TECHNICAL REPORT

TR-301

Architecture and Requirements for Fiber to the Distribution Point

Issue: 1
Issue Date: August 2015

Notice

The Broadband Forum is a non-profit corporation organized to create guidelines for broadband network system development and deployment. This Broadband Forum Technical Report has been approved by members of the Forum. This Broadband Forum Technical Report is not binding on the Broadband Forum, any of its members, or any developer or service provider. This Broadband Forum Technical Report is subject to change, but only with approval of members of the Forum. This Technical Report is copyrighted by the Broadband Forum, and all rights are reserved. Portions of this Technical Report may be copyrighted by Broadband Forum members.

The Forum draws attention to the fact that it is claimed that compliance with this Specification may involve the use of a patent ("IPR") concerning TR-301. The Forum takes no position concerning the evidence, validity or scope of this IPR.

The holder of this IPR has assured the Forum that it is willing to License all IPR it owns and any third party IPR it has the right to sublicense which might be infringed by any implementation of this Specification to the Forum and those Licensees (Members and non-Members alike) desiring to implement this Specification. Information may be obtained from:

British Telecommunications plc
PPC5A, BT Centre, 81Newgate Street, London, EC1A 7AJ, UK

Attention is also drawn to the possibility that some of the elements of this Specification may be the subject of IPR other than those identified above. The Forum shall not be responsible for identifying any or all such IPR.

THIS SPECIFICATION IS BEING OFFERED WITHOUT ANY WARRANTY WHATSOEVER, AND IN PARTICULAR, ANY WARRANTY OF NONINFRINGEMENT IS EXPRESSLY DISCLAIMED. ANY USE OF THIS SPECIFICATION SHALL BE MADE ENTIRELY AT THE IMPLEMENTER'S OWN RISK, AND NEITHER THE FORUM, NOR ANY OF ITS MEMBERS OR SUBMITTERS, SHALL HAVE ANY LIABILITY WHATSOEVER TO ANY IMPLEMENTER OR THIRD PARTY FOR ANY DAMAGES OF ANY NATURE WHATSOEVER, DIRECTLY OR INDIRECTLY, ARISING FROM THE USE OF THIS SPECIFICATION.

Broadband Forum Technical Reports may be copied, downloaded, stored on a server or otherwise re-distributed in their entirety only, and may not be modified without the advance written permission of the Broadband Forum.

The text of this notice must be included in all copies of this Broadband Forum Technical Report.

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

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

Editors	Michael Shaffer Dong Wei	Alcatel-Lucent Huawei Technologies
Fiber Access Network WG Chairs	Michael Shaffer Lin Wei	Alcatel-Lucent Huawei Technologies
End-to-End Architecture WG Chairs	David Allan David Thorne	Ericsson BT
End-to-End Architecture Vice Chair	Sven Ooghe	Alcatel- Lucent

TABLE OF CONTENTS

List of Figures

List of Tables

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.

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 High 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 MDU and 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 (IEEE802.3-2012 GbE, 10GbE [13])
- PON (ITU-T GPON [11] and XG-PON1 [12])

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 POTS and 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	*European Requirements for Reverse Powering of Remote Access Equipment*	ETSI	2014

[4]	TR-156 Issue 3	*Using GPON Access in the context of TR-101*	BBF	2012
[5]	TR-167 Issue 2	*GPON-fed TR-101 Ethernet Access Node*	BBF	2010
[6]	TR-178	*Multi-service Broadband Network Architecture and Nodal Requirements*	BBF	2014
[7]	RFC 6241	*Network Configuration Protocol (NETCONF)*	IETF	2011
[8]	RFC 6020	*YANG - A Data Modeling Language for the Network Configuration Protocol (NETCONF)*	IETF	2010
[9]	G.9701	*Fast access to user terminals (G.fast) - Physical layer specification*	ITU-T	2014
[10]	G.993.2	*Very high speed digital user line transceivers 2 (VDSL2)*	ITU-T	2011
[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 MultiService Architectures*	BBF	2008
[16]	RFC 3376	*Internet Group Management Protocol V3*	IETF	2002

2.3 Definitions and Acronyms

The following terminology is used in this Technical Report.

Complex DPU	A DPU that complies with all the requirements for Ethernet Access Nodes found in TR-101 and/or TR-178.
CPE	Customer Premises Equipment.
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	High Order Node. The first node upstream of the DPU.
Hybrid DPU	A DPU that supports both G.fast and VDSL2.
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.
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.
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.
WiFi	A generic name used to refer to all versions of IEEE 802.11

2.4 Abbreviations

This Technical Report uses the following abbreviations:

AC	Alternating Current
AFE	Analog Front End
AN	Access Node
ANCP	Access Node Control Protocol
ATA	Analog Terminal Adapter
BNG	Broadband Network Gateway
CBSU	Communications Based Start Up
CC	Continuity Check
CO	Central Office
DC	Direct Current
DECT	Digital Enhanced Cordless Telecommunications
DSLAM	Digital Subscriber Line Access Multiplexer
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
10GbE	10 Gigabit Ethernet
GPON	Gigabit Passive Optical Network
IA	Intermediate Agent
IGMP	Internet Group Management Protocol
IP	Internet Protocol
IPv4	Internet Protocol version 4
IPv6	Internet Protocol version 6
L2	Layer 2
MDSU	Metallic Detection Start Up
MDU	Multi-Dwelling Unit
MEP	Maintenance association End Point
MIP	Maintenance domain Intermediate Point
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
PON	Passive Optical Network
POTS	Plain Old Telephone Service
QLN	Quiet Line Noise
QoS	Quality of Service
RA	Relay Agent
RG	Residential Gateway
SFU	Single Family Unit
SFP	Small Form-factor Pluggable Transceiver
SSH	Secure Shell
TCP	Transmission Control Protocol
TLV	Type Length Value
TR	Technical Report
UC	Use Case
VLAN	Virtual Local Area Network
VoIP	Voice over IP
WG	Working Group
xDSL	Any Digital Subscriber Line Service

XGPON	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.

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 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 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 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 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 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-101/178, 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 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 DC voltage feed and DC signalling. Instead the voice service could be delivered as a derived service (VoIP), terminating in an ATA, 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 signalling 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 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 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 Uses 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 Case 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 (GPON, XGPON1) 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 Power Supply Equipment (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 which 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 fiber to the cabinet (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 High Order Node (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 Persistent Management Agent (PMA), which acts as a management proxy for the DPU when it was unpowered.
- Where the backhaul is PON-based, it must be possible to operate TR-156 ONTs and TR-167 ONUs on the same ODN as the new WT-301 DPUs. Ideally this would not require any changes to the 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 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.

