



Implementation Agreement

MEF 23.2.1

Models for Bandwidth Profiles with Token Sharing

Amendment to MEF 23.2

24 January 2017

Disclaimer

The information in this publication is freely available for reproduction and use by any recipient and is believed to be accurate as of its publication date. Such information is subject to change without notice and the MEF Forum (MEF) is not responsible for any errors. The MEF does not assume responsibility to update or correct any information in this publication. No representation or warranty, expressed or implied, is made by the MEF concerning the completeness, accuracy, or applicability of any information contained herein and no liability of any kind shall be assumed by the MEF as a result of reliance upon such information.

The information contained herein is intended to be used without modification by the recipient or user of this document. The MEF is not responsible or liable for any modifications to this document made by any other party.

The receipt or any use of this document or its contents does not in any way create, by implication or otherwise:

- a) any express or implied license or right to or under any patent, copyright, trademark or trade secret rights held or claimed by any MEF member which are or may be associated with the ideas, techniques, concepts or expressions contained herein; nor
- b) any warranty or representation that any MEF members will announce any product(s) and/or service(s) related thereto, or if such announcements are made, that such announced product(s) and/or service(s) embody any or all of the ideas, technologies, or concepts contained herein; nor
- c) any form of relationship between any MEF members and the recipient or user of this document.

Implementation or use of specific MEF standards or recommendations and MEF specifications will be voluntary, and no member shall be obliged to implement them by virtue of participation in the MEF Forum. The MEF is a non-profit international organization to enable the development and worldwide adoption of agile, assured and orchestrated network services. The MEF does not, expressly or otherwise, endorse or promote any specific products or services.

© The MEF Forum 2017. All Rights Reserved.

Table of Contents

1	LIST OF CONTRIBUTING MEMBERS.....	1
2	ABSTRACT	1
3	CHANGES TO SECTION 3 OF MEF 23.2 (TERMINOLOGY).....	2
4	CHANGES TO SECTION 8.6.1 OF MEF 23.2.....	2
5	NEW SECTION IN MEF 23.2	4
6	CHANGES TO SECTION 10 OF MEF 23.2 (REFERENCES)	22
7	NEW APPENDIX IN MEF 23.2	22

List of Figures

Figure A-1:	Bandwidth Profile Models - with and without token sharing in an Envelope	5
Figure A-2:	Concept of Dynamic Bandwidth Management	6
Figure A-3:	Generic Token Sharing Bandwidth Profile Model with three Flows.....	7
Figure A-4:	Example of a Customer Equipment Scheduler and Shaper connected to a MEF-defined service supporting token sharing Bandwidth Profiles.....	9
Figure A-5:	Token Flow for Model C/G/D.....	16
Figure A-6:	Token Flow for Model CX/G/R.....	18
Figure A-7:	Token Flow for Model CX/GY/R	20
Figure A-8:	Example of EVPL service used as access to IP-VPN Network	23
Figure A-9:	Token flow diagram for IP-VPN Use Case 1.....	24
Figure A-10:	Egress Bandwidth Profile with token sharing for two OVC End Points	27
Figure A-11:	Token flow diagram for Egress Bandwidth Profile Use Case	29
Figure A-12:	Simplified example of EVPL service used for Mobile Backhaul	31
Figure A-13:	Token flow diagram for the CX/GY/R Use Case	32
Figure A-14:	Service Model for Mobile Backhaul, 2-EVC Use Case.....	35
Figure A-15:	Token flow diagram for the 2 EVC Use Case.....	36

List of Tables

Table A-1:	Terms added to Table 1 of MEF 23.2.....	2
Table A-2:	Comparison of Bandwidth Profile Service Models.....	11
Table A-3:	Token Sharing Bandwidth Profile Models	15
Table A-4:	Bandwidth Profile Flow parameter requirements for Model C/G/D.....	17
Table A-5:	Bandwidth Profile Flow parameter requirements for Model CX/G/R	19
Table A-6:	Bandwidth Profile Flow parameter requirements for Model CX/GY/R	21
Table A-7:	Ingress Bandwidth Profile Flow Parameter Values for IP-VPN Use Case 1	25
Table A-8:	Ingress Bandwidth Profile Flow Parameter Values for IP-VPN Use Case 2	26
Table A-9:	PE shaping of IP traffic classes and mapping to CoS Label	28

Table A-10: Mapping of CoS Label per OVC to Egress Equivalence Class.....	28
Table A-11: Egress Bandwidth Profile Parameter Values.....	30
Table A-12: Ingress Bandwidth Profile Flow Parameter Values for the Mobile Backhaul Use Case	34
Table A-13: Ingress Bandwidth Profile Flow Parameter Values for Two-EVC Use Case	37

1 List of Contributing Members

albis-elcon

Bell Canada

Canoga Perkins

Ciena Corporation

Cisco Systems

Colt Technology Services

HFR, Inc.

PLDT Corp. Business Solutions

Telus

Verizon

Windstream

XO Communications

2 Abstract

This amendment makes the following changes to MEF 23.2¹:

- Changes are made to Section 8.6.1 of MEF 23.2 to generalize certain requirements to apply to CoS Names instead of just CoS Labels.
- A new normative section is added that defines a set of token sharing Bandwidth Profile models. These models allow for multiple Bandwidth Profile Flows in an Envelope thus enabling token sharing among the Bandwidth Profile Flows.
- New terms are added to the Terminology table of MEF 23.2
- A new appendix is added that describes use cases for the defined token sharing Bandwidth Profile models.

¹ Implementation Agreement, MEF 23.2, Carrier Ethernet Class of Service - Phase 3, August 2016.

3 Changes to Section 3 of MEF 23.2 (Terminology)

The following entries are added to the terminology table of MEF 23.2.

Term	Definition	Reference
Bandwidth Type	One of the classifiers of a token sharing Bandwidth Profile Model, based on the color(s) of traffic declared by the Bandwidth Profile for the set of Bandwidth Profile Flows in the Envelope. In this document, there are three types: <i>Committed (C)</i> , <i>Committed + Excess (CX)</i> or <i>Excess (X)</i> .	This document
MFS	Maximum Frame Size. In the context of an EVC, it refers to the value of the EVC Maximum Service Frame Size Service Attribute per MEF 10.3; and in the context of an OVC, it refers to the value of the OVC Maximum Frame Size Service Attribute, per MEF 26.2.	This document
Token Source	One of the classifiers of a token sharing Bandwidth Profile Model, based on the color of the token sources within the Envelope. In this document, there can be G (Green), GY (Green and Yellow) or Y (Yellow).	This document
Token Flow	One of the classifiers of a token sharing Bandwidth Profile Model based on the path that unused tokens use between Bandwidth Profile Flows in the Envelope. In this document, three types of token flow are defined: <i>Down</i> , <i>Across</i> and <i>Recirculate</i> .	This document
Token Sharing	A Bandwidth Profile model that allows for bandwidth sharing across two or more Bandwidth Profile Flows in an Envelope. Unused tokens from one Bandwidth Profile Flow can be made available to one or more other Bandwidth Profile Flows in the Envelope.	This document
Token Sharing Model	A specific set of requirements on the Bandwidth Profile parameter values applied to Bandwidth Profile Flows within an Envelope. The models are classified by: <i>Bandwidth Type</i> , <i>Token Source</i> and <i>Token Flow</i> .	This document
Unused token	With respect to a given Bandwidth Profile Flow, a token that either overflows a bucket for that Bandwidth Profile Flow or bypasses it because of the CIR_{max}^i or EIR_{max}^i value for that flow.	This document

Table A-1: Terms added to Table 1 of MEF 23.2

4 Changes to Section 8.6.1 of MEF 23.2

Subsection 8.6.1 in MEF 23.2 is changed, as follows. Changes are shown in underlined or struck-through text.

8.6.1 Ingress Bandwidth Profiles for CoS Names Labels

This IA does not mandate that an EVC or OVC with CoS Names Labels have ingress bandwidth profiles, however ~~in order to support the intended applications for those CoS Labels~~ this IA does impose constraints when ingress bandwidth profiles are present, in order to avoid unintended behavior of the bandwidth profiles.

- [R6] When an ingress Bandwidth Profile Flow is present for a CoS Name Label at an EI, the value of the CBS parameter for that flow **MUST** be either equal to zero or greater than or equal to:
- For an EVC, the EVC Maximum Service Frame Size as defined in MEF 10.3 [1]
 - For an OVC, the lower bound specified in Table 47 of MEF 26.2 [10]
- [R7] When an ingress Bandwidth Profile Flow is present for a CoS Name Label at an EI, the value of the EBS parameter for that flow **MUST** be either equal to zero or greater than or equal to:
- For an EVC, the EVC Maximum Service Frame Size as defined in MEF10.3 [1]
 - For an OVC, the lower bound specified in Table 47 of MEF26.2 [10]

[R6] [R9] and [R7] [R10] mean that setting CBS or EBS to a value greater than zero is necessary and sufficient to ensure that the BWP is capable of declaring ingress frames of any allowable size to be green or yellow, respectively. This is not meaningful in a practical sense, however, unless the bandwidth profile is also configured to allow tokens to be replenished as they are consumed.

- [R8] When an ingress Bandwidth Profile Flow is present for a CoS Name Label at an EI, and when the value of CBS is greater than zero, the other BWP parameters for that flow **MUST** be configured in a way that allows tokens to be added to the committed token bucket over time.
- [R9] When an ingress Bandwidth Profile Flow is present for a CoS Name Label at an EI, and when the value of EBS is greater than zero, the other BWP parameters for that flow **MUST** be configured in a way that allows tokens to be added to the excess token bucket over time.

The configurations for a Bandwidth Profile Flow that allow tokens to be added to a committed token bucket include:

- 1) $CIR_{max}^i > 0$ and $CIR^i > 0$, or
- 2) $CIR_{max}^i > 0$ and the Bandwidth Profile Flow is in an envelope where the next higher rank flow has $CF^i = 0$ and has a configuration that allows tokens to be added to its committed token bucket over time.

The bandwidth profile configurations that allow tokens to be added to an excess token bucket include:

- 1) $EIR_{max}^i > 0$ and $EIR^i > 0$, or

- 2) $EIR_{max}^i > 0$ and $CF^i = 1$ and the configuration allows tokens to be added to the committed bucket over time, or
- 3) $EIR_{max}^i > 0$ and the Bandwidth Profile Flow is in an envelope where the next higher rank flow has a configuration that allows tokens to be added to its excess token bucket over time.

CoS Label H is typically considered a “green-only” service, however it is allowed to be “green-yellow”. CoS Label M is typically considered a “green-yellow” service, however it is allowed to be “green-only”. CoS Label L is typically considered a “green-yellow” or “yellow-only” service, however it is allowed to be “green-only”. This is formalized in the following requirements:

- [R10] When an ingress Bandwidth Profile Flow is present for CoS Label H, it **MUST** have $CBS > 0$.
- [R11] When an ingress Bandwidth Profile Flow is present for CoS Label M, it **MUST** have $CBS > 0$.
- [R12] When an ingress Bandwidth Profile Flow is present for CoS Label L, it **MUST** have $CBS + EBS > 0$.

When a bandwidth profile specifies $CBS = 0$ for CoS Label L, the CoS Frame Set associated with the bandwidth profile will contain only yellow frames. Since there are no green frames in such a CoS Frame Set the Performance Objectives for CoS Label L in section 9.2 would not apply to any frames in the CoS Frame Set.

5 New Section in MEF 23.2

The following section is added as a new Section 10 in MEF 23.2 and the existing Section 10 is renumbered as Section 11.

10 Models for Bandwidth Profiles with Token Sharing

The term *token sharing* is used to define a Bandwidth Profile model that allows for bandwidth sharing across two or more Bandwidth Profile Flows in an Envelope, where unused tokens from one Bandwidth Profile Flow can be made available to one or more other Bandwidth Profile Flows in the Envelope.

