

TR-383

Common YANG Modules for Access Networks

Issue: 1 Amendment 2
Issue Date: December 2018

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 copyrighted by the Broadband Forum, and all rights are reserved. Portions of this Technical Report may be 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 IS EXPRESSLY DISCLAIMED. ANY USE OF THIS TECHNICAL REPORT SHALL BE MADE ENTIRELY AT THE IMPLEMENTER'S OWN RISK, AND NEITHER THE BROADBAND 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 TECHNICAL REPORT.

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.

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

Issue History

Issue Number	Approval Date	Publication Date	Issue Editor	Changes
1	8 May 2017	2 June 2017	Joey Boyd, ADTRAN Ludwig Pauwels, Nokia	Original
1 Amendment 1	13 June 2018	17 July 2018	Joey Boyd, ADTRAN Ludwig Pauwels, Nokia	Provide YANG model updates for Layer 2 Forwarding and QoS; publish initial model for Layer 2 Multicast Management; remove YANG models with dependencies on a draft revision of ietf-hardware.
1 Amendment 2	3 December 2018	3 December 2018	Joey Boyd, ADTRAN Ludwig Pauwels, Nokia	Add 'ethernet-like' abstract interface type.

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

Editors	Joey Boyd	ADTRAN
	Ludwig Pauwels	Nokia
YANG Modules Project Stream Leaders	Joey Boyd	ADTRAN
	Ken Kerpez	ASSIA
Common YANG Work Area Directors	Joey Boyd	ADTRAN
	Sven Ooghe	Nokia

Table of Contents

1	Purpose and Scope	9
1.1	Purpose.....	9
1.2	Scope.....	9
2	References and Terminology	10
2.1	Conventions.....	10
2.2	References	10
2.3	Definitions.....	11
2.4	Abbreviations	11
3	Technical Report Impact	12
3.1	Energy Efficiency.....	12
3.2	IPv6	12
3.3	Security	12
3.4	Privacy	12
4	Modules.....	13
4.1	DHCP	13
4.1.1	<i>bbf-l2-dhcpv4-relay</i>	13
4.1.2	<i>bbf-l2-dhcpv4-relay-forwarding</i>	13
4.1.3	<i>bbf-ldra</i>	13
4.2	Equipment	13
4.3	Ethernet	14
4.3.1	<i>bbf-ethernet-performance-management</i>	14
4.4	Layer 2 Forwarding.....	14
4.4.1	<i>bbf-l2-forwarding</i>	14
4.5	Interfaces	15
4.5.1	<i>bbf-interfaces-performance-management</i>	15
4.5.2	<i>bbf-interface-statistics-management</i>	16
4.5.3	<i>bbf-interface-usage</i>	16
4.5.4	<i>bbf-ptm</i>	16
4.6	PPPoE.....	16
4.6.1	<i>bbf-pppoe-intermediate-agent</i>	16
4.7	QoS.....	16
4.7.1	<i>bbf-qos-classifiers</i>	17
4.7.2	<i>bbf-qos-filters</i>	17
4.7.3	<i>bbf-qos-policies</i>	17
4.7.4	<i>bbf-qos-policies-sub-interfaces</i>	17
4.7.5	<i>bbf-qos-rate-control</i>	17
4.7.6	<i>bbf-qos-traffic-mngt</i>	18
4.7.7	<i>bbf-qos-enhanced-scheduling</i>	18
4.7.8	<i>bbf-qos-policer-envelope-profiles</i>	18
4.7.9	<i>bbf-policing-types</i>	18
4.7.10	<i>bbf-policing</i>	18
4.7.11	<i>bbf-shaping</i>	19
4.7.12	<i>bbf-qos-types</i>	19

4.8	sub-interfaces	19
4.8.1	<i>bbf-frame-classification</i>	19
4.8.2	<i>bbf-sub-interface-tagging</i>	19
4.8.3	<i>bbf-sub-interfaces</i>	19
4.9	subscribers.....	20
4.9.1	<i>bbf-subscriber-profiles</i>	20
4.9.2	<i>bbf-subscriber-types</i>	20
4.10	types	20
4.10.1	<i>bbf-dot1q-types</i>	20
4.10.2	<i>bbf-if-type</i>	20
4.10.3	<i>bbf-inet-types</i>	21
4.10.4	<i>bbf-yang-types</i>	21
4.11	common.....	21
4.11.1	<i>bbf-availability</i>	21
4.12	Layer 2 Multicast	21
4.12.1	<i>bbf-mgmd</i>	22
4.12.2	<i>bbf-mgmd-types</i>	23
5	Documentation	24
6	Dependencies on related YANG modules and Standards.....	25
7	Layer 2 Forwarding Data Model.....	26
7.1	Sub-interfaces.....	26
7.1.1	<i>Interface Usage</i>	27
7.2	Forwarders	27
7.2.1	<i>Forwarder Ports and Port Groups</i>	27
7.2.2	<i>Split Horizon Profiles</i>	28
7.2.3	<i>MAC Learning</i>	28
7.2.4	<i>Flooding</i>	28
8	Ethernet-like Interfaces	29

List of Figures

Figure 1 – YANG Data Model Relationships9
Figure 2 – Sub-interface Example.....27
Figure 3 – Forwarder Ports.....28

Executive Summary

This Technical Report defines YANG data models for the management of Broadband Forum specified access network equipment used across many deployment scenarios. Broadband Forum-specified access network equipment comprises Access Nodes and FTTdp DPUs. There is no assumption for BBF YANG modules to apply globally, e.g., to apply to access network equipment other than BBF Access Nodes and FTTdp DPUs, or to apply to core network equipment.

