STATELESS RESOURCE SHARING (SRS)

VIA PER PACKET VALUE (PPV) MARKING

(NEW SERIOUS METHOD FOR ENSURING QOS CHARACTERISTICS)

2016-05-25, IEEE P802.1Qcr, Budapest
Balázs Varga (A), Szilveszter Nádas, János Farkas

INTRODUCTION
MOTIVATIONS

- Trigger: 802.1 Qcr PAR
  - Asynchronous Traffic Shaping (ATS) (urgency based forwarding)
  - Extensions to the QoS related functions
- Targets:
  - Leverage the extensions introduced by ATS for more
  - Stateless Resource Sharing (SRS) is a delay aware resource sharing framework:
    - requires enhanced queueing and operations upon arrival, which are also needed for ATS
    - does not need stream awareness at the bottleneck (stateless)
  - Explore whether ATS math can be re-utilized
    - Can SRS and ATS be combined (create a stream unaware behavior with similar delay guarantees)?
    - Is it possible to apply “ATS math” for several congestion levels?

MARKING TRAFFIC
SHADE = DROP PREFERENCE

Transport prefers dark shades

- QoS encoding
  - Priority: color, e.g. blue (business as usual)
  - Drop preference: dark – bright (new stuff)

Packet Value

- Example:
  - VBR sources with 6 Mbps max traffic
  - 3 service class: Bronze, Silver, Gold
    (wanted BW ratio: 1:2:3 if congestion occurs)
- Bottleneck:
  - Drop based on shade of blue: the darker the color the more unlikely to be dropped
  - E.g. Bottleneck 6 Mbps:
    - Only the darkest blue gets through
    - 1Mbps Bronze, 2 Mbps Silver, 3 Mbps Gold
  - E.g. Bottleneck 12 Mbps:
    - The two darkest shades get through
    - 2Mbps Bronze, 4 Mbps Silver, 6 Mbps Gold

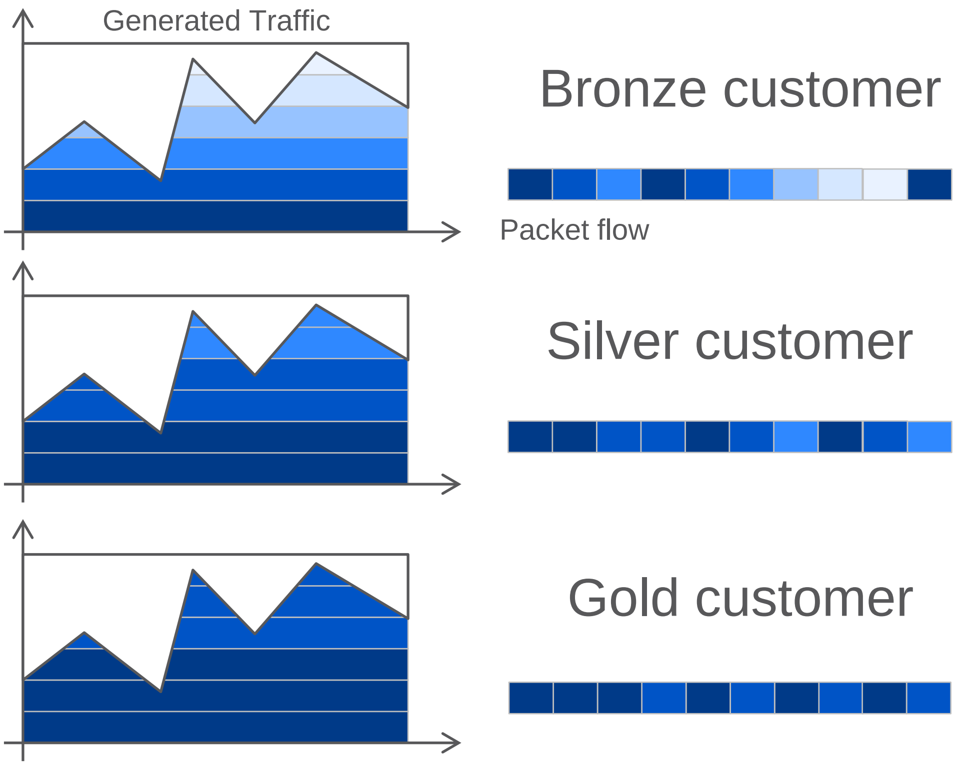

QUEUING
BASED ON COLOR SHADE

Darker packets kick off brighter packets

- New queuing rules
  (AQM: Active Queue Management)
  - FIFO queuing
    - avoid out-of-order delivery
  - Shade based dropping
    - Dark shades kick-off bright shades if buffer is full

