information option for DHCPv6 and DHCP messages is described in Appendix A. For DPU Discover, the option is included using the BBF enterprise-number value of 3561 and a SuboptionN-code of 192.

R-192 A DPU MUST include option 17 (Vendor-Specific Information) with Enterprise-number = 3561 and SuboptionN-code = 192 in all DHCPv6 DPU Discover messages.

R-193 A DPU MUST include option 125 (Vendor-Identifying Vendor-Specific Information) with Enterprise-number = 3561 and SuboptionN-code = 192 in all DHCP DPU Discover messages.

The TLVs used in the SuboptionN-data field for SuboptionN-code = 192 (DPU Discover) are shown in Table 16-4. Any Type value not listed in the table is reserved for future use.

| Type | Length | Value | Description |
|------|----------|-----------------|---|
| 00 | Variable | Registration ID | ID applied to the ONU for PON registration. |

Table 16-4 TLVs for DPU Discover

The Registration ID is described in Section 16.3.1.3. If a registration ID is configured in the DPU, it can be used by the network operator to facilitate mapping the DPU to a PMA.

R-194 If a DPU is configured with a registration ID, the DPU MUST include the Registration ID TLV in all DHCPv6 and DHCP DPU Discover messages.

16.5.2.2 PMA Offer message

PMA Offer is sent by a DHCPv6 or DHCP server to the DPU to provide reachability information for the PMA with which the DPU should connect. It can be formatted as DHCPv6 Reply, DHCPv6 Advertise, DHCPOFFER, or DHCPACK, depending on the type of server and the DHCP message type to which it is replying.

One or more PMA Information sub-options are encapsulated within the Vendor-specific Information option (option 17) in DHCPv6 messages and within the Vendor-Identifying Vendor-specific Information option (option 125) in DHCP messages, with the BBF Enterprise-number value of 3561 used to identify the vendor. The formatting used for the options 17 and 125 is described in Appendix A. Each instance of PMA Information is identified using the SuboptionN-code of 193.

R-195 A DHCPv6 server MUST include option 17 (Vendor-specific Information) with Enterprise-number = 3561 and at least one instance of SuboptionN-code = 193 (PMA Information) in all PMA Offer messages.

R-196 A DHCP server MUST include option 125 (Vendor-Identifying Vendor-specific Information) with Enterprise-number = 3561 and at least one instance of SuboptionN-code = 193 in all PMA Offer messages.

The TLVs used in the SuboptionN-data field for SuboptionN-code = 193 (PMA Information) are shown in Table 16-3. Any Type value not listed in the table is reserved for future use.

| Type | Length | Value | Description |
|------|-----------|---------------|---|
| 16 | 4 octets | IPv4 address | IPv4 address of the PMA being offered |
| 17 | 16 octets | IPv6 address | IPv6 address of the PMA being offered |
| 18 | 1 octet | PCP | Priority Code Point value to be used by the DPU for management traffic. |
| 19 | Variable | Domain name | The domain name of the PMA being offered. |
| 20 | 2 octets | TCP port name | Port number to be used by the DPU to open a TCP connection to the PMA being offered |

Table 16-5 TLVs for PMA Information

The TLVs in Table 16-5 are described below.

- IPv4 address, IPv6 address and Domain name each provide the reachability information in different forms to allow the DPU to establish a TCP connection with the PMA. Only one of these TLVs is typically provided.
- TCP port number provides the port number to be used by the DPU to open a TCP connection to the PMA being offered
- **PCP** provides the Priority Code Point value to be used by the DPU for all traffic sent to the PMA.

R-197 A PMA Information suboption MUST include either the IPv4 address, the IPv6 address or Domain name of the associated PMA.

R-273 A PMA Information suboption MAY include the port number to be used by the DPU to open a TCP connection to the PMA being offered. If the TCP port number is not included, it MUST default to port 4335.

R-274 If the PMA Information suboption includes the optional TLV for the TCP port number, it MUST NOT specify a port number assigned by IANA for a conflicting application.

For example, port number 4334 (assigned by IANA for NETCONF Call Home over SSH) cannot be specified in the TLV for NETCONF Call Home over TLS.