The models specified in this Technical Report are independent of any protocol.

1 Purpose and Scope

1.1 Purpose

This Technical Report defines YANG data models for the management of Broadband Forum specified access network equipment used across many deployment scenarios. Broadband Forum specified access network equipment comprises Access Nodes and FTTdp DPUs. There is no assumption for BBF YANG modules to apply globally, e.g., to apply to access network equipment other than BBF Access Nodes and FTTdp DPUs, or to apply to core network equipment.

The models specified in this Technical Report are independent of any protocol.

1.2 Scope

The data models defined by this Technical Report support the Broadband Forum requirements as applicable to access network equipment (e.g., BBF-specified Access Nodes and FTTdp DPUs) and form the set of core models which can be used for a multitude of other applications. It is intended that data models which are application specific can be built on, reference, and/or function alongside the common models.

The figure below provides a high level view of the functionality covered by this Technical Report (BBF YANG in green):

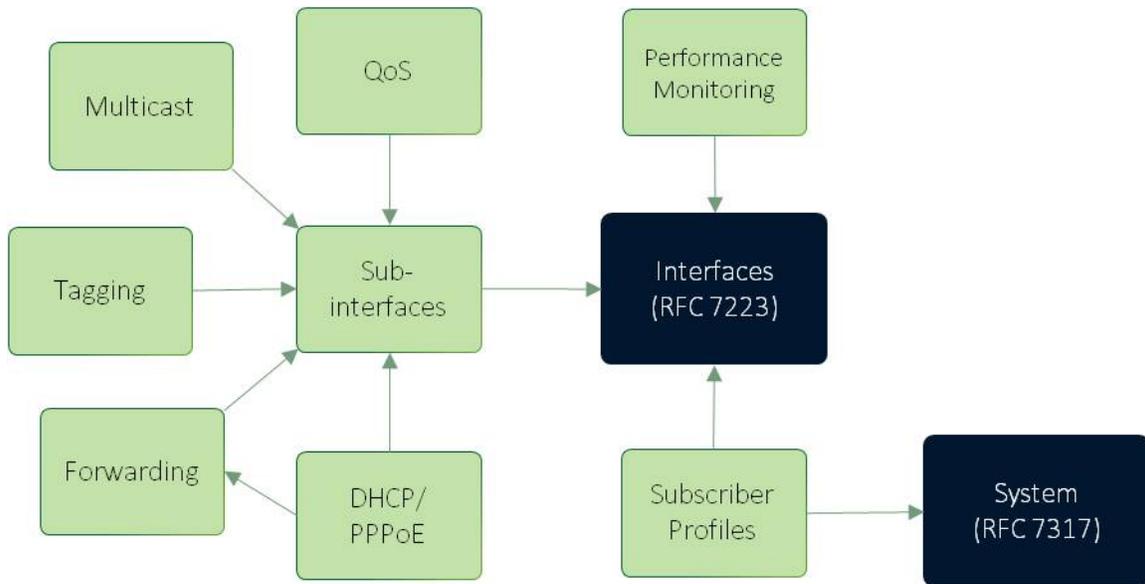


Figure 1 – YANG Data Model Relationships

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 in RFC 2119 [1].

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] RFC 2119	<i>Key words for use in RFCs to Indicate Requirement Levels</i>	IETF	1997
[2] RFC 7950	<i>The YANG 1.1 Data Modeling Language</i>	IETF	2016
[3] RFC 6991	<i>Common YANG Data Types</i>	IETF	2013
[4] RFC 7317	<i>A YANG Data Model for System Management</i>	IETF	2014

[5]	RFC 7223	<i>A YANG Data Model for Interface Management</i>	IETF	2014
[6]	RFC 8348	<i>A YANG Data Model for Hardware Management</i>	IETF	2018
[7]	RFC 3046	<i>DHCP Relay Agent Information Option</i>	IETF	2001
[8]	RFC 6221	<i>Lightweight DHCPv6 Relay Agent</i>	IETF	2011
[9]	TR-101i2	<i>Migration to Ethernet-Based Broadband Aggregation</i>	BBF	2011
[10]	TR-178	<i>Multi-service Broadband Network Architecture and Nodal Requirements</i>	BBF	2014
[11]	IEEE 802.3	<i>Ethernet Specification</i>	IEEE	2015
[12]	IEEE 802.1ag	<i>Connectivity Fault Management</i>	IEEE	2007
[13]	ITU-T Y.1731	<i>Operation, administration and maintenance (OAM) functions and mechanisms for Ethernet-based networks</i>	ITU-T	2015
[14]	MEF 35.1	<i>Service OAM Performance Monitoring Implementation Agreement</i>	MEF	2015
[15]	MEF 7.2	<i>Carrier Ethernet Management Information Model</i>	MEF	2013
[16]	MEF 10.3	<i>Ethernet Services Attributes Phase 3</i>	MEF	2013

2.3 Definitions

The following terminology is used throughout this Technical Report.

model	A data model.
module	A YANG module defines the hierarchy of data for the data model.
submodule	A YANG module may be broken up into a multiple submodules for ease of maintainability. The overall data model is comprised of a module and zero or more submodules.