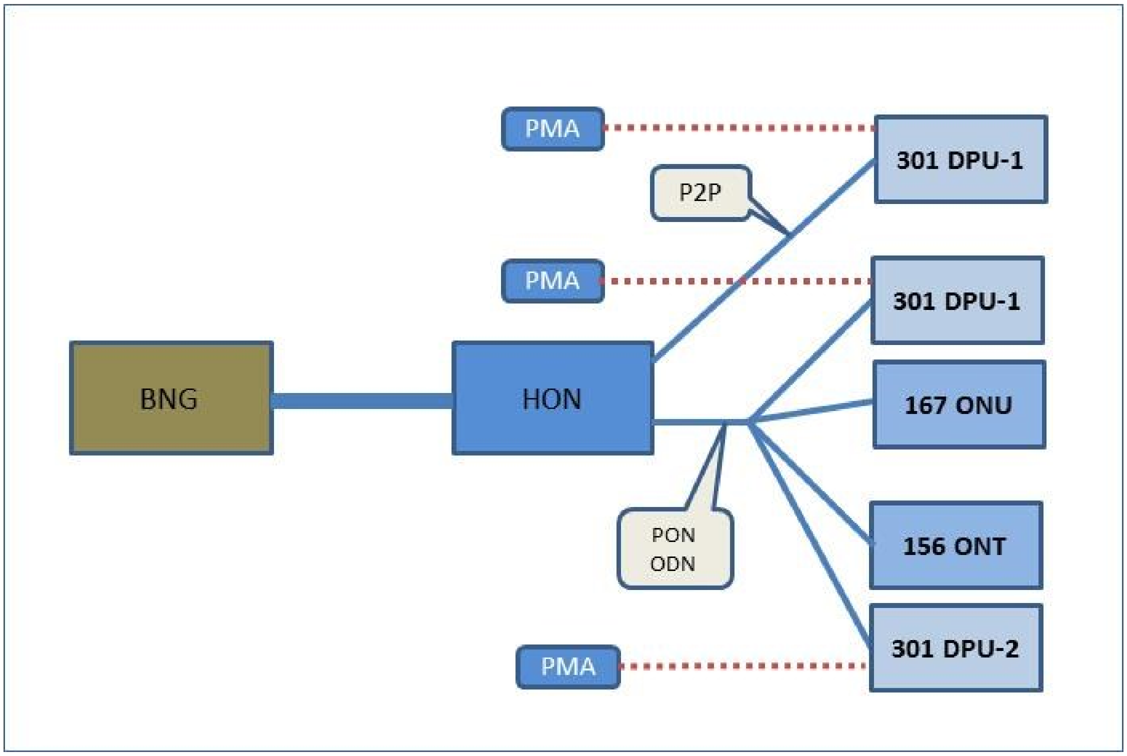

Figure 5-1 High level FTTdp deployment models and TR156/TR167 co-existence

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 (see 5.1 and 5.2 below) deployments. It is also likely that a DPU for Model 1 deployments could use an SFP to support both TR-167 compliant GPON and point-to-point Ethernet backhauls. This would then allow a single device to be used for both backhaul types, and for the migration between backhaul types after initial installation.

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

When point-to-point Ethernet or TR-167 [5] compliant GPON/XG-PON1 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 BNG. Additionally, N:1 VLANs are supported for the purpose of multicast delivery.

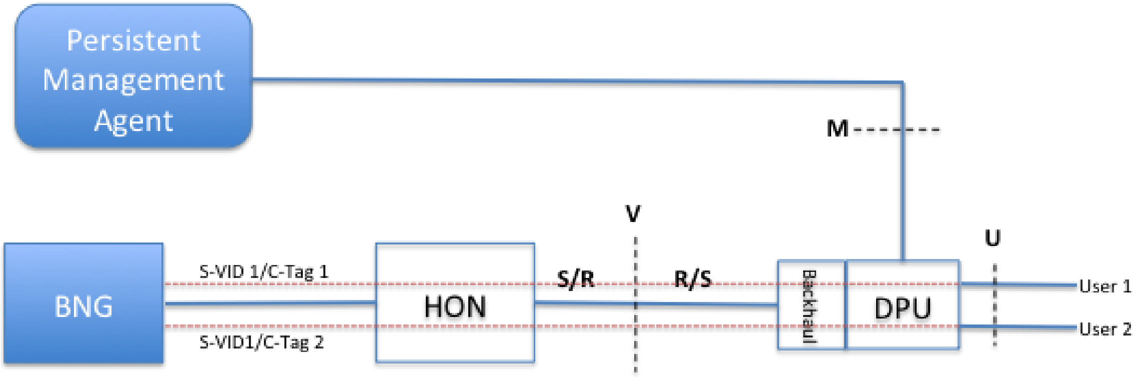

Figure 5-2 Deployment Model 1 (Pt-toPt Ethernet/TR-167 Backhaul)

For downstream multicast and broadcast frames, S-VLANs may be used to achieve efficiency in the use of GPON backhaul, The DPU provides 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 (ref: section 14).

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/XG-PON1 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 GEM ports that are unique to each virtual Ethernet Interface. The HON is able to perform 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.

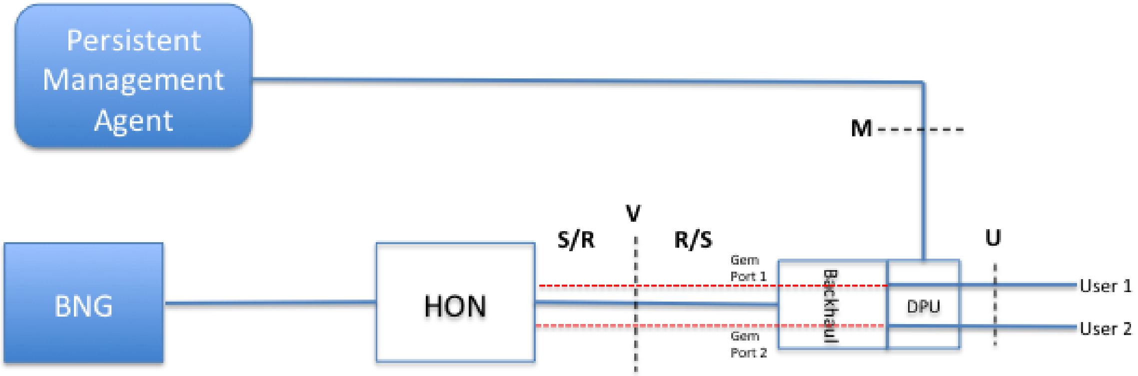

Figure 5-3 Deployment Model 2 (TR-156 Backhaul)

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, GPON, XG-PON1, or point-to-point Ethernet may be used.

ITU-T G.9701 and ETSI TS 101 548 are relevant to this BBF architecture document; the following diagram shows the generic view of references described by ITU-T and ETSI. It should be noted that, dependent on backhaul technology, different reference points might occur.

Figure 6-1 provides a reference model for the end-to-end FTTdp deployment. Note that the POTS and reverse power feed reference points are determined by ETSI.

