Frame Metering in 802.1Q

Version 01

Stephen Haddock

January 15, 2013

Overview

- Frame metering introduced in 802.1ad-2005 in conjunction with the DEI bit of the S-VLAN tag.
  - The DEI bit was introduced to allow “color” (drop eligibility) marking in Provider Networks.
    - Use of the DEI bit was extended to C-VLAN tags (replacing the old Canonical Format Indicator (CFI) bit) in 802.1Q-2011.
  - Frame metering was introduced to provide a framework for how the drop eligible parameter can get set and how it can be used.
    - The specification is minimal (!!!) to allow maximum flexibility for implementations to be optimized for different markets.
    - The primary motivation was to enable implementation Metro Ethernet Forum (MEF) Bandwidth Profiles, but to recognize that other bridging environments may have differing objectives/requirements for metering.

Specification in 802.1Q

- 6.9 Support of the EISS
  - Specifies “color marking” in VLAN tags.
    - Drop eligibility is encoded in the DEI bit and/or the PCP field.
    - Marking in Ingress frames can set the drop eligible parameter.
    - Drop eligible parameter used to mark egress frames.
- 8.6.5 Frame Metering
  - Specifies ingress meters that can set the drop eligible parameter associated with a frame.
    - Details on subsequent slides
- 8.6.7 Queue Management
  - Specifies that drop eligible frames may have a higher probability of discard than non-drop-eligible frames.

What is (and is not) specified

- The meters in 8.6.5 are ingress meters.
  - All frames going through a given meter come from the same reception Bridge Port.
  - The meters are located after all filtering (including active topology enforcement, ingress VID filtering, FDB filtering, and egress VID filtering). A frames is not counted by the meter if the bridge will not attempt to forward that frame.
  - The meters are located before the queues, so they are most naturally used for ingress rate policing.
  - 802.1Q does not specify any egress meters for egress rate policing, however it does specify scheduling algorithms (enhanced transmission selection, credit-based shaper) in 8.6.7.
- Meters are not required.
  - Zero or more meters per reception port are allowed.

What is (and is not) specified - 2

- Assignment of frames to meters is very flexible.
  - All frames from a given reception port go to one meter, or
  - may classify frames for separate meters based on any combination of DA, SA, VID, priority. (Classification based on other factors such as layer 3 parameters is not precluded, but not specifically mentioned because out of scope of 802.1Q.)
  - For example, Provider Bridges used to deliver MEF specified services, the assignment of frames to meters could be by port, by VID, or by VID and priority.
- The metering algorithm is not specified.
  - The MEF algorithm is referenced and recommended.
  - Having a single bit for drop eligibility allows two colors for frames that are forwarded (typically green and yellow). A three color metering algorithm is supported assuming frames with the third color (typically red) are always discarded by queue management.

What is (and is not) specified - 3

- What Queue Management (8.6.7) does with color indications is very loosely specified.
  - Implementer free to choose any queue management algorithm.
  - e.g. tail-drop, tail-drop with different color thresholds, Random Early Discard, Weighted Random Early Discard, …)
  - Only requirement is that drop eligible frames have a higher or equal probability of being discarded relative to non-drop eligible frames.

Color ≠ Priority

- Priority is an inherent characteristic of a frame, determined by the contents of the frame.
  - Different priorities can be mapped to different traffic classes, queued in different queues, and do not need order to be maintained between frames of different priorities.
- Color is not an inherent characteristic of a frame, and is determined based on the arrival time at a meter relative to the history of arrival times of previous frames.
  - Frames within a traffic class can be marked different colors but they should not be queued in different queues because order does need to be maintained.

The MEF Bandwidth Profile (metering) algorithm

MEF interpretation of color

MEF Bandwidth Profiles use three color indications.:

Green — "Committed" frames. Service Level Objectives such as frame loss rate, delay, delay variation, etc. are applicable to these frames. In theory, with proper policing at the edge of the network and proper allocation of buffer and bandwidth resources within the network, it is possible to guarantee lossless and timely delivery of all committed frames.