2.4 Abbreviations

This Technical Report uses the following abbreviations:

DEI	Drop Eligible Indicator
DHCP	Dynamic Host Configuration Protocol
LAG	Link Aggregation Group
PPPoE	Point-to-Point Protocol over Ethernet

3 Technical Report Impact

3.1 Energy Efficiency

TR-383 has no impact on energy efficiency.

3.2 IPv6

TR-383 includes YANG modules that support IPv6 deployments.

3.3 Security

TR-383 has no impact on security.

3.4 Privacy

Any issues regarding privacy are not affected by TR-383.

4 Modules

The YANG modules contained in TR-383 are briefly described here. These modules are published on GitHub at <https://github.com/BroadbandForum/yang/tree/master/standard>.

4.1 DHCP

There are three modules which are used to manage DHCP functionality. These modules can be found in the *networking* directory on Github.

4.1.1 bbf-l2-dhcpv4-relay

This module contains a collection of YANG definitions for supporting the Broadband Forum requirements on subscriber management via the DHCPv4 protocol as applicable to access network equipment. As such, this module is specific to access network equipment (e.g., BBF-specified Access Nodes and FTTdp DPUs).

4.1.2 bbf-l2-dhcpv4-relay-forwarding

This module contains a collection of YANG definitions for supporting the Broadband Forum requirements on subscriber management via the DHCPv4 protocol as applicable to access network equipment. As such, this module is specific to access network equipment (e.g., BBF-specified Access Nodes and FTTdp DPUs).

Specifically, this module augments bbf-l2-forwarding with subscriber management via the DHCPv4 protocol [7].

4.1.3 bbf-ldra

This module contains a collection of YANG definitions for supporting the Broadband Forum requirements on subscriber management via the DHCPv6 protocol as applicable to access network equipment. As such, this module is specific to access network equipment (e.g., BBF-specified Access Nodes and FTTdp DPUs).

This functionality is also known as a Lightweight DHCPv6 Relay Agent (LDRA) [8].

4.2 Equipment

The modules included in TR-383 which were based on an older revision of the IETF draft, draft-ietf-netmod-entity, have been removed from TR-383 Amendment 1. This is due to an overall architecture change in the IETF draft which rendered the BBF models obsolete. Updates are planned for a later publication.

4.3 Ethernet

There is one module which is specific to the management of Ethernet interfaces as defined by IEEE 802.3 [11]. This module can be found in the *interface* directory on Github.

4.3.1 bbf-ethernet-performance-management

This module contains a collection of YANG definitions for supporting the Broadband Forum requirements on Ethernet interface management as applicable to access network equipment. As such, this module is specific to access network equipment (e.g., BBF-specified Access Nodes and FTTdp DPUs).

Specifically, this module augments bbf-interface-performance-management with Ethernet specific counters.

4.4 Layer 2 Forwarding

There is one module which is comprised of 8 submodules which are used for the management of Layer 2 (L2) Forwarding. These modules and submodules can be found in the *networking* directory on Github.

4.4.1 bbf-l2-forwarding

This module contains a collection of YANG definitions for supporting the Broadband Forum requirements on layer 2 forwarding as applicable to access network equipment. As such, this module is specific to access network equipment (e.g., BBF-specified Access Nodes and FTTdp DPUs).

4.4.1.1 bbf-l2-forwarding-base

This submodule contains a collection of YANG definitions for defining the top level nodes for forwarding.

4.4.1.2 bbf-l2-forwarding-flooding-policies

This submodule contains a collection of YANG definitions for managing flooding policies.

Flooding policies define how the system forwards frames in case other forwarding mechanisms do not arrive at a forwarding decision.

4.4.1.3 bbf-l2-forwarding-forwarders

This submodule contains a collection of YANG definitions for managing forwarders.

A forwarder is used to forward traffic between two or more interfaces.

4.4.1.4 **bbf-l2-forwarding-forwarding-databases**

This submodule contains a collection of YANG definitions for managing forwarding databases.

A forwarding database contains the necessary information regarding the MAC addresses which are used in the forwarding decision.

4.4.1.5 **bbf-l2-mac-learning-control**

This submodule contains a collection of YANG definitions for managing MAC address learning constraints, i.e., to constrain MAC learning rules compared with the standard IEEE MAC learning.

4.4.1.6 **bbf-l2-mac-learning**

This submodule contains a collection of YANG definitions for managing MAC learning.

For a forwarder, it specifies the forwarding database to use for the specified forwarder. For an interface, it provides the ability to enable/disable MAC learning as well as specifies other parameters associated with MAC learning.

4.4.1.7 **bbf-l2-forwarding-split-horizon-profiles**

This submodule contains a collection of YANG definitions for managing split horizon profiles.

These profiles allow (or disallow) forwarding between various forwarder ports based on the underlying interface usage.

4.4.1.8 **bbf-l2-forwarding-shared-fdb**

This submodule contains a collection of YANG definitions for managing shared forwarding databases.

4.5 Interfaces

There are four modules which augment `ietf-interfaces` [5] with additional interface management. These modules can be found in the *interfaces* directory on Github.

4.5.1 **bbf-interfaces-performance-management**

This module contains a collection of YANG definitions for supporting the Broadband Forum requirements on interface management as applicable to access network equipment. As such, this module is specific to access network equipment (e.g., BBF-specified Access Nodes and FTTdp DPUs).