R-198 If a PMA Information suboption includes the Domain name TLV, the DPU MUST use the Domain name and ignore any IP address TLVs in the same PMA Information suboption.

R-199 If a PMA Information suboption includes the Domain name TLV, the Value field for the TLV MUST be encoded as defined in section 3.1 of RFC 1035 [24].

R-200 If a PMA Information suboption includes the Domain name TLV, the DPU MUST also acquire one or more DNS server addresses (via DHCP, DHCPv6, NDP, etc.).

R-201 The DPU MUST support DNS lookup over IPv4 and IPv6.

R-202 If the DPU needs to use a PCP value other than 0 for traffic sent to the PMA, the associated PMA Information sub-option MUST include the required PCP.

16.5.3 Discovery Process

The PMA discovery process starts when power is first applied to a DPU. The discovery process takes place as follows:

1. The DPU generates an IPv6 link-local address on its uplink interface and performs Duplicate Address Detection.
2. Using its link-local address as source address, the DPU sends a DPU Discover message formatted as DHCPv6 Information-request.
 - a. The DPU can also send the DPU Discover message formatted as DHCPDISCOVER and/or DHCHPv6 Solicit to solicit a routable address and/or to initiate discovery on IPv4 networks. It can also use SLAAC to generate a routable IPv6 address.
3. A DHCP or DHCPv6 server maps the DUID and/or registration ID encapsulated in the DPU Discover message to one or more PMAs.
4. The DHCP/DHCPv6 server responds with a PMA Offer message containing one PMA Information suboption for each PMA being offered. The PMA Information suboptions are listed under the Vendor-Specific Information option in the order in which the DPU should attempt to connect.
5. If the PMA Information suboption(s) in the PMA Offer contain Domain names, or if the PMA address(es) are IPv4 or IPv6 routable addresses, the DPU must complete the process of generating a corresponding routable address before proceeding to step 6. However, if the PMA IP address(es) contained in the PMA Offer are IPv6 link-local addresses, the DPU must proceed with step 6 using its own IPv6 link-local address.
6. The DPU performs the following steps when establishing a NETCONF connection with a PMA. In the steps below, the PMA is the NETCONF/TLS Client and the DPU is the NETCONF/TLS Server.
 - a. If the current PMA Information suboption contains a Domain name, the DPU performs a DNS lookup to resolve the domain name to an IP address.
 - b. The DPU opens a TCP connection to the offered PMA on the port as requested by the offered PMA (see section 16.5.2.2).
 - c. The PMA opens a TLS connection to the DPU as defined in RFC 7589.
 - d. The DPU and PMA each perform path certificate validation as per RFC 7589 and RFC 5280. The DPU ensures that the presented PMA certificate has a valid chain of trust to a pre-configured trust anchor and derives the PMA's NETCONF username from its certificate. The PMA ensures that the presented DPU certificate has a valid chain of trust to a pre-configured trust anchor and that the DPU's reference identifier matches the identifier pre-configured in the PMA. If either the PMA or the DPU cannot authenticate the other's identity, the connection is terminated.
 - e. After mutual authentication has completed and the TLS connection established, the DPU immediately starts the NETCONF server.
 - f. Once each node has authenticated the other, the PMA and the DPU can begin to exchange NETCONF messages sent as TLS application data, as specified in RFC 7589.
 - g. If the DPU cannot successfully establish a NETCONF connection with the PMA, it closes the TLS and/or TCP connections and attempts to establish a connection with the next offered PMA.

R-203 The DPU MUST attempt to establish a connection with each PMA offered in the order in which they are listed in the PMA Offer message.

R-204 If the DPU fails to connect to any of the offered PMAs, it MUST restart the discovery process beginning with the DPU Discovery message.

R-205 The DPU MUST NOT attempt to establish a connection with any more PMAs once it has successfully established a NETCONF connection with a PMA.