Bandwidth Profiles without token sharing apply independently to each Bandwidth Profile Flow (e.g., CoS Name) associated with an EVC or OVC, whereas a token sharing Bandwidth Profile can be used to share bandwidth between Bandwidth Profile Flows.

Figure A-1 below provides a simple, high-level comparison of the two models in the case of a Bandwidth Profile Flow per CoS Name.

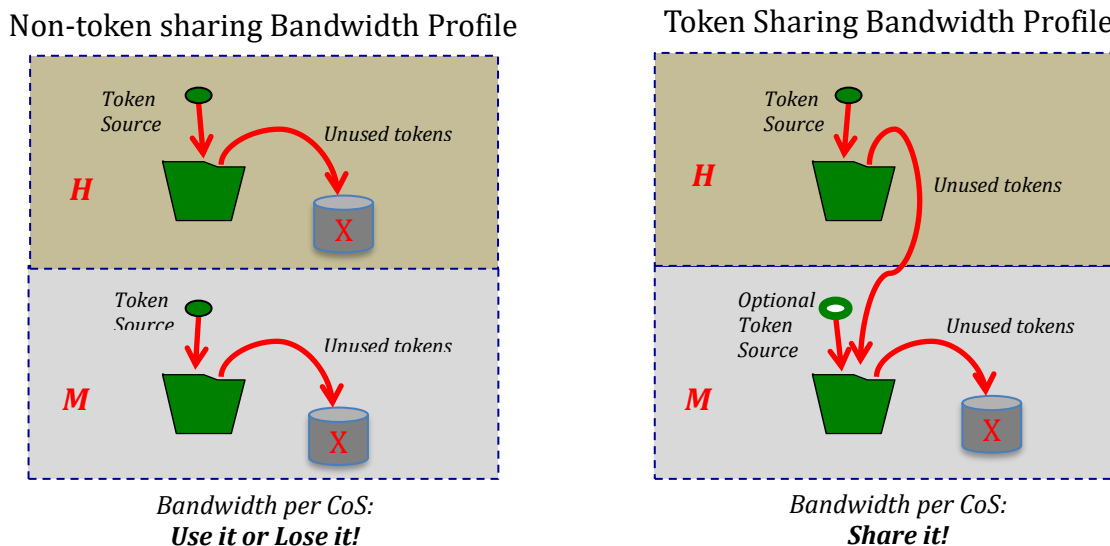


Figure A-1: Bandwidth Profile Models - with and without token sharing in an Envelope

In the non-token sharing Bandwidth Profile model, bandwidth is consumed separately at each Bandwidth Profile Flow (CoS Name). The unused bandwidth in the H CoS is lost, i.e., cannot be used by the M CoS. In the token sharing Bandwidth Profile model, any unused bandwidth in the H CoS is not necessarily lost; it can be used by traffic in the M CoS.

The Bandwidth Profile algorithm specified in MEF 10.3 [1] and MEF 26.2 [10] is very flexible and allows for many deployment models, i.e., there are lots of knobs, and therefore lots of possibilities (see Figure A-3). This document defines specific models for Bandwidth Profiles with token sharing that are each constrained to provide specific behaviors in the context of MEF defined services. The purpose of the defined models is to identify and give names to useful practical combinations of parameter values.

In practice, services without token sharing Bandwidth Profiles use one of the following approaches: a) all traffic is mapped to a single CoS Name; or b) traffic is divided into two, three or four CoS Names, each having a fixed bandwidth allocation. In scenario a), both delay and delay-variation sensitive traffic compete with bursty data traffic for a given bandwidth allocation -- all traffic is treated equally. The Operator or Service Provider will need to over-provision the CEN, resulting in higher cost for the network (and, presumably, for the Subscriber, as well). In scenario b), because of variable traffic mix, each CoS Name needs to support the peak bandwidth required for that CoS Name. This leads to provisioning more bandwidth than is needed for a given service, when the peak bandwidth requirements for each CoS Name are not coincident.

A Bandwidth Profile with token sharing enabled can be used to alleviate the issues identified above by dynamically allocating unused bandwidth among CoS Names. This functionality resides in the CEN, and hence does not require changes to the Subscriber's equipment.

This dynamic bandwidth allocation method is illustrated in Figure A-2 below:

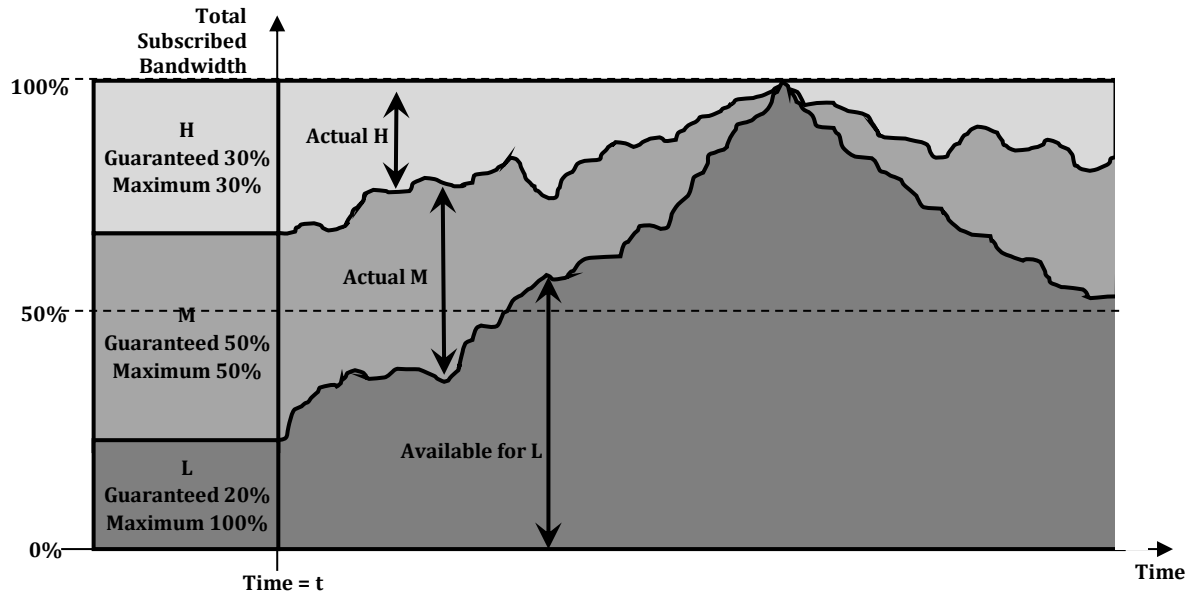


Figure A-2: Concept of Dynamic Bandwidth Management

As the traffic mix changes dynamically, more bandwidth can be made available for L, allowing for full utilization of the provisioned bandwidth. Such a bandwidth-sharing model can be used in many applications, including mobile backhaul and access to IP-VPN networks, both of which are described in Appendix I.

The Bandwidth Profiles with token sharing defined in this section are based on the general model defined in MEF 41 [1A], and the attributes, parameters and algorithms defined in MEF 10.3 [1] and MEF 26.2 [10]. They are intended to support the services defined in MEF 6.2 [16] and MEF 51 [17].

Token sharing Bandwidth Profiles can consist of many Bandwidth Profile Flows, with options for Green and Yellow bandwidth and with options for how tokens can be shared.

Note that in this document, whenever a Bandwidth Profile Flow parameter is used with the superscript i , e.g., CIR^i , it is shorthand for the Bandwidth Profile Flow parameter for the Bandwidth Profile Flow of rank i , $i = 1, \dots, n$, unless otherwise specified.

Figure A-3 below shows the generic token sharing Bandwidth Profile model, using an example of an Envelope with three Bandwidth Profile Flows. The figure shows all of the possible token sources, which could be used when CIR^i and/or EIR^i ($i=1, \dots, 3$) values are >0 . It also shows all of the potential token flow routes - potential token flow routes are shown using dotted gray lines, and are dependent on the CF^i parameter settings. This IA defines three specific token sharing models, which constrain the token flow in each model.




Figure A-3: Generic Token Sharing Bandwidth Profile Model with three Flows

The basic concept of a Bandwidth Profile is that ingress or egress frames consume tokens taken from a token bucket, where the number of tokens consumed is related to the length of the frame. In general, for a frame belonging to a particular Bandwidth Profile Flow, if a sufficient number of tokens are available in the corresponding Green token bucket, the frame is declared Green; otherwise if there are a sufficient number of tokens in the Yellow token bucket, the frame is declared Yellow; and if neither bucket has a sufficient number of tokens, the frame is declared Red. An exception is that in color-aware mode, frames marked as Yellow cannot consume tokens from the Green token bucket.

The token buckets are filled with tokens from one or more token sources that emit tokens at a predetermined rate. In a non-token-sharing Bandwidth Profile, unused tokens from the Green token bucket for a given Bandwidth Profile Flow can be used (optionally) to fill the Yellow token bucket for that Bandwidth Profile Flow, but otherwise any unused tokens are discarded; they cannot be used to fill the token buckets for any other Bandwidth Profile Flow. In a token-sharing Bandwidth Profile, unused tokens from the token buckets for one Bandwidth Profile Flow can be used to fill token

buckets for other Bandwidth Profile Flows. The effect of this is that if there is unused bandwidth from one Bandwidth Profile Flow, it can be used by frames belonging to another Bandwidth Profile Flow.

The size of each bucket, the rate of the token sources and the way that unused tokens are directed to different buckets are controlled by a number of parameters, as described below. See MEF 10.3 [1] and MEF 26.2 [10] for the complete parameter definitions.

- Multiple Bandwidth Profile Flows within one Envelope can share Green and/or Yellow tokens.
- Envelope and Rank (ER^i) is a Bandwidth Profile Flow parameter that consists of the Envelope ID and the rank, i ($i=1,\dots,n$), where n represents the highest rank flow in the Envelope and 1 represents the lowest rank flow in the Envelope. For example, $\langle XYZ,3 \rangle$ represents a Bandwidth Profile Flow that has a rank of 3 in Envelope XYZ .
- CBS^i and EBS^i are Bandwidth Profile Flow parameters that determine the depth of the Green and Yellow token buckets, respectively, for the flow with rank i ; i.e., the maximum number of tokens that could be available at any given time. When CBS^i is set to 0, the Bandwidth Profile for flow i has no ability to declare a frame Green. When EBS^i is set to 0, the Bandwidth Profile for flow i has no ability to declare a frame Yellow.
- CIR^i and EIR^i are Bandwidth Profile Flow parameters that (when set >0) are token sources, at rank i , for the Green and Yellow token buckets, respectively.
- The Coupling Flag (CF^i), $i=1,\dots,n$, is a Bandwidth Profile Flow parameter that is used to control the flow of unused Green tokens at rank i , as follows: $CF^i = 0$ results in unused Green tokens at rank i flowing downwards, to the Green token bucket of rank $i-1$; and $CF^i = 1$ results in unused Green tokens being converted to Yellow tokens that flow across to the Yellow token bucket of rank i .
- CIR^i_{max} and EIR^i_{max} are Bandwidth Profile Flow parameters at rank i that limit token flow into the Green and Yellow token buckets, respectively. A value of 0 means that no new tokens are added to the token bucket. An arbitrarily large value means that all unused tokens from higher ranks plus all tokens sourced at this rank can fill the token bucket.
- Unused Yellow tokens always flow downward; from higher to lower rank flows. Unused Yellow tokens from the lowest rank Bandwidth Profile Flow are discarded.
- Coupling Flag (CF^0), is an Envelope parameter that is used to control the flow of unused Green tokens at the bottom rank: $CF^0 = 0$ results in unused Green tokens being discarded; and $CF^0 = 1$ results in unused Green tokens being converted to Yellow tokens that flow up to the highest rank.