Specifically, this module reports performance management of statistics defined by the IETF `interfaces` data model, `ietf-interfaces` (RFC 7223) [5].

4.5.2 bbf-interface-statistics-management

This module contains a collection of YANG definitions for supporting the Broadband Forum requirements on interface management as applicable to access network equipment. As such, this module is specific to access network equipment (e.g., BBF-specified Access Nodes and FTTdp DPUs).

Specifically, this module augments ietf-interfaces [5] with a reset action for statistics.

4.5.3 bbf-interface-usage

This module contains a collection of YANG definitions for supporting the Broadband Forum requirements on interface management as applicable to access network equipment. As such, this module is specific to access network equipment (e.g., BBF-specified Access Nodes and FTTdp DPUs).

Specifically, this module contains a collection of YANG definitions defining how interfaces are used.

4.5.4 bbf-ptm

This module contains a collection of YANG definitions for supporting the Broadband Forum requirements on interface management as applicable to access network equipment. As such, this module is specific to access network equipment (e.g., BBF-specified Access Nodes and FTTdp DPUs).

Specifically, this module augments the IETF interfaces data model, ietf-interfaces (RFC 7223), with nodes for managing Packet Transfer Mode (PTM) interfaces.

4.6 PPPoE

There is one module for managing Point-to-Point Protocol over Ethernet. This module can be found in the *networking* directory on Github.

4.6.1 bbf-pppoe-intermediate-agent

This module contains a collection of YANG definitions for supporting the Broadband Forum requirements on subscriber management via the PPPoE protocol as applicable to access network equipment. As such, this module is specific to access network equipment (e.g., BBF-specified access Nodes and FTTdp DPUs).

4.7 QoS

There are twelve modules which provide management of Quality of Service (QoS). These modules can be found in the *networking* directory on Github.

4.7.1 bbf-qos-classifiers

This module contains a collection of YANG definitions for supporting the Broadband Forum requirements on the management of Quality of Service (QoS) as applicable to access network equipment. As such, this module is specific to access network equipment (e.g., BBF-specified Access Nodes and FTTdp DPUs).

Specifically, this module contains a collection of classifiers used by policies applied to packets.

4.7.2 bbf-qos-filters

This module contains a collection of YANG definitions for supporting the Broadband Forum requirements on the management of Quality of Service (QoS) as applicable to access network equipment. As such, this module is specific to access network equipment (e.g., BBF-specified Access Nodes and FTTdp DPUs).

Specifically, this module contains definitions of filter criteria.

4.7.3 bbf-qos-policies

This module contains a collection of YANG definitions for supporting the Broadband Forum requirements on the management of Quality of Service (QoS) as applicable to access network equipment. As such, this module is specific to access network equipment (e.g., BBF-specified Access Nodes and FTTdp DPUs).

Specifically, this module contains a collection of policies applied to packets.

4.7.4 bbf-qos-policies-sub-interfaces

This module contains a collection of YANG definitions for supporting the Broadband Forum requirements on the management of Quality of Service (QoS) as applicable to access network equipment. As such, this module is specific to access network equipment (e.g., BBF-specified Access Nodes and FTTdp DPUs).

Specifically, this module contains augments to sub-interfaces to support policies applied to packets.

4.7.5 bbf-qos-rate-control

This module contains a collection of YANG definitions for supporting the Broadband Forum requirements on the management of Quality of Service (QoS) as applicable to access network equipment. As such, this module is specific to access network equipment (e.g., BBF-specified Access Nodes and FTTdp DPUs).

Specifically, this module contains a possibility to rate limit a number of frames.

4.7.6 bbf-qos-traffic-mngt

This module contains a collection of YANG definitions for supporting the Broadband Forum requirements on the management of Quality of Service (QoS) as applicable to access network equipment. As such, this module is specific to access network equipment (e.g., BBF-specified Access Nodes and FTTdp DPUs).

Specifically, this module contains a collection of functions for QoS traffic management (TM).

4.7.7 bbf-qos-enhanced-scheduling

This module contains a collection of YANG definitions for supporting the Broadband Forum requirements on the management of Quality of Service (QoS) as applicable to access network equipment. As such, this module is specific to access network equipment (e.g., BBF-specified Access Nodes and FTTdp DPUs).

Specifically, this module contains a collection of enhanced traffic scheduling configuration.

4.7.8 bbf-qos-policer-envelope-profiles

This module contains a collection of YANG definitions for supporting the Broadband Forum requirements on the management of Quality of Service (QoS) as applicable to access network equipment. As such, this module is specific to access network equipment (e.g., BBF-specified Access Nodes and FTTdp DPUs).

Specifically, this module contains management of Policer (also known as Bandwidth Profile) Envelope Profiles as described in MEF 10.3 [16].

4.7.9 bbf-policing-types

This module contains a collection of YANG definitions for supporting the Broadband Forum requirements on the management of Quality of Service (QoS) as applicable to access network equipment. As such, this module is specific to access network equipment (e.g., BBF-specified Access Nodes and FTTdp DPUs).

Specifically, this module contains common types used for management of Policers.

