

Proposal for Supporting OpenFlow Management on VOLTHA by Using Profile Based Definition of Unsupported Access Technology Specific Service Attributes

Author: Shaun Missett/Kim Kempf/Saurav Das
E-mail: shaun.missett@radisys.com/kim.kempf@radisys.com/saurav.das@opennetworking.org
Owner: Shaun Missett

002-xxxx-000n
 Revision 0.7
 September 21, 2018

Note: To verify signoff, see the on-line repository.

Name	Title/Department
Prakash Siva	Prakash.Siva@radisys.com - VP Technology and Strategy, Tech BD
Prashant Sharma	Prashant.Sharma@radisys.com - Architect, Tech BD
Shad Ansari	shad@opennetworking.org - ONF
Jonathan Hart	jono@opennetworking.org - ONF

—
Approvers:

Reviewers:

Preface

On-line Document Repositories

- Approved versions of document are stored in Agile under the Radisys document number
- Interim released versions of document are stored here: \\TBD
- Pointers to documents and additional copies are also located in the collaboration site: <Provide link here>

Revision History

All revisions that start with zero (such as 0.5, 0.65, and the like) are preliminary drafts for internal review only.

Table 1 Revision History

No.	Date	Author	Description
0.1	02/02/2018	Shaun Missett/Kim Kempf	Initial draft
0.2	02/09/2018	Shaun Missett/Saurav Das/ONF	OpenFlow Updates/comments Update Table of Contents
0.3	02/15/2018	Shaun Missett/Kim Kempf	Update Figures
0.4	02/16/2018	Shaun Missett	Update Terminology, VOLTHA versus vOLTHA OpenFlow versus Open Flow variants – remove Revision in footer as there are PDF conversion issues – may fix later.
0.5	03/19/2018	Shaun Missett	Clarify Instance Control Use cases – add support for PON Level Instance Control. Add support for Bidirectional/Uni-Down/Ui-Up attributes for Multicast support. Add attribute to allow explicit or automatic control of PON DBA additional bandwidth.
0.6	04/12/2018 04/17/2018	Shaun Missett	Output of VOLTHA 2Q2018 Technical Lockdown meeting review of 'Service Profile' proposal: <ol style="list-style-type: none"> 1) Rename to be 'Technology Profile' 2) Remove 'Fixed' (R_F) DBA Bandwidth Component from Profile and add the use of a third Open Flow Meter Band for specifying this. Incorporate other clarifications based on ONF and AT&T feedback.

No.	Date	Author	Description
0.7	09/12/2018	Shaun Missett	Clarify Meters apply to both Upstream and Downstream Flows

Contents

OVERVIEW	5
Document Overview	5
Product Overview	5
Terminology.....	5
Acronyms.....	6
1. ISSUES LIST.....	7
OPENFLOW ABSTRACTION ISSUE WITH MANAGING NON-ETHERNET ACCESS TECHNOLOGIES	8
Problem Statement.....	8
Solution Proposal.....	8
OpenFlow and PON MEF to Meter Band Conversion	13
xPON Technology Profile	15
Technology Profile Identifiers:.....	15
Common Instance Control Attributes:.....	15
Upstream Scheduler Attributes:.....	15
Upstream GEM Port configuration and scheduling attributes:	15
Downstream Aggregate Scheduler Attributes:.....	16
Downstream GEM Port configuration and scheduling attributes:.....	16
Service Creation Examples.....	17
Single T-CONT with Four GEM Ports Example (4QueueHybridProfileMap1).....	20
OpenFlow Mapping to Technology profile '4QueueHybridProfileMap1' Using Table ID24	
Single T-CONT with Eight GEM Ports Example ('8QueueStrictProfileMap1')	27
OpenFlow Mapping to Technology Profile '8QueueStrictProfileMap1' Using Table ID 30	
Documentation.....	32
A.1.1. Radisys Documents.....	33
A.1.2. Other Documents	33
A.1.3. Industry-Standard References	33

Figures

Figure 1 Adding Technology Profile as Part of VOLTHA.....	10
Figure 2 ONU and OLT Traffic Management.....	19

Tables

Table 1 Revision History	2
Table 2 DBA Inferred Additional Bandwidth Type	15

OVERVIEW

Document Overview

This document describes the issues with using Open Flow independent of Technology Specific attributes and provides a solution which does not require extensive development of Technology Specific management interfaces to support different Access Technologies.

Product Overview

The VOLTHA ASFvOLT16 Adapter supports the Edgecore/Accton OLT which is an OCP 16 Port XGS-PON vOLT design. The ASFvOLT16 design is based on Broadcom Maple PON MAC silicon supporting 16x XFP ports of XGS-PON or NG-PON2 (10Gb/10Gb) and four 100 gigabit QSFP28 Ethernet uplink ports. The design includes a Broadcom StrataDNX (Qumran) switch to bridge the PON ports to the four NNI uplink ports.

Terminology

<u>Term</u>	<u>Definition</u>
Datapath	The aggregation of components of an OpenFlow logical switch that are directly involved in traffic processing and forwarding. The datapath includes the pipeline of flow tables, the group table and the ports.
GEM Port	An abstraction on the G-PON transmission convergence (GTC) adaptation sublayer representing a logical connection associated with a specific client packet flow.
GEM	In this Document GEM and XGEM are used interchangeably as are GEM Port and XGEM Port, primarily using the GEM terminology.
OpenFlow Protocol	The protocol defined by OpenFlow Switch specification [1]. Also called OpenFlow Switch Protocol
Packet	A series of bytes comprising a header, a payload and optionally a trailer, in that order, and treated as a unit for purposes of processing and forwarding.
Pipeline	The set of linked flow tables that provide matching, forwarding, and packet modification in an OpenFlow switch
Tag	A header that can be inserted or removed from a packet via push and pop actions
XGEM Port	An abstraction in the TWDM TC service adaptation sublayer representing a logical connection associated with a specific client packet flow.

Acronyms

Alloc-ID	Allocation Identifier
DB	Data Base
DBA	Dynamic Bandwidth Assignment
GEM	Gigabit-capable passive optical network Encapsulation Method
MAC	Media Access Control
ODN	Optical Distribuion Network
OF	OpenFlow
OLT	Optical Line Termination
OMCC	ONU Management and Control Channel
OMCI	ONU Management and Control Interface
ONF	Open Networking Foundation
ONU	Optical Network Unit
PON	Passive Optical Network
QoS	Quality of service
R_A	PON traffic flow Assured Bandwidth component (G.989.3 Section 7.1)
R_F	PON traffic flow Fixed Bandwidth component (G.989.3 Section 7.1)
R_M	PON traffic flow Maximum Bandwidth component (G.989.3 Section 7.1)
χ_{AB}	PON traffic flow Ternary eligibility indicator for additional bandwidth assignment: {none, non-assured (NA), best-effort (BE)}; (G.989.3 Section 7.1)
RED	Random Early Discard
SDN	Software-defined networking
SDNC	SDN controller
SLA	Service level agreement
SP	Strict Priority
T-CONT	Transmission Container
TOR	Top of rack
VID	VLAN identifier
VLAN	Virtual local area network
002-xxxxx-000n	

VM	Virtual machine
VN	Virtual network
VNE	Virtual network element
WAN	Wide area network
WRED	Weighted Random Early Discard
WRR	Weighted Round Robin
XGEM	10-Gigabit-capable PON Encapsulation Method

ISSUES LIST

This section provides a list of issues that need to be resolved in order to complete the document. Additional are scattered throughout the document, flagged by “!!!!” or formatted as questions.