Note that the Token Request Offset (F^i)², not shown in Figure A-3, is a Bandwidth Profile Flow parameter that adjusts the number of tokens requested (consumed) for each EI frame for flow i . For example, at an ENNI, $F^i = 4$ might be used for an Ingress Bandwidth Profile Flow to result in 4 fewer tokens requested for each ENNI frame. See MEF 26.2 [10], Appendix G for examples of using the token request offset.

Note that each token sharing Bandwidth Profile model can be used for either an Ingress Bandwidth Profile or an Egress Bandwidth Profile, and is intended to be applied to a service at an External Interface that is either a UNI or an ENNI. Such service may be either an EVC service, as defined in MEF 6.2 [16] or an OVC service, as defined in MEF 51 [17]. The application of the Bandwidth Profile is dependent on the perspective of the Operator applying it, e.g., an Operator providing an OVC service may not be aware of the EVC service(s) supported by the OVC service.

In general, the token sharing Bandwidth Profile models defined in this section are intended to be used in conjunction with the common shapers that are available in a wide-range of equipment that connects to the CEN at a UNI or an ENNI. In the most common model, the Customer Edge (CE) shapes traffic per EVC (typically a VLAN), and schedules traffic from multiple queues³ within that EVC. Conceptually, in the CEN, multiple Bandwidth Profile Flows would map to a single Envelope (represented by an EVC or OVC), such that each queue in the CE would correlate with a Bandwidth Profile flow in the Envelope.

Figure A-4 below depicts an example of the basic components of such a device, and its relationship with the CEN, for a given EVC-based service.




Figure A-4: Example of a Customer Equipment Scheduler and Shaper connected to a MEF-defined service supporting token sharing Bandwidth Profiles

In the above example, the customer's equipment shapes the aggregate traffic associated with the EVC. The customer equipment scheduler and shaper are servicing three queues, each one dedicated to a different CoS Label carrying different applications. In this example,

² The F^i attribute is defined in MEF 26.2 for OVC End Points. There is no similar attribute in MEF 10.3 for EVC per UNI. As such, F^i is appropriate to use in the context of OVC services only, i.e., in the Service Provider/Operator context per MEF 26.2.

³ It is expected that applications would be mapped to the different queues, using a PCP mapping mechanism.

- CoS Label *H* is carrying traffic that is predictable and has relatively low bandwidth requirements (e.g., expedited traffic that is used for VoIP and control applications).
- CoS Label *M* is carrying traffic that is less predictable and has higher bandwidth requirements (e.g., assured traffic that is used for prioritized data and/or video applications).
- CoS Label *L* is carrying traffic that is unpredictable and has high bandwidth requirements at times (e.g., default traffic that is used for base Internet applications).

In this example, the *H* queue is serviced with strict priority (SP) over the other queues. There is normally no need for a rate limiter for this traffic, since the applications that map to this queue are typically well behaved, and the peak bandwidth requirements well understood. The *M* and *L* queues share the remaining bandwidth, and are serviced with a weighted round robin (WRR) scheduler, in this example giving four times as much weight for *M* than *L*.

Note that configuring the shaper in a way that closely matches the Ingress Bandwidth Profile parameters that have been agreed reduces the likelihood of frames being declared Red by an Ingress Bandwidth Profile, and hence discarded. Appendix G provides detailed analysis of burst sizes and shaper considerations.

Note that the above example is in the context of a Subscriber attaching Customer Equipment to a UNI. The same behaviors and benefits apply to Operator Equipment attaching to an ENNI.

Table A-2 compares characteristics of three different Ingress Bandwidth Profile service models for services requiring multiple classes of service: a) Ingress Bandwidth Profile per EVC/OVC; b) Ingress Bandwidth Profile per Class of Service Name without token sharing; and c) Ingress Bandwidth Profile per Class of Service Name with token sharing. A similar comparison can be made for Egress Bandwidth Profiles.

Bandwidth Profile Behavior	Bandwidth Profile Service Model		
	per EVC/OVC	per Class of Service Name, without token sharing	per Class of Service Name, with token sharing
During periods without higher rank class traffic, the bandwidth can be used by a lower rank class.	✓		✓
Each Class of Service Name can have its own Bandwidth Profile		✓	✓
Allows burst size to be limited per Class of Service Name		✓	✓
Traffic in a lower rank Class of Service Name cannot affect the amount of traffic declared Green in a higher rank Class of Service Name		✓	✓
Aligns with shapers that limit total rate while using scheduling algorithms to allocate the total rate among the Class of Service Names	✓		✓

Table A-2: Comparison of Bandwidth Profile Service Models

In summary, a Bandwidth Profile with token sharing can be used to share bandwidth among the Class of Service Names while ensuring that traffic in a lower rank class cannot cause discards in a higher rank class. In addition, simplified shapers used by the customer edge that limit the total rate, and schedule multiple flows within the total rate, can be supported effectively, because the shaped traffic will not be declared Red by the Bandwidth Profile algorithm.

10.1 Common Requirements for Token Sharing Bandwidth Profile Models

The following common requirements apply to each of the token sharing Bandwidth Profile models defined in this document, which are intended to support MEF 6.2 [16] and MEF 51 [17] services.

[R1A] For a token sharing Bandwidth Profile model, at an EI, the *Token Share* service attribute **MUST** be *Enabled*.

[R2A] For a token sharing Bandwidth Profile model, for an EVC at a UNI, an Ingress Bandwidth Profile Flow **MUST** be based on [R134] criterion 3 of MEF 10.3 [1].

Criterion 3 of MEF 10.3 [R134] essentially means that each Bandwidth Profile Flow in the Ingress Bandwidth Profile consists of ingress Service Frames with a given Class of Service Name for a given EVC at the UNI.

[R3A] For a token sharing Bandwidth Profile model, for an OVC End Point, an Ingress Bandwidth Profile Flow **MUST** be based on [R230] criterion 2 of MEF 26.2 [10].

Criterion 2 essentially means that each Bandwidth Profile Flow in the Ingress Bandwidth Profile consists of ingress EI Frames with a Class of Service Name for a given OVC End Point.

Note that due to requirements [R2A] and [R3A], each Bandwidth Profile Flow can be associated with only one CoS Name or CoS Label.

- [R4A]** For a token sharing Bandwidth Profile model with an Ingress Bandwidth Profile, the Bandwidth Profile Flows from a single EVC / OVC End Point that are based upon CoS Labels⁴ and are mapped to the Envelope **MUST** be ranked from highest to lowest as follows: {H+, H, M, L}.

The above requirement means that an Envelope can consist of Bandwidth Profile Flows from one or more EVCs / OVC End Points, and those Bandwidth Profile Flows can be based upon CoS Labels or other CoS Names. For Bandwidth Profile Flows that are based upon CoS Labels, the relative ranking for those Bandwidth Profile Flows for a given EVC or OVC End Point is what is important. For example, an Envelope can consist of five Bandwidth Profile Flows for a single EVC that are based upon CoS Names: H, Gold, M, Silver and L. In this case H needs to be higher rank than M, which is higher rank than L. Gold and Silver can use any rank.

The following common requirements relate to an Egress Bandwidth Profile Flow.

- [R5A]** For a token sharing Bandwidth Profile model, for an EVC at a UNI, an Egress Bandwidth Profile Flow **MUST** be based on [R134] criterion 6 of MEF 10.3 [1].

Requirement [R5A] means that an Egress Bandwidth Profile Flow consists of egress Service Frames with a given Egress Equivalence Class Name for the EVC at the UNI.

- [R6A]** For a token sharing Bandwidth Profile model, for an OVC End Point, an Egress Bandwidth Profile Flow **MUST** be based on [R230] criterion 5 of MEF 26.2 [10].

Requirement [R6A] means that an Egress Bandwidth Profile Flow consists of egress EI frames with a given Egress Equivalence Class Name for the OVC End Point at either a UNI or an ENNI.

The following requirements are necessary to ensure that there are at least some tokens available to each Bandwidth Profile Flow. When the term *MFS* is used in this document, in the context of an EVC, it refers to the value of the EVC Maximum Service Frame Size Service Attribute per MEF 10.3; and in the context of an OVC, it refers to the value of the OVC Maximum Frame Size Service Attribute, per MEF 26.2.

- [R7A]** For each Bandwidth Profile Flow i , $i = 1, \dots, n$, in an Envelope, at least one of the following inequalities **MUST** hold:

- $CBS^i \geq MFS$ or
- $EBS^i \geq MFS$.

⁴ Three CoS Labels (H, M and L) are defined in this document and one CoS Label (H+) is defined in MEF 22.2 [2A].

Since R6 and R7 require that CBS^i and EBS^i are either 0 or greater than or equal to the MFS, the effect of [R7A] is to prohibit a Bandwidth Profile Flow in which both CBS^i and EBS^i are 0. Such a flow would not be able to declare any frames as Green or Yellow, so all frames belonging to the flow would be declared Red.

[R8A] For each Bandwidth Profile Flow $i, i = 1, \dots, n$, in an Envelope, if $CBS^i > 0$, then CIR_{max}^i **MUST** be greater than 0.

[R9A] For each Bandwidth Profile Flow $i, i = 1, \dots, n$, in an Envelope, if $EBS^i > 0$, then EIR_{max}^i **MUST** be greater than 0.

[R8A] and [R9A] ensure that if a bucket is not 0-sized, then at least some tokens are allowed to flow into it. If this were not the case, the bucket would never receive any tokens and, after the initial number of tokens in the token bucket is consumed, no frames would be able to be declared Green or Yellow (as appropriate) by consuming tokens from it.

[R10A] For each Bandwidth Profile Flow $j, j = 1, \dots, n - 1$, in an Envelope, if $CBS^j \geq MFS$, then at least one of the following **MUST** hold:

- $CIR^j > 0$ or
- $CF^{j+1} = 0$.

[R11A] For two Bandwidth Profile Flows $i, i = 2, \dots, n$, and $j, j = 1, \dots, i - 1$, in the same Envelope, if $CBS^i = 0$, CBS^j **MUST** be 0.

[R11A] means that if there is a Yellow-only Bandwidth Profile Flow (i.e. a flow where there is no possibility of declaring frames to be Green), then all lower-ranked flows are also Yellow-only. [R10A] and [R11A] ensure that there is a flow of tokens that can reach the Green bucket for a given Bandwidth Profile Flow, unless they are all consumed by higher ranked flows: either there is a Green token source for this Bandwidth Profile Flow ($CIR^j > 0$), or tokens overflowing from the next-higher-ranked flow can flow down into this bucket ($CF^{j+1} = 0$). [R13A] ensures there will be some Green tokens at the highest rank.

[R12A] For each Bandwidth Profile Flow $j, j = 1, \dots, n - 1$, in an Envelope, if $EBS^j \geq MFS$, then at least one of the following **MUST** hold:

- $CF^k = 1$ for at least one $k \in \{j, \dots, n\}$, or
- $CF^0 = 1$, or
- $EIR^k > 0$ for at least one $k \in \{j, \dots, n\}$.

[R12A] ensures that there is a flow of tokens that can reach the Yellow bucket for a given flow. As unused Yellow tokens always flow down to the next lower ranked Yellow bucket, it is sufficient to ensure that there is a flow of tokens that can reach this or some higher rank; i.e. either there is a

Yellow token source for this or a higher-ranked flow ($EIR^k > 0$, for $k \geq j$); or unused Green tokens from the lowest rank are recirculated around to become Yellow tokens for the highest rank ($CF^0 = 1$); or unused Green tokens at this or some higher rank are used to fill the Yellow bucket at that rank ($CF^k = 1$ for some $k \geq j$).

[R13A] For the Bandwidth Profile Flow with the highest rank, n , in an Envelope, all of the following inequalities **MUST** hold:

- $CBS^n \geq MFS$, and
- $CIR^n \geq CIR_{max}^n > 0$.

[R13A] ensures that the highest rank is capable of declaring frames to be Green. This requires that there is a Green token source at that rank ($CIR^n > 0$), that tokens are allowed to enter the Green bucket ($CIR_{max}^n > 0$) and that the bucket is not 0-sized ($CBS^n \geq MFS$).