16.6 Management VLAN

All discovery traffic and DPU-PMA traffic uses a predefined management VID in the range from 0 (P-tagged) to 4094 at the DPU interface. If the network uses a different VLAN for management traffic to and from the PMA, it is the responsibility of the higher order node (HON) to which the DPU connects to perform tag operations that place the upstream and downstream DPU-PMA traffic on the correct VLANs. In the case of GPON being used as the DPU backhaul, the mapping of frames to the correct VLAN may take place in the backhaul function of the DPU by command from the HON. The PMA Information sub-option within the PMA Offer message contains a PCP field that provides the priority value to be used by the DPU for PMA traffic. The Higher Order Node (HON) should perform any VLAN tag manipulations necessary to convert between the management VLAN used by the DPU and the management VLAN used by the network.

R-206 The DPU MUST support use of a predefined VID in the range of 0 - 4094 for all discovery messages and all traffic sent to the PMA.

R-207 The DPU MUST support use of a predefined VID p-bit value for all discovery messages and all traffic sent to the PMA.

R-208 If the PMA Information contains a PCP TLV, the DPU MUST use the associated PCP value for DPU-PMA traffic.

16.7 Management Frame Handling

The processing of management frames within the DPU must take account of the correct identification and queuing of these frames as they are forwarded from the DPU to higher level management systems.

16.7.1 Management Frames In DPUs With Integrated Backhauls

DPUs with an integrated backhaul, have a single host processor that is used to manage both the backhaul and the PFFF. All Model 2 DPUs have backhauls with are integrated with the PFFF. Some Model 1 DPUs may have backhauls that are integrated with the PFFFs.

All management frames are created by software executing on the host processor with the TCID of the management frames under control of the software that creates the frames. When a management frame is created, the software uses the predefined VID and p-bit along with any VID and p-bit value provisioned on the backhaul to create the TCID that is used for forwarding the management frame toward the HON. Once the frame is created, the software places the frame into the appropriate upstream queue at the backhaul interface based on the contents of the TCID.

Example

A predefined VID of 6 and a p-bit value of 2 is used by the DPU. The backhaul is provisioned to translate all frames with a VID of 6 to a VID of 50 and p-bit value of 7 at the virtual interface between the backhaul and PFFF. The software on the DPU uses the combination of the predefined VID and the provisioned translation at the backhaul-PFFF virtual interface and creates all management frames with a VID of 50 and a p-bit value of 7. The frames are then placed directly into the upstream queue associated with p-bit value 7.

16.7.2 Management Frames In DPUs With Non-Integrated Backhauls

DPUs with a non-integrated backhaul, have separate host processors dedicated to the backhaul and PFFF management functions. Only Model 1 DPUs are expected to have non-integrated backhaul and these will typically take the form of a DPU with a pluggable backhaul like an SFP.

All management frames are created by software executing on the host processor responsible for management of the PFFF. When a management frame is created, the software uses the predefined VID and p-bit to create the TCID that is used for forwarding the management frame toward the HON. In the case of a pluggable backhaul, the software then places the frame into the appropriate upstream queue at the backhaul plug interface based on the contents of the TCID.

After the management frame is received across the interface from the PFFF, any tag/p-bit translation provisioned on the backhaul is performed and the frame is queued to the backhaul uplink interface based on the contents of the TCID.

16.8 Time Management

For the logging of events and alarms, the DPU needs to maintain the local time. This includes time zone and daylight savings time offsets since the PMA and DPU may be in different time zones or daylight savings time regions.

R-209 The DPU MUST support the setting and querying of the current system time zone by the PMA.

R-210 The DPU MUST support the setting and querying of the daylight saving time function by the PMA.

R-211 The DPU MUST support an interface to set the time format between the PMA and the DPU (UTC local time).

R-212 The DPU MUST support sending an alarm to the PMA with event log information using the DPU time local time or UTC time.

The Network Time Protocol (NTP) is widely used to synchronize computer clocks in the Internet. NTP version 4 (NTPv4) described in RFC 5905, is backwards compatible with NTP version 3 (NTPv3), described in RFC 1305, as well as previous versions of the protocol. NTPv4 includes a modified protocol header to accommodate the Internet Protocol version 6 address family. It includes a dynamic server discovery scheme, so that in many cases, specific server configuration is not required.

R-213 The DPU MUST support an NTPv4 client.