Yellow — "Excess" frames. Service Level Objectives are not applicable to these frames. Excess frames are delivered on a "best effort" basis.

Red — "Non-conformant" frames that are always discarded.

MEF Bandwidth Profile Algorithm

- Is a Two Rate Three Color algorithm.
  - Can be configured as a Single Rate Three Color algorithm, or as a Single Rate Two Color algorithm.
- Operates in “Color Aware” or “Color Blind” mode.
  - Color Aware: Algorithm considers the color indication of incoming frames. Incoming frames without a color indication get a default color prior to entering the meter. Frames are never “promoted”; An incoming yellow frame is never changed to green.
  - Color Blind: Algorithm ignores the color indication (if any) of incoming frames. Effectively all incoming frames are assumed to start out green.
- Is a Token Bucket algorithm.
  - Guarantees that over any time interval T, the amount of data declared “conformant” is less than or equal to the Information Rate times T plus the Bucket Size.

Token Bucket Basics

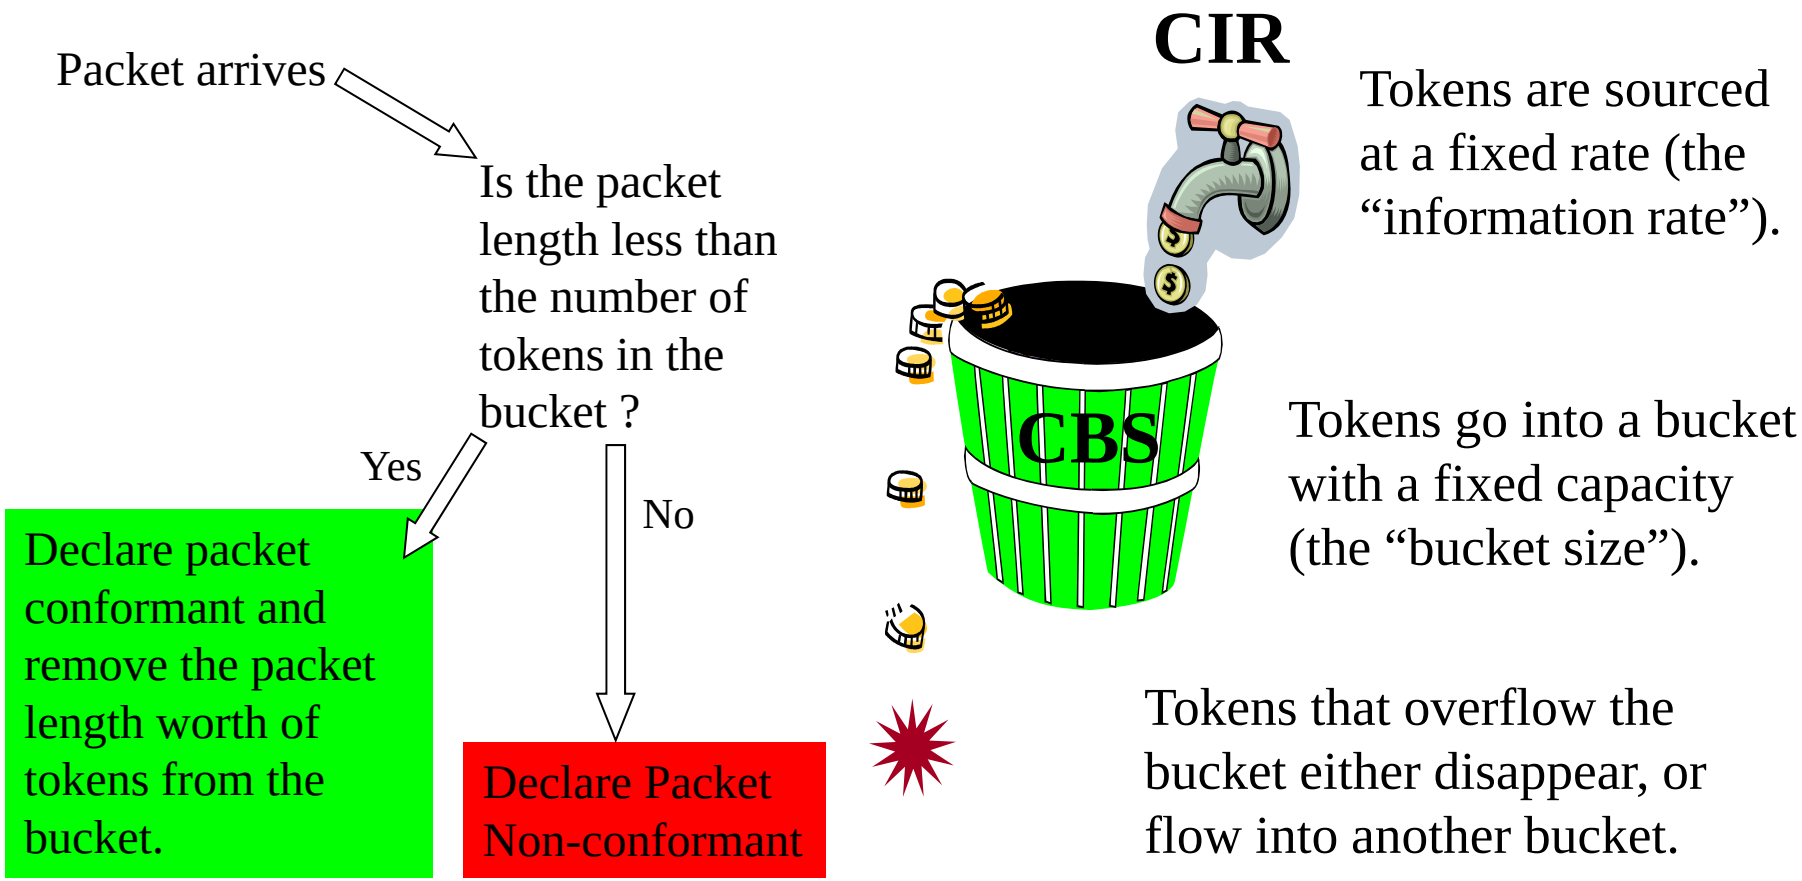

A Single Rate Two Color Meter/Marker

Single Rate Three Color Marker

CIR

(Overflow falls into Yellow bucket)

(Overflow lost)

Single Rate Three Color Marker

Color Blind Mode (CB)