1. [Issue # N.] RESOLUTION → Resolution of Issue #N.

OPENFLOW ABSTRACTION ISSUE WITH MANAGING NON-ETHERNET ACCESS TECHNOLOGIES

Problem Statement

Currently OpenFlow can program Match/Action rules for unidirectional Service Flows but cannot specify Queue Scheduling Policies or Access Technology Specific attributes:

For xPON:

- T-CONT and GEM Port allocations
- T-CONT Upstream Bandwidth Scheduling Attributes(G.989.3 Section 7)
- GEM Port Scheduling Policy

Similarly for Access Technologies such as EPON or G.fast there is the same problem.

Note that OpenFlow can additionally specify Meters and/or an output Queue for a given Flow on a given interface.

The goal would be to eliminate the need for adding additional Technology Specific Data Models (such as Yang) for creating each and every Subscriber Flow on each and every specific Access Technology directly on VOLTHA so that VOLTHA can remain a generic hardware abstraction for OLT based services using OpenFlow.

Solution Proposal

Instead of supporting technology specific Service Data Models in VOLTHA, and implementing the Local and Northbound interfaces required for Service creation, it is proposed that VOLTHA will support a set of 'Technology Profiles' which describe technology specific attributes required to implement Subscriber Services on an OpenFlow managed Logical Switch overlaid upon an OLT or other technology specific platform.

Each 'Technology Profile' can be an XML encoded(TBD) technology specific Record with an OpenFlow Flow Table ID as part of the Record. Each Service Flow will reference a Flow Table with a Table ID containing a default Flow Entry which will 'Match All' Flows, with No Actions and this Flow Table ID will be used to Reference a Technology Profile with the Access Technology Specific information contained within it, examples are provided at the end of the document.

It is proposed that optionally:

- 1) A Local User Interface will allow creation of each Technology Profile Record,
- 2) The Technology Profile Records can be created via a Northbound interface to an SDN Controller,
- 3) The Records can be downloaded via Secure File Transfer from a Server.

Also initially, the Technology Profile Records can be Locally Stored as Fixed Default Entries to Speed development.

Each Subscriber Record within the SADIS Application used by ONOS and the OpenFlow Manager can be updated to indicate a Flow Table ID to use as the Reference to a Technology Profile known to VOLTHA.

A reserved Range of Flow Table IDs could be used for Technology Profile reference (e.g 64 - 255) ONOS Applications will manage the Flow Table ID to Technology Profile reference. Any Service Flow without a Flow Table ID Reference for a

Service Flow will default to use a System Default Profile – this Profile can be configured after system startup and stored in non-volatile memory. The Flow Table ID range 0-63 can be used for normal OpenFlow packet processing.