16.9 PMA Aggregator

The PMAA provides an aggregation function between the individual PMAs and the NMS/OSS. This has a number of benefits, namely:

- Reducing the number of management channels that need to be setup, monitored and maintained i.e., 1 per PMAA rather than 1 per PMA.
- Reducing the number of IP addresses that the NMS needs to be aware of
- Reducing the maintenance burden on the NMS, for example:
 - The PMAA can manage the scheduling of firmware upgrades across all its associated PMAs
 - Only a single copy of the new firmware needs to be downloaded to the PMAA.

16.9.1 PMAA Location

Many of the benefits of the PMAA stem from it being co-located with its PMAs, but this is not a requirement, and the PMAA can be in an entirely different physical location.

While the PMAA and its PMAs can be physically separated, they are indivisible from a functional perspective. The interface between PMAA and PMA is NOT exposed, and is therefore not a defined point of interoperability. This means that both the PMAA and the PMA need to come from the same vendor.

16.9.2 Northbound Interface

The northbound interface between the PMAA and the NMS is beyond the scope of this specification.

16.9.3 PMAA Security & Scalability

While security is an issue for any management system, the PMAA does not of itself introduce any new security requirements.

The whole point of the PMAA is to increase the management system scalability, but this is implementation dependent and not subject to any quantitative formal requirements.

16.9.4 PMAA and PMA Addressing

At the southbound interface, each PMA has a one-to-one relationship with a DPU. Each PMA communicates with its respective DPU via a TCP socket identified by a unique combination of IP address and TCP port number. The PMA/PMAA architecture may be designed in different ways. In each case below, the DPU communicates with the IP address used by the PMA for southbound communication. The configuration of IP addresses in the PMAA and the PMAs is out of scope for this specification.

- PMAs may be fully contained within a PMAA (see Figure 16-5). The PMAA may be implemented as a standalone device or as a function within a device such as a switch, an OLT, or a network management server. The PMAs communicate with their respective DPUs using a common network interface and are differentiated by the combination of IP address (which may also be the address used by the PMAA for northbound communication) and TCP port number.
- PMAs may be implemented separately from the PMAA, which aggregates the northbound traffic to and from the NMS (See Figure 16-6). Each PMA has its own IP address to communicate with its DPU. The PMAs may be implemented in the same device as each other and as the PMAA, or in different devices that may be distributed in the network. Details of the PMA/PMAA interface are out of scope.

9.

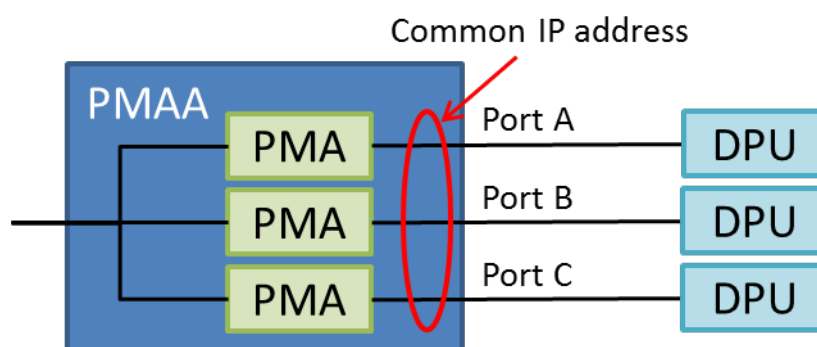


Figure 16-5 PMAs Sharing a Common Address




Figure 16-6 PMAs With Unique Addresses

16.10 Software image Management

R-214 The DPU MUST be able to store at least 2 software images.

17 Operations and Maintenance

17.1 DPU Installation

Service provider technicians rather than users install DPUs. One can envision installation teams installing all the DPUs required to cover a service area as a focused activity weeks in advance of the activation of service. This means that the DPU needs to be installed without disruption to existing service and be tested for correct operation without the presence of customer premises equipment. The DPU should also support the ability to be swapped out with a different unit should the initial tests fail. Ideally, an installer should be able to install a DPU without requiring support from personnel at the management head end.

17.1.1 DPU Installation Requirements

R-215 The DPU MUST support connection of a handheld tester, which can emulate a modem and reverse power feed to any user port of the DPU to facilitate installation, commissioning, and trouble shooting

R-216 The DPU MUST be capable of providing system status to its installer prior to contacting its PMA.