10.2 Classification of Token Sharing Bandwidth Profile Models

The following subsections define some specific token sharing Bandwidth Profile models. In each model, the number of Bandwidth Profile Flows is generalized from 1 to n , where n is the highest rank flow. The models defined in this document are intended for use when there are two or more Bandwidth Profile Flows in an Envelope. The classification is based on terms that are described below:

Bandwidth Type describes the possible declarations by the Bandwidth Profile algorithm for the set of flows in the Envelope.

- *C* (for committed) means that the Bandwidth Profile algorithm can only declare a frame as Green or Red, i.e. $CBS^i \geq MFS$ and $EBS^i = 0$ for all Bandwidth Profile Flows.
- *CX* (for committed and excess) means that the Bandwidth Profile algorithm can declare a frame as Green, Yellow or Red, i.e. $CBS^i \geq MFS$ for at least some Bandwidth Profile Flows and $EBS^i \geq MFS$ for at least some Bandwidth Profile Flows. Note that in an Envelope, a single Bandwidth Profile Flow could have frames declared Green, Yellow, or Red, or one Bandwidth Profile Flow could have frames declared only Green or Red and another Bandwidth Profile Flow could have frames declared only Yellow or Red.
- *X* (for excess) means that the Bandwidth Profile algorithm can only declare a frame as Yellow or Red, i.e. $CBS^i = 0$ and $EBS^i \geq MFS$ for all Bandwidth Profile Flows.

Token Source describes the color of the token sources.

- *G* means that there are *Green* token sources only, i.e., $CIR^n > 0$, and CIR^i may be > 0 for other Bandwidth Profile Flows. In addition, $EIR^i = 0$ for all Bandwidth Profile Flows.

- *GY* means that there are *Green* and *Yellow* token sources, i.e., $CIR^n > 0$, and CIR^i may be > 0 for other Bandwidth Profile Flows. In addition, $EIR^i > 0$ for at least one of the Bandwidth Profile Flows.
- *Y* means that there are *Yellow* token sources only, i.e., $EIR^n > 0$, and EIR^i may be > 0 for other Bandwidth Profile Flows. In addition, $CIR^i = 0$ for all Bandwidth Profile Flows. Note that models using Yellow-only token sources do not conform with [R13A], but are included for completeness.

Token Flow describes the flow of tokens in the system.

- *D* means the token flow is *down*, which means that $CF^0 = 0$ for the Envelope and that for Bandwidth Profile Flow $i, i = 1, \dots, n, CF^i = 0$ for each of the Bandwidth Profile Flows - all tokens flow down, and there is no conversion of Green tokens to Yellow tokens.
- *A* means that unused Green tokens flow *across* to the Yellow token bucket at one or more of the ranks, i.e., for Bandwidth Profile Flow $i, i = 1, \dots, n, CF^i = 1$ for at least one of the Bandwidth Profile Flows (which requires $CF^0 = 0$)⁵. Other Bandwidth Profile Flows may have CF^i set to 0.
- *R* means that unused Green tokens at the bottom rank are converted to Yellow tokens and *recirculate* - feeding those Yellow tokens to the highest rank Bandwidth Profile Flow, i.e., $CF^0 = 1$ for the Envelope (which requires for Bandwidth Profile Flow $i, i = 1, \dots, n$, that $CF^i = 0$ for each of the Bandwidth Profile Flows).

Table A-3 below lists specific token sharing Bandwidth Profile models.

Bandwidth Type	Token Source	Token Flow	Model Name
C	G	D	C/G/D
CX	G	R	CX/G/R
CX	G	A	CX/G/A
CX	GY	R	CX/GY/R
CX	GY	A	CX/GY/A
CX	GY	D	CX/GY/D
X	Y	D	X/Y/D

Table A-3: Token Sharing Bandwidth Profile Models

Table note: Given the three-classifier model, only the models with meaningful classifier combinations are shown in the above table. The possible models that are not shown are because of conflict between two (or more) of the classifiers. One example of such conflict is when *Bandwidth Type* is *C* and Token Flow *A* or *R* is used. A second example is when *Bandwidth Type* is *C* and Token Source of *GY* is used. In these examples, any Yellow tokens would be wasted, since EBS for each of the Bandwidth Profile Flows is set to 0. Another example is when *Bandwidth Type* is *CX*,

⁵ See MEF 10.3 [R150] and MEF 26.2 [R247].

Token Source is G and Token Flow is D . In this case, there is no possibility of feeding Yellow tokens to any of the Bandwidth Profile Flows.

Of the seven models shown in Table A-3, three specific token sharing Bandwidth Profile models are normatively defined in this document: $C/G/D$, $CX/G/R$ and $CX/GY/R$. These are described in the following subsections. Note that model $X/Y/D$, dark gray shaded row above, is shown for completeness in Table A-3, but cannot be normatively defined, as it does not conform with [R13A]. The remaining three models ($CX/G/A$, $CX/GY/A$ and $CX/GY/D$), gray shaded rows above, are for further study.

10.3 Model $C/G/D$

This is the simplest model in concept, since EBS is 0 for each Bandwidth Profile Flow. Green tokens are shared between flows and unused Green tokens at the bottom rank are discarded. One or more Green token sources ($CIR > 0$) are used, and the token flow is down.

Figure A-5 below depicts the token flow for Model $C/G/D$, using an example of an Envelope with three Bandwidth Profile Flows.

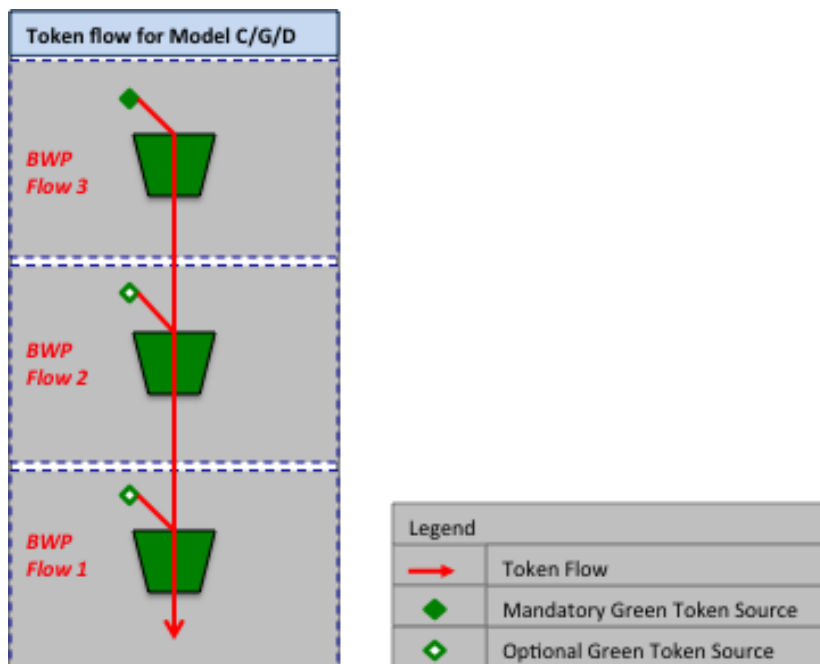


Figure A-5: Token Flow for Model $C/G/D$

A frame can be declared Green or Red by the Bandwidth Profile algorithm.

[R14A] For Model $C/G/D$, CF^0 for the Envelope **MUST** be 0.

Note that Model $C/G/D$ requires $CF^0 = 0$, although it doesn't really matter to the behavior which value is used, since no frames can be declared Yellow by the Bandwidth Profile algorithm.

[R15A] For Model C/G/D, the Bandwidth Profile Flow parameter values and constraints as specified in Table A-4 below **MUST** apply for each bandwidth profile flow i , $i = 1, \dots, n$.

Bandwidth Profile Flow Parameter, for $i=1, \dots, n$	Requirements for Model C/G/D
CM^i	No additional constraints from MEF 10.3 or MEF 26.2
CF^i	$CF^i = 0$ for all flows
CIR^i	No additional requirements beyond those specified in Section 10.1. Note that [R13A] requires that $CIR^n > 0$.
CBS^i	$CBS^i \geq MFS$ for all flows. See Appendix H for guidance on selection of CBS^i values
CIR^i_{max}	No additional requirements beyond those specified in Section 10.1. Note that [R8A] requires $CIR^i_{max} > 0$ for all flows in this model, since $CBS^i > MFS$ for all flows.
EIR^i	$EIR^i = 0$ for all flows
EBS^i	$EBS^i = 0$ for all flows.
EIR^i_{max}	$EIR^i_{max} = 0$ for all flows
ER^i	Requirement per [R4A]
F^i	No additional constraints from MEF 26.2

Table A-4: Bandwidth Profile Flow parameter requirements for Model C/G/D

The highest rank Bandwidth Profile Flow has a Green token source ($CIR^n > 0$) and each of the other flows may or may not have a Green token source.

In the C/G/D model, unused Green tokens from one or more higher ranked Bandwidth Profile Flows can be used by lower ranked flow i when $CIR^i_{max} > CIR^i$. Consequently, this model has token sharing. If $CIR^i_{max} > \sum CIR^j$ for $j=i, \dots, n$, then flow i could use all of the unused Green tokens from higher ranks plus the Green tokens sourced at rank i when the only traffic is in flow i .

This simple model may be useful for general application. One specific application - access to IP-VPN services - is described in Appendix I.1.

10.4 Model CX/G/R

In this model, the *Bandwidth Type* is CX, i.e., *committed and excess*. Only *Green* tokens are sourced. These are shared between the flows, with unused Green tokens at the bottom rank converted to Yellow and fed to the upper rank's Yellow token bucket. One or more Green token sources are used (i.e., $CIR^i > 0$ for at least one rank i , $i = 1, \dots, n$), and the token flow is *recirculate*.

Figure A-6 below depicts the token flow for Model CX/G/R, using an example of an Envelope with three Bandwidth Profile Flows.

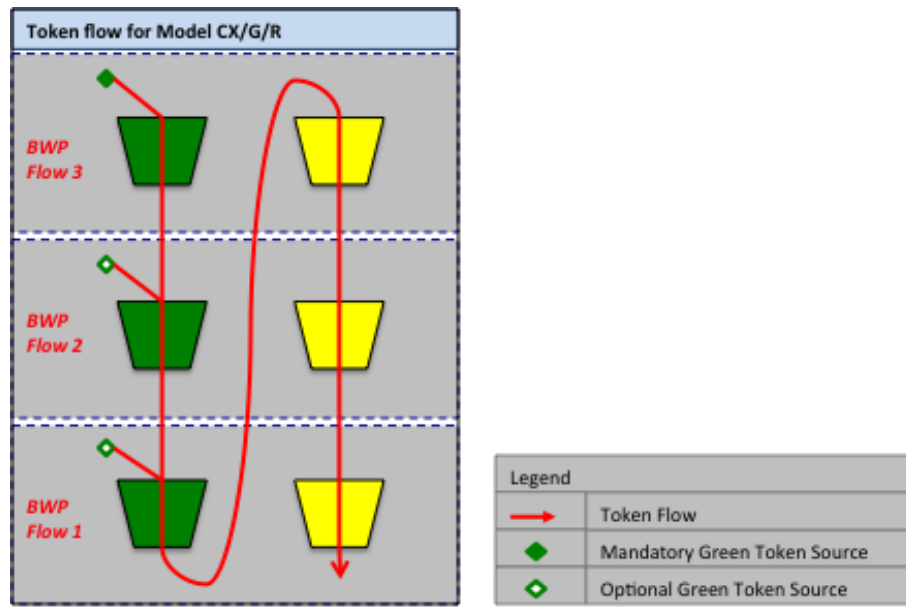


Figure A-6: Token Flow for Model CX/G/R

The key difference between this model and Model C/G/D is that Green tokens that are lost in Model C/G/D are converted to Yellow tokens in the CX/G/R model. This means that unused bandwidth from a low-priority Bandwidth Profile Flow can be used by a higher-priority one that is transmitting at or above its CIR_{max} , while still guaranteeing that some bandwidth is available for the low-priority flow when both flows are transmitting at or above CIR_{max} . Traffic in the higher-priority flow that consumes the recirculated tokens (i.e., traffic above CIR_{max}) is declared Yellow, so in the case of an Ingress Bandwidth Profile, the SLS does not apply to it.

