

MEF

Technical Specification

MEF 41

Generic Token Bucket Algorithm

October 2013

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 Metro Ethernet 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 company which are or may be associated with the ideas, techniques, concepts or expressions contained herein; nor
- b) any warranty or representation that any MEF member companies 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 member companies and the recipient or user of this document.

Implementation or use of specific Metro Ethernet standards or recommendations and MEF specifications will be voluntary, and no company shall be obliged to implement them by virtue of participation in the Metro Ethernet Forum. The MEF is a non-profit international organization accelerating industry cooperation on Metro Ethernet technology. The MEF does not, expressly or otherwise, endorse or promote any specific products or services.

© The Metro Ethernet Forum 2013. All Rights Reserved.

Table of Contents

| | | |
|--|---|-----------|
| 1. | List of Contributing Members | 1 |
| 2. | Abstract | 1 |
| 3. | Terminology and Acronyms | 1 |
| 4. | Scope | 3 |
| 5. | Compliance Levels | 3 |
| 6. | Introduction | 3 |
| 7. | Token Requests and Token Request Flows | 4 |
| 8. | Algorithm Parameters | 5 |
| 8.1 | General Parameters | 5 |
| 8.2 | TRF Parameters | 5 |
| 9. | Token Count Updating | 7 |
| 10. | Color Determination | 9 |
| 10.1 | Sequence of Color Determination | 9 |
| 10.2 | Color Determination Logic..... | 9 |
| 11. | References | 10 |
| Appendix A Examples of Referencing the GTBA | | 10 |
| A.1 | Bandwidth Profile at the UNI Example..... | 10 |
| A.2 | Bandwidth Profile at the ENNI Example | 13 |

List of Figures

| | | |
|----------|--|----|
| Figure 1 | – Color Determination Logic for a TRF Green Token Request | 4 |
| Figure 2 | – Color Determination Logic for a TRF Yellow Token Request | 4 |
| Figure 3 | – Generic Token Bucket Algorithm Token Flows | 7 |
| Figure 4 | – The Color Determination Logic | 10 |
| Figure 5 | – Active/Standby EVCs..... | 11 |

List of Tables

| | | |
|---------|--|----|
| Table 1 | – Terminology and Acronyms | 3 |
| Table 2 | – Bandwidth Profile Parameter Values | 11 |
| Table 3 | – Token Request Flow Parameter Values | 12 |
| Table 4 | – Setting the Value of c for TRF i | 12 |
| Table 5 | – Token Request Flow Parameter Values | 13 |

1. List of Contributing Members

The following members of the MEF participated in the development of this document and have requested to be included in this list.

| | |
|-------------------|---------------------------------|
| at&t | EXFO Inc. |
| BTI Systems Inc. | Pulse Communications (Pulsecom) |
| Ciena Corporation | Verizon |

2. Abstract

This document specifies a Generic Token Bucket Algorithm (GTBA) based on the Bandwidth Profile Algorithm specified in MEF 10.3.[1] It is expected that other documents can reference this specification for describing new Bandwidth Profiles and methods for regulating the use of other resources.

3. Terminology and Acronyms

This section defines the terms used in this document. In many cases, the normative definitions to terms are found in other documents. In these cases, the third column is used to provide the reference that is controlling, in other MEF or external documents.

| Term | Definition | Source |
|---------------------------------------|---|---------------|
| Bandwidth Profile | A characterization of the lengths and arrival times for Frames at a reference point. | MEF 10.3 [2] |
| CF^0 | A GTBA parameter that controls overflow of rank 1 Green tokens to rank n Yellow tokens. | This document |
| CF^i | A GTBA parameter that controls the destination of overflow TRF i Green tokens. | This document |
| Generic Token Bucket Algorithm | A method for determining a color for each element in a sequence of Token Requests. | This document |
| Color | The value assigned to each Token Request by the GTBA. Can be one of three possible values; Green, Yellow, or Red. | This document |
| GTBA | Generic Token Bucket Algorithm | This document |
| GTR | Green Token Rate | This document |
| GTR^i | A GTBA parameter that represents the replenishing rate of new Green tokens for TRF i | This document |