Note: The DPU should be able to be replaced quickly in the event of failure, with minimal customer down-time – this will have implications for the connection method.

17.1.2 DPU Startup With POTS From Exchange/Cabinet

This section applies to DPUs that support the deployment scenarios with POTS from the CO/exchange or cabinet present (refer to options 2 and 3 in Table 1 of ETSI 101 548-1 [3]). The high level start-up sequence for a reverse powered DPU with POTS from the CO/exchange or cabinet is depicted in Figure 17-1.

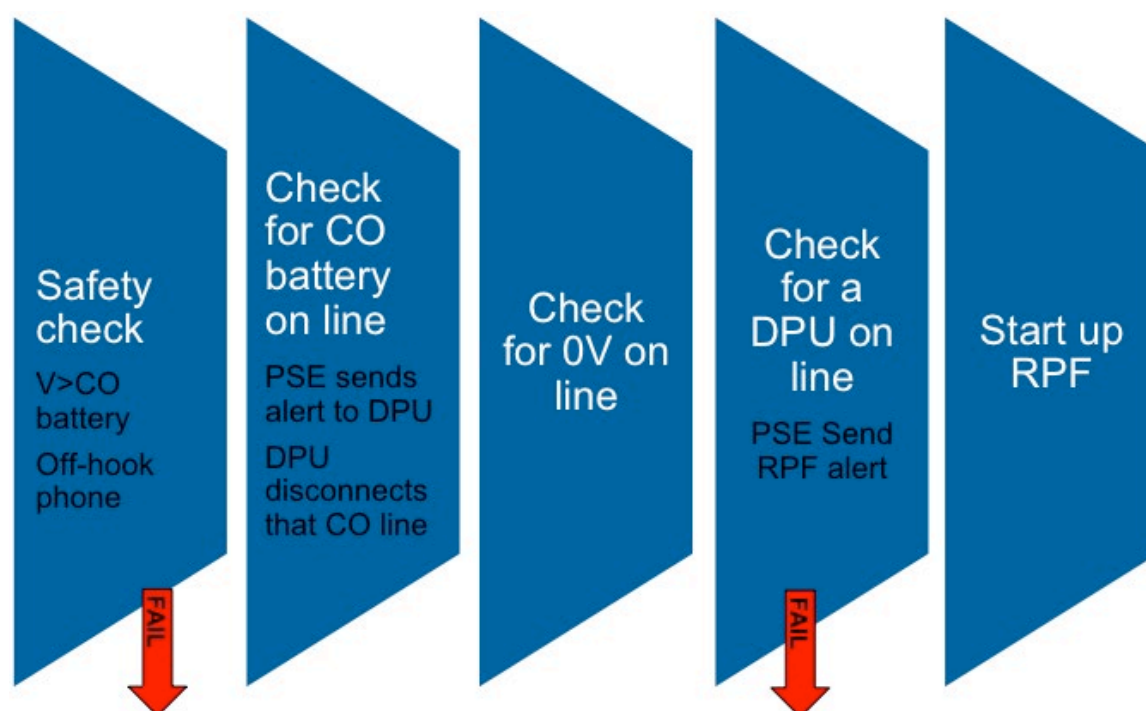


Figure 17-1 DPU Startup with Exchange/Cabinet POTS

DPUs and PSEs deployed in scenarios that include POTS delivered from the CO/exchange or cabinet must comply with the following requirements.

- R-217** A DPU MUST support initial installation such that the direct metallic path from the CO/Exchange or cabinet to each customer on that DP is retained, and all voice-band and DSL-based broadband services can continue to be delivered to those customers on their original pair, and without additional impairment.
- R-218** If the DPU has ports for incoming CO/exchange or cabinet copper lines, the DPU MUST support initial installation such that the DC is blocked from the CO/cabinet
- R-219** The DPU MUST support the optional PRP protocol as defined in Section 6.2.5 of ETSI TS 101 548-1 [3].
- R-220** Upon detection of POTS disconnect trigger SR1, SR2, SR3, SRany on a given pair, as defined in ETSI TS 101 548-1 [3], the DPU MUST disconnect any exchange copper connection to that pair, subject to any start-up mode override set by network management.
- R-221** The CO/exchange or cabinet connection MUST be able to be reinstated via network management.
- R-222** The DPU MUST NOT reconnect any CO/exchange or cabinet line automatically after a power cycle.
- R-223** The PSE MUST support the optional PRP protocol as defined in clause 6.2.5 of ETSI TS 101 548-1 [3].