4.7.10 bbf-policing

This module contains a collection of YANG definitions for supporting the Broadband Forum requirements on the management of Quality of Service (QoS) as applicable to access network equipment. As such, this module is specific to access network equipment (e.g., BBF-specified Access Nodes and FTTdp DPUs).

Specifically, this module contains a collection of type and data definitions for use of some policer types that act on an individual flow.

4.7.11 bbf-shaping

This module contains a collection of YANG definitions for supporting the Broadband Forum requirements on the management of Quality of Service (QoS) as applicable to access network equipment. As such, this module is specific to access network equipment (e.g., BBF-specified Access Nodes and FTTdp DPUs).

Specifically, this module contains shaping configuration data.

4.7.12 bbf-qos-types

This module contains a collection of YANG definitions for supporting the Broadband Forum requirements on the management of Quality of Service (QoS) as applicable to access network equipment. As such, this module is specific to access network equipment (e.g., BBF-specified Access Nodes and FTTdp DPUs).

Specifically, this module contains type definitions used in multiple QoS modules.

4.8 sub-interfaces

There are three modules which provide management definitions for sub-interfaces. These modules can be found in the *interfaces* directory on Github.

4.8.1 bbf-frame-classification

This module contains a collection of YANG definitions for supporting the Broadband Forum requirements on frame classification as applicable to access network equipment. As such, this module is specific to access network equipment (e.g., BBF-specified Access Nodes and FTTdp DPUs).

Specifically, this module contains reusable groupings defined for frame classification.

4.8.2 bbf-sub-interface-tagging

This module contains a collection of YANG definitions for supporting the Broadband Forum requirements on the management of sub-interfaces as applicable to access network equipment. As such, this module is specific to access network equipment (e.g., BBF-specified Access Nodes and FTTdp DPUs).

Specifically, this module defines the ingress and egress tagging of a VLAN sub-interface.

4.8.3 bbf-sub-interfaces

This module contains a collection of YANG definitions for supporting the Broadband Forum requirements on the management of sub-interfaces as applicable to access network equipment. As

such, this module is specific to access network equipment (e.g., BBF-specified Access Nodes and FTTdp DPUs).

4.9 subscribers

There are two modules which provide management of subscriber related functionality. These modules can be found in the *networking* directory on Github.

4.9.1 bbf-subscriber-profiles

This module contains a collection of YANG definitions for supporting the Broadband Forum requirements on the management of subscribers as applicable to access network equipment. As such, this module is specific to access network equipment (e.g., BBF-specified Access Nodes and FTTdp DPUs).

Specifically, this module defines subscriber profiles.

4.9.2 bbf-subscriber-types

This module contains a collection of YANG definitions for supporting the Broadband Forum requirements on the management of subscribers as applicable to access network equipment. As such, this module is specific to access network equipment (e.g., BBF-specified Access Nodes and FTTdp DPUs).

Specifically, this module defines common types associated with subscribers and subscriber protocols.

4.10 types

There are four modules which provide reusable type definitions for use across all BBF YANG models. These modules can be found in the *common* directory on Github.

4.10.1 bbf-dot1q-types

This module contains a collection of YANG definitions for supporting the Broadband Forum requirements on reusable data types as applicable to access network equipment. As such, this module is specific to access network equipment (e.g., BBF-specified Access Nodes and FTTdp DPUs).

Specifically, this module defines common types for support of some IEEE standards.

4.10.2 bbf-if-type

This module contains a collection of YANG definitions for supporting the Broadband Forum requirements on reusable data types as applicable to access network equipment. As such, this

module is specific to access network equipment (e.g., BBF-specified Access Nodes and FTTdp DPUs).

Specifically, this module defines interface types that are needed for BBF applications but are not defined in *iana-if-type*.

4.10.3 bbf-inet-types

This module contains a collection of YANG definitions for supporting the Broadband Forum requirements on reusable data types as applicable to access network equipment. As such, this module is specific to access network equipment (e.g., BBF-specified Access Nodes and FTTdp DPUs).

Specifically, this module defines generally useful derived YANG data types for Internet addresses and related things added by BBF on top of the ones defined in IETF.

4.10.4 bbf-yang-types

This module contains a collection of YANG definitions for supporting the Broadband Forum requirements on reusable data types as applicable to access network equipment. As such, this module is specific to access network equipment (e.g., BBF-specified Access Nodes and FTTdp DPUs).

Specifically, this module defines common types used throughout BBF data models.

4.11 common

There is one module which provides support for common requirements for use across all BBF YANG models. This module can be found in the *common* directory on Github.

4.11.1 bbf-availability

This module contains a collection of YANG definitions for supporting the Broadband Forum requirements on the general availability of specific resources as applicable to access network equipment. As such, this module is specific to access network equipment (e.g., BBF-specified Access Nodes and FTTdp DPUs).

4.12 Layer 2 Multicast

There are two modules, one of which is comprised of 5 submodules which are used for the management of Layer 2 (L2) Multicast. These modules and submodules can be found in the *networking* directory on Github.

4.12.1 bbf-mgmd

This module contains a collection of YANG definitions for supporting the Broadband Forum requirements on layer 2 multicast management as applicable to access network equipment. As such, this module is specific to access network equipment (e.g., BBF-specified Access Nodes and FTTdp DPUs).

Specifically, this module contains a collection of YANG definitions for supporting the Multicast Group Membership Discovery (MGMD) Protocols. In particular, it describes objects used for managing the Internet Group Management Protocol (IGMP) and the Multicast Listener Discovery (MLD) protocol in systems that act as a multicast proxy, snooper, or a snooper with proxy reporting.