- IMPACT (behavior in case of congestion):
  - Bright shades are dropped
  - Darker shades are preferred
  - The more severe the congestion is the darker shades get dropped

- Example (theoretical solution):
  - No congestion
    - New packet queued

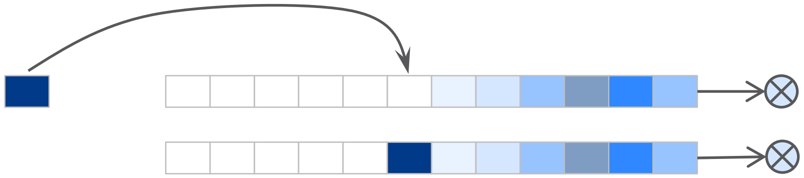

  - Congestion
    - Drop the packet with the brightest shade
      (if no brighter packet → drop new packet)
    - Shift packets
    - Put new packet at queue end

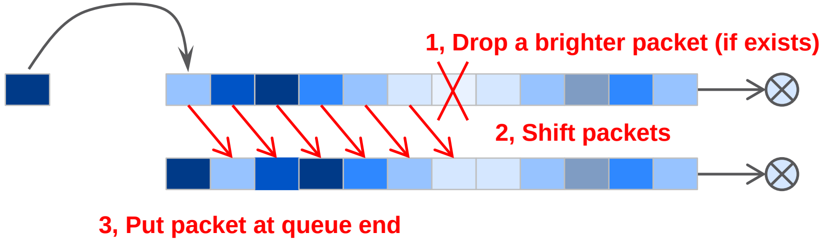

CONGESTION
IMPACT ON TRAFFIC

Congestion level defines what shades can get through

- Congestion results in packet drop:
  - Darker packets kick-off brighter ones
  - Based on the shade: The brighter the shade, the more likely to be dropped during congestion
- Congestion level:
  - Can be defined by a „congestion shade” (All brighter packets dropped and all darker packets are forwarded)

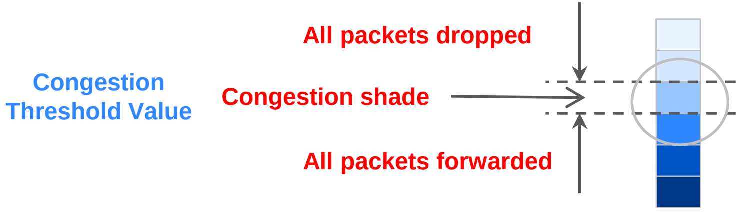

- Severity of the congestion:
  - Depends on the actual traffic over-load
  - Depends on shade distribution of actual traffic

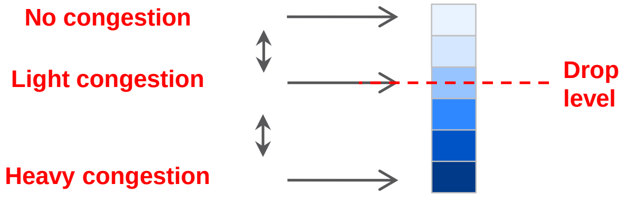

- Throughput can be designed in advance
  - By a ”-shading schema” that defines the pre-defined BW ratio between flows
  - No traffic situation specific pre-configuration is needed at the congestion point !!! :
  - BW share is automatically controlled by the shades of packets participating in the congestion situation

Note: Explicit feedback is possible to sources if needed what level of congestion they face during transport

CONGESTION
EXPECTED THROUGHPUT

Congestion level defines what shades can get through

- 3 types of groups can be defined based on what happens with packets:
  - forwarded: darker than congestion shade
  - congestion shade: may be forwarded, may be dropped
  - dropped: brighter than congestion shade