| Term | Definition | Source |
|---------------------------|---|---------------|
| GTR_{\max}^i | A GTBA parameter that represents the limit on the sum of the replenishing rate of Green tokens and the overflow Green tokens for TRF i . | This document |
| GTV | Green Token Volume | This document |
| GTV^i | A GTBA parameter that represents the upper limit on the Green token bucket count for TRF i . | This document |
| n | A positive integer representing the number of TRFs. | This document |
| Rank | The order in which Token Sharing would occur among Token Request Flows. Any Token Sharing would occur beginning with the highest ranking TRF and continuing in rank order to the lowest ranking TRF. | This document |
| Token Request | A 4-tuple of the form $\langle l, t, c, r \rangle$ where l is a positive number representing the number of tokens requested, t is a number representing the time of occurrence of the token request, c is the type of tokens requested with value Green, Yellow, or Red, and r is a positive integer that both identifies the TRF and its rank for Token Sharing. | This document |
| Token Request Flow | A sequence of Token Requests where each Token Request in the TRF has the same value of r . | This document |
| TRF | Token Request Flow | This document |
| YTR | Yellow Token Rate | This document |
| YTR^i | A GTBA parameter that represents the replenishing rate of new Yellow tokens for TRF i . | This document |
| YTR_{\max}^i | A GTBA parameter that represents the limit on the sum of the replenishing rate of Yellow tokens and the overflow Yellow tokens for TRF i . | This document |
| YTV | Yellow Token Volume | This document |
| YTV^i | A GTBA parameter that represents the upper limit on the Yellow token bucket count for TRF i . | This document |

Table 1 – Terminology and Acronyms**4. Scope**

This document specifies a Generic Token Bucket Algorithm (GTBA) that can be used for regulating the use of a resource. A prime example is a Bandwidth Profile where the resource in question is bandwidth. The GTBA specified includes a token sharing schema that is identical to that specified in Section 10 of MEF 10.3. [2]

The goal of this document is to allow documents that need a token bucket algorithm to reference this document. This will avoid having to specify the complex logic in the referencing document.

Implementation considerations for the GTBA are beyond the scope of this document. Implementation considerations when this algorithm is used for defining a Bandwidth Profile can be found in Appendix C.3 of MEF 10.3. [2]

Section 7 introduces the concepts of Token Request and Token Request Flow.

Section 8 describes the parameters that control the operation of the GTBA.

Section 9 lays out the equations for updating token counts including token sharing among Token Request Flows.

Section 10 specifies the logic for the Color Determination of each Token Request.

Appendix A contains examples of referencing this document to specify a Bandwidth Profile.

5. Compliance Levels

The key words "MUST", "MUST NOT", "REQUIRED", "SHALL", "SHALL NOT", "SHOULD", "SHOULD NOT", "RECOMMENDED", "MAY", and "OPTIONAL" in this document are to be interpreted as described in RFC 2119 [1]. All key words must be in upper case, bold text.

Items that are **REQUIRED** (contain the words **MUST** or **MUST NOT**) are labeled as [Rx] for required. Items that are **RECOMMENDED** (contain the words **SHOULD** or **SHOULD NOT**) are labeled as [Dx] for desirable. Items that are **OPTIONAL** (contain the words **MAY** or **OPTIONAL**) are labeled as [Ox] for optional.

6. Introduction

The Generic Token Bucket Algorithm (GTBA) is a method for determining the color for each element in a sequence of Token Requests. A Token Request includes the number of tokens requested and the type of tokens requested (Green, Yellow, or Red). The result of each such determination is a one of three colors: Green, Yellow, or Red.

The sequence of Token Requests is composed of one or more subsequences of Token Requests. Such a subsequence is called a Token Request Flow (TRF). The GTBA maintains a token bucket

of Green tokens and a token bucket of Yellow tokens for each TRF. These token buckets have upper bounds on their token counts.