4.12.1.1 bbf-mgmd-configuration-interface-to-host

This module contains a collection of YANG definitions for supporting the Multicast Group Membership Discovery (MGMD) Protocols. In particular, it describes configuration objects used for managing the Internet Group Management Protocol (IGMP) and the Multicast Listener Discovery (MLD) protocol on interfaces that connect the system to multicast hosts.

4.12.1.2 bbf-mgmd-configuration-interface-to-router

This module contains a collection of YANG definitions for supporting the Multicast Group Membership Discovery (MGMD) Protocols. In particular, it describes configuration objects used for managing the Internet Group Management Protocol (IGMP) and the Multicast Listener Discovery (MLD) protocol on interfaces that connect to multicast routers.

4.12.1.3 bbf-mgmd-configuration-multicast-snoop

This module contains a collection of YANG definitions for supporting the Multicast Group Membership Discovery (MGMD) Protocols. In particular, it describes configuration objects used for managing the Internet Group Management Protocol (IGMP) and the Multicast Listener Discovery (MLD) protocol in case the system acts as a snooper.

4.12.1.4 bbf-mgmd-operational-interface-to-host

This module contains a collection of YANG definitions for supporting the Multicast Group Membership Discovery (MGMD) Protocols. In particular, it describes state objects used for managing the Internet Group Management Protocol (IGMP) and the Multicast Listener Discovery (MLD) protocol on interfaces that connect a system to multicast hosts.

4.12.1.5 bbf-mgmd-operational-interface-to-router

This module contains a collection of YANG definitions for supporting the Multicast Group Membership Discovery (MGMD) Protocols. In particular, it describes state objects used for

managing the Internet Group Management Protocol (IGMP) and the Multicast Listener Discovery (MLD) protocol on interfaces that connect a system to multicast routers.

4.12.2 bbf-mgmd-types

This module contains a collection of YANG definitions for supporting the Broadband Forum requirements on layer 2 multicast management as applicable to access network equipment. As such, this module is specific to access network equipment (e.g., BBF-specified Access Nodes and FTTdp DPUs).

Specifically, this module contains a collection of YANG type and feature definitions for use in modules supporting the Multicast Group Membership Discovery (MGMD) Protocols. In particular, it describes objects used for managing the Internet Group Management Protocol (IGMP) and the Multicast Listener Discovery (MLD) protocol.

5 Documentation

There are “README.md” files; these are short text files giving brief descriptions of the contents of the directories they are in.

Documentation for each module can be found in the *docs* folder of the corresponding directory, e.g., *networking*. There are two types of documents per each top level module.

- *.tree: Provides a tree diagram of the module.
- *.xml: Provides an XML schema representation of the module

Each .tree and .xml file are named according to each model's module name along with –full to indicate it the full tree and not just the nodes defined in a given module (in the case of submodule support), e.g., bbf-dhcpv4-relay-full.tree and bbf-dhcpv4-relay-full.xml

6 Dependencies on related YANG modules and Standards

TR-383 is based on YANG 1.1 (RFC 7950 [2]).

The following YANG modules are used by TR-383:

- ietf-interfaces [5]
- ietf-yang-types [3]
- ietf-inet-types [3]
- ietf-system [4]

7 Layer 2 Forwarding Data Model

The intent of this section is to provide some general information regarding the usage of the layer 2 forwarding data model. It is not possible to describe every possible application which would use the model but rather it provides the theory behind the model and illustrates some general use cases.

7.1 Sub-interfaces

Before traffic can be forwarded, it must first be classified to determine what to forward, where to forward and how to manipulate the packet if so desired. The concept of a vlan-sub-interface has been introduced to provide a VLAN interface which can be used as the source or destination interface of a forwarding decision. Each vlan-sub-interface classifies traffic from a particular lower layer interface into a forwarder. This classification consists of a set of rules specified using match criteria on packet fields (e.g., VLAN-ID, p-bit). The lower layer interface can be either a non-aggregated physical or logical interface (e.g., Ethernet), an aggregation of physical or logical interfaces (e.g., LAG) or can be another vlan-sub-interface.

A vlan-sub-interface is created each time a new forwarding context is required (e.g., 1:1 VLAN). Each vlan-sub-interface can then have multiple rules associated with it if different classification results in the same forwarding decision. For example, one rule can catch frames tagged with a particular VLAN-ID, a second rule can catch untagged frames, and a third rule can catch priority-tagged frames. The second and third rules in this example cover the concept of a port default VLAN.

As stated above, multiple vlan-sub-interfaces can refer to the same lower layer interface in order to provide multiple traffic classifications based on different, but potentially overlapping, match criteria. In order to provide deterministic classification, each rule is given a priority. The scope of the priority is over all rules defined within all vlan-sub-interfaces referring to same lower layer interface. A packet ingressing the lower layer interface would then be compared to each rule starting with the highest priority rule and proceeding to the lowest priority rule. If a match occurs, the packet is processed accordingly. If no match occurs, the packet is dropped.

The figure below shows how two vlan-sub-interfaces are associated with the same physical interface classifying traffic for two different forwarding decisions.

