

Bandwidth Profile

A characterization of the lengths and arrival times for Service Frames at a reference point
(MEF 10.3 p8)

A Bandwidth Profile (BWP) enforces limits on bandwidth utilization according to the [Service Level Specification \(SLS\)](#) that has been agreed upon by the [Subscriber](#) and the [Service Provider](#) as part of the [Service Level Agreement \(SLA\)](#). It can be thought of as enforcing the long term average guaranteed bandwidth (CIR) and excess bandwidth (EIR) allowed by the service.

CIR (Committed Information Rate) is the bit rate for which the SP provides performance guarantees in terms of performance attributes for the service.

EIR (Excess Information Rate) is the additional bit-rate that the subscriber can utilize and for which traffic can probably pass through the [CEN](#) as long as there is no congestion.

Note that the total rate, sometimes known as PIR (Peak Information Rate), is the sum of CIR and EIR.

Enforcement of a bandwidth profile is an operation called **policing**. The policer indicates a frame's compliance or non-compliance to a bandwidth profile using one of three "color" markings.

Color Markings

Green

Frames within the CIR / CBS compliance level. These frames are subject to the SLS.

Yellow

Frames exceeding the CIR/CBS but are within the EIR/EBS. These frames are delivered as "best effort" without being subject to the SLS. The CEN may drop some or all of these frames based on congestion conditions in the network.

Red

Frames that do not conform to either the CIR/CBS or the EIR/EBS are dropped.

Applying Bandwidth Profiles

Bandwidth Profile can be applied:-

- Per UNI

? Unknown Attachment

- Per EVC

? Unknown Attachment

- Per CoS ID per EVC

? Unknown Attachment

Ingress and Egress

There are two types of bandwidth profiles:

Status

DRAFT

Source(s) and Reference(s)

MEF 10.3 |

Contributor(s)

Reviewer(s)

- Ingress Bandwidth Profile – limits the rates of frames entering the CEN.
- Egress Bandwidth Profile – limits the rate of frames egressing the CEN and thus protecting overload towards the egress CE. Egress Bandwidth profiles are usually not used with Ethernet Line services since those services can be policed at the ingress. They are more commonly used for Ethernet LAN services where multiple UNIs can be sourcing traffic destined for a single UNI.

At a given UNI, at most one ingress BWP can be applied and at most one egress BWP can be applied to a given service frame. This means, for example, that if a BWP is applied to the UNI, then no other bandwidth profiles (in the same direction) can be applied at that UNI.

Egress Bandwidth Profiles (per EVC)

? Unknown Attachment

The Bandwidth Profile Algorithm

The MEF bandwidth profile is a two rate, three color marker (trTCM) that modeled using a dual token bucket algorithm. One bucket, referred to as the 'Committed' or 'C' bucket, is used to determine CIR-conformant, in-profile Service Frames while a second bucket, referred to as the 'Excess' or 'E' bucket, is used to determine EIR-conformant, excess Service Frames.

Token Buckets

? Unknown Attachment

Each token bucket is a bucket of byte 'tokens'. Initially, each token bucket is full of tokens. Each bucket can hold the allowed burst size, CBS, for the Committed bucket and EBS for the Excess bucket. The buckets are refilled by adding tokens at the respective rates (CIR/8 for the Committed bucket and EIR/8 for the Excess bucket). When a bucket is full, no additional tokens are added (but see the discussion about **CF** below).

The conformance for each Service Frame that passes through the policer (on ingress or egress) is determined by comparing the number of bytes in the Service Frame to the number of tokens in the buckets. If the Committed Bucket has at least as many tokens as the number of bytes in the frame, then the frame is **green** and those tokens are removed from the Committed Bucket. If not, then the Excess Bucket is compared. If there are enough tokens in the Excess bucket, then the frame is **yellow** and those tokens are removed from the Excess Bucket. Otherwise the frame is **red**. This is shown in the following diagram.

? Unknown Attachment

? Unknown Attachment

Color Blind Algorithm Skeleton

```


If (Service Frame length < C-Bucket tokens)
    {declare green; remove tokens from C-Bucket}

Else if (Service Frame length < E-Bucket tokens)
    {declare yellow; remove tokens from E-Bucket}

Else
    {declare red}
  
```


Example

100 Mbps UNI services 3 EVCs. Each EVC supports CIR of 30 Mbps and EIR of 20 Mbps.



The subscriber is guaranteed 30 Mbps per service. The sum of CIRs must be \leq UNI speed. However, the Subscriber cannot use the excess bandwidth simultaneously for the 3 services, or else the Subscriber would be trying to send 150 Mbps over a 100 Mbps link.