A frame can be declared Green, Yellow or Red by the Bandwidth Profile algorithm.

[R16A] For Model CX/G/R, CF^0 for the Envelope **MUST** be 1.

[R17A] For Model CX/G/R, the Bandwidth Profile Flow parameter values and constraints as specified in Table A-5 below **MUST** apply for each bandwidth profile flow i , $i = 1, \dots, n$.

Bandwidth Profile Flow Parameter, for $i=1,\dots,n$	Requirements for Model CX/G/R
CM^i	No additional constraints from MEF 10.3 or MEF 26.2
CF^i	$CF^i = 0$ for all flows.
CIR^i	No additional requirements beyond those specified in Section 10.1. Note that [R13A] requires that $CIR^n > 0$.
CBS^i	No additional requirements beyond those specified in Section 10.1. Note that [R13A] requires that $CBS^n \geq MFS$. See Appendix H for guidance on selection of CBS^i values.
CIR^i_{max}	No additional requirements beyond those specified in Section 10.1. Note that [R13A] requires that $CIR^n_{max} > 0$ and that [R8A] requires $CIR^i_{max} > 0$ for flows with $CBS^i > MFS$.
EIR^i	$EIR^i = 0$ for all flows.
EBS^i	$EBS^i \geq MFS$ for at least one of the flows. See Appendix H for guidance on selection of EBS values.
EIR^i_{max}	No additional requirements beyond those specified in Section 10.1. Note that [R9A] requires $EIR^i_{max} > 0$ for flows with $EBS^i > MFS$.
ER^i	Requirement per [R4A]
F^i	No additional constraints from MEF 26.2

Table A-5: Bandwidth Profile Flow parameter requirements for Model CX/G/R

The highest ranked Bandwidth Profile Flow has a Green token source and each of the other flows may or may not have a Green token source.

In the CX/G/R model, unused Green tokens from one or more higher ranked Bandwidth Profile Flows can be used by lower ranked flow i when $CIR^i_{max} > CIR^i$. Consequently, this model has token sharing. Note that if $CIR^i_{max} \geq \sum CIR^j$ for $j=i,\dots,n$, then flow i could use all of the unused Green tokens from higher ranks plus the Green tokens sourced at rank i when the only traffic is in flow i . In addition, if $EIR^i_{max} \geq \sum CIR^k$, $k=1,\dots,i-1$, then flow i could use all of the Yellow tokens converted from unused Green tokens at ranks lower than i .

The CX/G/R model may be useful for general application. One specific application - an Egress Bandwidth Profile for two OVC End Points at a UNI, providing access to an IP-VPN service - is described in Appendix I.2.

10.5 Model CX/GY/R

This model supports flows that have both Green and Yellow token sources.

Figure A-7 below depicts the token flow for Model CX/GY/R, using an example of an Envelope with three Bandwidth Profile Flows.




Figure A-7: Token Flow for Model CX/GY/R

When applied to an Ingress Bandwidth Profile, the main difference from Model CX/G/R is that even when all flows are transmitting at the maximum rate to which the SLS can apply, additional traffic is permitted for some flows, but the SLS does not apply to this additional traffic.

For Model CX/GY/R, a frame can be declared Green, Yellow or Red by the Bandwidth Profile algorithm.

[R18A] For Model CX/GY/R, CF^0 for the Envelope **MUST** be 1.

Note that this requirement ensures that unused Green tokens at the bottom rank are converted to Yellow and used to provide an additional feed of Yellow tokens. Setting $CF^0 = 1$ also results in $CF = 0$ for each Bandwidth Profile Flow (see [R150] in MEF 10.3 [1] and [R247] in MEF 26.2 [10]).

[R19A] For Model CX/GY/R, the Bandwidth Profile Flow parameter values and constraints as specified in Table A-6 below **MUST** apply for each bandwidth profile flow i , $i = 1, \dots, n$.

Bandwidth Profile Flow Parameter, for $i=1,\dots,n$	Requirements for Model CX/GY/R
CM^i	No additional constraints from MEF 10.3 or MEF 26.2
CF^i	$CF^i = 0$, for all flows
CIR^i	No additional requirements beyond those specified in Section 10.1. Note that [R13A] requires that $CIR^n > 0$
CBS^i	No additional requirements beyond those specified in Section 10.1. Note that [R13A] requires that $CBS^n \geq MFS$. See Appendix H for guidance on selection of CBS values
CIR^i_{max}	No additional requirements beyond those specified in Section 10.1. Note that [R13A] requires that $CIR^n_{max} > 0$ and that [R8A] requires $CIR^i_{max} > 0$ for flows with $CBS^i > MFS$
EIR^i	$EIR^i > 0$ for at least one of the flows
EBS^i	$EBS^i \geq MFS$ for each of the Bandwidth Profile Flows with $EIR^i > 0$. See Appendix H for guidance on selection of EBS^i values.
EIR^i_{max}	No additional requirements beyond those specified in Section 10.1. Note that [R9A] requires $EIR^i_{max} > 0$ for flows with $EBS^i > MFS$.
ER^i	Requirement per [R4A]
F^i	No additional constraints from MEF 26.2

Table A-6: Bandwidth Profile Flow parameter requirements for Model CX/GY/R

The highest ranked Bandwidth Profile Flow has a Green token source and each of the other flows may or may not have a Green token source. In addition, at least one of the flows must have a Yellow token source.

In the CX/GY/R model, unused Green tokens from one or more higher ranked Bandwidth Profile Flows can be used by lower ranked flow i when $CIR^i_{max} > CIR^i$. Consequently, this model has token sharing. Note that if $CIR^i_{max} \geq \sum CIR^j$ for $j=i,\dots,n$, then flow i could use all of the unused Green tokens from higher ranks plus the Green tokens sourced at rank i when the only traffic is in flow i . In addition, if $EIR^i_{max} \geq \sum CIR^j$ ($j = 1,\dots,i-1$) + $\sum EIR^k$ ($k = i,\dots,n$), then flow i could use all of the Yellow tokens converted from unused Green tokens at ranks lower than i plus the Yellow tokens sourced at rank i and above when the only traffic is in flow i .

The CX/GY/R model may be useful for general application. One specific application - mobile backhaul - is described in Appendix I.3.

10.6 Scalability Considerations

The maximum number of ranks that can be supported by an implementation of the token sharing Bandwidth Profile models described above is limited, and it depends on the chosen model, and the line rate.

In a regular (i.e. without token sharing) Bandwidth Profile implementation, each frame is associated with a single Bandwidth Profile Flow, and the token counts for that specific flow must be updated for each processed frame. In a token sharing implementation, a frame is associated with a specific Bandwidth Profile Flow and an Envelope. When a frame is processed, token counts must be updated not only for a specific Bandwidth Profile Flow, but for all flows in the envelope. Moreover, for the token sharing models with token recirculation (such as CX/G/R and CX/GY/R) there may be twice as many token counts to be updated for each processed frame, versus the model without token recirculation.

While implementation specific optimizations are possible, the amount of processing needed per frame increases linearly with the number of ranks in the Envelope, and it is higher for models with token recirculation. At the same time, however, the available processing time per frame is limited, and it depends on the line rate – the higher the line rate, the less time available for per-frame processing.

Hence, when defining services using token sharing Bandwidth Profiles, one should take into account the maximum number of ranks that can be practically implemented, for specific line rates and token sharing models.

6 Changes to Section 10 of MEF 23.2 (References)

The section number for References in MEF 23.2 is changed from 10 to 11. The following references are to be added to the reference list in MEF 23.2.

[1A] MEF Forum MEF 41, *Generic Token Bucket Algorithm*, October 2013.

[2A] MEF Forum MEF 22.2, *Mobile Backhaul Phase 3*, January 2016.

7 New Appendix in MEF 23.2

The following section is added as a new Appendix I in MEF 23.2.

Appendix I Token Sharing Bandwidth Profile Use Cases

This appendix describes use cases for the token sharing Bandwidth Profile models defined in Section 10. In these use cases, depending on parameter values, Bandwidth Profile Flows may get a minimum guaranteed bandwidth, and may be able to get more than the guaranteed minimum bandwidth depending on whether the higher rank flows have unused tokens. There are two cases where a Bandwidth Profile Flow may have unused tokens: a) the ingress traffic rate is lower than the rate of new tokens (causing the bucket to fill); or b) the rate of new tokens exceeds CIR^i_{max} .

In all of the following use cases, the *CBS* and *EBS* values used are aligned with the suggestions in Appendix H, unless there is specific reason for other values. Of course, in practice, Operators may use different *CBS* and/or *EBS* values.

I.1 Use Cases: IP-VPN Access with Token Sharing Model C/G/D

This section describes two use cases for Token Sharing Model C/G/D. Figure A-8 below shows an example of an EVPL service connecting a remote user to an IP-VPN network.

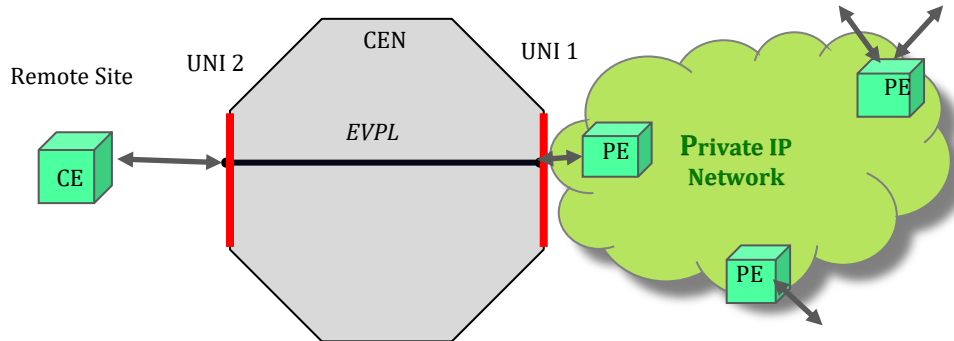


Figure A-8: Example of EVPL service used as access to IP-VPN Network

In this example, there are three Ingress Bandwidth Profile Flows, one for each of the CoS Labels {H,M,L} supported by the EVPL service. These three Bandwidth Profile Flows are contained in a single Envelope, whose Envelope ID = XYZ. The Bandwidth Profile Flows are 'Green only', i.e., the *EBS* value for each is set to 0. Model C/G/D is used. Unused Green tokens flow down, from H to M, and from M to L.

There are two use cases for Model C/G/D described below. These use cases vary in CIR^i and CIR^i_{max} values. The first use case allows a high-rank Bandwidth Profile Flow to consume all of the tokens such that lower ranked flows might not receive any unused tokens and hence become starved for new tokens. The second use case constrains the token use for the highest rank flow, ensuring there are always some Green tokens that flow to the lower ranks.

Use Case 1 - IP-VPN Access with Token Sharing Model C/G/D

This use case involves an Ingress Bandwidth Profile using token sharing for an EVC at UNI_2, per Figure A-8. The Subscriber and Service Provider agree to a CIR^3 value of 100 Mbps (for rank 3), which is the only token source for the Envelope. The specification also includes a CIR^i_{max} value of 100 Mbps for each of the ranks. In this service, if rank 3 needs all of the tokens, it can get them. Any unused tokens from rank 3 are made available to rank 2, which can use all of the available tokens if needed. Any unused tokens from rank 2 are then made available to rank 1, which can use all of the available tokens, if needed.