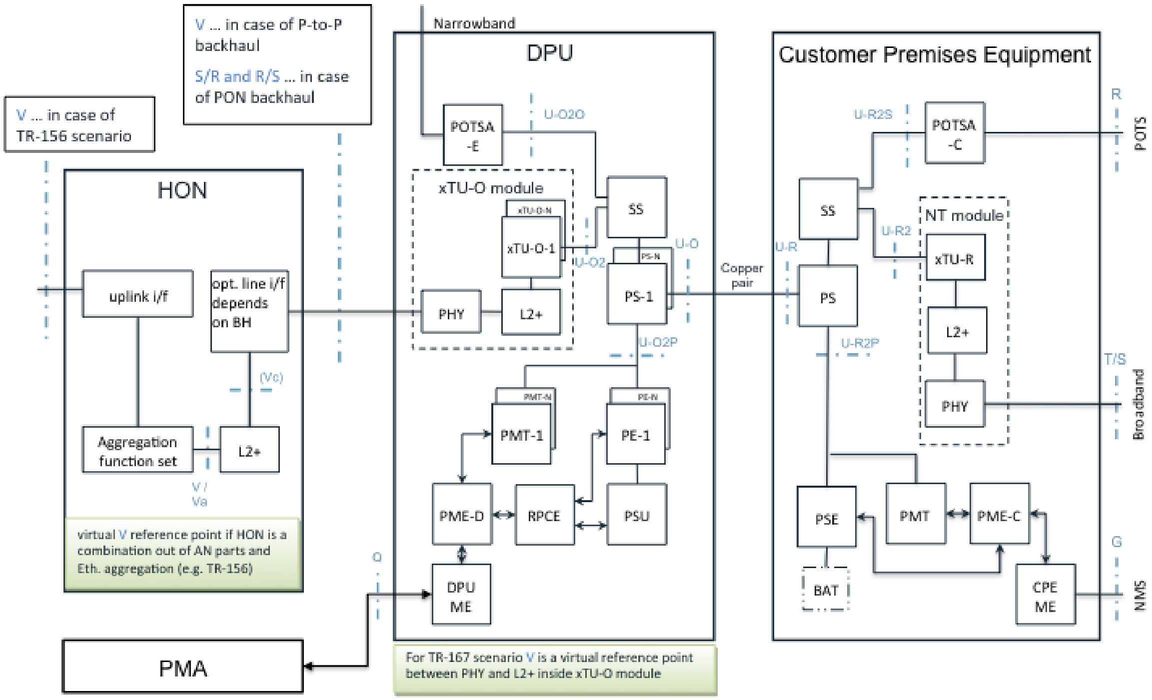

Figure 6-1 FTTdp Functional Reference Model

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.

Some service provider deployments require the support of baseband telephony over the same twisted pair as reverse power feed *within the customer premises* while others do not. In recognition of these different deployment scenarios, two RPF types are considered here. Type 1 uses a communications based startup method (CBSU) which is based on a simple message exchange between PSE and DPU. Type 2 uses a metallic detection start up method (MDSU) based on detection of the resistive signature located in the DPU. The requirements in this section apply to both types. Detailed specifications for both types are defined in [3].

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 A DPU MUST take a roughly equal share of power (as measured at the DPU) from all connected, powered-on users when operating in full power mode.

R-14 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-15 The DPU MUST support sending the appropriate Dying Gasp code plus the identifier of the affected line, if applicable, from the list below to its PMA immediately prior to shutting down:

- Unknown/Undefined
- Last reverse power feed to DPU disappeared with no fault indication from the its PSE
- Safe temperature of DPU exceeded
- Last error code received from a PSE

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

R-16 The DPU and PMA MUST maintain the following current RPF status for each of its user lines:

- Line is powered by PSE on mains
- Line is powered by PSE on battery power (default state is false)
- NT Module on battery power. If the NT Module and PSE are integrated this state is equal to the PSE on battery state.

R-17 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:

- Loss of power (valid in both Mains and Battery states)
- Unprotected off-hook phone
- Excess current demand

R-18 When the DPU is powered up, it MUST respond to periodic keep-alive messages from the PMA.

R-19 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-20 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-21 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-22 The DPU MUST be able to be configured on a per port basis to NOT draw power from a line when it is in PSE-on-battery state, if power is available from other sources.

R-23 The DPU MUST notify the PMA of the receipt of a Dying Gasp from a user line.

R-24 The DPU start up protocol MUST operate irrespective of the presence of MELT-P signatures encountered in FTTdp deployments.

R-25 Type 1 (CBSU) and Type 2 (MDSU) based DPU reverse powering functionality MUST comply with ETSI TS 101 548 [3].

R-26 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:

- Keep the line in service
- Move the line to low-power mode
- Shutdown the transceiver, i.e. take the line out of service
- Support emergency services only

R-27 The DPU MUST be able to automatically put a provisioned port in-service when reverse powering is detected on that line.

R-28 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 20ms per feeding line at a repetition rate of 20 seconds.

R-29 A DPU SHOULD prevent electrical noise and 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-30 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.

8.1.1 DPUs With RCR

The following requirements apply to DPUs that support RCR.

R-31 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-32 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-33 The PSE of a single active line MUST be able to power the DPU it is connected to in both mains-powered and battery-powered operation.

R-34 The PSE MUST comply with one of the power classes defined in ETSI TS 101 548 [3].

R-35 The PSE MUST send a Dying Gasp indication to the DPU immediately before it removes power from the line.

R-36 The PSE MUST be able to start up a DPU from an unpowered state.

R-37 The PSE MUST promptly remove power from a line upon the detection of a fault.

R-38 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 a 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-39 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-40 Type 1 (CBSU) and Type 2 (MDSU) user side reverse powering functionality MUST comply with ETSI TS 101 548 [3].

R-41 Type 1 and Type 2 PSE MUST comply with one or more of the RPF classes specified in ETSI TS 101 548 [3]

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, 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.

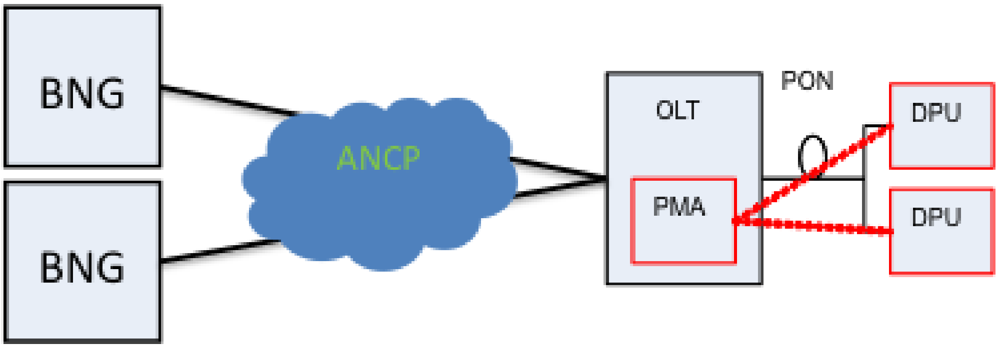

Figure 9-1 ANCP With PMA In The HON

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

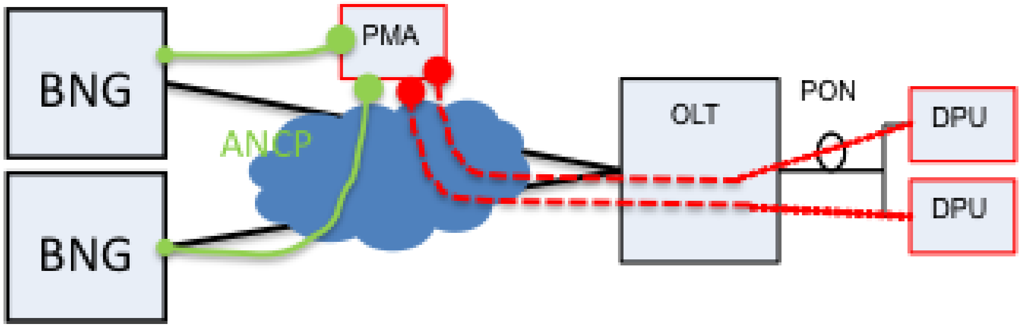

Figure 9-2 ANCP With The PMA External To The HON

R-42 A DPU MUST support at least one of the following customer facing copper drop technologies:
- G.fast
- VDSL2

R-43 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-44 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-45 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-46 The DPU MUST be able to report the attainable 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-47 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 attainable NDR on any port changes by more than a configurable threshold.