Figure 1 shows an informal description of the Color Determination logic for a TRF token request for l Green tokens in pseudo code.¹ Figure 2 shows an informal description of the Color Determination logic for a TRF token request for l Yellow tokens in pseudo code.¹ All requests for Red tokens are determined to be Red.

```

if( $l \leq$  Green token bucket count)
    {Color Determination = Green;
    Green token bucket count = Green token bucket count - 1;}
else if( $l \leq$  Yellow token bucket count)
    {Color Determination = Yellow;
    Yellow token bucket count = Yellow token bucket count - 1;}
else
    {Color Determination = Red;}

```

Figure 1 – Color Determination Logic for a TRF Green Token Request

```

if( $l \leq$  Yellow token bucket count)
    {Color Determination = Yellow;
    Yellow token bucket count = Yellow token bucket count - 1;}
else
    {Color Determination = Red;}

```

Figure 2 – Color Determination Logic for a TRF Yellow Token Request

The structures of a Token Request and a TRF are described in Section 7. There are several parameters whose values control the Green and Yellow token bucket counts as detailed in Section 8. In addition, the token bucket counts for one TRF can be modified by the addition of tokens not used by other TRFs. This is referred to Token Sharing. Section 9 specifies how the token bucket counts are determined including Token Sharing. Section 10 contains the specification of the logic for the Color Determination of token requests.

One of the uses of the GTBA is for controlling bandwidth use. In such a case, each token corresponds to a byte of transmission and each Token Request is derived from the length of a frame or packet at a reference point. Appendix A shows examples of how to use the GTBA in specifying a Bandwidth Profile.

7. Token Requests and Token Request Flows

A Token Request consists of a 4-tuple of the form $\langle l, t, c, r \rangle$ where:

- l is a positive number representing the number of tokens requested,
- t is a number representing the time of occurrence of the token request,
- c is the type of tokens requested with value Green, Yellow, or Red, and

¹ See Section 10 for the normative logic specification.

- r is a positive integer that both identifies the TRF and its rank for Token Sharing.

Note that the algorithm operates with three colors: Green, Yellow, and Red. Consequently, the color of a Token Request, c , is allowed to take on the values Green, Yellow, or Red. And, the color declaration can be Green, Yellow, or Red.

The disposition of the requested resource based on the color declaration is to be specified by the referencing document. For example, a Bandwidth Profile would discard a frame whose color is declared to be Red.

A TRF is the sequence of Token Requests where each Token Request in the TRF have the same value of r . For example, $\langle l_j, t_j, c_j, 3 \rangle, j = 0, 1, 2, \dots$ could represent TRF 3 that, by definition, has rank 3.

The order in which Token Requests have their color determined by the GTBA is very important and is specified in Section 10.1.

8. Algorithm Parameters

There are two kinds of parameters that control the behavior of GTBA: General Parameters and TRF Parameters as described below.

8.1 General Parameters

There are two General Parameters:

- n a positive integer representing the number of TRFs and
- CF^0 can be 0 or 1 and controls overflow of rank 1 Green tokens to rank n Yellow tokens.²

[R1] Values of n and CF^0 **MUST** be specified for each instance of a GTBA.

[R2] When $n = 1$, CF^0 **MUST** have the value of 0.

8.2 TRF Parameters

Let the n TRFs be ranked $1, 2, \dots, n$ where n is the highest rank. Then for each $i \in \{1, 2, \dots, n\}$, there are the following TRF Parameters:

- GTR^i a non-negative number in units of tokens per unit time that represents the replenishing rate of new Green tokens for TRF i ,

² See Figure 3 and Section 9 for the details of how CF^0 is used.

- GTR_{\max}^i a non-negative number in units of tokens per unit time that represents the limit on the sum of the replenishing rate of Green tokens and the overflow Green tokens for TRF i ,
- GTV^i a non-negative number in units of tokens that represents the upper limit on the Green token bucket count for TRF i ,
- YTR^i a non-negative number in units of tokens per unit time that represents the replenishing rate of new Yellow tokens for TRF i ,
- YTR_{\max}^i a non-negative number in units of tokens per unit time that represents the limit on the sum of the replenishing rate of Yellow tokens and the overflow Yellow tokens for TRF i ,
- YTV^i a non-negative number in units of tokens that represents the upper limit on the Yellow token bucket count for TRF i , and
- CF^i can be 0 or 1 and controls the destination of overflow TRF i Green tokens.