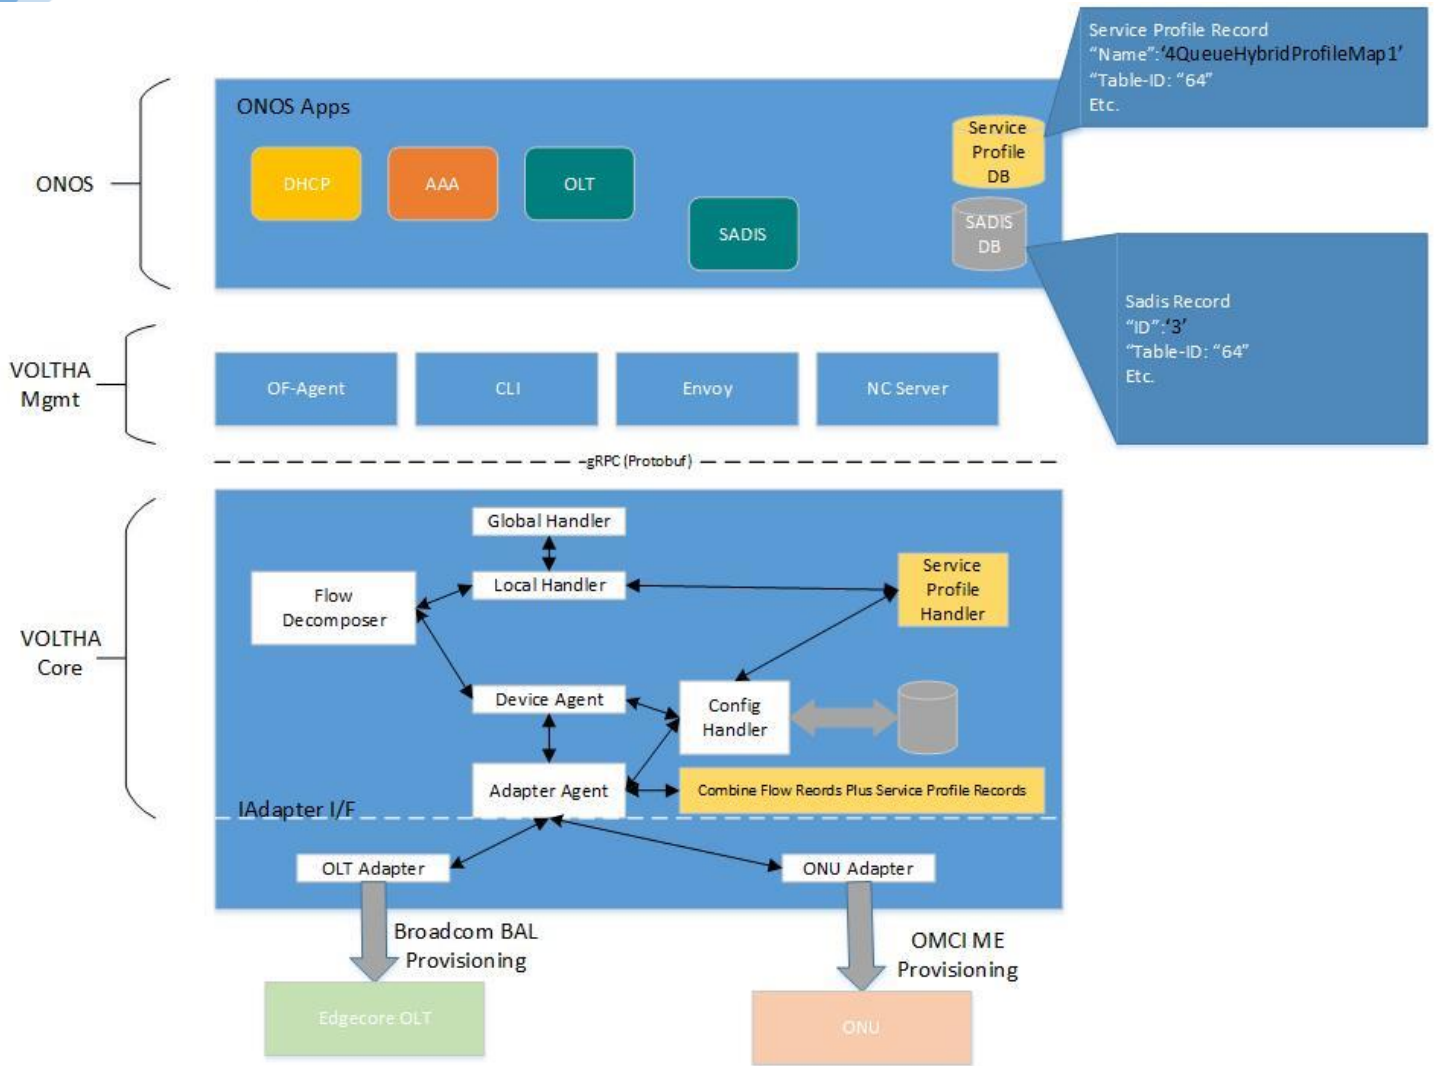


Figure 1 Adding Technology Profile as Part of VOLTHA

One area of concern from a scalability perspective is the number of Technology Profiles required since the Flow Table ID field is 8 bits wide. This concern is mitigated somewhat by the fact that currently we use only two Flow Table IDs for OLT based Services (0 and 1), and the fact that the expectation would be that for the same type of service across an Operators subscriber base, a Single Technology Profile can be used.

Generally Operators use three models of Service configuration:

- 1) Every attribute is defined for each Subscriber service and explicitly configured each time.
- 2) A Profile is defined for each Subscriber Service and applied per Subscriber
- 3) A Profile is used but Overrides are defined for Service attributes which can be variable between Subscribers with the same Service e.g. bandwidth attributes.

Since we want to support a small number of Technology Profiles it would be more scalable to separate out any bandwidth specific Flow information and use OpenFlow Meters to carry that.

If a Service bandwidth requirement changes then the OpenFlow service definition Meter Bands will be updated and the modified bandwidth parameters downloaded independent of the Access Specific Technology Profile and the per Subscriber Bandwidth can change independent of the referenced Technology Profile/Flow Table ID.

Business Services (MEF 10.3) are usually defined in terms of the following attributes:

- CIR – Committed Information Rate
- CBS - Committed Burst Size
- EIR - Excess Information Rate
- EBS - Excess Burst Size

A Residential Service Flow may be defined in terms of PIR and PBS and sometimes in terms of Policy based scheduling models defined by Priority and Weight.

- PIR – Peak Information Rate
- PBS - Peak Burst Size

For Access Technologies that use a Downstream Broadcast method with a Multi-Access Upstream scheme where Upstream Transmission Opportunities are allocated either based on Demand or a 'Fixed' Configuration a third type of Bandwidth can be assigned which represents the 'Fixed'/Unsolicited bandwidth which we will term 'Guaranteed' Information Rate (GIR).

Where CIR, EIR and PIR can be considered 'Work Conserving' i.e. Upstream Transmission Opportunities can be granted based on Demand with no waste, the GIR Upstream Transmission opportunities are allocated based on 'Fixed' Configuration independent of Demand and if not used by the Subscriber they are allocated for are 'lost' and not available for others to use.

For ITU based PON Systems the GIR Bandwidth represents the 'Fixed' F_R component of Upstream Bandwidth (See ITU G.989.3, Section 7) and for IEEE EPON systems it could represent Upstream Bandwidth configured for Unsolicited Grant Service (See Cable Labs DPoE-SP-MULPIv2.0-I13-180228).

A Flow is a unidirectional construct and there will be both an Upstream and Downstream component of the Technology Profile. Only Upstream Flows will program Guaranteed Information Rate bandwidth, whereas Downstream Flows are limited to CIR/EIR/PIR bandwidth.

The parameters CIR/CBS/EIR/EBS can be used for both policing and shaping (MEF 10.3 A.3) both Upstream and Downstream, GIR will also be applied for Policing and Shaping purposes but only in the Upstream direction – the GIR and CIR components are combined for the CIR part of a two rate three color policer.

For the Access nodes the need is primarily for shaping to reduce loss and optimize TCP throughput.

OPENFLOW AND PON MEF TO METER BAND CONVERSION

The OpenFlow 'Meter' Table will be used to convey the per Flow Bandwidth Profile Parameters of either the Upstream or Downstream traffic depending on the directionality of the Flow(See Section 5.7 of the OpenFlow 1.3.4 specification):

Per the OpenFlow standard the Meter Table consists of Meter Entries and each Meter Entry contains:

- Meter Identifier: a 32 bit unsigned integer uniquely identifying the meter
- Meter Bands: an unordered list of meter bands, where each meter band specifies the rate of the band and the way to process the packet
- Counters: updated when packets are processed by a meter

Each Meter Band is identified by its Rate and contains:

- Band Type : This field represents Optional behaviours, 'drop' (1) or 'dscp remark' (2) (Note: we would use option 2 set to not change the DSCP value)
- Rate: used by the Meter to select the Meter Band, and defines the lowest rate at which the band can apply(Kb/s)
- Burst: Burst Size in Kbits
- Counters: updated when packets are processed by a meter – See Table 5 OpenFlow Spec 1.3.4
- Band Type Specific Arguments: Increment to DSCP Field (zero increment in this case)

In order to support the Upstream PON DBA model with OpenFlow Meters one Meter Band in a Meter will represent the CIR (Upstream:G.989.3 Section 7, R_A :Assured bandwidth) and CBS attributes and a second Meter Band will represent the CIR+EIR = PIR(Peak Information Rate) value (Upstream:G.989.3 Section 7, R_M :Maximum bandwidth) and the CBS+EBS = PBS.

Where 'Fixed' Bandwidth (F_R) is required a third Meter Band will be used in the Upstream direction only. The Third Meter Band will Specify a Rate with a '0' Burst capability – this will identify the GIR bandwidth component.