R-48 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-49 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-50 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).

9.2 DPU Backhaul Physical Interface Requirements

R-51 A DPU MUST support at least one of the following backhaul types:
- GPON [11]
- Gigabit Optical Ethernet [13]
- 10 Gigabit Optical Ethernet [13]
- XG-PON1 [12]

10 Traffic Management and 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 of 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. The PFFF then removes or translates tags prior to their transmission to the user. 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.

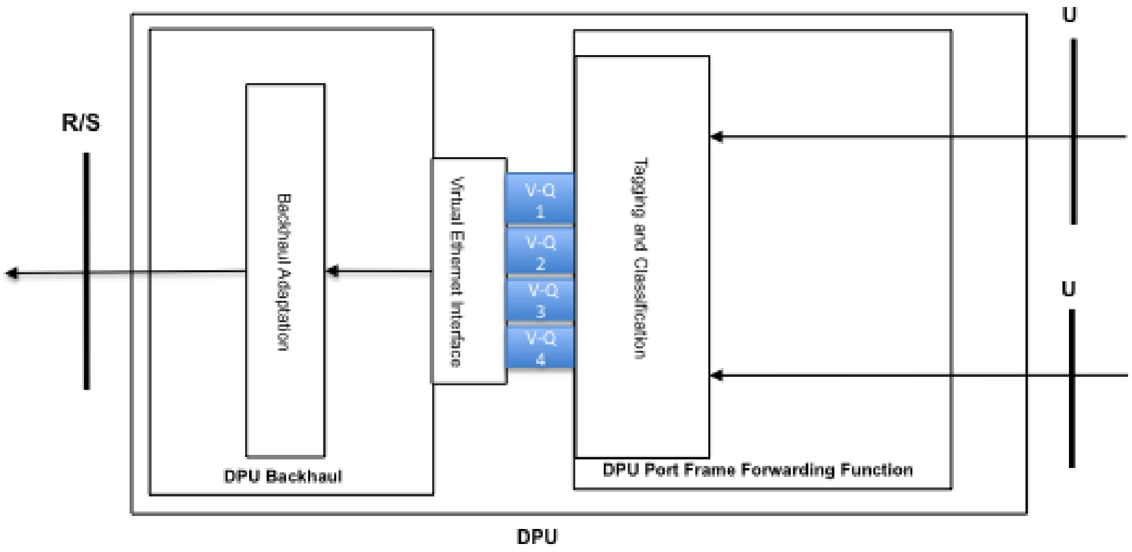

Figure 10-1 Deployment Model 1 DPU Upstream Frame Handling

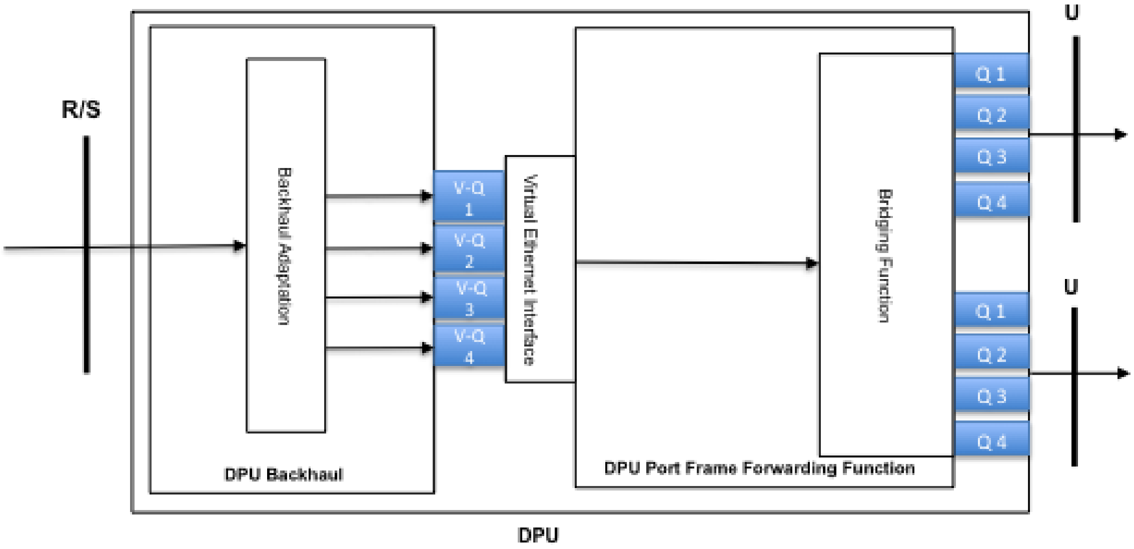

Figure 10-2 Deployment Model 1 Downstream Frame Handling

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.

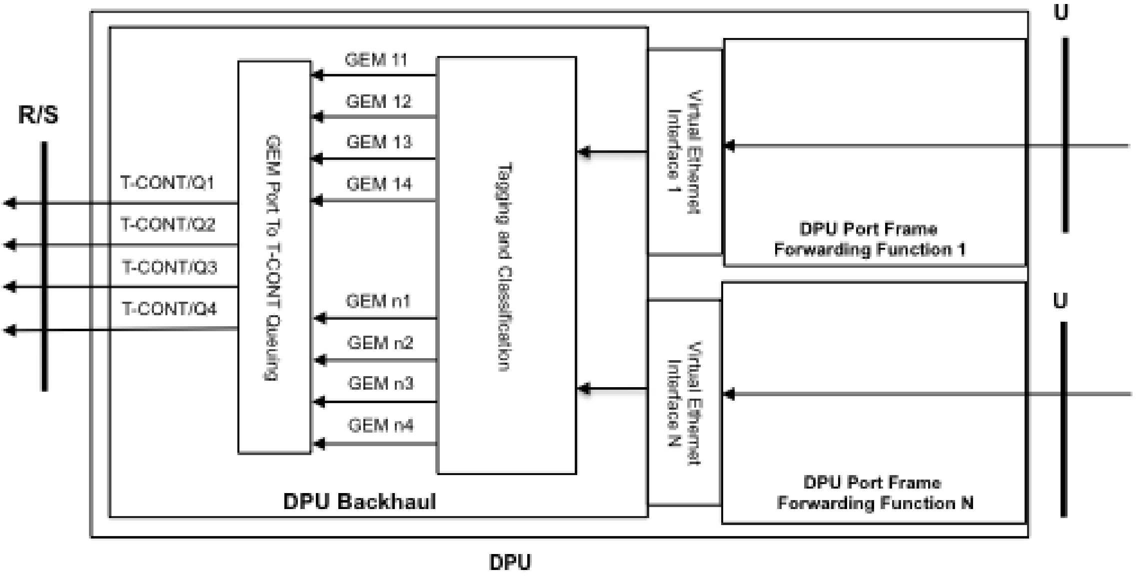

Figure 10-3 Deployment Model 2 DPU Upstream Frame Handling

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.