[R3] If $CF^0 = 1$, then CF^i **MUST** equal 0 for $i = 1, 2, \dots, n$.

Figure 3 shows a conceptual diagram of the token flows with three TRFs.

- The green and yellow trapezoids represent the Green and Yellow token bucket for each TRF.
- The solid arrows labeled GTR^i or YTR^i represent a token source flowing into each token bucket.
- The dashed arrows show where tokens that overflow each token bucket go when the number of tokens in the token bucket reaches the capacity (GTV^i or YTV^i).
- The circles represent decision points in the path of the overflow tokens that are controlled by CF^i .
- The X's represent points at which tokens are discarded (i.e., not added to any of the token buckets).
- GTR_{\max}^i and YTR_{\max}^i are not shown on the figure.

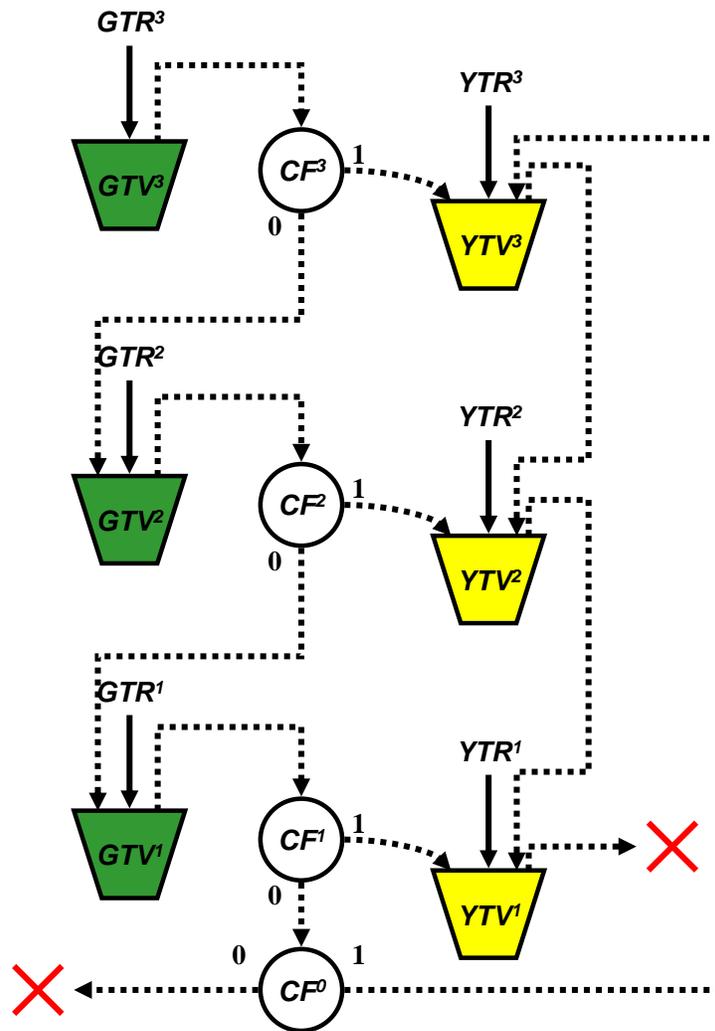


Figure 3 – Generic Token Bucket Algorithm Token Flows

9. Token Count Updating

This Section specifies how the token bucket counts for each TRF are calculated. $C_G^i(t)$ represents the Green token bucket count at time t for TRF i , $i = 1, 2, \dots, n$ and $C_Y^i(t)$ represents the Yellow token bucket count at time t for TRF i , $i = 1, 2, \dots, n$.

$T_G^i(t_1, t_2)$ and $T_Y^i(t_1, t_2)$ represent the maximum number of tokens that might be added to the Green and Yellow token bucket counts, respectively, for TRF i over the time interval t_1 to t_2 , for $i = 1, 2, \dots, n$.