17.2 CPE Installation

Users perform CPE installation in an FTTdp deployment. These installations require the FTTdp network to support auto-configuration and remote management such that service provider personnel are not required to visit the DPU or user site. Auto-configuration and remote management are assumed to include:

- The ability of the DPU to automatically detect the presence of a new CPE and to report its presence to its PMA, and to enable services on that CPE according to pre-provisioned attributes.
- The ability to remotely control the service state of ports on a DPU.

In addition to auto-configuration and remote management, the FTTdp network may support the ability to reconfigure the copper loop such that it is physically disconnected from the CO and connected to the DPU without a site visit. This also requires the ability to perform the reverse action when the CPE is removed. This capability is referred to as Remote Copper Reconfiguration (RCR).

17.2.1 CPE Installation Requirements

- R-224** A DPU MUST support the migration of individual customers on that DP to a DPU based service.
- R-225** Requirement R-224 MUST be able to be done remotely e.g., via a management system.
- R-226** Requirement R-224 MUST NOT require a visit to the DPU.
- R-227** The DPU MUST support the remotely activated metallic disconnection of a customer's individual pair from the CO/exchange or cabinet as part of the migration of that customer to a DPU based service.
- R-228** The DPU MUST support remotely activated reversion of individual customers on that DP to a CO/Exchange or cabinet-based service.

Appendix A – DHCPv6 and DHCP Vendor Specific Option formatting

This Appendix describes how BBF-specific information fields used for DPU-PMA discovery are mapped to a DHCPv6 or a DHCP message.

A.1 DHCPv6 option 17 formatting

The DPU uses Vendor-specific Information option 17 to provide DPU Discover information regarding its certifications and supported features over DHCPv6 for discovery. The DHCPv6 server uses the same option to provide one or more PMA Offers, including certification, supported features, and other data, to the DPU. The format for the Vendor-specific Information option is shown below and described in RFC3315. Multiple instances of this option can be included in a DHCPv6 message with each instance containing a unique enterprise-number.

1 1 1 1 1 1 1 1 1 2 2 2 2 2 2 2 2 3 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


| | |
|----------------------------|----------------|
| Option-code = 17 | Option-len |
| Enterprise-number = 0x0DE9 | |
| Suboption1-code | Suboption1-len |
| Suboption1-data | |
| ... | |
| Suboption2-code | Suboption2-len |
| Suboption2-data | |
| ... | |
| ... | |

Optional

| Field | Length | Description |
|-------------------|----------|--|
| Option-code | 2 octets | Vendor-specific Information (17) |
| Option-len | 2 octets | Total length of all following option data in octets. This value is exclusive of the option-code and option-len octets. |
| Enterprise-number | 4 octets | The vendor's 32-bit Enterprise Number as registered with IANA. The Broadband Forum value is 3561 (0x0DE9). |
| SuboptionN-code | 2 octets | For DPU-PMA discovery, the type of discovery message. 192: DPU Discover, sent by the DPU to request discovery information. Only one instance of this suboption type may occur within a DHCPv6 message. 193: PMA Information, sent by a DHCPv6 server to provide PMA reachability information to a DPU. One instance of this suboption type may occur for each PMA being offered. |
| SuboptionN-len | 2 octets | The length of the SuboptionN-data field in octets. |
| SuboptionN-data | Variable | The SuboptionN-data field contains information specific to the SuboptionN-code field. See Table A1-2 for formatting. |

Table A1-1 DHCPv6 Option 17 fields

The suboptionN-data fields are formatted as a series of Type/Length/Values (TLVs). The TLVs have the format shown below.