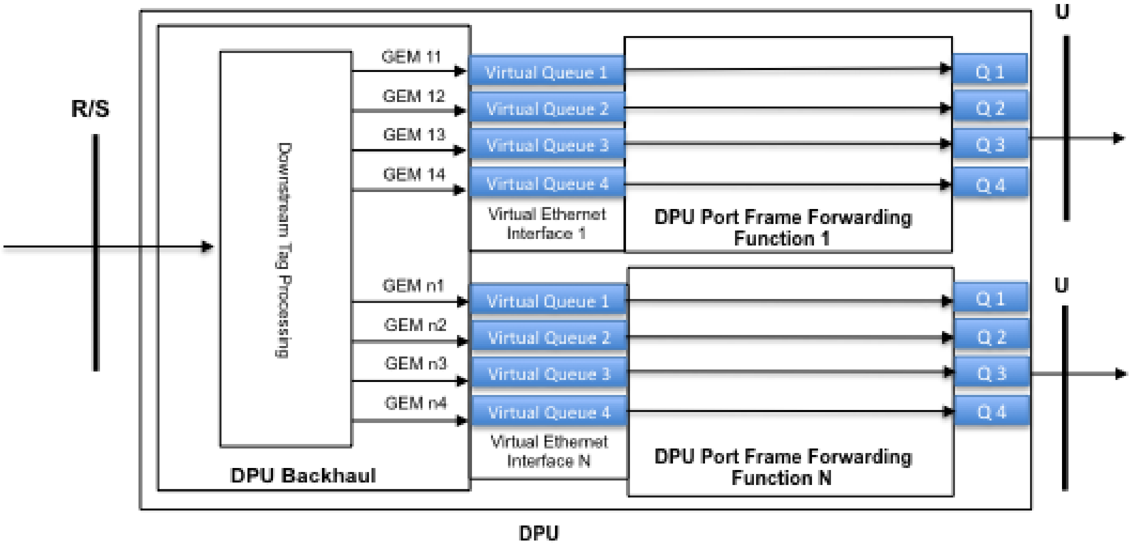

Figure 10-4 Deployment Model 2 Downstream Frame Handling

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-52 In the downstream direction, the DPU MUST support 4 strict priority queues per user port.

R-53 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-54 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-55 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-56 The DPU MUST support configuration of the actions required by R-54 and R-55 by the PMA.

11 VLAN Handling

Deployment Model 1 (ref: 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 a 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 a 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 (ref: 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-57 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-58 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-59 The DPU PFFF MUST support Ethernet tag addition, removal, or translation based on an arbitrary combination of: user port, VID, and received P-bit markings.

R-60 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-61 The DPU PFFF MUST support Ethernet tag addition, removal, or translation based on EtherType.

R-62 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-63 The DPU PFFF MUST support both 0x8100 (C-tag) and 0x88A8 (S-tag) Ethertype values.

R-64 The DPU MUST NOT alter any VLAN tags beyond the outer two, and MUST treat any additional tags as part of the payload.

R-65 The DPU PFFF MUST perform any necessary VID and P-bit manipulations before performing the mapping into upstream queues.

R-66 The DPU PFFF MUST perform the mapping of frames into downstream queues prior to any necessary VID and P-bit manipulations.

R-67 The DPU PFFF MUST support multiple P-bit values being used in the same VLAN.

R-68 The DPU PFFF MUST NOT prevent multiple VLANs from using the same P-bits.

R-69 The DPU PFFF MUST support at least 4 simultaneously active VLANs per user port.

R-70 The DPU PFFF MUST support at least 16 simultaneous tagging operation rules per user port.

11.2 Deployment Model 2 DPU VLAN Requirements

R-71 The DPU PFFF MUST transparently forward upstream frames from the user port to the corresponding virtual Ethernet interface.

R-72 The DPU PFFF MUST transparently forward downstream frames from the virtual Ethernet interface to the user port.

R-73 The DPU Backhaul MUST support at least 4 simultaneously active VLANs per user port.

R-74 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

R-75 The DPU PFFF MUST support enabling and disabling IGMP snooping on a per user port basis.

R-76 The DPU PFFF MUST support processing IGMP packets on a per VLAN basis.

R-77 The DPU PFFF MUST support the identification and processing of user-initiated IGMP messages. When this function is disabled on a port and/or VLAN, these messages MUST be transparently forwarded.

R-78 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-79 The transparent snooping function MUST be able to snoop the multicast source IP address and destination IP group address in IGMP packets, and set the corresponding MAC group address filters as specified in R-80.

R-80 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-81 The transparent snooping function MUST be able to translate the upstream VLAN of IGMP packets to the configured VLAN pertaining to that 1:n multicast VLAN

R-82 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-83 The DPU PFFF MUST support matching groups conveyed by IGMP messages to a provisioned list of multicast groups corresponding to a multicast VLAN associated with the receiving user port.

R-84 When no match is found by R-83, the IGMP message MUST be either forwarded or dropped, based on configuration.

R-85 When there is a match found by R-83, the IGMP message MUST be forwarded within a multicast VLAN, and frames from the matching group forwarded to the requesting port and VLAN.

R-86 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 IGMP 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 normally be sealed units, and may be located in fairly inaccessible locations. Therefore fairly comprehensive 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-87 The DPU MUST support configuring a MEP on its backhaul interface.

R-88 The DPU MUST support configuring a MEP on each of its end-user facing ports.

R-89 The DPU SHOULD support configuring a MIP on each of its end-user facing ports.

R-90 The Maintenance Domain level of each MEP and MIP MUST be configurable.

R-91 The DPU MUST support IEEE802.1ag and Y.1731 Loopback and Link Trace.

R-92 The DPU MUST support IEEE802.1ag and Y.1731 CC.

R-93 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-94 The DPU MUST support IEEE802.3ah OAM clause 57.2.9 active mode.

R-95 The DPU MUST support configuring IEEE802.3ah loopback (enable and disable) on the CPE via each end-user facing interface,.

R-96 The DPU MUST forward looped traffic from the CPE to its WAN interface on the normal data path.

13.2 CPE OAM Requirements

R-97 The G.fast NT Module/CPE MUST support configuring IEEE802.3ah Passive mode.

R-98 The G.fast NT Module/CPE MUST be able to configure Loopback (enable/disable) when requested by an Active Mode DPU.

R-99 The G.fast NT Module/CPE SHOULD support 802.1ag/Y.1731.

R-100 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/ONUID/Slot/Port[:VLAN-ID]

The resulting syntax is different from the syntax used by DPUs in a Model 1 / TR-167 deployment (ref Figure 14-2). 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.