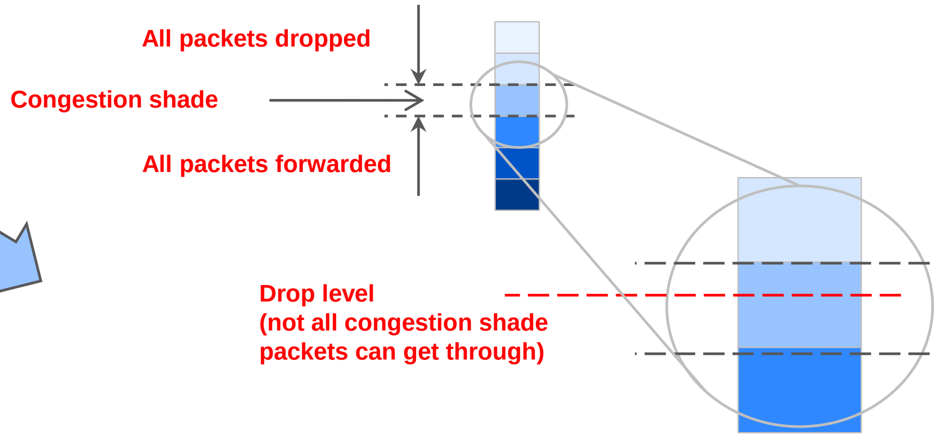

- Throughput of a source depends on
  - ratio of packets darker than "congestion shade"

  AND
  - what happens with "congestion shade" packets

- Within the Congestion shade:
  - No ratio between sources can be ensured
  - Situation depends on source parameters (e.g. how aggressive the source is, etc.)
  - Same type of sources result in equal share of BW within "congestion shade"

MARKING SCHEMA
SETTING THE RULES

Marking schema defines BW share in advance

- Shading schema
  - The shading rules on packets of a flow/source/etc.
  - Grant per shade defines how much traffic can use a given shade
  - For any congestion scenario the (resulting) BW of participating flows can be calculated in advance
- Implementation (theoretical)
  - Many options, e.g. Multiple Leaky Buckets