$A_G^i(t_1, t_2)$ and $A_Y^i(t_1, t_2)$ represent the number of tokens actually added to the Green and Yellow token bucket counts, respectively, for TRF i over the time interval t_1 to t_2 , for $i = 1, 2, \dots, n$.

$O_G^i(t_1, t_2)$ and $O_Y^i(t_1, t_2)$ represent the number of tokens that overflow the Green and Yellow token buckets, respectively, for TRF i over the time interval t_1 to t_2 , for $i = 1, 2, \dots, n$. Tokens are said to overflow when the number of tokens to be added to a token bucket count would cause the count to exceed the upper limit on the token bucket count.

$S_G^i(t_1, t_2)$ and $S_Y^i(t_1, t_2)$ represent the number of tokens available to be shared to the Green and Yellow token bucket counts of TRF $i-1$, respectively, from TRF i over the time interval t_1 to t_2 for $i = 2, \dots, n, n+1$.

The maximum number of tokens that might be added to Green token bucket count, $T_G^i(t_1, t_2)$, includes Green tokens sourced at the rate GTR^i over the time interval and any Green tokens shared from the next higher rank TRF. The maximum number of tokens that might be added to the Yellow token bucket count, $T_Y^i(t_1, t_2)$, includes Yellow tokens sourced at the rate YTR^i over the time interval, any Yellow tokens shared from the higher rank TRF, and any overflow Green tokens from the Green token bucket that are allowed by CF^i . Note that for the Green token bucket count, TRF n has no tokens shared from a higher rank since there are no TRFs with a rank higher than n . Therefore $S_G^{n+1}(t_1, t_2) = 0$. For the Yellow token bucket count there are no tokens shared from a rank higher than n , however, depending on CF^0 , there may be tokens shared from the overflow of the Green token bucket at rank 1. Therefore $S_Y^{n+1}(t_1, t_2) = CF^0 \times O_G^1(t_1, t_2)$.

$$T_G^i(t_1, t_2) = GTR^i \times (t_2 - t_1) + S_G^{i+1}(t_1, t_2) \text{ for } i = 1, 2, \dots, n$$

$$T_Y^i(t_1, t_2) = YTR^i \times (t_2 - t_1) + S_Y^{i+1}(t_1, t_2) + CF^i \times O_G^i(t_1, t_2) \text{ for } i = 1, 2, \dots, n$$

Note that [R3] mandates that $CF^n = 0$ if $CF^0 = 1$ in the equation for $T_Y^n(t_1, t_2)$ and thus

$$T_Y^n(t_1, t_2) = \begin{cases} YTR^n \times (t_2 - t_1) + CF^n \times O_G^n(t_1, t_2), & \text{if } CF^0 = 0 \\ YTR^n \times (t_2 - t_1) + CF^0 \times O_G^1(t_1, t_2), & \text{if } CF^0 = 1 \end{cases}$$

The number of tokens actually added to a token bucket count is some or all of the tokens available to be added, limited by GTV^i and YTV^i and by the maximum rate at which tokens are allowed to be added (GTR_{\max}^i and YTR_{\max}^i).

$$A_G^i(t_1, t_2) = \min \{ T_G^i(t_1, t_2), GTV^i - C_G^i(t_1), GTR_{\max}^i \times (t_2 - t_1) \} \text{ for } i = 1, 2, \dots, n$$

$$A_Y^i(t_1, t_2) = \min \{ T_Y^i(t_1, t_2), YTV^i - C_Y^i(t_1), YTR_{\max}^i \times (t_2 - t_1) \} \text{ for } i = 1, 2, \dots, n$$

The number of tokens that overflow each token bucket is the number of available tokens that are not actually added to the token bucket count.

$$O_G^i(t_1, t_2) = T_G^i(t_1, t_2) - A_G^i(t_1, t_2) \text{ for } i = 1, 2, \dots, n$$

$$O_Y^i(t_1, t_2) = T_Y^i(t_1, t_2) - A_Y^i(t_1, t_2) \text{ for } i = 1, 2, \dots, n$$

The number of tokens available to be shared from TRF i to the next lower rank, TRF $i-1$, is the number of tokens that overflow the Green and Yellow token buckets and, in the case of the Green token bucket count, are not made available to the Yellow token count by CF^i .