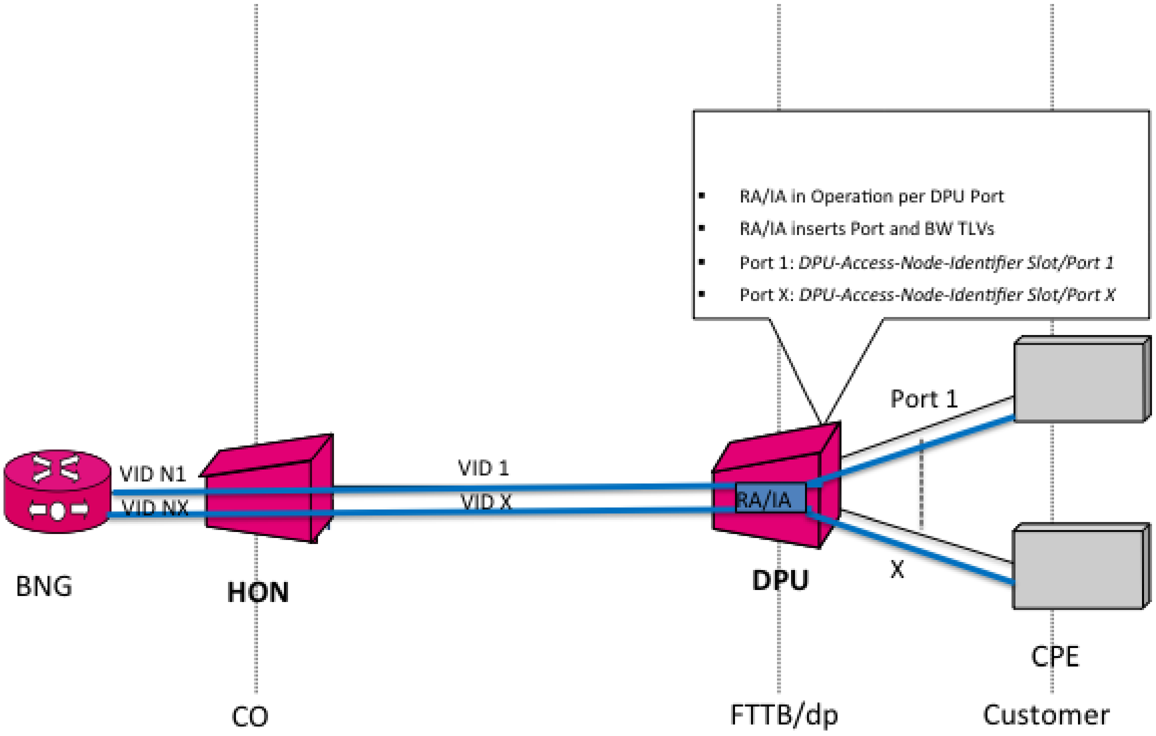

Figure 14-1 RA/IA in a Model 1 DPU

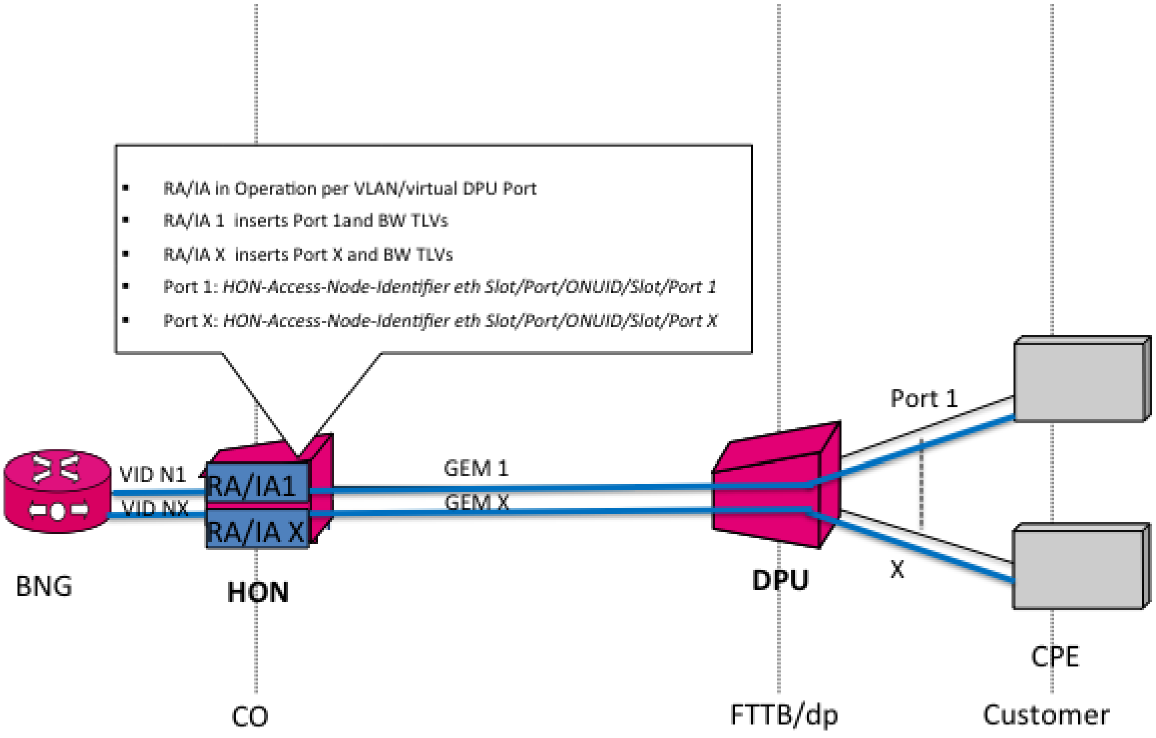

Figure 14-2 RA/IA In The HON (Model 2 DPU)

14.1 RA/IA Operation Requirements

R-101 DPUs in a Model 1 deployment MUST comply with section 5.4.7 of TR-178.

R-102 HONs in a Model 2 deployment MUST comply with section 5.4.7 of TR-178.

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-103 The RA/IA function MUST support the Sub TLVs in Table 14-1 G.fast Sub-TLVs

		G.fast G.9701 and G.997.2		
Pos.	Corresponding Sub TLV	Message Type	Information	Reference
1	0x0081/0x0082	Net Data Rate (NDR) Up- and Downstream	Reports the net data rate Upstream and Downstream	ITU-T G.997.2 Section 7.11.1.1
2	0x0085/0x0086	Attainable net Data rate (ATTNDR)	Reports the attainable net data rate Up- and Downstream	ITU-T G.997.2 Section 7.11.2.1
3	TBD	ETR	Reports expected throughput Up and Downstream	ITU-T G.997.2 Section 7.11.1.2
4		ATTETR	Reports attainable expected throughput Up and Downstream	ITU-T G.997.2 Section 7.11.2.2

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-104 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-105 The DPU MUST support counters for the total number of bytes received and bytes transmitted on its backhaul interface.

R-106 The DPU MUST support counters for the total number of bytes received and bytes transmitted on each activated user port.

R-107 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-108 The DPU MUST support the retrieval of the previous 15 minute interval counts upon request from the PMA.

R-109 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-110 The DPU MUST use the same 15 minute interval for all its counters.

R-111 The DPU MUST support the synchronization of its 15 minute interval with the PMA.

R-112 The DPU MUST support enabling and disabling of individual counters upon request from the PMA

R-113 The DPU MUST be able to include periodically reported but not binned data in the 15 minute count report.

R-114 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.

15.2 On Demand Diagnostics

R-115 The DPU MUST report, on demand from the PMA, the identity of the mains powered lines from which it is drawing power.

R-116 The DPU MUST report, on demand from the PMA, the DPU average power consumption over a known period of time.

R-117 The DPU MUST report, on demand from the PMA, the lines that are contributing to the DPU powering below a configurable percentage of the mean power being delivered to the DPU.

Note: This requirement is expressed by the following logic:

- mean_power = Sum(power delivered by mains powered lines) / $n_{\text{mains\_powered\_lines}}$
- If (line_power < mean_power x percentage_threshold) then include in the report

R-118 The DPU MUST be able to measure QLN and Hlog on a per provisioned line basis.

R-119 The DPU MUST be able to determine downstream NDR, ATT NDR, and margin for each provisioned line.

R-120 The DPU MUST be able to determine upstream NDR, ATT NDR, and margin for each provisioned line.

R-121 The DPU MUST be able to detect a short circuit between the 2 legs of any drop pair.

R-122 The DPU MUST be able to detect an open circuit on either leg of any drop pair.

R-123 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-124 The DPU MUST be able to make a coarse measurement of DC voltage on any drop pair.