Figure A-9 shows a simplified token flow diagram for Use Case 1. The red line depicts the general flow of tokens. The solid black lines with arrows show the specific unused token flow route for this use case, determined by the Coupling Flag values and the CIR_{max} values for each rank.

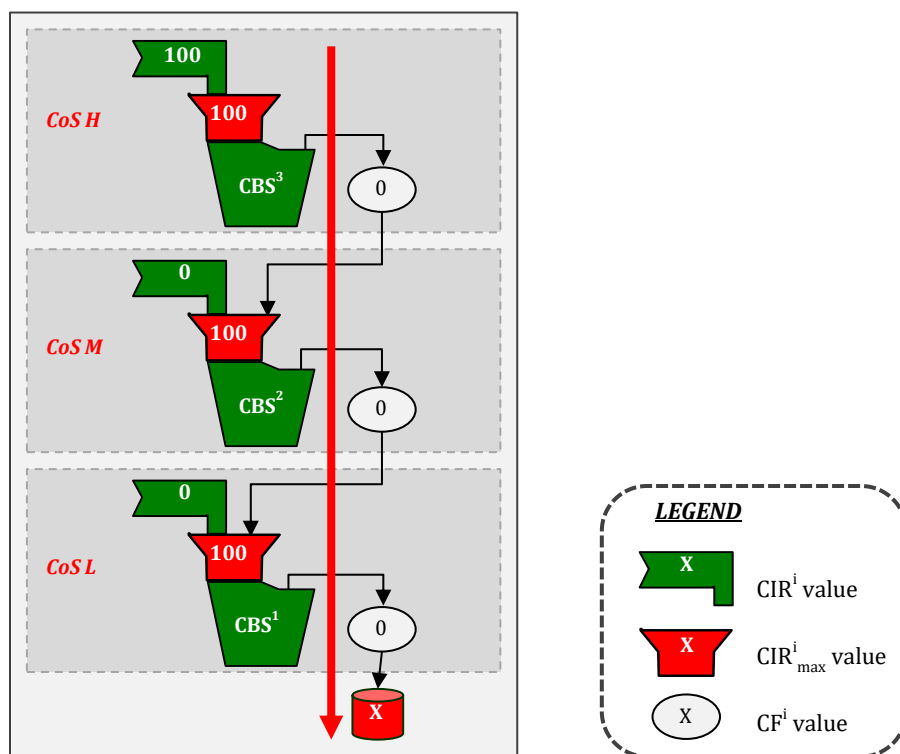


Figure A-9: Token flow diagram for IP-VPN Use Case 1

One question normally asked for this use case is: if all traffic can be sent into the best performing CoS Label, e.g., CoS Label H, what would motivate the Subscriber to distribute the load across multiple CoS Labels? Consider the following scenario: a Service Provider offers a usage-based billing model over an agreed time period with the Subscriber, and one that is sensitive to Class of Service Labels (i.e., usage⁶ in CoS Label H is more expensive per byte than usage in CoS Label M, which is more expensive per byte than usage in CoS Label L). In such a scenario, the Subscriber could be motivated to distribute the traffic across the classes, with the applications really requiring the best performance getting mapped to CoS Label H, and other applications getting mapped to less costly CoS Labels.

Table A-7 below shows the Ingress Bandwidth Profile Parameter values for Envelope XYZ that are agreed by the Service Provider and Subscriber for this service.

⁶ Usage in this context is defined as the number of bytes declared Green by the Bandwidth Profile in time period, T.

Model C/G/D Use Case 1 - Ingress Bandwidth Profile Flow Parameter values			
Envelope Parameters: $CF^0 = 0$ and Envelope ID = XYZ			
Bandwidth Profile Flow Parameters	CoS Label H	CoS Label M	CoS Label L
CM^i {color-aware, color-blind}	color-blind	color-blind	color-blind
CF^i {0,1}	0	0	0
CIR^i (Mbps)	100	0	0
CBS^i (B)	8*MFS	24*MFS	24*MFS
CIR^i_{max} (Mbps)	100	100	100
EIR^i (Mbps)	0	0	0
EBS^i (B)	0	0	0
EIR^i_{max} (Mbps)	0	0	0
ER^i	<XYZ,3>	<XYZ,2>	<XYZ,1>
F^i	0	0	0

Table A-7: Ingress Bandwidth Profile Flow Parameter Values for IP-VPN Use Case 1

There is only a single token source at the top rank, $CIR^3 = 100$ Mbps. Both CIR^2 and CIR^1 are 0. Since this is a 'Green-only' service, there are no Yellow tokens in this model, thus EIR^i , EIR^i_{max} and EBS^i are 0 for all flows. Unused Green tokens at the bottom rank are discarded.

Note that CIR^i_{max} for each Bandwidth Profile Flow is equal to CIR^3 , i.e., CIR^i_{max} for each flow is configured such that each flow can potentially consume all generated tokens. For a CE using a weighted round robin or strict priority queuing scheme, with aggregate traffic shaped at 100 Mbps, the Ingress Bandwidth Profile of this use case will not drop traffic.

Use Case 2 - IP-VPN Access with Token Sharing Model C/G/D with constrained maximum bandwidth for the Bandwidth Profile Flow at rank 3

This use case is different from C/G/D Use Case 1 in that the Bandwidth Profile Flows are more constrained and that two Green token sources are used.

- Rank 3 (CoS Label H) has a CIR^3 value of 20 Mbps and a CIR^3_{max} value of 20 Mbps.
- Rank 2 (CoS Label M) has a CIR^2 value of 80 Mbps and a CIR^2_{max} value of 100 Mbps.
- Rank 1 (CoS Label L) has a CIR^1 value of 0 Mbps and a CIR^1_{max} value of 100 Mbps.

Rank 3 is guaranteed to get up to 20 Mbps. Rank 2 is guaranteed to get at least 80 Mbps, and can get up to the full 100 Mbps when rank 3 is underutilized. Rank 1 is not guaranteed any bandwidth, but can get up to the full 100 Mbps when ranks 2 and 3 are both underutilized.

Table A-8 below shows the Ingress Bandwidth Profile Parameter values that are agreed between the Subscriber and Service Provider for the service.

Model C/G/D Use Case 2 - Ingress Bandwidth Profile Flow Parameter values			
Envelope Parameters: $CF^0 = 0$ and Envelope ID = XYZ			
Bandwidth Profile Flow Parameters	CoS Label H	CoS Label M	CoS Label L
CM^i {color-aware, color-blind}	color-blind	color-blind	color-blind
CF^i {0,1}	0	0	0
CIR^i (Mbps)	20	80	0
CBS^i (B)	8*MFS	24*MFS	24*MFS
CIR^i_{max} (Mbps)	20	100	100
EIR^i (Mbps)	0	0	0
EBS^i (B)	0	0	0
EIR^i_{max} (Mbps)	0	0	0
ER^i	<XYZ,3>	<XYZ,2>	<XYZ,1>
F^i	0	0	0

Table A-8: Ingress Bandwidth Profile Flow Parameter Values for IP-VPN Use Case 2

Similar to the first use case, since this is a 'Green only' service, i.e., there are no Yellow tokens in this model, and thus EIR^i , EIR^i_{max} and EBS^i are 0 for all flows. Unused Green tokens at the bottom rank are discarded.

Subscribers may prefer this multiple token source approach when Service Providers are billing by CIR^i value per CoS⁷.

This scenario is useful for cases where IP level QoS configurations constrain the amount of bandwidth per Bandwidth Profile Flow, e.g., VoIP applications use CoS Label H and can use, at most, a relatively small percentage of the aggregate bandwidth. Applications mapped to CoS Labels M and L can use up to full bandwidth when rank 3 has unused tokens to share.

Note that the Ingress Bandwidth Profile of Use Case 2 will not drop traffic for a CE that: a) allocates a maximum bandwidth of 20 Mbps to CoS Label H (either by configuring a rate limiter/shaper or just knowing that the applications mapped to that class will not exceed the bandwidth) and gives it strict priority; and b) uses a weighted round robin or strict priority scheduler for the remaining CoS Labels; and c) ensures that the shaped rate for the EVC does not exceed 100 Mbps.

I.2 Use Case: Egress Bandwidth Profile with Token Sharing Model CX/G/R

This use case describes a scenario using token sharing Model CX/G/R for an Egress Bandwidth Profile applied to two OVC End Points at a UNI.

Figure A-10 below depicts the example using two Access E-Line services connecting UNI_1 with two different ENNIs (ENNI_A and ENNI_B). There is assumed to be an IP-VPN network behind (to the right of) the ENNIs, and the PEs shown are Provider Edge devices providing access to the IP-VPN service.

⁷ The behavior of this use case would be the same with only a single token source at rank 3, with no change to the values of the other Bandwidth Profile parameters.

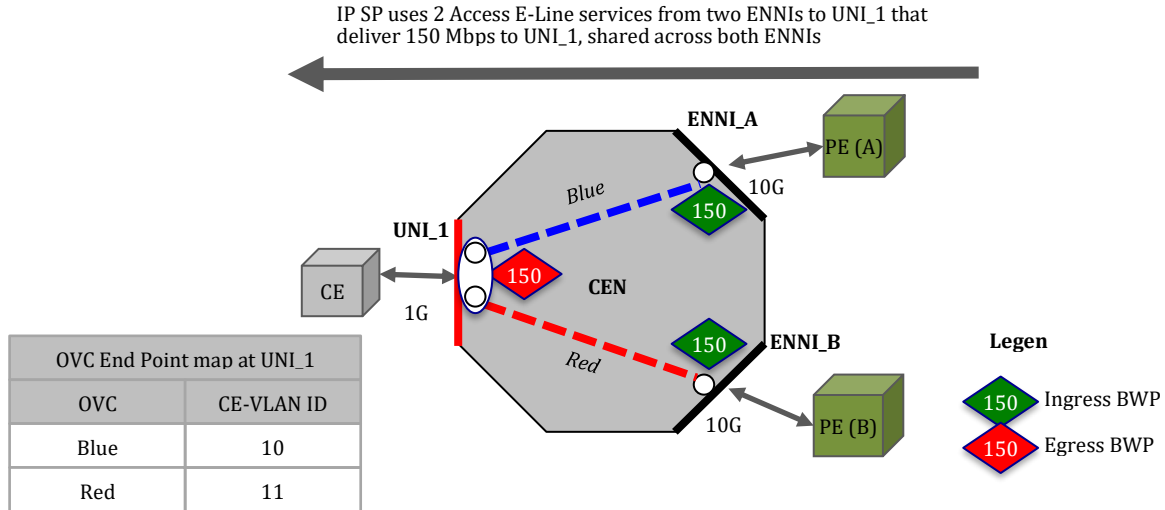


Figure A-10: Egress Bandwidth Profile with token sharing for two OVC End Points

The assumptions for this use case are as follows:

- The IP Service Provider (IP SP) orders access services from the CEN Operator to connect from the ENNIs to the UNI shown in Figure A-10. The IP SP wants to deliver a maximum of 150 Mbps to UNI_1 via the two OVCs: the Blue OVC connects UNI_1 to ENNI_A and the Red OVC connects UNI_1 to ENNI_B. Both OVCs may be carrying traffic concurrently, or, at times, a single OVC may be carrying all the traffic.
- At any given time, the CEN Operator does not know how much traffic will be arriving at ENNI_A or ENNI_B.
- The Ingress Bandwidth Profile for the OVC End Point at each ENNI is configured using the C/G/D token sharing model, with $CIR^3 = 150$ Mbps and CIR_{max}^i for each Bandwidth Profile Flow = 150 Mbps. This allows up to 150 Mbps into the CEN at each ENNI, ensuring that either OVC can carry the full load at any given period of time.
- The IP SP has four IP traffic classes that map to three Carrier Ethernet CoS Labels at the OVC End Points, as follows:
 - The Expedited Forwarding (EF) class maps to CoS Label H.
 - The two Assured Forwarding classes (AF4 and AF2) are both mapped to CoS Label M.
 - The Default Forwarding (DF) class is mapped to CoS Label L.
- Each PE is configured to shape traffic to the ENNI for the given OVC, and scheduling is done within for the traffic classes. To support the IP-VPN service, the SP requires 50 Mbps

to be reserved for CoS Label H traffic at UNI_1, with the remaining available bandwidth allocated to CoS Labels M and L, preferentially weighted 2:3. See Table A-9 below.