The CBS and PBS Meter Band Burst values MUST be > 0.

Downstream the MEF bandwidth model is used as applied to the OpenFlow Meter Band programming where one meter band represents CIR and CBS, another the PIR(CIR+EIR) and PBS(CBS+EBS).

Bandwidth Rules:

In order to use OpenFlow Meters for PON Scheduling the following conventions apply:

1) Where No Meter is specified for a Flow:

- **Upstream Only:**The Flow scheduling bandwidth can be defined by a combination of { ($R_M = \text{PON Rate Max}$), $\chi_{AB(BE)}$ }.
- If the Extended PON Bandwidth Model is Enabled then Bandwidth is assigned based on Priority and Weight as specified in the Technology Profile.
- **Downstream Only:** Packet Transmission based on Flow Priority and Queue Weight. No PIR Specified so no rate limit.

2) Where a Single Meter Band is specified in a Meter and Burst Size > 0:

- The Meter Band will represent the PIR(Peak Information Rate) value (Upstream Only:G.989.3 Section 7, **R_M :Maximum bandwidth**) and the PBS(Peak Burst Size).
 - **Upstream Only:**The PON scheduling bandwidth can be defined by a combination of $\{R_M, \chi_{AB(BE)}\}$.
 - **Downstream Only:** Packet Transmission based on Flow Priority and Queue Weight. Flow Rate limited to PIR Specified.

3) Where a Single Meter Band is specified in a Meter and Burst Size = 0:

- One Meter Band with Burst == 0 MUST represent the GIR (Upstream Only:G.989.3 Section 7, **R_F :Fixed bandwidth**).
 - **Upstream Only:**The PON scheduling bandwidth can be defined by a combination of $\{R_F, \chi_{AB(=None)}\}$.
 - **Downstream Only:** Not applicable to Downstream. If CIR is desired there must be two Meter Bands specified for Downstream Flows.

4) Where Two Meter Bands are specified in a Meter (Burst Size == Don't Care):

- One Meter Band will represent the CIR (Upstream:G.989.3 Section 7, **R_A :Assured bandwidth, Downstream: CIR**) and CBS attributes
- A second Meter Band will represent the CIR+EIR = PIR(Peak Information Rate) value (Upstream Only:G.989.3 Section 7, **R_M :Maximum bandwidth**) and the CBS+EBS = PBS(Peak Burst Size):
 - If PIR > CIR: the flow is AF with two rates; { Upstream: $R_A, R_M, \chi_{AB(NA)}$ }, {Downstream:PIR}.
 - If PIR == CIR: the flow is AF with CIR component only; { Upstream: $R_A = R_M, \chi_{AB(None)}$ }, {Downstream:PIR==CIR}.
- Note: CIR <= PIR, so does not matter for order of Band specification

5) Where Three Meter Bands are specified in a Meter (Upstream Only):

- One Meter Band with Burst == 0 MUST represent the GIR (Upstream Only:G.989.3 Section 7, **R_F :Fixed bandwidth**).
- A Second Meter Band will represent the CIR (Upstream Only:G.989.3 Section 7, **R_A :Assured bandwidth, $\chi_{AB(NA)}$**) and CBS attributes (CIR > 0 and CIR<PIR).
 - If CIR == 0 and PIR > 0 then **R_A :Assured bandwidth = 0 and $\chi_{AB} = BE$** .
 - If CIR == 0 and PIR == 0 then this is a GIR only Meter: **R_F :Fixed bandwidth, $\chi_{AB} = None$** .
- A Third Meter Band will represent the CIR+EIR = PIR(Peak Information Rate) value (Upstream Only:G.989.3 Section 7, **R_M :Maximum bandwidth**) and the CBS+EBS = PBS:
 - If PIR > CIR: the flow is Type 5 T-CONT with three rates; { Upstream Only: $R_F, R_A, R_M, \chi_{AB(NA)}$ }.
 - If PIR == CIR: the flow is Type 5 T-CONT with two rates; { Upstream Only: $R_F, R_A = R_M, \chi_{AB(None)}$ }.

Note1: χ_{AB} is inferred from the specification and use of $\{R_F, R_A, R_M\}$. However the Technology Profile allows the χ_{AB} Attribute to be overridden by selecting a specific value as opposed to 'auto'.

Note2: An Upstream Meter Band with a Burst Size == 0 can only be specified when there is a Single Meter Band (i.e. only 'Fixed' Bandwidth) or there are Three Meter Bands (i.e. a combination of R_F, R_A, R_M).

Note3: A Downstream Meter supports a Maximum of Two Meter Bands (CIR and PIR) – three Meter Bands are not applicable to Downstream Flows.

R_F	R_A	R_M	$R_A = R_M$	$R_A < R_M$	Inferred χ_{AB}
x					None
	x	x	x		None
	x	x		x	Non-Assured
x	x	x		x	Non-Assured
		x			Best Effort
x		x			Best Effort
x	x	x	x		None

Table 2 DBA Inferred Additional Bandwidth Type

xPON Technology Profile

This section describes an xPON Technology Specific Technology Profile:

TECHNOLOGY PROFILE IDENTIFIERS:

Name: 'Technology Profile Name'

OpenFlow Table ID; Used by the OpenFlow Service definition to Reference the Technology Profile

Profile Type: xPON

Profile Version: 1.0

COMMON INSTANCE CONTROL ATTRIBUTES:

- Single Instance per ONU or multi-instance
- Single Instance per UNI or multi-instance (Not Applicable if Single Instance per ONU is TRUE)
- Number of GEM Port-ID (Queues)

UPSTREAM SCHEDULER ATTRIBUTES:

- T-CONT; Additional Bandwidth {None | NA | BE | Auto}
- T-CONT; xPON Priority (DBA Scheduling Priority for best-effort bandwidth assignment)
- T-CONT; xPON Weight (DBA Scheduling weight for best-effort bandwidth assignment)
- Queue Scheduling Policy; {WRR | Strict Priority | Hybrid}

UPSTREAM GEM PORT CONFIGURATION AND SCHEDULING ATTRIBUTES:

- Per GEM Port Attributes:
 - a) P-Bit Map to GEM Port
 - b) AES Encryption
 - c) GEM Port Scheduling Policy; {WRR | Strict Priority}
 - d) GEM Port (Queue) Priority
 - e) GEM Port (Queue) Weight
 - f) Queue Discard Policy; {Tail-drop | WTail-drop | RED | WRED}
 - g) Discard Control Attributes:
 - i. Min Threshold

- ii. Max Threshold
- iii. Max Probability

DOWNSTREAM AGGREGATE SCHEDULER ATTRIBUTES:

- Priority (Scheduling Priority)
- Weight (Scheduling weight)
- Queue Scheduling Policy; {WRR | Strict Priority | Hybrid}

DOWNSTREAM GEM PORT CONFIGURATION AND SCHEDULING ATTRIBUTES:

- Per GEM Port Attributes:
 - a) P-Bit Map to GEM Port
 - b) AES Encryption
 - c) GEM Port (Queue)Scheduling Policy; {WRR | Strict Priority}
 - d) GEM Port (Queue) Priority
 - e) GEM Port (Queue) Weight
 - f) Queue Discard Policy; {Tail-drop | WTail-drop | RED | WRED}
 - g) Discard Control Attributes:
 - iv. Min Threshold
 - v. Max Threshold
 - vi. Max Probability

Service Creation Examples

The following sections detail specific Service creation examples and a proposed set of 'Built-in' Technology Profiles specifically for xPON, but this can be extended with more Service Specific examples with other access technologies.