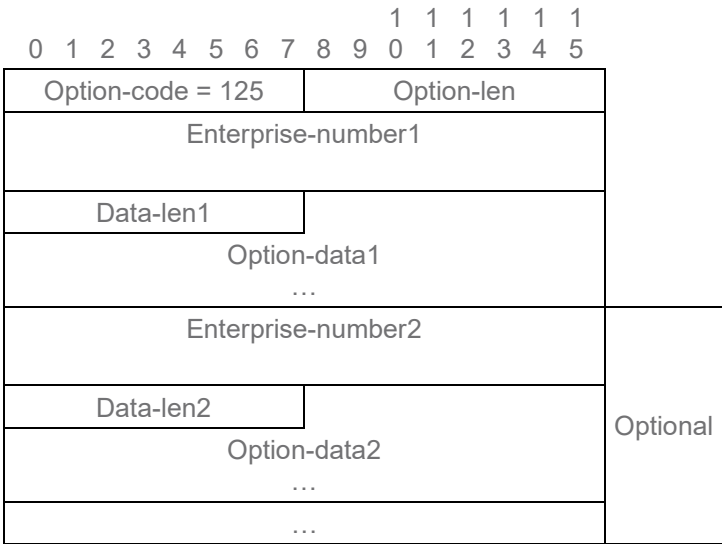


| Field | Length | Description |
|---------|----------|--|
| TypeN | 2 octets | The TypeN field identifies the type of data contained in the ValueN field. |
| LengthN | 2 octets | The LengthN field contains the length of the ValueN field in octets. |
| ValueN | Variable | The Value field is zero or more octets and contains information specific to the Type. The format and length of the Value field are determined by the Type and Length fields. |

Table A1-2 Format of BBF-specific TLV fields for DHCPv6 Option 17

A.2 DHCP (IPv4) option formatting

The DPU uses Vendor-Identifying Vendor-specific Information option 125 to provide DPU Discover information regarding its certifications and supported features over DHCP for discovery. The DHCP server uses the same option to provide one or more PMA Offers, including certification, supported features, and other data, to the DPU. The format for the Vendor-Identifying Vendor-specific option is shown below and described in RFC3925. The option may also be divided into a series of smaller options as defined in RFC3396.



| Field | Length | Description |
|--------------------|----------|--|
| Option-code | 1 octet | Vendor-Identifying Vendor-specific Information (125) |
| Option-len | 1 octet | Total length of all following option data in octets. This value is exclusive of the option-code and option-len octets. |
| Enterprise-numberN | 4 octets | The vendor's 32-bit Enterprise Number as registered with IANA. The Broadband Forum value is 3561 (0x0DE9). |
| Data-lenN | 1 octet | Length of Option-dataN field in octets |

| | | |
|--------------|----------|--|
| Option-dataN | Variable | Vendor-specific option data. When Enterprise-numberN = 3561 and used for DPU-PMA discovery, see Table A2-2 for formatting. |
|--------------|----------|--|

Table A2-1 DHCP Option 125 fields

When used with the Broadband Forum enterprise-number for DPU-PMA discovery, the option-data field for option 125 is formatted as one or more suboptions as shown in Table A2-2.

| | | | | | | | | | | | | | | | |
|-----------------|---|---|---|---|---|---|---|---|----------------|---|---|---|---|---|---|
| 0 | 1 | 2 | 3 | 4 | 5 | 6 | 7 | 8 | 9 | 0 | 1 | 2 | 3 | 4 | 5 |
| Suboption1-code | | | | | | | | | Suboption1-len | | | | | | |
| Suboption1-data | | | | | | | | | | | | | | | |
| ... | | | | | | | | | | | | | | | |
| Suboption2-code | | | | | | | | | Suboption2-len | | | | | | |
| Suboption2-data | | | | | | | | | | | | | | | |
| ... | | | | | | | | | | | | | | | |
| ... | | | | | | | | | | | | | | | |

| Field | Length | Description |
|-----------------|----------|--|
| SuboptionN-code | 1 octet | For DPU-PMA discovery, the type of discovery message. 192: DPU Discover, sent by the DPU to request discovery information. Only one instance of this suboption type may occur within a DHCP message. 193: PMA Information, sent by a DHCP server to provide PMA reachability information to a DPU. One instance of this suboption type may occur for each PMA being offered. |
| SuboptionN-len | 1 octet | The length of the SuboptionN-data field in octets. |
| SuboptionN-data | Variable | The SuboptionN-data field contains information specific to the SuboptionN-code field. See Table A2-3 for formatting. |