$$S_G^i(t_1, t_2) = (1 - CF^i) \times O_G^i(t_1, t_2) \text{ for } i = 2, 3, \dots, n.$$

$$S_Y^i(t_1, t_2) = O_Y^i(t_1, t_2) \text{ for } i = 2, 3, \dots, n.$$

10. Color Determination

10.1 Sequence of Color Determination

The Color Determination logic specified in Section 10.2 operates on a sequence of Token Requests $\langle l_k, t_k, c_k, r_k \rangle, k = 0, 1, \dots$ that is constructed from the n TRFs. A document that references this algorithm needs to unambiguously specify how $\langle l_k, t_k, c_k, r_k \rangle, k = 0, 1, \dots$ is derived from the n TRFs. This sequence can be tailored in any way so long as [R4] is satisfied.

[R4] $t_k \leq t_{k+1}$ **MUST** hold for all $k = 0, 1, \dots$

10.2 Color Determination Logic

The Color Determination logic for each Token Request in $\langle l_k, t_k, c_k, r_k \rangle, k = 0, 1, \dots$ is specified in this Section.

[R5] The Color Determination for $\langle l_k, t_k, c_k, r_k \rangle, k = 0, 1, \dots$ **MUST** be as specified in Figure 4 where $C_G^i(t_0) = GTV^i$ and $C_Y^i(t_0) = YTV^i$.

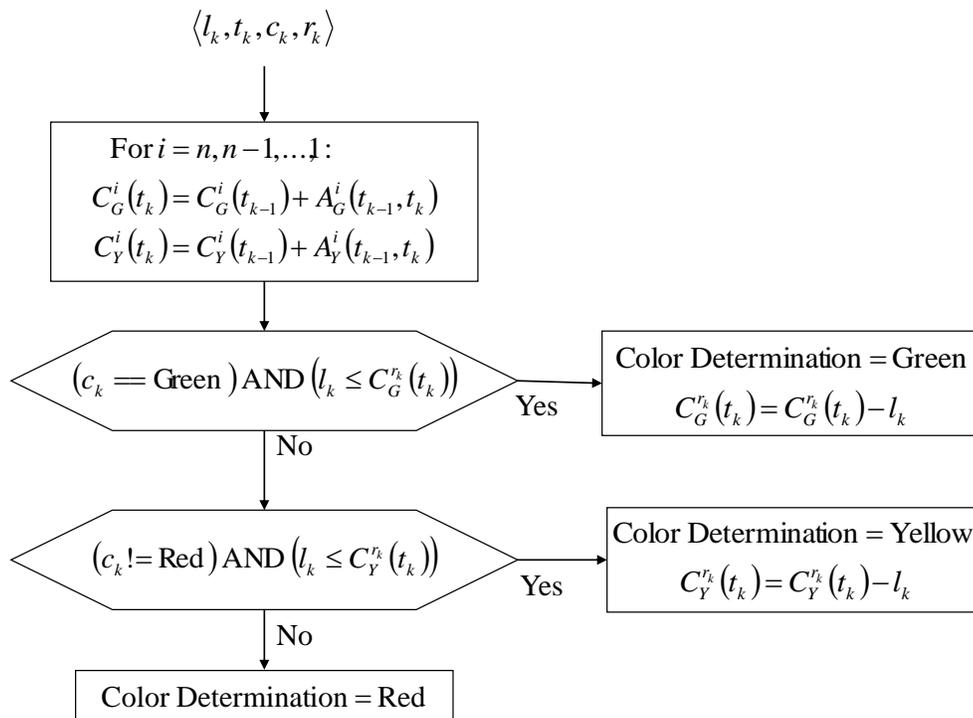


Figure 4 – The Color Determination Logic

11. References

- [1] S. Bradner, RFC 2119, *Key words for use in RFCs to Indicate Requirement Levels*, March 1997.
- [2] Metro Ethernet Forum MEF 10.3, *Ethernet Services Attributes Phase 3*, October 2013.
- [3] Metro Ethernet Forum MEF 26.1, *External Network Network Interface (ENNI) – Phase 2*, January 2012.