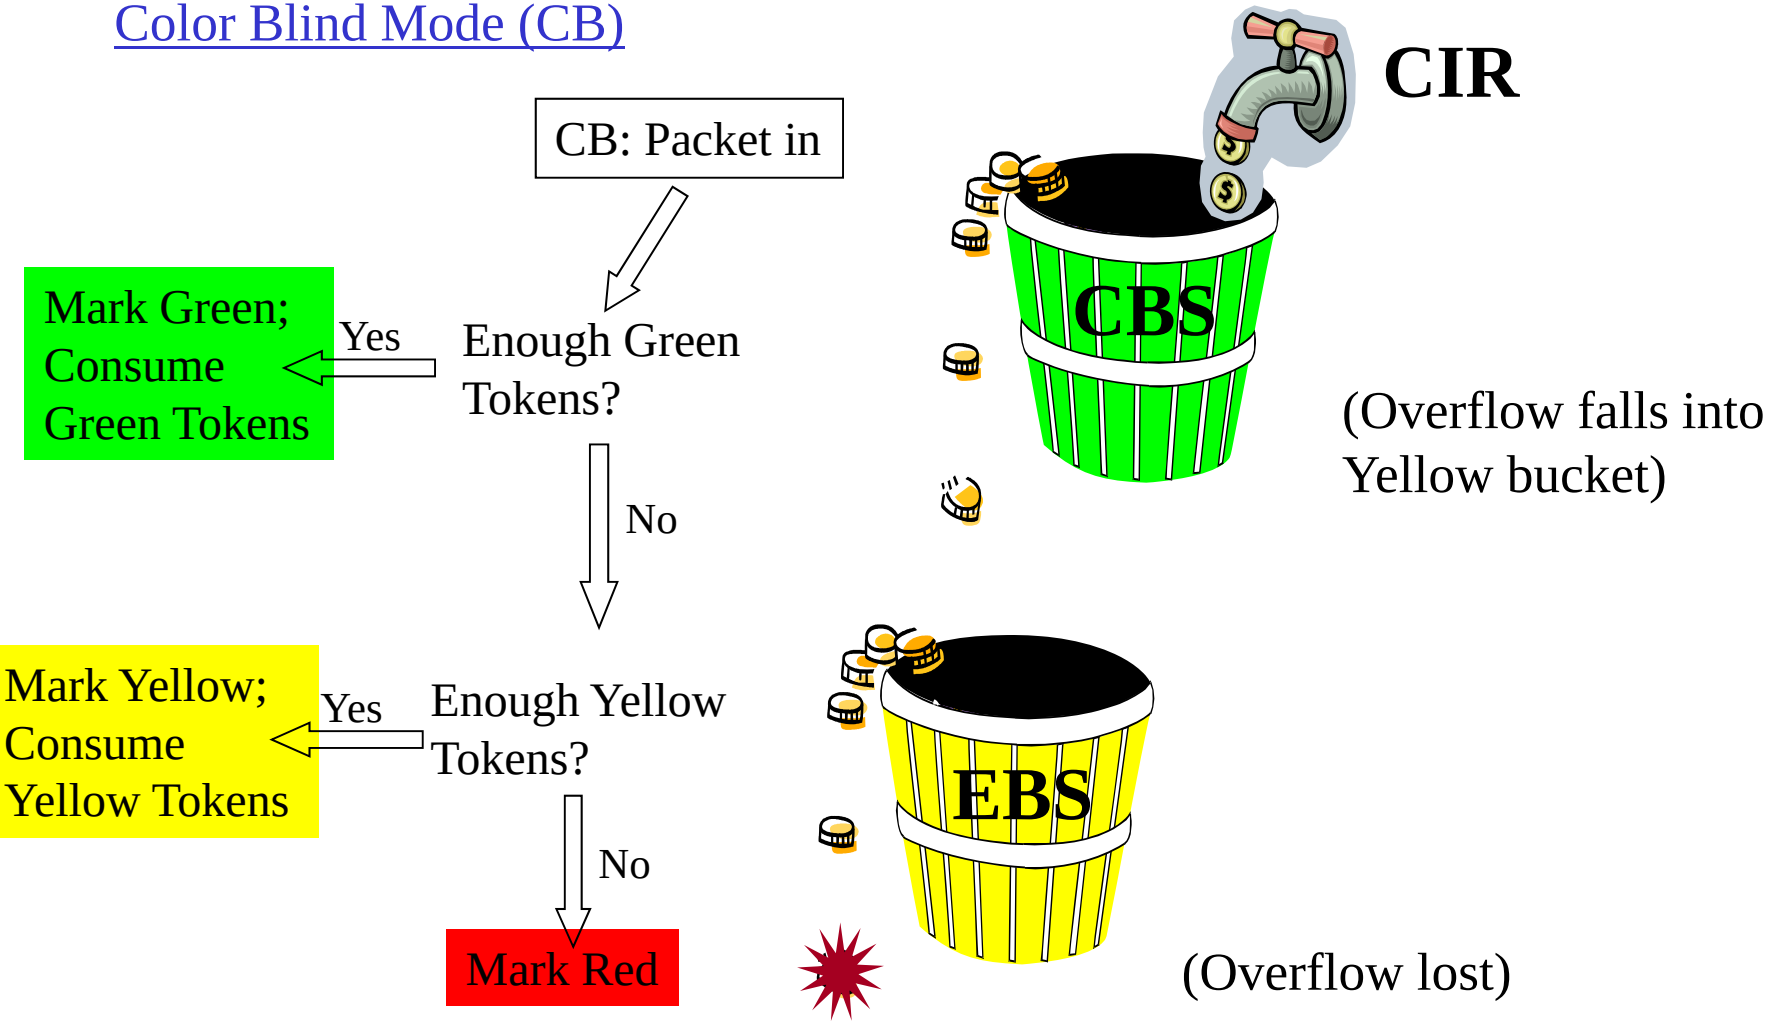

Single Rate Three Color Marker

Color Aware Mode (CA)

CA: Green Packet in

Mark Green; Consume Green Tokens

Yes

Enough Green Tokens?

No

CA: Yellow Packet in

Mark Yellow; Consume Yellow Tokens

Yes

Enough Yellow Tokens?

No

CA: Red Packet in

Mark Red

CIR

(Overflow falls into Yellow bucket)

(Overflow lost)

Single Rate Three Color Marker (RFC 2697)