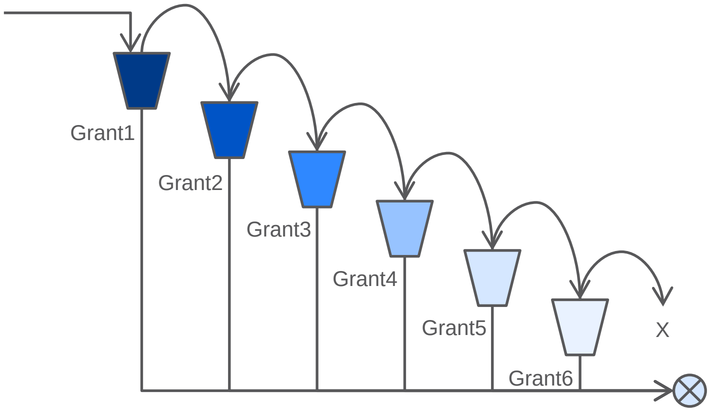

Grant per shade

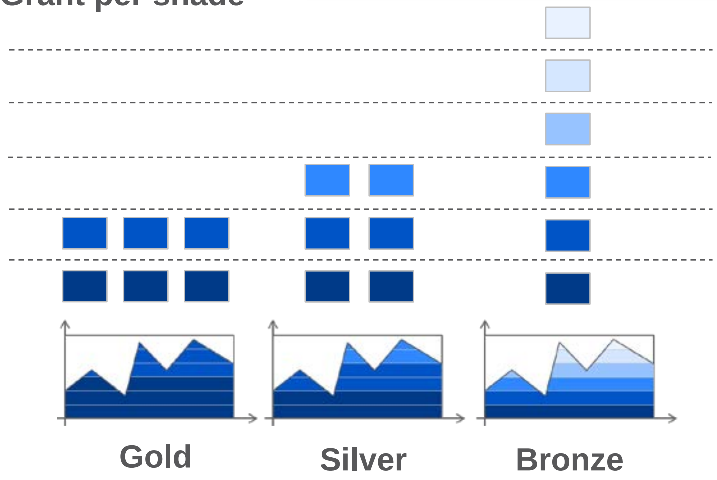

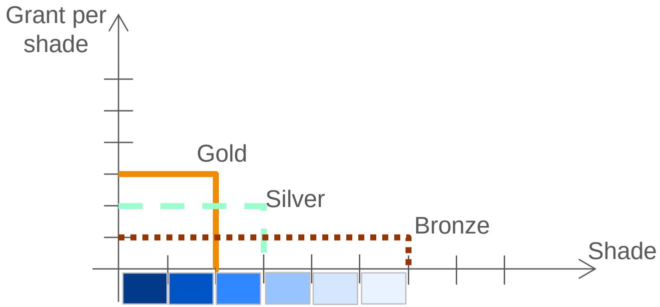

SUMMARY
UNIQUE CHARACTERISTICS

SRS via PPV at-a-glance

- SRS essentials:
  - Share of available BW between flows is encoded in the packets (i.e. shade of a packet)
  - Shade based dropping in queues does not need traffic situation specific pre-configuration at the congestion point(s)
  - BW share is automatically controlled by the shades of packets participating in congestion situation
  - Accuracy of BW share depends on the number of shades used
    (Note: predefined ratio exactly ensured if drop level = border between shades)
  - Explicit feedback is possible for systems that need congestion notification
    (i.e. what level of congestion they will face during transport is known from the congestion shade)

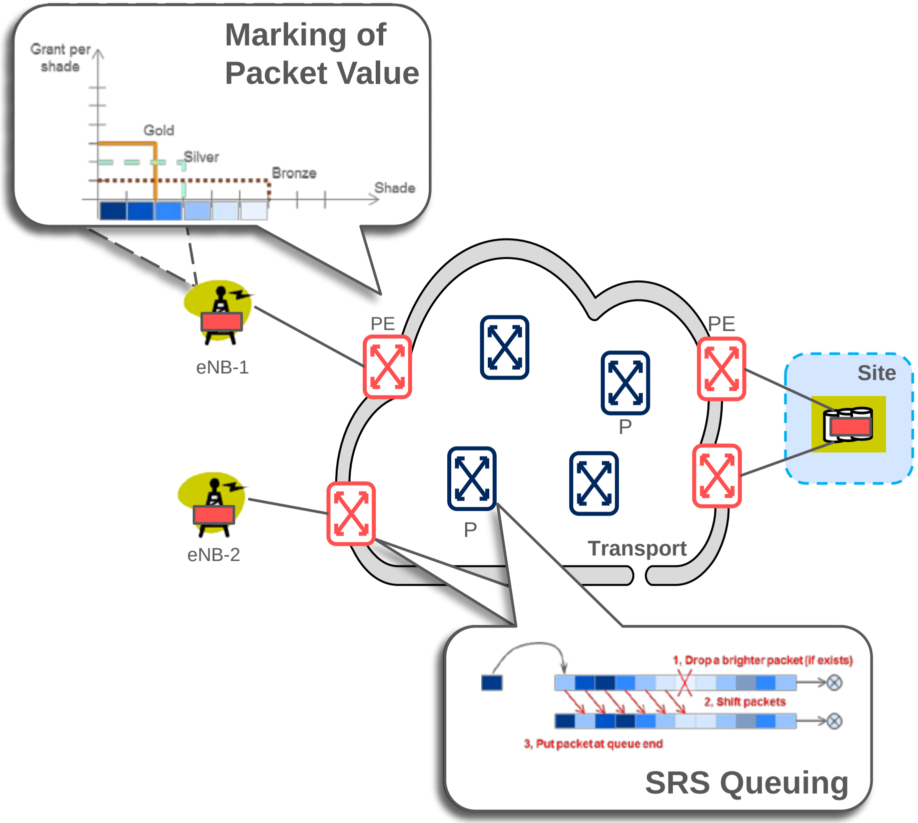

USER GUIDE
SRS HINTS (ONE EXAMPLE)

Introducing SRS in real networks with minimal effort

- SRS
  - can be used in addition to existing PQ, WFQ, CBWFQ, etc.
  - can be used to replace existing queuing techniques
- Service Provider target
  - Maximize network throughput
  - Relative fairness between customer groups
- SRS can be used to achieve these targets without complete change of SP’s network design, dimensioning, etc.

- Example
  - SP owns a multiservice network (e.g. Internet service, Business VPN, IPTV, VoIP)
  - SP has already introduced 4 QoS classes in its network (best-effort [0], prio-data [1], video [4] and voice [5])
  - SP intends to distinguish service classes within the predefined QoS classes
  - Dimensioning ensures that QoS classes of 1, 4 and 5 are not congested in the network → no need for SRS in those QoS classes
  - Best-effort traffic is highly elastic and causes congestion in the network regularly. Overdimensioning is not an option for the SP.
  - SRS can be used to control congestion in best-effort QoS class ONLY

USER GUIDE
SRS HINTS (ONE EXAMPLE)