IP VPN traffic class	Allocation	CoS Label
EF	50 Mbps	H
AF4	10%	M
AF2	30%	
DF	60%	L

Table A-9: PE shaping of IP traffic classes and mapping to CoS Label

- The SP would like the Egress Bandwidth Profile at the UNI to support the shaped egress bandwidth from a PE (e.g., shaped to VLAN, with strict priority for EF, and the AF and DF classes using WRR scheduler, with a higher weighting for DF).

At UNI_1, the IP SP and Operator agree to six Egress Equivalence Classes. Table A-10 maps the CoS Labels for each OVC to an Egress Equivalence Class.

OVC	CoS Label	Egress Equivalence Class
Blue	H	Blue-H
Red	H	Red-H
Blue	M	Blue-M
Red	M	Red-M
Blue	L	Blue-L
Red	L	Red-L

Table A-10: Mapping of CoS Label per OVC to Egress Equivalence Class

In the remainder of this section, "CoS Label H traffic" refers to egress Service Frames at the UNI that belong to either EEC_Blue-H or EEC_Red-H. The same convention is used for CoS Label M and CoS Label L.

The IP SP and Operator agree on a Bandwidth Profile Flow for each Egress Equivalence Class. The Envelope {XYZ} consists of these six Bandwidth Profile Flows with the flows for H ranked above the flows for M and the flows for M ranked above the flows for L.

Figure A-11 shows the token flow diagram for this use case. The red line depicts the general flow of tokens. The solid black lines with arrows show the specific unused token flow route for this use case, determined by the Coupling Flag values and the CIR_{max} values for each rank.




Figure A-11: Token flow diagram for Egress Bandwidth Profile Use Case

Table A-11 below shows the Egress Bandwidth Profile Parameter values that are agreed between the IP SP and the CEN Operator.

Model CX/G/R Use Case - Egress Bandwidth Profile Flow Parameter values						
Envelope Parameters: $CF^0 = 1$ and Envelope ID = XYZ						
Bandwidth Profile Flow Parameters	Egress Equivalence Class					
	Blue-H	Red-H	Blue-M	Red-M	Blue-L	Red-L
CM^i {color-aware, color-blind}	color-blind	color-blind	color-blind	color-blind	color-blind	color-blind
CF^i {0,1}	0	0	0	0	0	0
CIR^i (Mbps)	50	0	40	0	60	0
CBS^i (B)	8*MFS	8*MFS	24*MFS	24*MFS	24*MFS	24*MFS
CIR^i_{max} (Mbps)	50	50	40	40	60	60
EIR^i (Mbps)	0	0	0	0	0	0
EBS^i (B)	0	0	24*MFS	24*MFS	24*MFS	24*MFS
EIR^i_{max} (Mbps)	0	0	110	110	90	90
ER^i	<XYZ,6>	<XYZ,5>	<XYZ,4>	<XYZ,3>	<XYZ,2>	<XYZ,1>
F^i	0	0	0	0	0	0

Table A-11: Egress Bandwidth Profile Parameter Values

As Bandwidth Profile Flows for the Blue OVC have a higher rank than the Bandwidth Profile Flows for the Red OVC, the Blue OVC could take all of the available bandwidth for a given CoS Label.

With these parameter values, up to 150 Mbps can be delivered to UNI_1 across the two OVCs.

CoS Label H can have up to 50 Mbps in egress traffic at UNI_1, from either the Blue OVC or the Red OVC, or shared across both. Similarly, CoS Label M can have up to 40 Mbps and CoS Label L can have up to 60 Mbps. This addresses the desire to maintain a 2:3 ratio between CoS Labels M and L during times when all flows are sending traffic at their maximum rate.

When there is little or no traffic in a higher ranked Bandwidth Profile Flow, its available bandwidth is re-allocated to the lower ranked flows.

When all traffic is in just one CoS Label, M or L, then the Operator will deliver this traffic up to 150 Mbps due to use of Yellow tokens for these flows; i.e., the sum of the CIR^i_{max} and EIR^i_{max} values add up to 150 Mbps for a given Bandwidth Profile Flow i . This bandwidth can be shared across both OVCs or utilized fully by just one of them.

When there is no traffic in CoS Label H, and traffic is running on both OVCs for CoS Labels M and L, then the Operator can deliver the sum of CoS Label M and L traffic up to the 150 Mbps. In this scenario, the Operator will deliver Blue-M and Red-M traffic up to 40 Mbps each, Blue-L traffic up to 60 Mbps and Red-L traffic up to 10 Mbps. This results in an M:L ratio of 8:7, which is acceptable to the IP SP for this scenario.

I.3 Use Case: Mobile Backhaul with Token Sharing Model CX/GY/R

This appendix describes a Mobile Backhaul use case using Model CX/GY/R for an Ingress Bandwidth Profile for a single EVC at the UNI.

Figure A-12 below shows an example of an EVPL service connecting a cell site UNI to an aggregation site UNI.

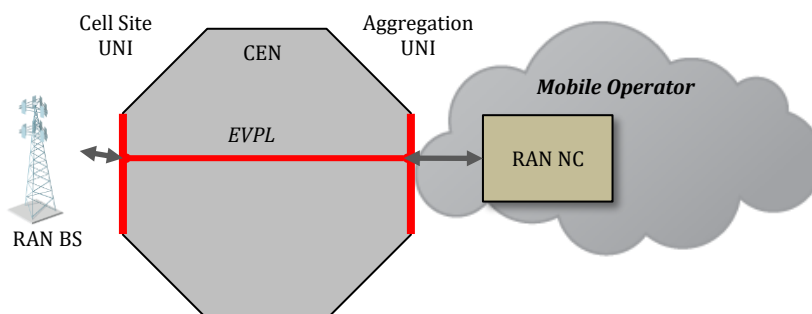


Figure A-12: Simplified example of EVPL service used for Mobile Backhaul

The assumptions for this use case are as follows:

- A 3-sector eNodeB is used at the RAN BS (Radio Access Network Base Station), where each 20 MHz radio (cell per sector) enables data rates up to 150 Mbps in the downlink (DL) and 50 Mbps in the uplink (UL). Note that because of the DL/UL asymmetry, and the symmetry of Ethernet services used for Mobile Backhaul in most deployments, the shaping arrangement at the RAN BS cell site is not as critical as at the RAN NC (Network Controller) site.
- DL peak rate/load per sector = 150 Mbps (PHY layer)
- DL average rate/load per sector - depends on the cell/sector size⁸
- Air interface overhead: ~2 %
- Backhaul Bandwidth overhead - at the UNI, a Service Frame with MAC Client Data that includes encapsulation overhead⁹ of ~25%
- One sector has peak load and the other two sectors have average load each (for this example, we use an average of 45 Mbps). DL total BW for all three sectors = 150 + 45 + 45 = 240 Mbps, of which 235 Mbps is user payload (240 - 2%).

⁸ Average BW per sector may range between 25-75 Mbps (PHY).

⁹ MAC Client Data overhead includes protocols stacked on the customer's data frame to carry the frame between the RAN BS and the RAN NC. These protocols typically include GTP-U, UDP, IPv4, ETH and, optionally, IP Sec. 25% is a calculation based on a distribution of small, medium and large size packets.

Thus, to cover the RAN BS peak load, backhaul bandwidth demand with the overhead = 235 Mbps * 1.25 = 293 Mbps. This can be shared between *CIR* and *EIR* in a variety of ways; it is up to the Mobile Operator to decide and work out with the Service Provider. In this case, it is assumed that the Mobile Operator chooses a value of *CIR* such that under normal load all traffic is committed and hence guaranteed under the SLA. Under peak load more traffic may be carried, but is not subject to the SLA. In this use case, the average load for three sectors is: 45 Mbps * 3 * 1.25 = 169 Mbps. Assuming *CIR*/*EIR* granularity of 50 Mbps, the Mobile Operator chooses: $CIR^n = 200$ Mbps and $EIR^n = 100$ Mbps. The value of CIR^n will cover normal operating conditions. The value of EIR^n covers the RAN BS peak load periods¹⁰.

Figure A-13 below shows an example of the token flow for three CoS Labels for this model. The red line depicts the general flow of tokens. The solid black lines with arrows show the specific unused token flow route for this use case, determined by the Coupling Flag values and the CIR_{max} values for each rank.




Figure A-13: Token flow diagram for the CX/GY/R Use Case

¹⁰ The "one-peak, two average" principle used in this example is just one option; alternatively, one may choose, for example, "one-peak" or "all-average".

In this example, there are three Bandwidth Profile Flows, one for each of the CoS Labels {H,M,L} supported by the EVPL instance. These three Bandwidth Profile Flows are contained in a single Envelope, whose Envelope ID = XYZ. Frames with CoS Label H can only be declared Green or Red, i.e., the EBS value is set to 0. Frames with CoS Label M and L can be declared Green, Yellow, or Red. CoS Label H is the source of both the Green tokens and the Yellow tokens (the Yellow tokens could also be sourced by CoS Label M, since the EBS for CoS Label H is set to 0). Unused tokens flow down, from H to M, and from M to L. Note that the Yellow token flow consists of EIR^i at the top rank and the unused Green tokens from the bottom rank that are converted to Yellow for the top rank.

In general, Model CX/GY/R is useful for cases where the bandwidth requirements for applications mapped to CoS Label H is predictably less than the bandwidth for the Envelope, and the bandwidth requirements for applications mapped to CoS Labels M and L are not predictable. In addition, the Mobile Operator may require some committed bandwidth for each Bandwidth Profile Flow. One example of this may be where CoS Label H is carrying VoIP and control traffic for a mobile backhaul application, CoS Label M is carrying high priority data/video traffic and CoS Label L is carrying base Internet traffic. In this example, it is assumed that CoS Label M will require less than full aggregate CIR^i bandwidth most of the time, but may at times burst up to that level. This is why an additional amount of Yellow bandwidth is provided. When there is no traffic in CoS Labels H & M, then CoS Label L would be able to use 300 Mbps.

In this use case, a CIR^3 value of 200 Mbps and an EIR^3 value of 100 Mbps are specified. These are the only token sources in the Envelope. In addition, a CIR^3_{max} value of 90 Mbps and an EIR^3_{max} value of 0 are specified. For the lower ranked Bandwidth Profile Flows, $CIR^2_{max} = 150$ Mbps and $EIR^2_{max} = 150$ Mbps, which allows rank 2 to get up to 300 Mbps, when the only traffic is in the Bandwidth Profile Flow for rank 2; and $CIR^1_{max} = 200$ Mbps and $EIR^1_{max} = 100$ Mbps, which allows rank 1 to get up to 300 Mbps, when the only traffic is in the Bandwidth Profile Flow for rank 1.

Table A-12 below shows the Bandwidth Profile Parameter values that are configured for the service, as a result of agreement between the Service Provider and the Mobile Operator.