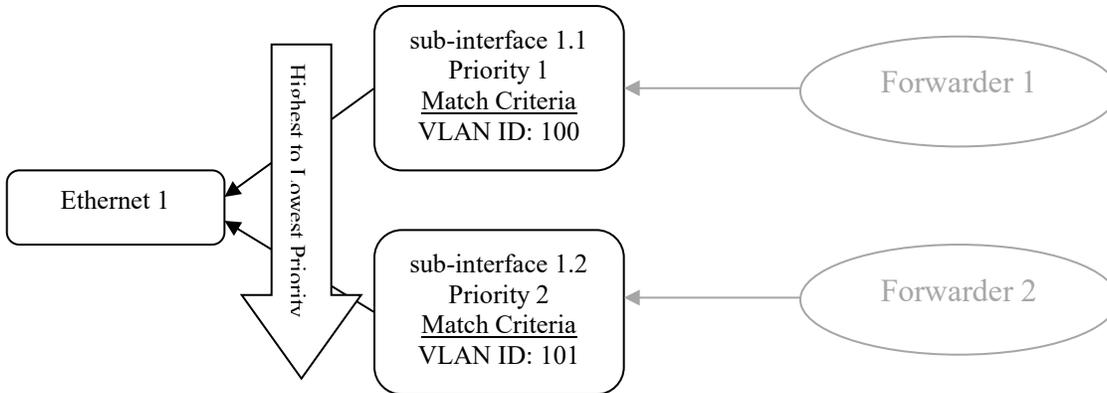


Figure 2 – Sub-interface Example

In addition to classification of traffic, the sub-interface also contains rules for any ingress or egress actions to take on each matched packet. These actions include pushing or popping tags, rewrite of p-bits or rewrite of Drop Eligible Indication (DEI) bits.

7.1.1 Interface Usage

For the case of N:1 or N:M VLAN forwarding, the role each interface plays in the network is important to determine how traffic flow is managed. For example, in the context of an access node, traffic ingressing a user port should not normally be forwarded to another user port. Certain mechanisms to be discussed later will be used to enforce this restriction. First, however, the way in which an interface is used must be explicitly known. For this the interface usage must be configured either by the user or by the system if the usage is already known. The 3 types of interface usage are:

- user port: The interface connects an access node to a user.
- network port: The interface connects an access node to a network.
- subtended-node port: The interface connects an access node to another access node.

7.2 Forwarders

Once traffic has been classified and possibly manipulated, it needs to be forwarded appropriately to another sub-interface. A forwarder is used to determine how traffic is routed between two or more forwarder ports each of which is associated with a sub-interface. This forwarder can be used to handle 1:1 VLAN, N:1 VLAN and N:M VLAN applications.

7.2.1 Forwarder Ports and Port Groups

Each forwarder port is associated with a sub-interface whose underlying interface is either a user port, a network port or a subtended node port. Forwarder ports with similar forwarding characteristics can be placed into forwarding groups and referenced collectively when configuring the forwarder.

Below shows the relationships between a forwarder, its forwarder ports, and the referenced sub-interfaces and their lower layer interfaces.

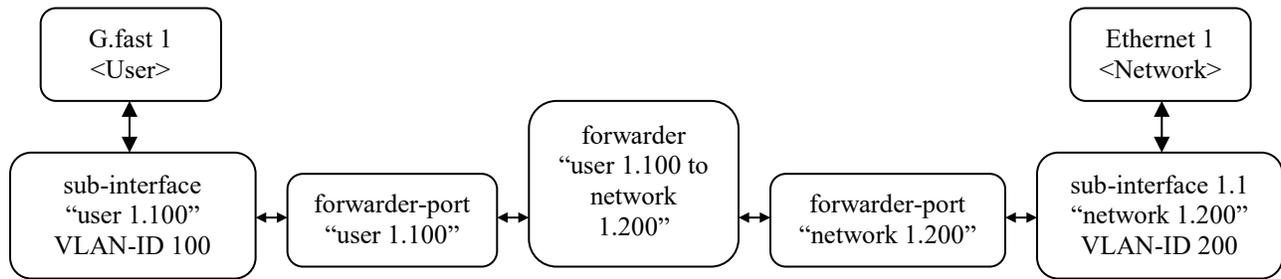


Figure 3 – Forwarder Ports

In the simplest use case of a 1:1 VLAN, this is all the forwarder needs to forward traffic between two sub-interfaces. The sub-interfaces determine which packets will be forwarded and how they will be manipulated. The forwarder just provides the means to associate the sub-interfaces.

7.2.2 Split Horizon Profiles

Once the interface usage is configured, a split horizon profile can be created and applied per forwarder to configure how traffic is forwarded between the various types of interfaces. Each profile specifies the usage of the ingress interface then lists the usages for which egress of the packet is not allowed from the ingress interface. For example, a profile could specify that for an ingress interface that is a user port, it is not allowed to send traffic to interfaces that are also user ports.

7.2.3 MAC Learning

In addition to the usage of an interface, the source and destination MAC addresses are key to making correct forwarding decisions for N:1 and N:M VLAN forwarding. Each forwarder contains configuration which determines how and if MAC source addresses are learned. It can also prevent traffic from being forwarded if it contains a certain MAC source address.

Once a MAC source address is learned, it is installed in the forwarding database for a given forwarder port. When a packet egresses a forwarder port, its MAC destination address is compared to the addresses in each of the other forwarder port's forwarding database to determine where the packet should be forwarded.