- Marking schema
  - How many color-shades?
    The more color-shades are used, the more precisely the targeted ratio is achievable at any congestion scenario
  - What BW increments between color-shades?
    Depends on the source set (service portfolio)
  - In this example:
    - All sources: 6 Mbps (peak rate)
    - Color-shade per 1 Mbps BW increments
    - 3 service classes (BW ratio 1:2:3)

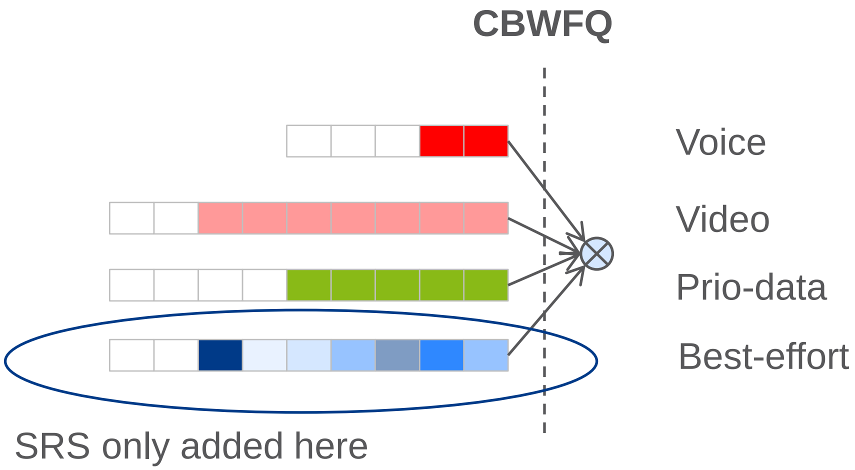

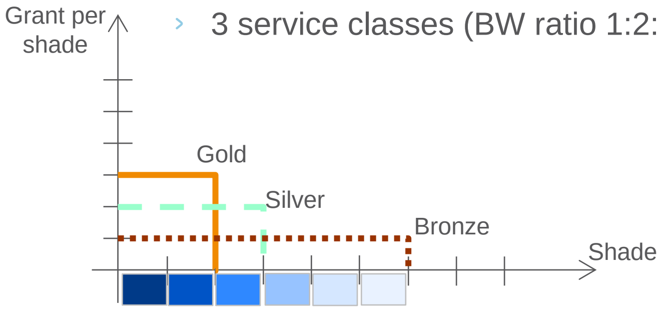

FURTHER USE-CASES
WHAT ELSE CAN BE ACHIEVED BY SRS?

- Other targets
  - Delay Class: Dealing with delay bounded traffic
    - Adding further marking to packets
    - The Packet Value (PV) is a scalar, with preferably many bits to represent it. The larger the value space the more precision resource sharing can be expressed and enforced.
    - The other value attached to packets at the edges is the Delay Class. Unlike the PV that is scalar, we allow only a few distinct Delay Classes. Nodes shall transmit packets in various Delay Classes according to operator-defined maximum delay.
    - Packet Values and Delay Classes are orthogonal.
  - Feedback on congestion level
    - Congestion Threshold Value represents the congestion situation
  - Load balancing
    - Congestion Threshold Value on the alternatives should be the same. If not then moving flows from a more congested alternative to a less congested one will result in the transmission of extra PV.
  - Resource Balancing
    - In case of resource balancing it is not the traffic that is balanced across two given resources, but the resources are balanced under two groups of traffic flows.
  - Network Value maximalization
    - load balancing and resource balancing solutions can be generalized to the network level

IMPLEMENTATION ALTERNATIVES

- Ideal way-of-working
  - Congestion
    - Kick-off the packet with the brightest shade (if no brighter packet → drop new packet)
    - Shift packets behind dropped packet
    - Put new packet at queue end