R-125 The DPU MUST be able to detect the presence of AC mains on any drop pair and report it to the PMA.

R-126 The DPU MUST be able to report 'safe to touch' on demand in the absence of AC mains on any drop pair.

Note: R-125 and R-126 may depend on the availability of local ground connection.

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/XG-PON1 and point-to-point Ethernet. While these uplink technologies all include their own 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 interface. 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.

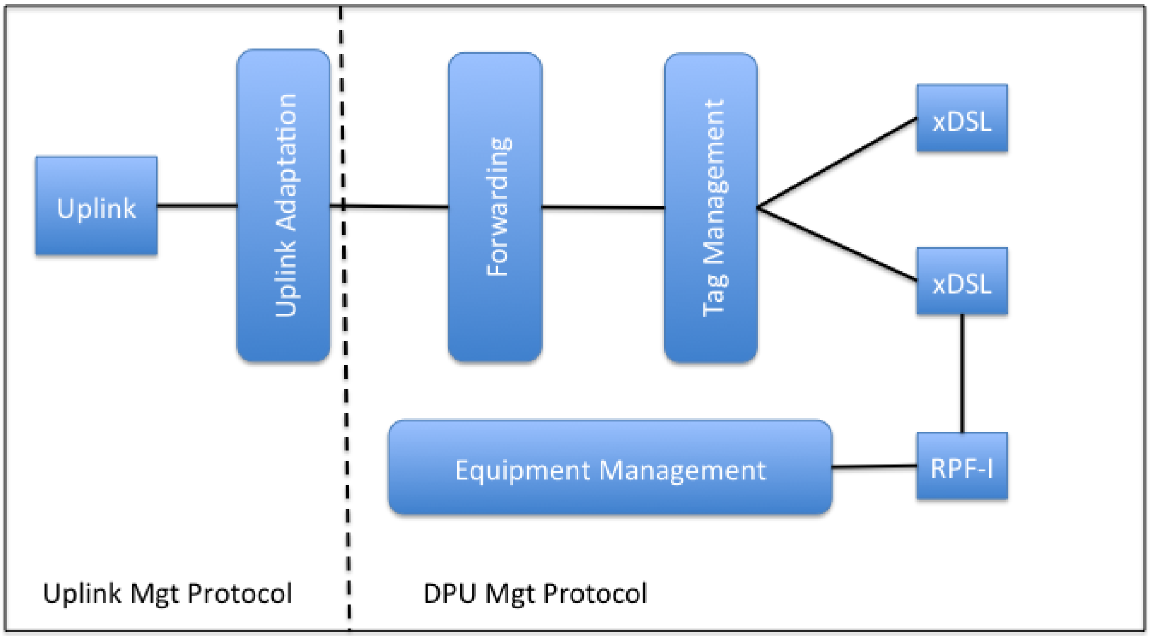

Figure 16-1 : Model 1 DPU Management Domains

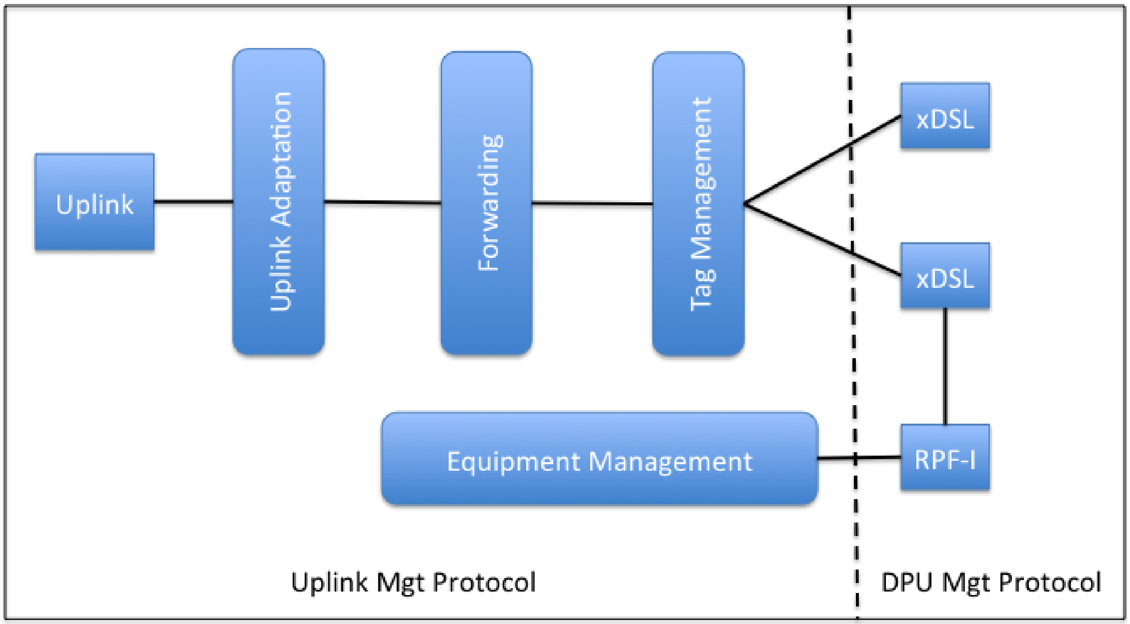

Figure 16-2 Model 2 DPU Management Domains

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 Persistent Management Agent (PMA).

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

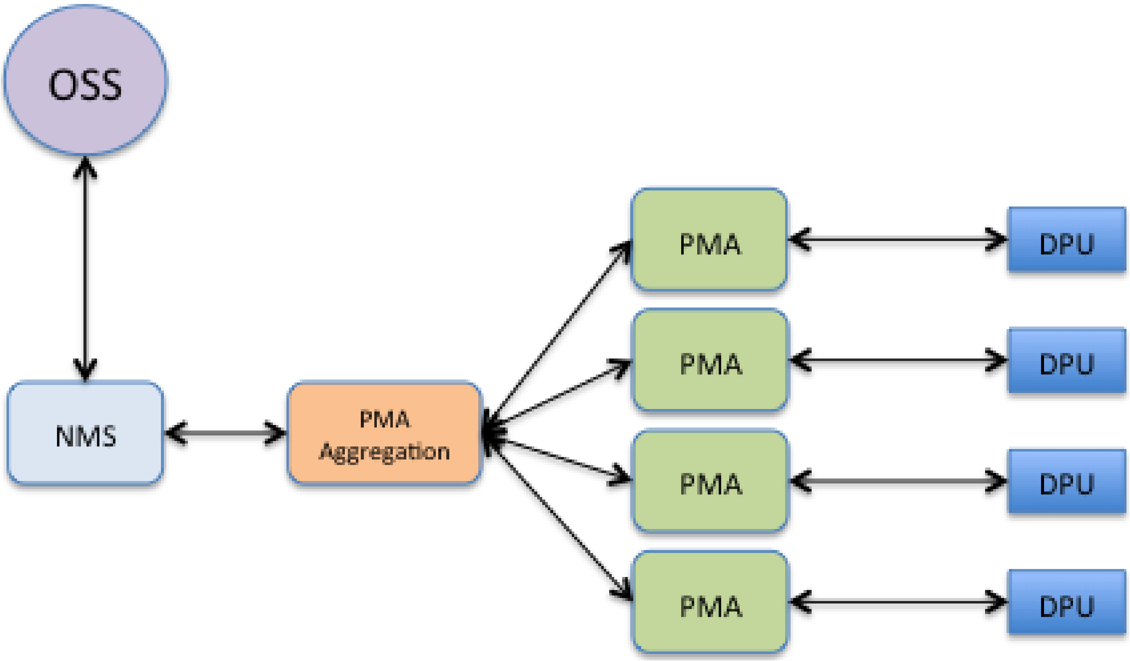

Figure 16-3 : DPU Management Architecture

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.