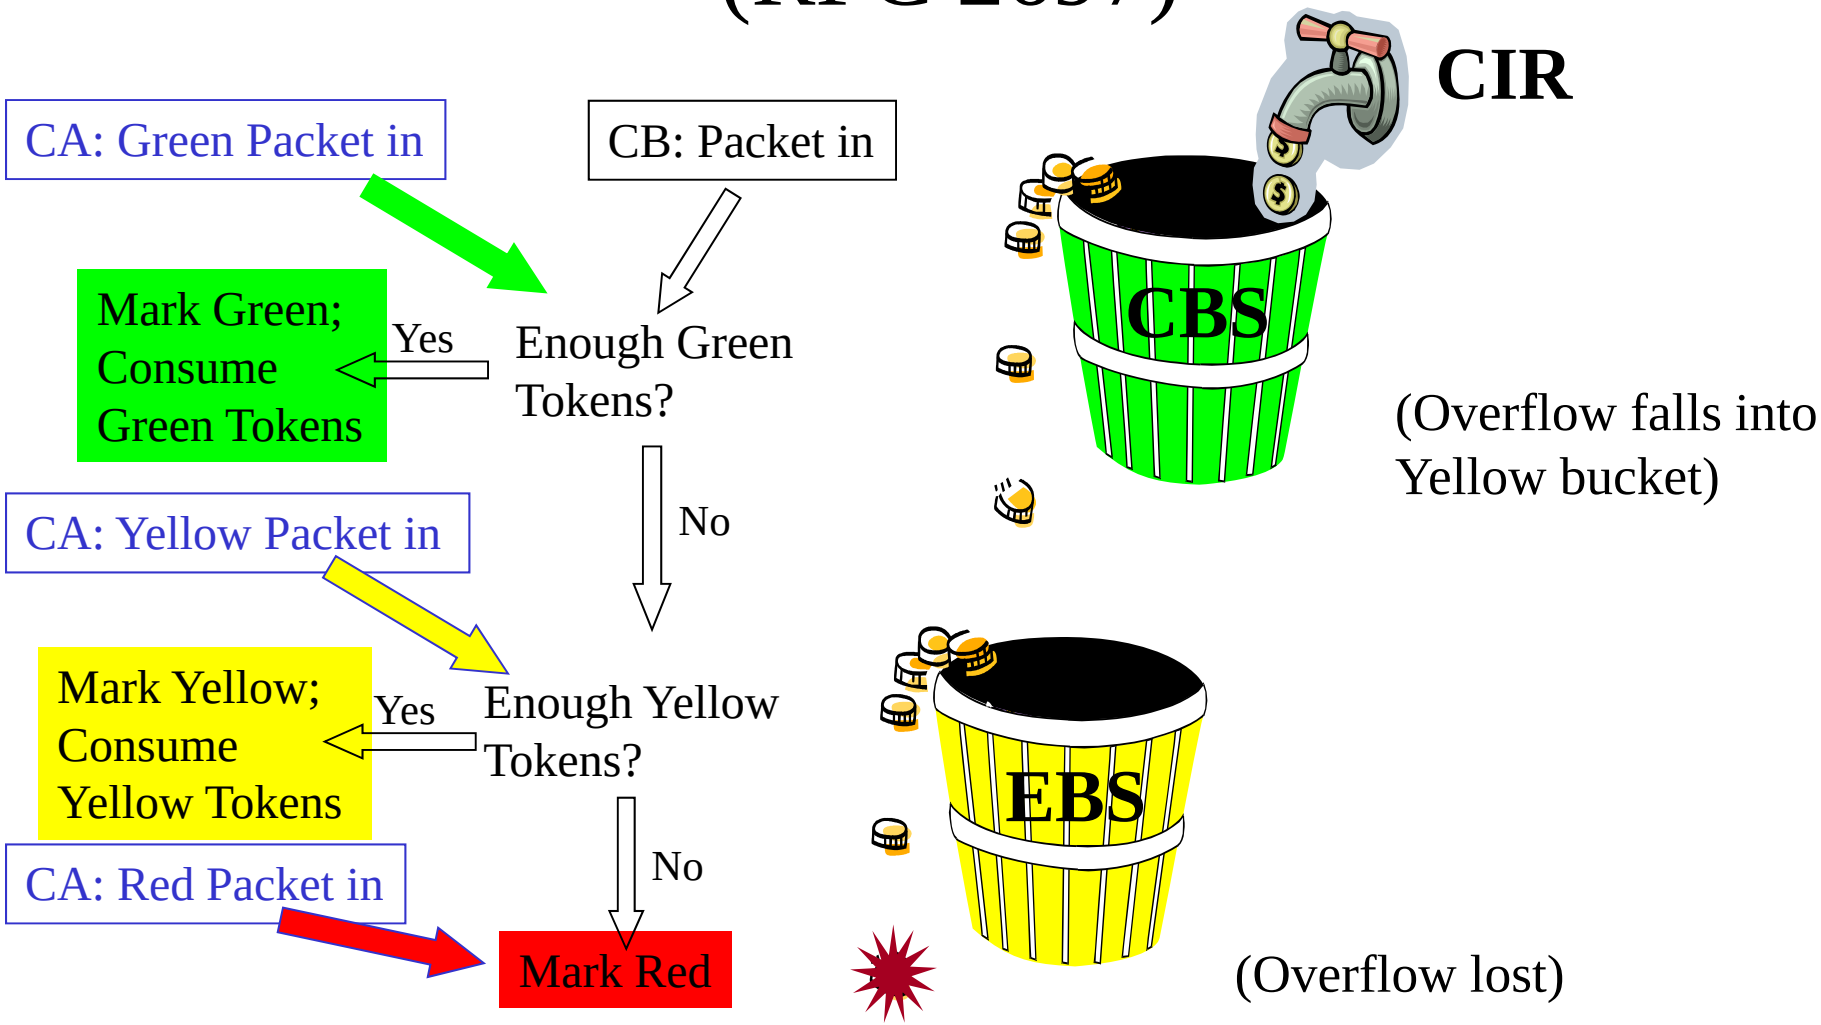

Properties of SRTCM

1. Only green packets consume green tokens;
   Only yellow packets consume yellow tokens.
   - Yellow traffic cannot interfere with green traffic.
2. Overflow from the Committed bucket goes into the Excess bucket.
   - Color blind mode: CBS can strictly control bursts of committed traffic, while excessively bursty traffic is accommodated on a best effort basis.
   - Color aware mode: In the absence of green traffic, the full rate is available to yellow traffic.
   - Since tokens only go into Excess bucket when the Committed bucket is full, any conformant green traffic can be accommodated at any time.