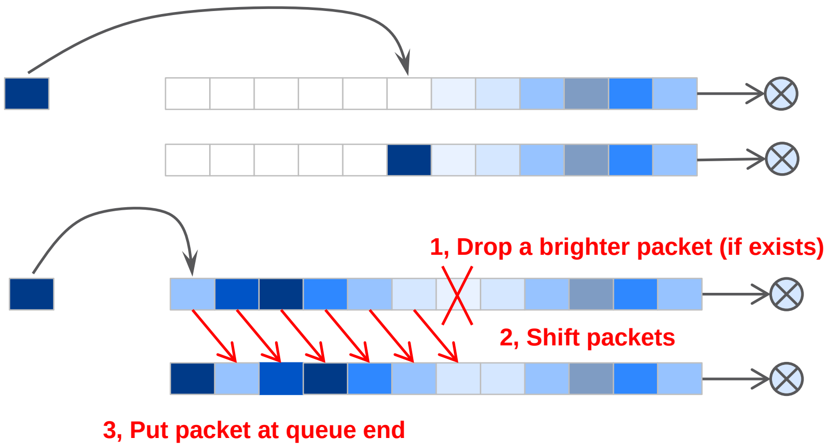

- Achieving same characteristics, but by packet dropping upon packet's arrival
  - Alternative1 – Multi-level drop thresholds
    - A way to reuse existing AQM algorithms is to quantize Packet Values to a few levels (e.g. DSCPs) and configure Weighted RED [13] profiles for each level. This can approximate the behavior of the algorithms above on existing network equipment.
  - Alternative2 – Packet drop probability function
    - Define a packet drop probability function with its inputs being the queue length and the PV of the incoming packet. More generally, it is possible to adapt, e.g., the PIE algorithm to also take into account the PV of the incoming packets when determining the packet drop probability.

ALTERNATIVES
SOME SIMULATION RESULTS

- Method:
  - Ideal SRS implemented
- Scenario
  - 3 types of traffic sources:
    - Gold (TCP),
    - Silver (TCP) and
    - Background (UDP)
  - Traffic shares the same buffer
- Target:
  - Protect Gold and Silver TCP flows from Background UDP flows, while achieving their desired BW share