Figure 1 - Application of Bandwidth Profiles



Frames are marked with their color so that subsequent components of the carrier network can know how to handle them. Although there are several ways to mark the frames, the most common are:

- Green frames are not marked at all. Their lack of a marking indicates that they are green.
- Yellow frames are marked either by setting the Discard Eligibility Indicator (DEI) bit in the S-VLAN tag or by setting a particular IP DSCP value.
- Red frames are dropped so there is no need to mark them.

The bandwidth profile has two additional parameters, the **Color Mode, CM**, and the **Coupling Flag, CF**.

The **Color Mode** can have values **color-blind** and **color-aware**. The color mode is relevant only if the frame has traversed another policer before this one. In **color-blind** mode, each frame starts green and is then marked only according to operation of this policer. In **color-aware** mode, the color marking of the previous policer is the starting point for this frame—an incoming green frame can be green, yellow, or red, but an incoming yellow frame can only be yellow or red.

The **Coupling Flag** indicates whether tokens that would overflow the C-bucket should be added to the E-bucket. If **CF=1**, then a token that would overflow the C-bucket (i.e. a token that would be added except that the C-bucket is full) is, in effect, added to the E-bucket (unless it is full). If **CF=0** then this doesn't happen. CF has a negligible effect when the policing is **color-blind**. When **CF=1** and when operating in **color-aware** mode, more yellow Service Frames are allowed into the service provider's network.

When **CF=0**, the long term average bit rate of Service Frames that are declared Yellow is bounded by EIR. When **CF=1**, the long term average bit rate of Service Frames that are declared Yellow is bounded by CIR + EIR depending on the volume of the offered Service Frames that are declared Green. In both cases the burst size of the Service Frames that are declared Yellow is bounded by EBS.

The diagram at the right shows the dynamics of the token buckets. The buckets start full. Then there is a burst of six frames at a high rate. The first two frames can be taken from the C-bucket so they are green. The C-bucket doesn't have enough tokens for frame 3 so it is taken from the yellow bucket (and marked yellow). Frame 4 can be taken from the green bucket but frame 5 is taken from the yellow bucket. There aren't enough tokens in either bucket for frame 6, so it is red and discarded. At the end of the burst, both buckets are heavily depleted—although not quite as much as shown since they are being refreshed while these frames were coming in (but the assumption in this example is that the frames are coming in much faster than the buckets are being refreshed).

Effect of Each Parameter Setting

Setting the appropriate values for CIR and EIR is often relatively easy. Since CIR is a guaranteed bandwidth it should (usually) be set to exactly the agreed value or a little more. For example, if a customer contracts for 20 Mbps guaranteed service with it may be wiser to set CIR to be 21 or 22 Mbps to avoid occasional slips. Since EIR is a *best efforts* value, it can be set to the agreed on value.

Setting the burst size is a bit more complicated. The requirements of the application must be taken into account. For non-bursty traffic such as circuit emulation and other traffic that needs low [frame delay variation](#), smaller burst sizes are appropriate (maybe 2*Maximum Frame Size). For TCP/IP traffic, TCP windowing creates larger bursts, so a larger burst size is needed to achieve reasonable throughput. Note that CBS and EBS **must** be at least as large as the maximum frame size (MFS).

There are two special cases:

- EIR=0, EBS=0, CIR>0, CBS>0 – If no excess bandwidth is needed
- CIR=0, CBS=0, EIR>0, EBS>0 – If no guaranteed bandwidth is needed, all traffic is best efforts (yellow)

If a rate is 0 its bucket should be 0. If a rate is >0, its bucket **must** be MFS.

Applications

Consider an enterprise with 3 locations: An HQ and 2 branch offices as shown in figure 1 at right.

The enterprise asks for EP-LAN service of 100 Mbps guaranteed. An EVC Ingress BWP could be set to provide

CIR=100 Mbps

CBS=50,000 Bytes

EIR=0

EBS=0

CM=0

CF=0

If the enterprise would like specific delay guarantees to its VoIP application, then the SP may offer 2 CoS ID, letting the CE perform the priority marking. The VoIP service is assumed to require 10 Mbps. High marked with PCP=5 and Low marked with PCP=0. All other PCP values could be blocked.

Attribute	High CoS	Low CoS
CIR	10 Mbps	90 Mbps
CBS	3,000 bytes (very few bursts on VoIP)	50,000 bytes
EIR	0	10 Mbps (this enables using the VoIP 8W for the low priority traffic when VoIP is not being used)
EBS	0	50,000 bytes
CM	0	0
CF	0	0

Now two ingress BWPs should be set - one per each CoS ID.

Example(s)

Related and Further Reading

[Bandwidth Profiles for Ethernet Services](#)

Categories