Generically all xPON Services Profiles will tend to fall into the following categories:

- Single T-CONT with Single GEM Port
- Single T-CONT with Two GEM Ports
- Single T-CONT with Four GEM Ports
- Single T-CONT with Eight GEM Ports

Operators will vary as to whether the T-CONT is unique per ONU (a PON Class of Service - e.g. TR-156 4+1 model), per UNI per ONU or per Service per UNI per ONU – this usage can be controlled per Technology Profile via the attributes:

- T-CONT; Single Instance per ONU or multi-instance
- T-CONT; Single Instance per UNI or multi-instance

If the model is a single instance per ONU then the Second Attribute 'instance per UNI' control is not used. In this case all Service Flows referencing this Technology Profile will share a Single T-CONT (This is the TR156 4+1 model) where 4 Service T-CONTs are created per ONU and each T-CONT represent a distinct PON Class of Service. There will be 4 Technology Profiles required in the TR156 4+1 case each representing a T-CONT Instance per ONU and a single GEM Port Instance per UNI.

Where a T-CONT is shared per ONU then having unique instances of GEM Ports per UNI allows rate limiting the GEM Flows on the ONU such that any given service cannot access the Total available T-CONT Bandwidth since the T-CONT would be assigned the Sum of the Upstream Bandwidths for each of the Flows referencing it.

If the T-CONT per ONU is Multi-Instance instead of Single Instance then each Service Flow referencing the Technology Profile will have its's own instance of a T-CONT governed by the second T-CONT attribute the Per 'UNI instance' field.

If the second T-CONT field indicates Single Instance per UNI there can be a single shared T-CONT per UNI for all Service flows referencing the Profile. Each Service Flow in this case will share the same set of GEM Ports per UNI, and the T-CONT allocated the sum of the Upstream Bandwidths of the Flows referencing it.

If the second T-CONT field indicates Multi-instance then each Service Flow reference has a unique T-CONT Instance as well as unique GEM Port Instances.

The next TCONT field provides an override to the Meter Inferred Additional Bandwidth Field:

- TCONT: Additional Bandwidth {None | NA | BE | Auto}

'Auto' Signals that the Inferred 'Additional Bandwidth' field should be used whereas any other value indicates that the Inferred value (See:

Table 2 DBA Inferred Additional Bandwidth Type) should be overwritten.

Since the Technology Profile is allocated to a certain Class of Traffic the T-CONT Priority and Weight are also part of the Technology Profile:

- T-CONT; xPON Priority (DBA Scheduling Priority – Extended Model)
- T-CONT; xPON Weight (DBA Scheduling weight – Extended Model)

Where Additional Bandwidth Assignment is managed via the Priority and Weight Mechanism these two DBA Scheduling Attributes are used.

The T-CONT Queue (GEM Port) Scheduling policy is defined in the Technology Profile:

- Queue Scheduling Policy; {WRR | Strict Priority | Hybrid}

Each Technology Profile will additionally contain information about the number of GEM Ports (Queues) assigned to the service and for each GEM Port the PBits assigned (used to create a GEM Port Mapper ME via OMCI) , whether AES encryption is enabled, the GEM Port Scheduling behavior, its T-CONT Priority Queue Assignment, Weight and Discard Policy.

Operators may select a Technology Profile with a Single T-CONT and 4 GEM Ports but may differ in the PBit Mapping to the GEM Ports, AES, Scheduling and Discard Policies, requiring multiple instances of this Type of Technology Profile (unique OpenFlow Table ID). Once an Operator has selected a Type of Technology Profile then OpenFlow is used to select the Match Action Rules on Ingress and the Bandwidth Profile applied.

The T-CONT and its' Queue(s) are shown pictorially in the following diagram, Figure 2:

Note: In the Downstream Data Path the T-CONT plays no part in the packet handling, only the GEM Port-IDs – T-CONT is an Upstream Construct only used to assign Alloc-IDs for which upstream Grants are issued by the DBA Engine.




Figure 2 ONU and OLT Traffic Management

SINGLE T-CONT WITH FOUR GEM PORTS EXAMPLE (4QUEUEHYBRIDPROFILEMAP1)

The AT&T Open OMCI document describes an High Speed Internet Access (HSIA) Service with a Single T-CONT and Four Gem Ports (Queues):

<i>Service</i>	<i>Traffic</i>	<i>Priority</i>	<i>PBIT</i>	<i>Scheduling</i>	<i>Upstream</i>	<i>Downstream</i>	<i>Traffic Type</i>
<i>GEM</i>	<i>Class</i>	<i>Queue</i>		<i>Policy</i>	<i>Weight</i>	<i>Weight</i>	
<i>GEM4</i>	TC6	Q1	6,7	Strict Priority	N/A	N/A	NC, H1 (voice)
<i>GEM3</i>	TC4	Q2	5	Strict Priority	N/A	N/A	EF (igmp)
<i>GEM2</i>	TC2	Q3	1,3,4	WRR	75%	90%	H2, L1 (vod/icc)
<i>GEM1</i>	TC0	Q4	0,2	WRR	25%	10%	BE,AF (HSIA)

The 'Built-in' Technology Profile which can be used for an instance of this Service is described by the **'4QueueHybridProfileMap1'** Technology Profile:

Name: '4QueueHybridProfileMap1'

OpenFlow Table ID; **64**

Profile Type: xPON

Profile Version: 1.0

Upstream Aggregate Scheduler Attributes:

- T-CONT; Single Instance per ONU or multi-instance: **Multi-instance**
- T-CONT; Single Instance per UNI or multi-instance: **Single-instance**
- T-CONT; Additional Bandwidth: **Auto**
- T-CONT; xPON Priority (DBA Scheduling Priority): **0**
- T-CONT; xPON Weight (DBA Scheduling weight): **0**
- Queue Scheduling Policy; {WRR | Strict Priority | Hybrid}: **Hybrid**

Upstream GEM Port configuration and scheduling attributes:

- GEM Ports (Queues) per Service: **4**

- **GEM Port 1 Attributes:**
 - i. P-Bit Map to GEM Port: **PBit Mask (00000101), PBits 0 and 2**
 - ii. AES Encryption: **True**
 - iii. GEM Port Scheduling Policy; {WRR | Strict Priority}: **WRR**
 - iv. GEM Port (Queue) Priority: **0**
 - v. GEM Port (Queue) Weight: **25%**
 - vi. Queue Discard Policy; {Tail-drop | WTail-drop | RED | WRED}: **Tail-drop**
 - a. Min Threshold:**0 (N/A)**
 - b. Max Threshold:**0 (N/A)**
 - c. Max Probability:**0 (N/A)**
- **GEM Port 2 Attributes:**
 - i. P-Bit Map to GEM Port: **PBit Mask (00011010), PBits 1, 3 and 4**
 - ii. AES Encryption: **True**
 - iii. GEM Port Scheduling Policy; {WRR | Strict Priority}: **WRR**
 - iv. GEM Port (Queue) Priority: **1**
 - v. GEM Port (Queue) Weight: **75%**
 - vi. Queue Discard Policy; {Tail-drop | WTail-drop | RED | WRED}: **Tail-drop**
 - a. Min Threshold:**0 (N/A)**
 - b. Max Threshold:**0 (N/A)**
 - c. Max Probability:**0 (N/A)**
- **GEM Port 3 Attributes:**
 - i. P-Bit Map to GEM Port: **PBit Mask (00100000), PBit 5**
 - ii. AES Encryption: **True**
 - iii. GEM Port Scheduling Policy; {WRR | Strict Priority}: **Strict Priority**
 - iv. GEM Port (Queue) Priority: **2**
 - v. GEM Port (Queue) Weight: **0**
 - vi. Queue Discard Policy; {Tail-drop | WTail-drop | RED | WRED}: **Tail-drop**
 - a. Min Threshold:**0 (N/A)**
 - b. Max Threshold:**0 (N/A)**
 - c. Max Probability:**0 (N/A)**
- **GEM Port 4 Attributes:**
 - i. P-Bit Map to GEM Port: **PBit Mask (11000000), PBits 6 and 7**
 - ii. AES Encryption: **True**
 - iii. GEM Port Scheduling Policy; {WRR | Strict Priority}: **Strict Priority**
 - iv. GEM Port (Queue) Priority: **3**
 - v. GEM Port (Queue) Weight: **0**
 - vi. Queue Discard Policy; {Tail-drop | WTail-drop | RED | WRED}: **Tail-drop**
 - a. Min Threshold:**0 (N/A)**
 - b. Max Threshold:**0 (N/A)**
 - c. Max Probability:**0 (N/A)**

Downstream Aggregate Scheduler Attributes:

- Single Instance per ONU or multi-instance: **Multi-instance**
- Single Instance per UNI or multi-instance: **Single-instance**
- Priority (Scheduling Priority) : **0**

- Weight (Scheduling weight) : 0
- Queue Scheduling Policy; {WRR | Strict Priority | Hybrid}: Hybrid

Downstream GEM Port configuration and scheduling attributes:

- GEM Ports (Queues) per Service: 4
- GEM Port 1 Attributes:
 - P-Bit Map to GEM Port: PBit Mask (00000101), PBits 0 and 2
 - AES Encryption: True
 - GEM Port Scheduling Policy; {WRR | Strict Priority}: WRR
 - GEM Port (Queue) Priority: 0
 - GEM Port (Queue) Weight: 10%
 - Queue Discard Policy; {Tail-drop | WTail-drop | RED | WRED}: Tail-drop
 - Min Threshold:0 (N/A)
 - Max Threshold:0 (N/A)
 - Max Probability:0 (N/A)
- GEM Port 2 Attributes:
 - P-Bit Map to GEM Port: PBit Mask (00011010), PBits 1, 3 and 4
 - AES Encryption: True
 - GEM Port Scheduling Policy; {WRR | Strict Priority}: WRR
 - GEM Port (Queue) Priority: 1
 - GEM Port (Queue) Weight: 90%
 - Queue Discard Policy; {Tail-drop | WTail-drop | RED | WRED}: Tail-drop
 - Min Threshold:0 (N/A)
 - Max Threshold:0 (N/A)
 - Max Probability:0 (N/A)
- GEM Port 3 Attributes:
 - P-Bit Map to GEM Port: PBit Mask (00100000), PBit 5
 - AES Encryption: True
 - GEM Port Scheduling Policy; {WRR | Strict Priority}: Strict Priority
 - GEM Port (Queue) Priority: 2
 - GEM Port (Queue) Weight: 0
 - Queue Discard Policy; {Tail-drop | WTail-drop | RED | WRED}: Tail-drop
 - Min Threshold:0 (N/A)
 - Max Threshold:0 (N/A)
 - Max Probability:0 (N/A)
- GEM Port 4 Attributes:
 - P-Bit Map to GEM Port: PBit Mask (11000000), PBits 6 and 7
 - AES Encryption: True
 - GEM Port Scheduling Policy; {WRR | Strict Priority}: Strict Priority
 - GEM Port (Queue) Priority: 3
 - GEM Port (Queue) Weight: 0
 - Queue Discard Policy; {Tail-drop | WTail-drop | RED | WRED}: Tail-drop
 - Min Threshold:0 (N/A)
 - Max Threshold:0 (N/A)
 - Max Probability:0 (N/A)

OpenFlow Mapping to Technology profile '4QueueHybridProfileMap1' Using Table ID

Since the '4QueueHybridProfileMap1' is a 'Built-in' Pre-defined Technology Profile it can be allocated a Table ID from a range 64 – 127; Profiles within this 'Built-in' range cannot be modified dynamically but new Technology Profiles can be added in the Table ID range 128 – 255.

The '4QueueHybridProfileMap1' Technology Profile is assigned the Table ID value of 64, and assume that the Service Flow associated with Table ID 64 is a HSIA and the Customer C-TAG to be used is 122, with an S-Tag of 2.

The Service Frame is passed Upstream to the ONU UNI and will Ingress the UNI as a Priority Tagged Frame (VLAN = 0) and Egress the NNI as a Double Tagged Frame, S-Tag/C-Tag of 2/122.

Downstream the Customer Traffic will Ingress the OLT NNI as a Double Tagged Frame with S-Tag/C-Tag of 2/122 and Egress the Customer UNI as a Priority Tagged Frame.

The Customer has signed up for a 1G Upstream and 1G Downstream Service.