Appendix A Examples of Referencing the GTBA

This Appendix provides examples of how this document can be referenced for the purposes of specifying Bandwidth Profiles.

A.1 Bandwidth Profile at the UNI Example

This example involves an active and a standby EVC at UNI A as illustrated in Figure 5 where the two EVCs share 40×10^6 bps of committed bandwidth and 8×10^6 bps of excess bandwidth. Bandwidth Profile Flows are based on Criterion 2 of [R157] of MEF 10.3 [2]. This means that ingress Service Frames mapped to EVC 1 constitute one Bandwidth Profile Flow while ingress Service Frames mapped to EVC 2 constitute a second Bandwidth Profile Flow. These two

Bandwidth Profile Flows are mapped to a single Envelope (called Env) with EVC 2 having rank 2 and EVC 1 having rank 1.

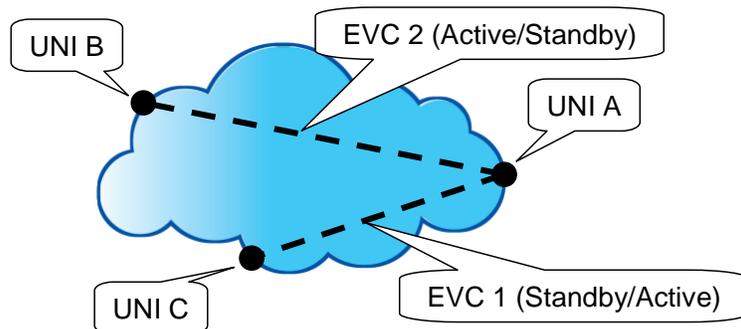


Figure 5 – Active/Standby EVCs

The Bandwidth Profile Parameter $CF^0 = 0$ and Table 2 contains the other Bandwidth Profile parameter values for this example.

| EVC 1 Bandwidth Profile Flow | | EVC 2 Bandwidth Profile Flow | |
|------------------------------|-----------------------|------------------------------|----------------------|
| Parameter | Value | Parameter | Value |
| CIR^1 | 8000 bps ³ | CIR^2 | 40×10^6 bps |
| CBS^1 | 15,000 bytes | CBS^2 | 15,000 bytes |
| CIR_{max}^1 | ∞ bps | CIR_{max}^2 | ∞ bps |
| EIR^1 | 0 bps | EIR^2 | 8×10^6 bps |
| EBS^1 | 15,000 bytes | EBS^2 | 15,000 bytes |
| EIR_{max}^1 | ∞ bps | EIR_{max}^2 | ∞ bps |
| CF^1 | 0 | CF^2 | 0 |
| CM^1 | color-aware | CM^2 | color-aware |
| ER^1 | <Env, 1> | ER^2 | <Env, 2> |

Table 2 – Bandwidth Profile Parameter Values

To reference the GTBA, each Bandwidth Profile Flow is associated with a TRF with the parameter values as shown in Table 3. Note that bits per second are converted to tokens per second by dividing by 8 and bytes are directly converted to tokens.

³ This value is meant to accommodate routing protocols on the backup EVC.

| EVC 1 Token Request Flow | | EVC 2 Token Request Flow | |
|--------------------------|----------------------------|--------------------------|-----------------------------------|
| Parameter | Value | Parameter | Value |
| GTR^1 | 1000 tokens per second | GTR^2 | 5×10^6 tokens per second |
| GTV^1 | 15,000 tokens | GTV^2 | 15,000 tokens |
| GTR_{\max}^1 | ∞ tokens per second | GTR_{\max}^2 | ∞ tokens per second |
| YTR^1 | 0 tokens per second | YTR^2 | 1×10^6 tokens per second |
| YTV^1 | 15,000 tokens | YTV^2 | 15,000 tokens |
| YTR_{\max}^1 | ∞ tokens per second | YTR_{\max}^2 | ∞ tokens per second |
| CF^1 | 0 | CF^2 | 0 |