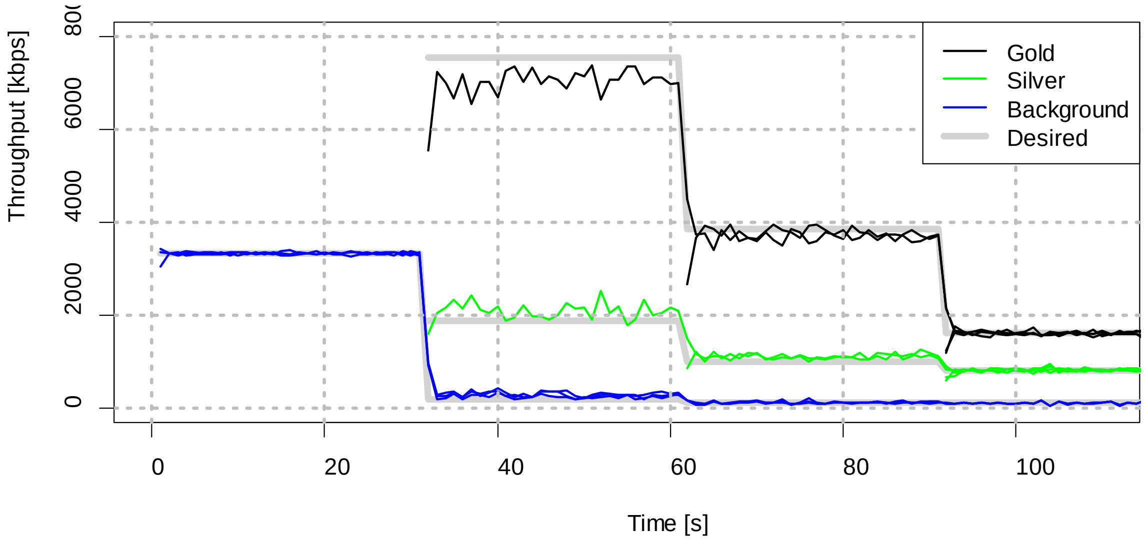

ALTERNATIVE APPROXIMATION – WEIGHTED RED

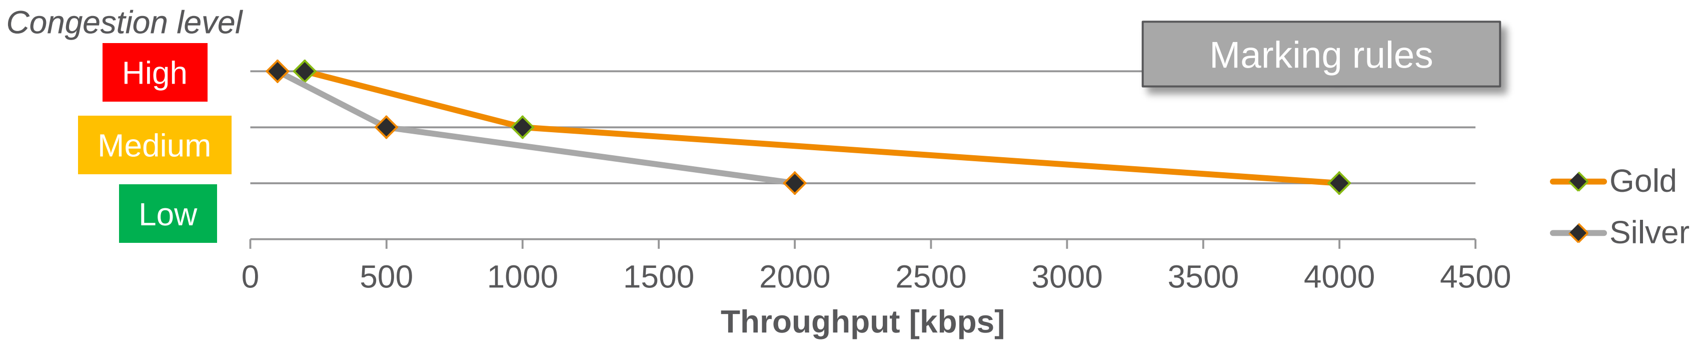

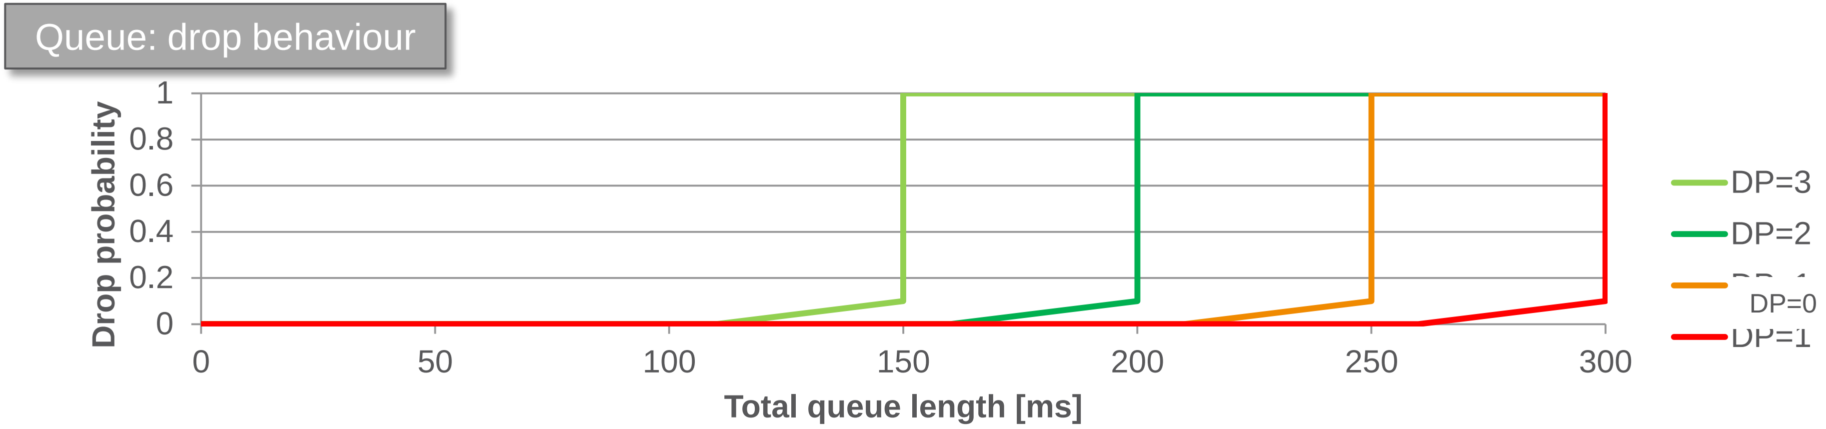

ALTERNATIVE SIMULATION – WEIGHTED RED

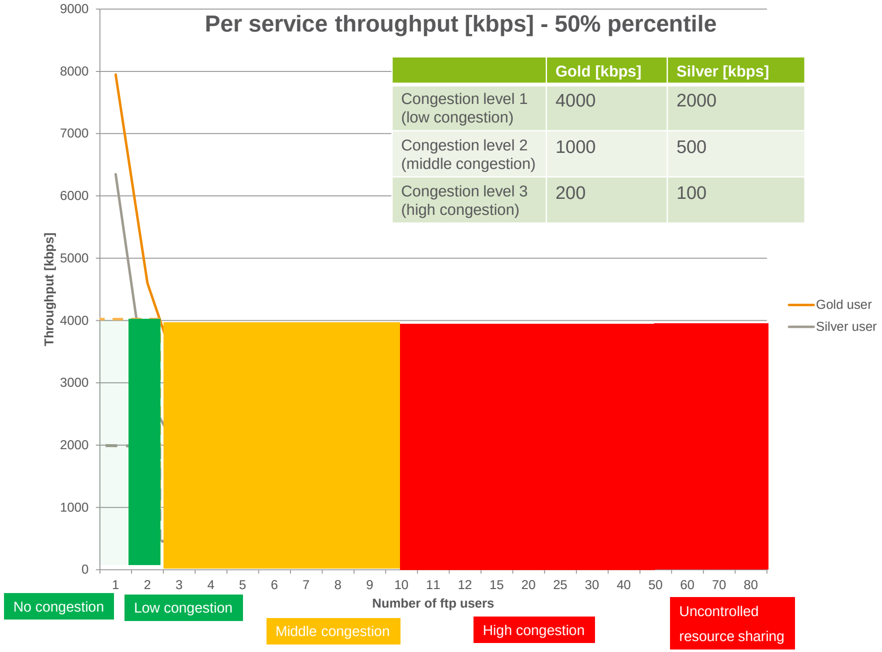

	Gold [kbps]	Silver [kbps]
Congestion level 1 (low congestion)	4000	2000
Congestion level 2 (middle congestion)	1000	500
Congestion level 3 (high congestion)	200	100

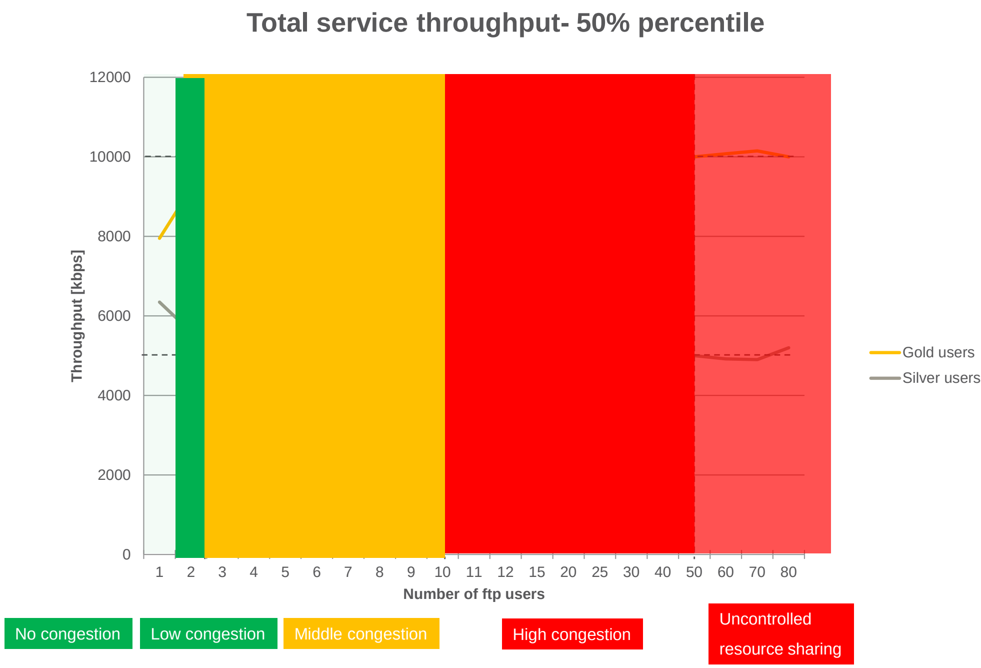

CONCLUSION
PROPOSAL

- P802.1Qcr ATS is about to introduce an enhanced queuing scheme to 802.1Q bridging
- The new queueing scheme could be leveraged for stateless resource sharing with a number of benefits
  - Marking remains at network edge
  - Bandwidth share is automatically controlled
  - No need for pre-configuration at the congestion point(s)
  - Explicit feedback is possible for systems that need congestion notification
  - Universal, simple and flexible control via shading-schemes (adapts to any provider preferences)

ERICSSON