OpenFlow will create a Meter in the Meter Table for the Customer HSIA **Downstream** Service Flow: Single Meter Band:

Band Type : NA(No actions on the Packet when Band Exceeded)

Rate: 1Gb/s

Burst: 384 Kbits

OpenFlow will create a Meter in the Meter Table for the Customer HSIA **Upstream** Service Flow : Single Meter Band:

Band Type : NA(No actions on the Packet when Band Exceeded)

Rate: 1Gb/s

Burst: 384 Kbits

Note that, port 129 is NNI and port 21 is UNI. Also 2/122 are s-tag and c-tag for the given Customer.

OpenFlow Customer 2/122 Downstream Messages:

Table 64 (The Technology Profile Reference) contains a Default Rule of Match-All, No Action.

```
ADDED, bytes=0, packets=0, table=0, priority=1000, selector=[IN_PORT:129, METADATA:15, VLAN_VID:2],  
treatment=[immediate=[VLAN_POP], transition=TABLE:1]
```

```
ADDED, bytes=0, packets=0, table=1, priority=1000, selector=[IN_PORT:129, VLAN_VID:122],  
treatment=[immediate=[VLAN_ID:0], deferred=[OUTPUT:21], meter=ID, transition=TABLE:64]
```

In “table 1” we say that if the frame customer C-tag 122 is coming in from NNI port 129, we want to change the vlan to be Priority Tagged VLAN 0 and associate the meter (ID) with this flow, following which we want to go to “table 64” for the rest of the Technology profile, and we will defer the output action (to port 21) once the packet has gone through table 64. Since the Default flows in table 64 are Match-All and No Action, the pipeline execution will stop and the accumulated action-set will be executed.

Note: The ONU will automatically Reverse the Upstream Tagging Action for Downstream Frames without needing to be explicitly configured.

OpenFlow Customer 2/122 Upstream Messages:

```
ADDED, bytes=0, packets=0, table=0, priority=1000, selector=[IN_PORT:21, VLAN_VID:0],  
treatment=[immediate=[VLAN_ID:122], transition=TABLE:1]
```

```
ADDED, bytes=0, packets=0, table=1, priority=1000, selector=[IN_PORT:21, VLAN_VID:122],  
treatment=[immediate=[VLAN_PUSH:vlan, VLAN_ID:2], deferred=[OUTPUT:129], meter=ID,  
transition=TABLE:64]
```

SINGLE T-CONT WITH EIGHT GEM PORTS EXAMPLE ('8QUEUESTRICTPROFILEMAP1')

The DT Service Presentation document describes a High Speed Internet Access (HSIA) Service with a Single T-CONT and Eight Gem Ports (Queues):

Service GEM	Traffic Class	Priority Queue	PBIT	Scheduling Policy	Upstream Weight	Downstream Weight	Traffic Type
GEM8	TC7	Q0	7	Strict Priority	N/A	N/A	?
GEM7	TC6	Q1	6	Strict Priority	N/A	N/A	?
GEM6	TC5	Q2	5	Strict Priority	N/A	N/A	?
GEM5	TC4	Q3	4	Strict Priority	N/A	N/A	?
GEM4	TC3	Q4	3	Strict Priority	N/A	N/A	?
GEM3	TC2	Q5	2	Strict Priority	N/A	N/A	?
GEM2	TC1	Q6	1	Strict Priority	N/A	N/A	?
GEM1	TC0	Q7	0	Strict Priority	N/A	N/A	?

The 'Built-in' Technology Profile which can be used for an instance of this Service is described by the **'8QueueStrictProfileMap1' Technology Profile:**

Name: '8QueueStrictProfileMap1'

OpenFlow Table ID; **65**

Profile Type: xPON

Profile Version: 1.0

Upstream Aggregate Scheduler Attributes:

- T-CONT; Single Instance per ONU or multi-instance: **Multi-instance**
- T-CONT; Single Instance per UNI or multi-instance: **Single-instance**
- T-CONT; Additional Bandwidth: **Auto**
- T-CONT; xPON Priority (DBA Scheduling Priority): **0**
- T-CONT; xPON Weight (DBA Scheduling weight): **0**

- Queue Scheduling Policy; {WRR | Strict Priority | Hybrid}: **Strict Priority**

Upstream and Downstream (identical) GEM Port configuration and scheduling attributes:

- GEM Ports (Queues) per Service: **8**
- **GEM Port 1 Attributes:**
 - P-Bit Map to GEM Port: **PBit Mask (00000001), PBit 0**
 - AES Encryption: **True**
 - GEM Port Scheduling Policy; {WRR | Strict Priority}: **Strict Priority**
 - GEM Port (Queue) Priority: **0**
 - GEM Port (Queue) Weight: **0**
 - Queue Discard Policy; {Tail-drop | WTail-drop | RED | WRED}: **Tail-drop**
 - Min Threshold:**0 (N/A)**
 - Max Threshold:**0 (N/A)**
 - Max Probability:**0 (N/A)**
- **GEM Port 2 Attributes:**
 - P-Bit Map to GEM Port: **PBit Mask (00000010), PBit 1**
 - AES Encryption: **True**
 - GEM Port Scheduling Policy; {WRR | Strict Priority}: **Strict Priority**
 - GEM Port (Queue) Priority: **1**
 - GEM Port (Queue) Weight: **0**
 - Queue Discard Policy; {Tail-drop | WTail-drop | RED | WRED}: **Tail-drop**
 - Min Threshold:**0 (N/A)**
 - Max Threshold:**0 (N/A)**
 - Max Probability:**0 (N/A)**
- **GEM Port 3 Attributes:**
 - P-Bit Map to GEM Port: **PBit Mask (00000100), PBit 2**
 - AES Encryption: **True**
 - GEM Port Scheduling Policy; {WRR | Strict Priority}: **Strict Priority**
 - GEM Port (Queue) Priority: **2**
 - GEM Port (Queue) Weight: **0**
 - Queue Discard Policy; {Tail-drop | WTail-drop | RED | WRED}: **Tail-drop**
 - Min Threshold:**0 (N/A)**
 - Max Threshold:**0 (N/A)**
 - Max Probability:**0 (N/A)**
- **GEM Port 4 Attributes:**
 - P-Bit Map to GEM Port: **PBit Mask (00001000), PBit 3**
 - AES Encryption: **True**
 - GEM Port Scheduling Policy; {WRR | Strict Priority}: **Strict Priority**
 - GEM Port (Queue) Priority: **3**
 - GEM Port (Queue) Weight: **0**
 - Queue Discard Policy; {Tail-drop | WTail-drop | RED | WRED}: **Tail-drop**
 - Min Threshold:**0 (N/A)**
 - Max Threshold:**0 (N/A)**
 - Max Probability:**0 (N/A)**
- **GEM Port 5 Attributes:**