7.2.4 Flooding

In the case that the destination MAC address is not found in any forwarding database, it may be desired to flood the packet to all appropriate forwarder ports. To determine when and how this flooding occurs, a flooding policy profile can be created and associated with each forwarder. Each policy can be specified for a particular interface usage (e.g., user port) and/or a specific destination MAC address. It then assigns an appropriate action of either discarding the packet or flooding it to all interfaces of specified usage(s). For example, a forwarder may be configured to flood all packets with an unknown MAC address coming from a network port to all user ports.

8 Ethernet-like Interfaces

There are several instances in the Common YANG modules where the interface list from `ietf-interfaces`[5] is augmented with a constraint on the types of interfaces to which the augmented nodes apply. This constraint limits the interface types to those which transport Ethernet frames. The following is an example of a constrained augmentation.

```
augment '/if:interfaces/if:interface' {
  when
    "derived-from-or-self(if:type, 'ianaift:ethernetCsmacd') or
     derived-from-or-self(if:type, 'ianaift:ieee8023adLag') or
     derived-from-or-self(if:type, 'ianaift:ptm') or
     derived-from-or-self(if:type, 'bbfift:vlan-sub-interface')" {
  description
    "Interfaces that can have QoS policy profiles assigned.";
}
```

The augmentation to add a QoS policy reference to an interface is constrained to interfaces which are of one of four types or derived from those types. See RFC 7950[2] for the full definition and usage of the `derived-from-or-self()` function.

Similarly, there are nodes which are references to an interface whose type is also constrained to those which transport Ethernet frames.

```
leaf interface {
  type if:interface-ref;
  must
    "derived-from-or-self(
     /if:interfaces/if:interface[if:name = current()]
     /if:type, 'ianaift:ethernetCsmacd')
     or
     derived-from-or-self(
     /if:interfaces/if:interface[if:name = current()]
     /if:type, 'ianaift:ieee8023adLag')
     or
     derived-from-or-self(
     /if:interfaces/if:interface[if:name = current()]
     /if:type, 'ianaift:ptm')
     or
     derived-from-or-self(
     /if:interfaces/if:interface[if:name = current()]
     /if:type, 'bbfift:sub-interface') ";
  mandatory true;
  description
    "References the lower-layer interface.";
}
```

In this example, the reference to the interface must be one of four types or derived from those types.

While this approach works well, it does not lend itself to extensibility when new interface types are defined either by the Broadband Forum, a vendor, an operator, or some other organization which is defining YANG data models. In order for these augments and must statements to be applicable to those interface types, either the new type or types need to be derived from one of these existing types or the new ones have to be added to the modeled constraints. This presents a challenge of keeping these models aligned and may not even be possible depending on the source of the newly defined interface type.

One solution that has been introduced is the creation of an abstract Ethernet type from which new interface types can be defined.

```
identity ethernet-like {
  base bbf-interface-type;
  description
    "An abstract identity defining a class of interfaces which
    represents a logical interface transporting Ethernet frames,
    i.e. frames with a destination and source MAC address, an
    Ethernet type or length field, and a payload. This
    'interface type' is intended only to be used to define
    constraints against a class of interfaces each of which have
    their 'type' derived from this identity (as well as potentially
    others). At no time should this identity be used as the 'type'
    for an interface.";
}
```

This abstract type is added to the constraints.

Updated augment example:

```
augment '/if:interfaces/if:interface' {
  when
    "derived-from-or-self(if:type, 'ianaift:ethernetCsmacd') or
    derived-from-or-self(if:type, 'ianaift:ieee8023adLag') or
    derived-from-or-self(if:type, 'ianaift:ptm') or
    derived-from-or-self(if:type, 'bbfift:vlan-sub-interface') or
    derived-from(if:type, 'bbfift:ethernet-like')" {
    description
      "Interfaces that can have QoS policy profiles assigned.";
  }
}
```

Updated must statement example:

```
leaf interface {
  type if:interface-ref;
  must
    "derived-from-or-self(
      /if:interfaces/if:interface[if:name = current()]
      /if:type, 'ianaift:ethernetCsmacd')"
```

```

    or
    derived-from-or-self(
      /if:interfaces/if:interface[if:name = current()]
      /if:type, 'ianaift:ieee8023adLag')
    or
    derived-from-or-self(
      /if:interfaces/if:interface[if:name = current()]
      /if:type, 'ianaift:ptm')
    or
    derived-from-or-self(
      /if:interfaces/if:interface[if:name = current()]
      /if:type, 'bbfift:sub-interface')
    or
    derived-from(
      /if:interfaces/if:interface[if:name = current()]
      /if:type, 'bbfift:ethernet-like') ";

    mandatory true;
    description
      "References the lower-layer interface.";
  }

```

The use of ‘derived-from[2]’ stems from the identity’s definition which states it is an abstract identity which is not to be used as an actual interface type.

The method of using this abstract type is to add it as a base identity[2] to any Ethernet type definition which satisfies the definition of ‘ethernet-like’. For example,

```

identity new-ethernet-type {
  base bbfift:Ethernet-like;
  description
    "A new Ethernet type.";
}

```

By utilizing the abstract interface type, the Common YANG modules which define these constraints no longer have to be updated when a new Ethernet interface type is created.

End of Broadband Forum Technical Report TR-383