Use Case for Model CX/GY/R - Ingress Bandwidth Profile Flow Parameter values			
Envelope Parameters: $CF^0 = 1$ and Envelope ID = XYZ			
Bandwidth Profile Flow Parameters	CoS Label H	CoS Label M	CoS Label L
CM^i {color-aware, color-blind}	color-blind	color-blind	color-blind
CF^i {0,1}	0	0	0
CIR^i (Mbps)	200	0	0
CBS^i (B)	8*MFS	24*MFS	24*MFS
CIR^i_{max} (Mbps)	90	150	200
EIR^i (Mbps)	100	0	0
EBS^i (B)	0	24*MFS	24*MFS
EIR^i_{max} (Mbps)	0	150	100
ER^i	<XYZ,3>	<XYZ,2>	<XYZ,1>
F^i	0	0	0

Table A-12: Ingress Bandwidth Profile Flow Parameter Values for the Mobile Backhaul Use Case

In this example, CoS Label H can take, at most, 45% of the CIR^3 bandwidth ($CIR^3_{max} = 90$ Mbps). In normal times, this Bandwidth Profile Flow takes much less than 45%, but it is critical that frames are not dropped, so 45% is a safe level for even the worst times. Due to token sharing Bandwidth Profiles, unused Green tokens at H can be shared in the lower rank Bandwidth Profile Flows. Similarly, CoS Label M can take, at most, 75% of the CIR bandwidth for Green traffic ($CIR^i_{max} = 150$ Mbps). The above parameter values have been chosen to allow full peak bandwidth in CoS Label M or in CoS Label L if the other Bandwidth Profile Flows in the Envelope make all of their tokens available for sharing. Note that SLS guarantees do not apply to Yellow frames.

The parameter values (specifically, CIR_{max} and EIR_{max}) for each Bandwidth Profile Flow in this example have been carefully chosen to ensure that the simplified shaper described in Section 10 could be used effectively. Specifically, CIR_{max} and EIR_{max} for CoS Label M, as well as for CoS Label L, are set so that the sum is at least equal to the shaping rate. Use of other values, e.g., where the sum of CIR_{max} and EIR_{max} is less than the shaping rate, may result in the need for a hierarchical scheduler and shaper in the Mobile Operator's equipment.

The above example just describes one type of behavior, but Model CX/GY/R is flexible and is able to support a wide range of behaviors.

I.4 Use Case: Mobile Backhaul, Two-EVC Use Case

This use case describes a mobile backhaul application using two diversely routed EVCs from a cell site to connect to an aggregation site. The Mobile Operator requires a total of 200 Mbps bandwidth shared by the EVCs, each of which has three Bandwidth Profile Flows. The Ingress Bandwidth Profile at the cell site UNI (UNI_1) is applied to the six Bandwidth Profile Flows in the Envelope, whose Envelope ID = XYZ. It supports either an active-standby or an active-active application.

Figure A-14 below shows the simplified service model for this use case.

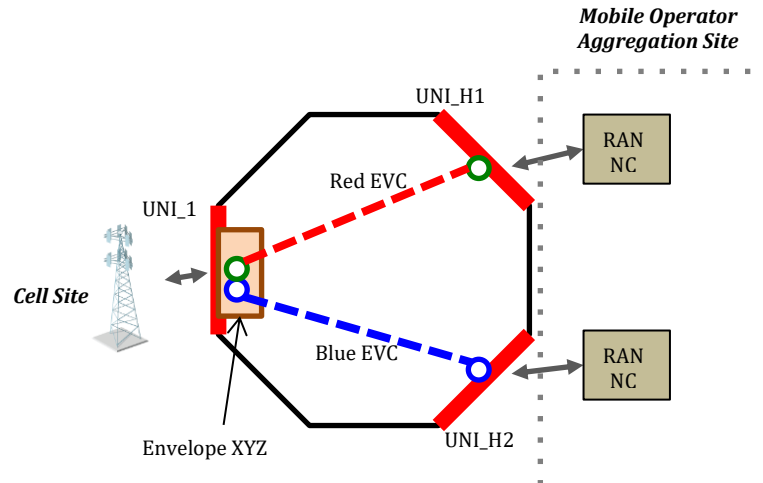


Figure A-14: Service Model for Mobile Backhaul, 2-EVC Use Case

This example is summarized, as follows:

- The Mobile Operator may use this service in an active-standby scenario, where the majority of the traffic is handled by one EVC, with only hand-shaking traffic on the standby EVC. Or, it can be used in an active-active scenario, where one or more flows in each EVC are carrying traffic. It is assumed that the Mobile Operator is in control of this arrangement and that the CEN Operator does not have to know whether active-active or active-standby is being used.
- The Bandwidth Profile Flows consist of three CoS Labels for each EVC {H, M, L} in Envelope XYZ at UNI_1. Frames with CoS Label H can only be declared Green or Red. Frames with CoS Label M can be declared Green, Yellow, or Red. Frames with CoS Label L can only be declared Yellow or Red. The Bandwidth Profile Flows are ranked from highest to lowest, as follows: {Red-H, Blue-H, Red-M, Blue-M, Red-L, Blue-L}.
- Two Green token sources are used to provide the required committed bandwidth for the SLS, and one Yellow token source is used to provide additional bandwidth to satisfy demands during peak time periods. Unused Green tokens are converted to Yellow to feed the highest rank to provide additional bandwidth for each Bandwidth Profile Flow during times when one or more flows are carrying no traffic.
- The CE shapes to the aggregate $CIR + EIR$ for the Envelope, and schedules within.

The token flow diagram for this use case is shown in Figure A-15 below. The red line depicts the general flow of unused tokens. The solid black lines with arrows show the specific unused token flow route for this use case, determined by the Coupling Flag values and the CIR_{max} values for each rank.

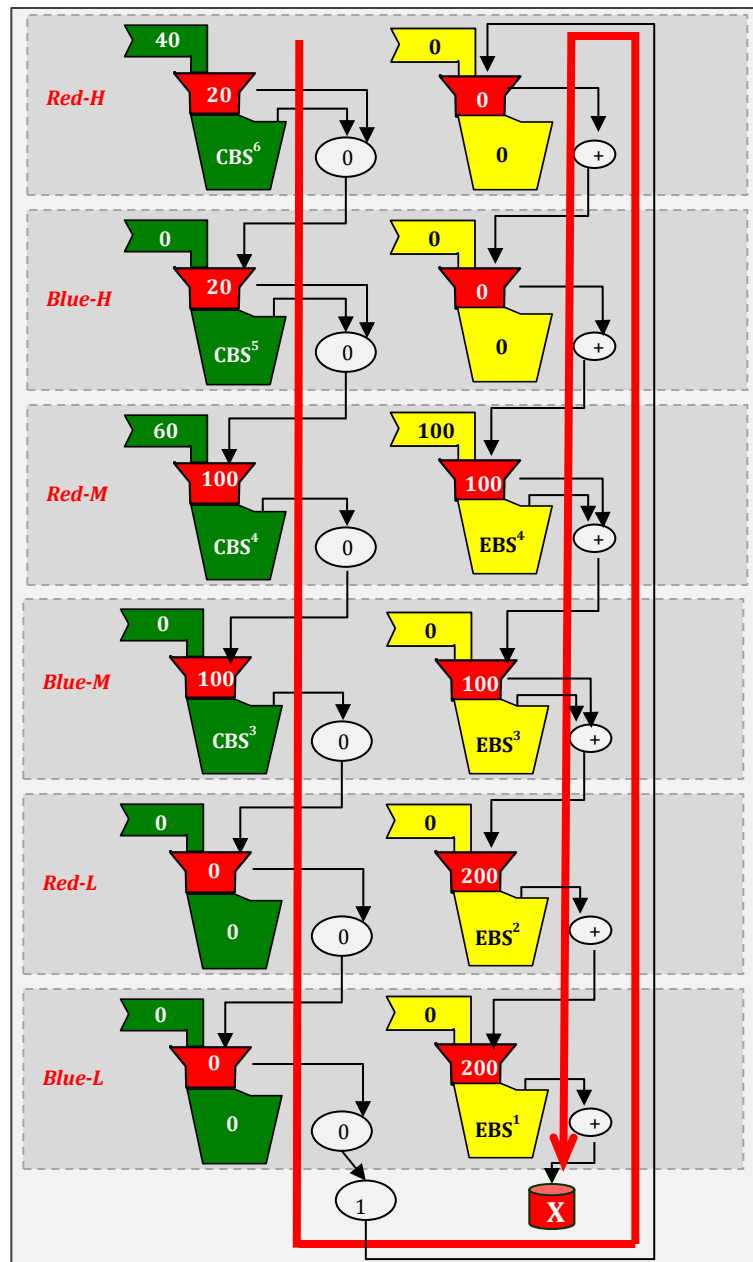


Figure A-15: Token flow diagram for the 2 EVC Use Case

In this example, Green tokens are sourced at ranks 6 and 4, and Yellow tokens are sourced at rank 4 (other arrangements are possible).

Green token flow: Rank 6 (Red-H) sources Green tokens at a rate of 40 Mbps and has a CIR_{max}^6 value of 20 Mbps, allowing it to take half the bandwidth. Any unused Green tokens from rank 6 flow down to rank 5 (Blue-H). Rank 5 has a CIR_{max}^5 value of 20 Mbps, allowing it to take only up to 20 Mbps. So, CoS Label H for each EVC is provided 20 Mbps. Any unused Green tokens from

ranks 5 and 6 flow down to rank 4. Rank 4 (Red-M) has a Green token source of 60 Mbps, and a CIR_{max}^4 value of 100 Mbps. Thus, rank 4 is provided at least 60 Mbps, and can get up to 100 Mbps of committed bandwidth if ranks 6 and 5 have unused tokens. Any unused Green tokens from rank 4 flow down to rank 3. In this example, rank 3 can get up to 100 Mbps, depending on the available tokens flowing down from the higher ranks. All of the unused Green tokens from rank 3 flow down to ranks 2 and then 1, neither of which can use Green tokens ($CBS^i=0$ and $CIR_{max}^i = 0$ for ranks 2 and 1). These unused Green tokens are then converted to Yellow tokens for feeding the Yellow token buckets.

Yellow token flow: The Yellow tokens are sourced by the unused Green tokens at the bottom of rank 1 and the Yellow token source at rank 4 (Red-M), which is set at 100 Mbps. EIR_{max}^4 is set to 100 Mbps, allowing it to utilize all of its Yellow tokens. EIR_{max}^3 is also set to 100 Mbps, allowing rank 3 to fully utilize any Yellow tokens flowing down from rank 4. Any unused Yellow tokens in rank 3 flow down to rank 2 (Red-L) and then to rank 1 (Blue-L). These ranks do not have separate token sources and rely completely on the unused tokens from rank 3. Both ranks 2 and 1 have an EIR_{max}^i value of 200 Mbps, allowing either to get the full 200 Mbps bandwidth when the higher ranks make the tokens available.

Table A-13 below shows the Ingress Bandwidth Profile Parameter values that are agreed between the Service Provider and Subscriber for this service.

Model CX/GY/R Use Case - Ingress Bandwidth Profile Flow Parameter values						
Envelope Parameters: $CF^0 = 1$ and Envelope ID = XYZ						
Bandwidth Profile Flow Parameters	Red EVC, CoS Label H	Blue EVC CoS Label H	Red EVC CoS Label M	Blue EVC CoS Label M	Red EVC CoS Label L	Blue EVC CoS Label L
CM^i {color-aware, color-blind}	color-blind	color-blind	color-blind	color-blind	color-blind	color-blind
CF^i {0,1}	0	0	0	0	0	0
CIR^i (Mbps)	40	0	60	0	0	0
CBS^i (B)	8*MFS	8*MFS	24*MFS	24*MFS	0	0
CIR_{max}^i (Mbps)	20	20	100	100	0	0
EIR^i (Mbps)	0	0	100	0	0	0
EBS^i (B)	0	0	24*MFS	24*MFS	24*MFS	24*MFS
EIR_{max}^i (Mbps)	0	0	100	100	200	200
ER^i	<XYZ,6>	<XYZ,5>	<XYZ,4>	<XYZ,3>	<XYZ,2>	<XYZ,1>
F^i	0	0	0	0	0	0

Table A-13: Ingress Bandwidth Profile Flow Parameter Values for Two-EVC Use Case

In summary, this use case supports bandwidth sharing among six Bandwidth Profile Flows across two EVCs. Both active-active and active-standby scenarios can be supported.