- i. P-Bit Map to GEM Port: **PBit Mask (00010000), PBit 4**
- ii. AES Encryption: **True**
- iii. GEM Port Scheduling Policy; {WRR | Strict Priority}: **Strict Priority**
- iv. GEM Port (Queue) Priority: **4**
- v. GEM Port (Queue) Weight: **0**
- vi. Queue Discard Policy; {Tail-drop | WTail-drop | RED | WRED}: **Tail-drop**
 - a. Min Threshold:**0 (N/A)**
 - b. Max Threshold:**0 (N/A)**
 - c. Max Probability:**0 (N/A)**
- **GEM Port 6 Attributes:**
 - i. P-Bit Map to GEM Port: **PBit Mask (00100000), PBit 5**
 - ii. AES Encryption: **True**
 - iii. GEM Port Scheduling Policy; {WRR | Strict Priority}: **Strict Priority**
 - iv. GEM Port (Queue) Priority: **5**
 - v. GEM Port (Queue) Weight: **0**
 - vi. Queue Discard Policy; {Tail-drop | WTail-drop | RED | WRED}: **Tail-drop**
 - a. Min Threshold:**0 (N/A)**
 - b. Max Threshold:**0 (N/A)**
 - c. Max Probability:**0 (N/A)**
- **GEM Port 7 Attributes:**
 - i. P-Bit Map to GEM Port: **PBit Mask (01000000), PBit 6**
 - ii. AES Encryption: **True**
 - iii. GEM Port Scheduling Policy; {WRR | Strict Priority}: **Strict Priority**
 - iv. GEM Port (Queue) Priority: **6**
 - v. GEM Port (Queue) Weight: **0**
 - vi. Queue Discard Policy; {Tail-drop | WTail-drop | RED | WRED}: **Tail-drop**
 - a. Min Threshold:**0 (N/A)**
 - b. Max Threshold:**0 (N/A)**
 - c. Max Probability:**0 (N/A)**
- **GEM Port 8 Attributes:**
 - i. P-Bit Map to GEM Port: **PBit Mask (10000000), PBit 7**
 - ii. AES Encryption: **True**
 - iii. GEM Port Scheduling Policy; {WRR | Strict Priority}: **Strict Priority**
 - iv. GEM Port (Queue) Priority: **7**
 - v. GEM Port (Queue) Weight: **0**
 - vi. Queue Discard Policy; {Tail-drop | WTail-drop | RED | WRED}: **Tail-drop**
 - a. Min Threshold:**0 (N/A)**
 - b. Max Threshold:**0 (N/A)**
 - c. Max Probability:**0 (N/A)**

OpenFlow Mapping to Technology Profile '8QueueStrictProfileMap1' Using Table ID

Since the '8QueueStrictProfileMap1' is a 'Built-in' Pre-defined Technology Profile it can be allocated a Table ID from a range (TBD) 64 – 127; Profiles within this 'Built-in' range cannot be modified dynamically but new Technology Profiles can be added in the Table ID range 128 – 255.

Assume the '8QueueStrictProfileMap1' is assigned the Table ID value of 65, and assume that the Service Flow associated with Table ID 65 is a HSIA – the DT HSIA model requires the UNI Ingress frames to be Tagged (could be more than one Tag) and the ONU must pass through the Frame Unchanged with the OLT adding a Single S-Tag before Egressing the NNI Upstream. Assume the Customer is identified by the S-Tag of 1000.

Downstream the Customer Traffic will Ingress the OLT NNI as a Tagged Frame (more than Single Tagged) with S-Tag of 1000, which will be removed and Egress the Customer UNI again as a Tagged Frame.

The Customer has signed up for a PIR of 1G Upstream and 1G Downstream Service, with CIR of 100 Mb/s in both directions.

OpenFlow will create a Meter in the Meter Table for the Customer HSIA **Downstream** Service Flow: Two Meter Bands:

Band 1:

Band Type : NA(No actions on the Packet when Band Exceeded)

Rate: 1Gb/s

Burst: 384 Kbits

Band 2:

Band Type : NA(No actions on the Packet when Band Exceeded)

Rate: 100 Mb/s

Burst: 384 Kbits

OpenFlow will create a Meter in the Meter Table for the Customer HSIA **Upstream** Service Flow: Two Meter Bands:

Band 1:

Band Type : NA(No actions on the Packet when Band Exceeded)

Rate: 1Gb/s

Burst: 384 Kbits

Band 2:

Band Type : NA(No actions on the Packet when Band Exceeded)

Rate: 100 Mb/s

Burst: 384 Kbits

Note that, port 129 is NNI and port 21 is UNI. Also s-tag = 1000.

OpenFlow Customer S-Tag=1000 Downstream Messages:

```
ADDED, bytes=0, packets=0, table=0, priority=1000, selector=[IN_PORT:129, METADATA:15,
VLAN_VID:1000], treatment=[defer=[VLAN_POP], transition=TABLE:1]
```

```
ADDED, bytes=0, packets=0, table=1, priority=1000, selector=[IN_PORT:129, VLAN_VID:1000],
treatment=[deferred=[OUTPUT:21], meter=ID, transition=TABLE:65]
```

OpenFlow Customer S-Tag=1000 Upstream Messages:

```
ADDED, bytes=0, packets=0, table=0, priority=1000, selector=[IN_PORT:21, VLAN_VID: OFFVID_PRESENT],
treatment=[immediate=[], deferred=[], transition=TABLE:1]
```

```
ADDED, bytes=0, packets=0, table=0, priority=500, selector=[IN_PORT:21], treatment=[immediate=[],
deferred=[]]
```

Tagged Frames coming in at UNI port 21 will get forwarded ([OFFVID_PRESENT](#)). All other packets (untagged) on port 21 will match the lower priority flow, and get dropped (no actions, no goto-table).

```
ADDED, bytes=0, packets=0, table=1, priority=1000, selector=[IN_PORT:21, VLAN_VID: OFFVID_PRESENT],
treatment=[immediate=[VLAN_PUSH:QinQ, VLAN_ID:1000], deferred=[OUTPUT:129], meter=ID,
transition=TABLE:65]
```

Note: If you wanted the TPID to be 0x88a8, then the VLAN_PUSH action would have the tpid 0x88a8 and the ONOS representation would be VLAN_PUSH:QinQ, if tpid 0x8100 then represented in the ONOS flow as 'vlan'.

Documentation

n document.

A. References

A.1. Related Documents

A.1.1. Radisys Documents

A.1.2. Other Documents

A.1.3. Industry-Standard References

1. OpenFlow Switch Specification, Version 1.3.5, March 26, 2015
2. OpenFlow Management and Configuration Protocol, OF-CONFIG, Version 1.2, 2014
3. ITU G.989 Series; 40-Gigabit-capable passive optical networks (NGPON2)
4. ITU G.984 Series; Gigabit-capable passive optical networks (G-PON)