Two ~~Single~~ Rate Three Color Marker

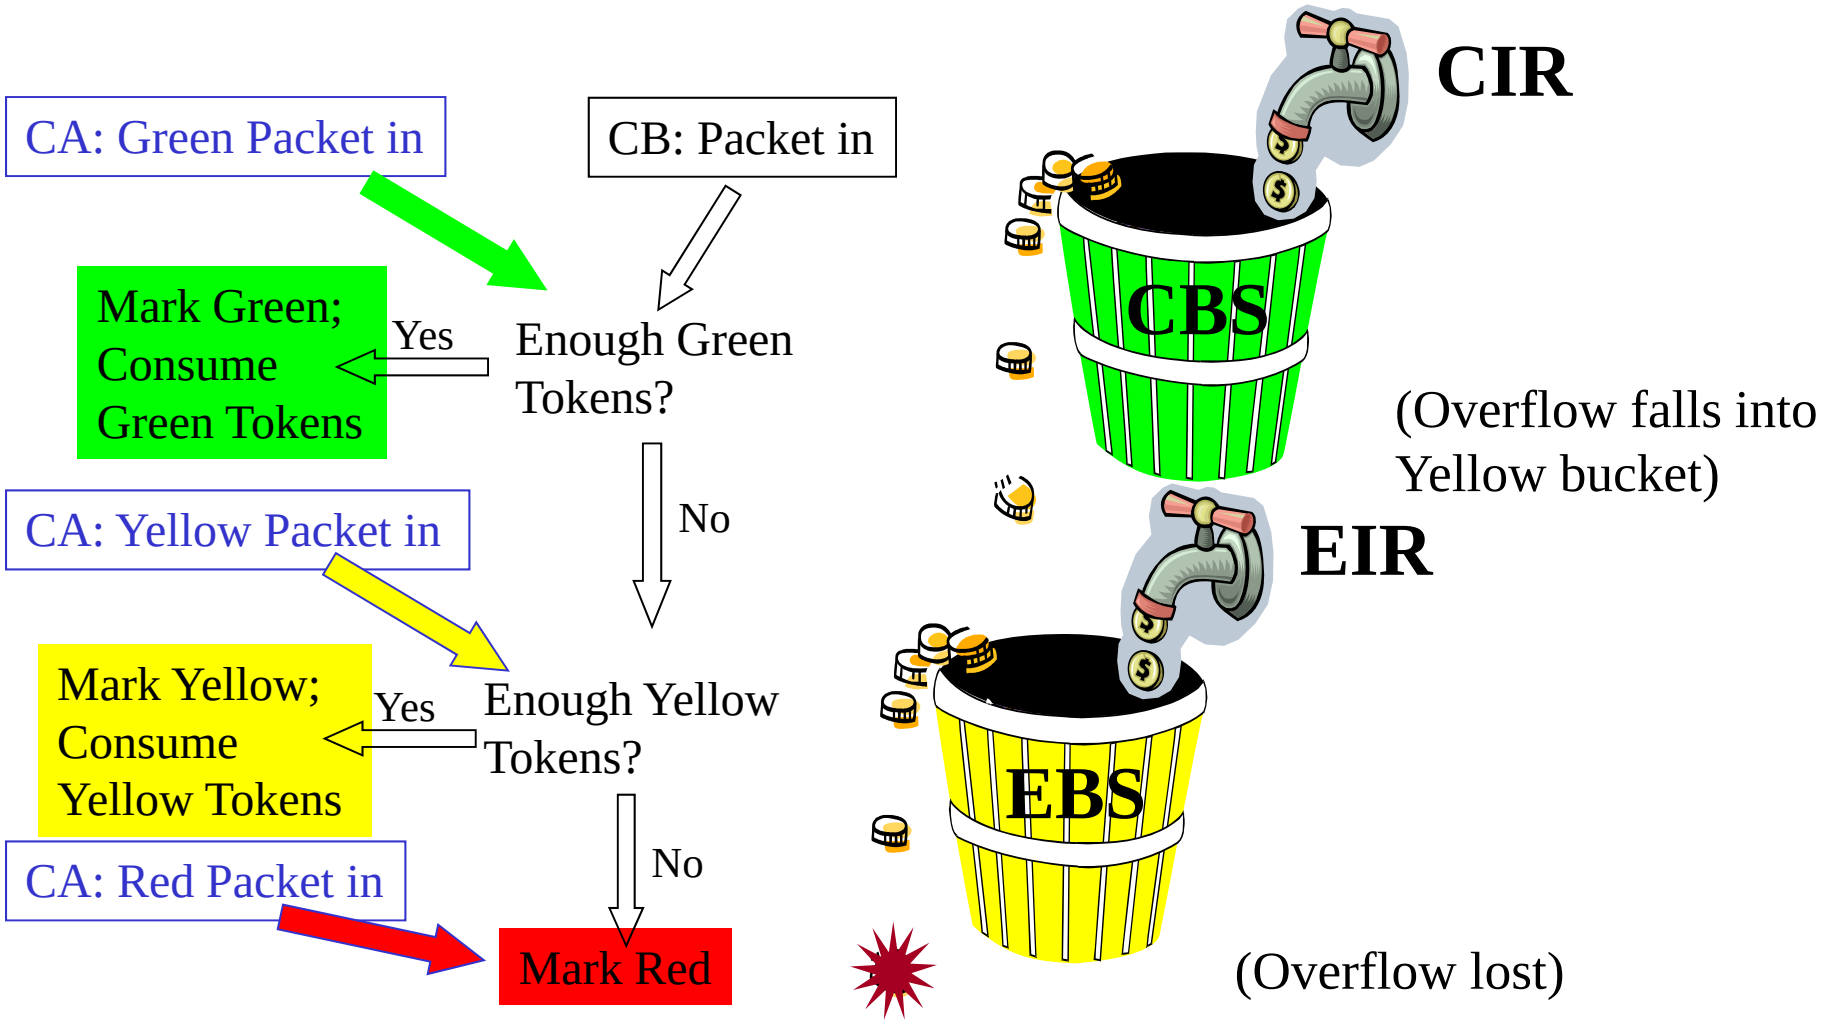

Varieties of TRTCM

- RFC 2698:
  - Uses “peak” token bucket instead of “excess” token bucket.
  - PIR = CIR+EIR and PBS = CBS+EBS.
  - Frames declared green consume tokens from both the committed and peak buckets; frames declared yellow consume tokens only from the peak bucket.
  - Burst of yellow frames can deplete the peak bucket, so subsequent green are declared red even though there are tokens available in the committed bucket. Most people consider this a significant bug.
- RFC 4115:
  - Uses committed/excess bucket model (not committed/peak) so resolves the RFC 2698 bug.
  - Does not allow tokens overflowing the committed bucket to go into the yellow bucket. Therefore cannot be configured as a single rate three color marker, and does not allow yellow traffic to utilize unused green bandwidth.
- MEF Bandwidth Profile Algorithm:
  - Uses committed/excess bucket model.
  - Configuration variable (“coupling flag”) allows tokens overflowing the committed bucket to go into the yellow bucket or to simply disappear.

TRTCM Comparison

Feature	Significance	RFC 2697	RFC 2698	RFC 4115	MEF
All incoming green packets conformant to CIR/CBS profile are guaranteed to be marked green.	Makes it possible to provision a network that guarantees a level of service for green traffic.	Yes	No	Yes	Yes
Yellow traffic may reclaim unused green bandwidth when green token bucket is full.	Enables higher utilization of available network bandwidth, without sacrificing ability to guarantee service for green traffic.	Yes	Yes	No	Yes
Allows yellow burst limit To be set > or < green limit	Simple, Flexible configuration	Yes	No	Yes	Yes
2rate3color marker	Simple, Flexible configuration	No	Yes	Yes	Yes
1rate3color marker when EIR = PIR-CIR = 0	Simple, Flexible configuration	Yes	Yes	No	Yes
1rate2color marker when EBS = PBS-CBS = 0	Simple, Flexible configuration	Yes	Yes	Yes	Yes

Thank You

.