Table 3 – Token Request Flow Parameter Values

Note that there is no parameter corresponding to the Bandwidth Profile Color Mode parameter CM . The value of CM^i determines the value of c in the Token Requests, $\langle l, t, c, r \rangle$, for TRF i as per Table 4.

| Value of CM^i | Value of c |
|-----------------|----------------------------|
| color-aware | Color of the Service Frame |
| color-blind | Green |

Table 4 – Setting the Value of c for TRF i

The final step is specifying the Token Requests in each TRF. For TRF 1, each time an ingress Service Frame that is mapped to EVC 1 arrives at UNI A, a Token Request $\langle l, t, c, r \rangle$ is specified where:

- l = the length of the Service Frame in bytes,
- t = the arrival time of the Service Frame,
- c = the color indicated by the Color Identifier of the Service Frame or Green if there is no Color Identifier, and
- $r = 1$.

TRF 1 is then $\langle l_j, t_j, c_j, 1 \rangle, j = 0, 1, \dots$ where $t_{j+1} > t_j, j = 0, 1, \dots$.⁴

For TRF 2, the Token Requests are based on ingress Service Frames mapped to EVC 2 at UNI A where:

- l = the length of the Service Frame in bytes,
- t = the arrival time of the Service Frame,

⁴ Since two Service Frames cannot arrive at the same time at a UNI, the inequality is strict.

- c = the color indicated by the Color Identifier of the Service Frame or Green if there is no Color Identifier, and
- $r = 2$.

TRF 2 is then $\langle l_j, t_j, c_j, 2 \rangle, j = 0, 1, \dots$ where $t_{j+1} > t_j, j = 0, 1, \dots$.⁴

Note that if CM^1 equaled color-blind, then c would be set to Green for all TRF 1 Token Requests and similarly for TRF 2.

With the above specification of Token Request Flows, the Bandwidth Profile declaration for an ingress Service Frame mapped to either EVC 1 or EVC 2 at UNI A is equal to the Color Determination of the corresponding Token Request.

A.2 Bandwidth Profile at the ENNI Example

This example shows how an ingress Bandwidth Profile per Class of Service Identifier at the ENNI as described in MEF 26.1 [3] can be specified using the GTBA. We assume that the ingress Bandwidth Profile is applied to all ingress ENNI Frames that map to the OVC in question and have PCP = 4. We further assume that the DEI is used to indicate the color of each ENNI Frame. The parameters for this Bandwidth Profile are $\langle CIR, CBS, EIR, EBS, CF, CM \rangle$.⁵

In this case there is only one TRF and the parameters are set as shown in Table 5.

| Parameter | Value |
|---------------|----------------------------|
| GTR^1 | $CIR/8$ tokens per second |
| GTV^1 | CBS tokens |
| GTR_{max}^1 | ∞ tokens per second |
| YTR^1 | $EIR/8$ tokens per second |
| YTV^1 | EBS tokens |
| YTR_{max}^1 | ∞ tokens per second |
| CF^1 | CF |

Table 5 – Token Request Flow Parameter Values

The final step is specifying the Token Requests in the TRF. Each time an ingress ENNI Frame that is mapped to the OVC and has PCP = 4 arrives at the ENNI, a Token Request $\langle l, t, c, r \rangle$ is specified where:

- l = the length of the ENNI Frame in bytes,

⁵ The Bandwidth Profile Algorithm in MEF 26.1 [3] does not include token sharing and thus only these six parameters are needed to specify the Bandwidth Profile. A future revision of MEF 26.1 might introduce token sharing in which case this specification could be used as a reference in the revision.

- t = the arrival time of the ENNI Frame,
- c = Green if DEI = 0 and Yellow if DEI = 1, and
- $r = 1$.

The TRF is then $\langle l_j, t_j, c_j, 1 \rangle, j = 0, 1, \dots$ where $t_{j+1} > t_j, j = 0, 1, \dots$.⁶

With the above specification of Token Request Flows, the Bandwidth Profile declaration for each ENNI Frame that is mapped to the OVC and that has PCP = 4 is equal to the Color Determination of the corresponding Token Request.

⁶ Since two ENNI Frames cannot arrive at the same time at an ENNI, the inequality is strict.