Table A2-2 BBF-specific Option-data field

When used with the Broadband Forum enterprise-number for DPU-PMA discovery, the SuboptionN-data field is formatted as a series of Type/Length/Values (TLVs). The TLVs have the format shown in Table A-5. This formatting is similar to that for the TLV fields for Options 17 as defined in section A.1, except that the Type and Length fields are only one octet in length.

| | | | | | | | | | | | | | | | | | |
|--|--------|---|---|---|---|---|---|---|---|---------|---|---|---|---|---|---|---|
| | 0 | 1 | 2 | 3 | 4 | 5 | 6 | 7 | 8 | 9 | 0 | 1 | 1 | 1 | 1 | 1 | 1 |
| | Type1 | | | | | | | | | Length1 | | | | | | | |
| | Value1 | | | | | | | | | | | | | | | | |
| | ... | | | | | | | | | | | | | | | | |
| | Type2 | | | | | | | | | Length2 | | | | | | | |
| | Value2 | | | | | | | | | | | | | | | | |
| | ... | | | | | | | | | | | | | | | | |

| Field | Length | Description |
|---------|----------|--|
| TypeN | 1 octet | The TypeN field identifies the type of data contained in the ValueN field. |
| LengthN | 1 octet | The LengthN field contains the length of the ValueN field in octets. |
| ValueN | Variable | The Value field is zero or more octets and contains information specific to the Type. The format and length of the Value field are determined by the Type and Length fields. |

Table A2-3 Format of BBF-specific TLV fields for DHCP Option 125

Appendix B Certificate Management

This section contains informative certificate management clarifications and requirement scope limitations in TR-301 Issue 2 [32].

B.1 Certificate Management

TR-301 Issue 2 [32] Section 16.1 describes the DPU management architecture. Between the DPU and PMA, the use of NETCONF over TLS is mandated as the communication transport for management information. NETCONF over TLS provides transaction confidentiality, data integrity, and requires certificate-based mutual authentication between the DPU and PMA.

Per RFC 7589, mutual authentication between TLS client and server have to be assured before a NETCONF session is established. Both peers use X.509 certificate path validation (refer to RFC5280 section 6) to verify the integrity of the certificate presented by the peer. The presented X.509 certificate may also be considered valid if it matches one obtained by another trusted mechanism, such as certificate pinning which uses a locally configured certificate fingerprint.

RFC 7589 defines the process by which the NETCONF client (PMA) validates the NETCONF server's (DPU) identity. The server presents an identifier in its certificate and the client checks its configured reference identifiers against the server presented identifier. For the DPU, the reference identifiers are a 2-tuple manufacturer and serial numbers (both of which are encoded in the IdevIDs certificate subject field).

As part of the authentication process referenced by RFC 6241, the transport protocol provides the NETCONF server with an authenticated NETCONF client identity (or username) whose permissions are known to the server. The access permissions for the authenticated client are then enforced by the server for the remainder of the NETCONF session.

RFC 7589 defines the algorithm used within the server to derive a NETCONF username from a certificate when TLS is used as the transport protocol. The server maintains an ordered list of mappings of certificates to NETCONF usernames. Each entry in the list contains: a certificate fingerprint used for matching; a map type defining how the username is derived from the matching certificate; and optional auxiliary data containing the username if the map type is 'specified.'

To support interoperability in cases where the DPU is not pre-configured with information about the PMA, the server can use the map type 'specified' to provide a username to NETCONF, resulting in a default set of access permissions which allows the NETCONF client in the PMA to perform the functions required of it. Any other certificate mappings or access permissions beyond the default set are outside the scope of this document.

Certificate management functions that do not directly impact DPU- PMA management interface interoperability are considered out of scope in TR-301 Issue 2 [32]. In many cases, the requirements associated with these functions are expected to vary from one network operator to another. Examples of these functions include:

- Certificate fields other than those specified in this document.
- Requirements for certificate authorities.
- Methods and infrastructure requirements pertaining to distribution or renewal of certificates, revocation of certificates, or certificate revocation lists.

End of Broadband Forum Technical Report TR-301