16.2 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 G.fast 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 very many more than existing access nodes such as DSLAMs. PMA aggregation therefore still help 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.

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

Although the DPU and PMA communicate using NETCONF over SSH 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 use self-assigned link-local addresses.

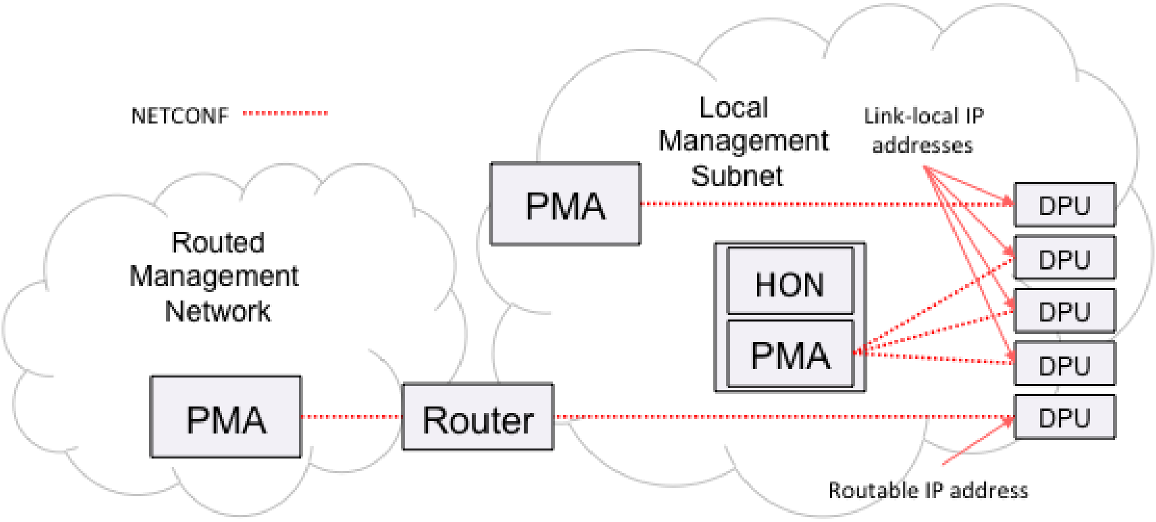

Figure 16-4 : Management Architecture Applied to Routable and Non-Routable IP Address Domains

In other cases, the PMA may be in a separate subnet 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.

16.4 PMA Concepts

The fundamental purpose of the PMA is to allow the 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 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.

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 is the logical management protocol interworking point for DPU management. While there is a single management protocol and data model used between the PMA and DPU, the northbound management interface for the PMA is intentionally left undefined so that it may support various deployed protocols like SNMP, TL1, or CORBA.

16.5 DPU-PMA Connectivity Monitoring

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.

The management protocol between the PMA and DPU is NETCONF. This runs over SSH which has an inherent keep-alive mechanism. There is no need to specify an additional PMA-DPU management keep-alive.

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-127 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

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). 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.

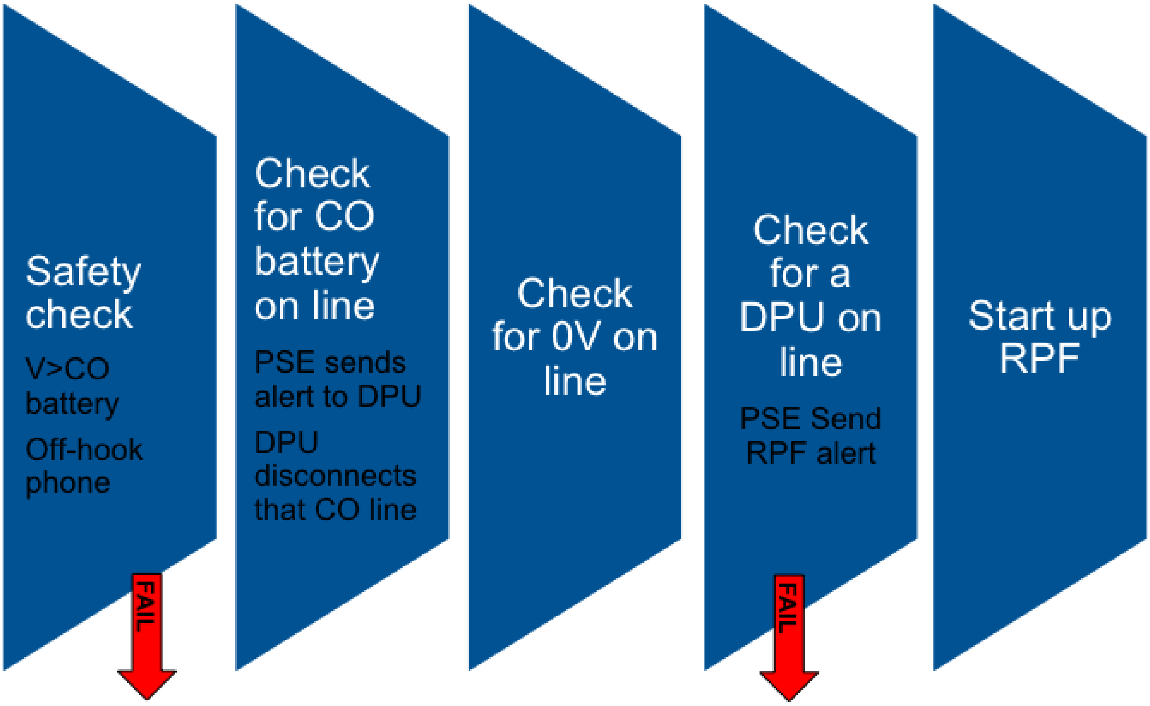

Figure 17-1 DPU Startup with Exchange/Cabinet POTS

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

R-128 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-129 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-130 The DPU MUST support receiving any of the ETSI defined RPF 'alerts' on every copper drop pair.

R-131 Upon detection of an RPF 'alert' on a given pair, the DPU MUST disconnect any exchange copper connection to that pair, subject to any start-up mode override set by network management.

R-132 The DPU MUST be able to scavenge enough power (e.g. from one or more connected exchange lines), to receive the RPF 'alert' on any customer pair and then disconnect that pair from the CO/exchange or cabinet.

R-133 The DPU power required by this start-up procedure MUST NOT result in line seizure.

R-134 The CO/exchange or cabinet connection MUST be able to be reinstated via network management.

R-135 The DPU MUST NOT reconnect any CO/exchange or cabinet line automatically after a power cycle.

R-136 The PSE MUST be able to send at least one of the 'alert' signals defined in ETSI TS 101 548

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-137 A DPU MUST support the migration of individual customers on that DP to a DPU based service.

R-138 Requirement R-137 MUST be able to be done remotely e.g. via a management system.

R-139 Requirement R-137 MUST NOT require a visit to the DPU.

R-140 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-141 The DPU MUST support remotely activated reversion of individual customers on that DP to a CO/Exchange or cabinet-based service.

End of Broadband Forum Technical Report TR-301