
cisco Live!
Let's go
#CiscoLive

The bridge to possible

Catalyst 9000 Switching QoS Deep Dive

Kenny Lei
Technical Marketing Engineer (TME)
BRKENS-2096

Cisco Webex App

Questions?

Use Cisco Webex App to chat
with the speaker after the session

How

1. Find this session in the Cisco Live Mobile App
2. Click “Join the Discussion”
3. Install the Webex App or go directly to the Webex space
4. Enter messages/questions in the Webex space

Webex spaces will be moderated by the speaker until June 9, 2023.

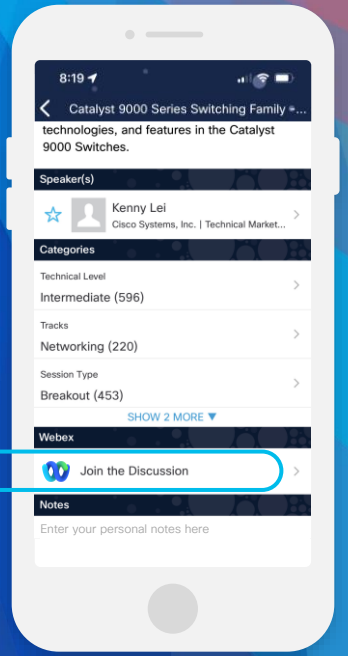

https://ciscolive.ciscoevents.com/ciscolivebot/#BRKENS-2096

Agenda

- QoS Overview
- UADP QoS
- Silicon One Q200 QoS
- Config migration examples

Overview

Do we need QoS?

User Experience

Guaranteeing voice quality	Bandwidth Savvy Business Applications	protect network infrastructure to deal with abnormal events
Video Quality	de-prioritizing non-business applications	protecting the control planes

QoS helps define the latency priority for your traffic packets

The QoS Toolset

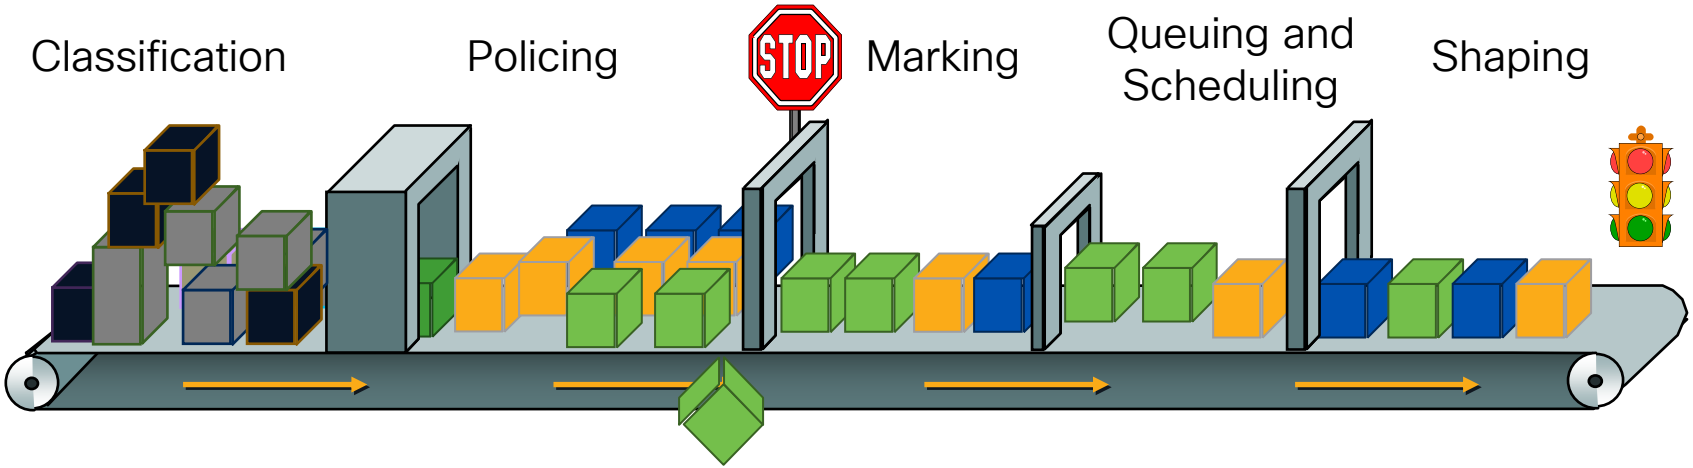

Identify and Split Traffic into Different Classes

Discard Misbehaving Traffic to Maintain Network Integrity

Mark Traffic According to Behavior and Business Policies

Prioritize, Protect and Isolate Traffic Based on Markings

Control Bursts and Conform Traffic

QoS Terminology

Term	Explanation
Trust	Retain the packet markings as it is
Classification	Identify packet priority and place it into different classes
Marking	Change the tags (priority) on the packets
Policing	Limit the traffic to specified rate. Excess traffic can either be dropped or assigned lower priority
Shaping	Limit the traffic to specified rate. Excess traffic will be queued and buffered.
Queueing	Process the packet into separate queues
Buffering	Storage for packets to be queued

Modular QoS CLI (MQC)

class-map
What traffic do we care about?

policy-map
What actions do we take on the classes?

service-policy
Where do we apply the policy?

```
class-map match-any Voice
  match dscp ef
class-map match-any Video
  match dscp 34


Policy-map POLICY-QOS
  class Voice
    priority level 1
  class Video
    set dscp 10


interface x/y
  service-policy (input/output) POLICY-QOS
```

Catalyst 9000 family of Switches

UADP 2.0x

Catalyst® 9200CX Series

Catalyst 9300X models

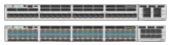

Catalyst 9200 Series

Catalyst 9300 Series

Catalyst 9400 Series With Sup-1

Catalyst 9500 Series

UADP 3.0x

Catalyst 9600 Series with Sup-1

Catalyst 9400 Series with Sup2

Catalyst 9500 High performance Series

Silicon One Q200

Catalyst 9600 Series with Sup-2

Catalyst 9500X Series

Class-based Weighted Fair Queue (CBWFQ) QoS

Virtual Output Queue (VoQ) QoS

UADP QoS (CBWFQ)

QoS Fundamental Actions in UADP

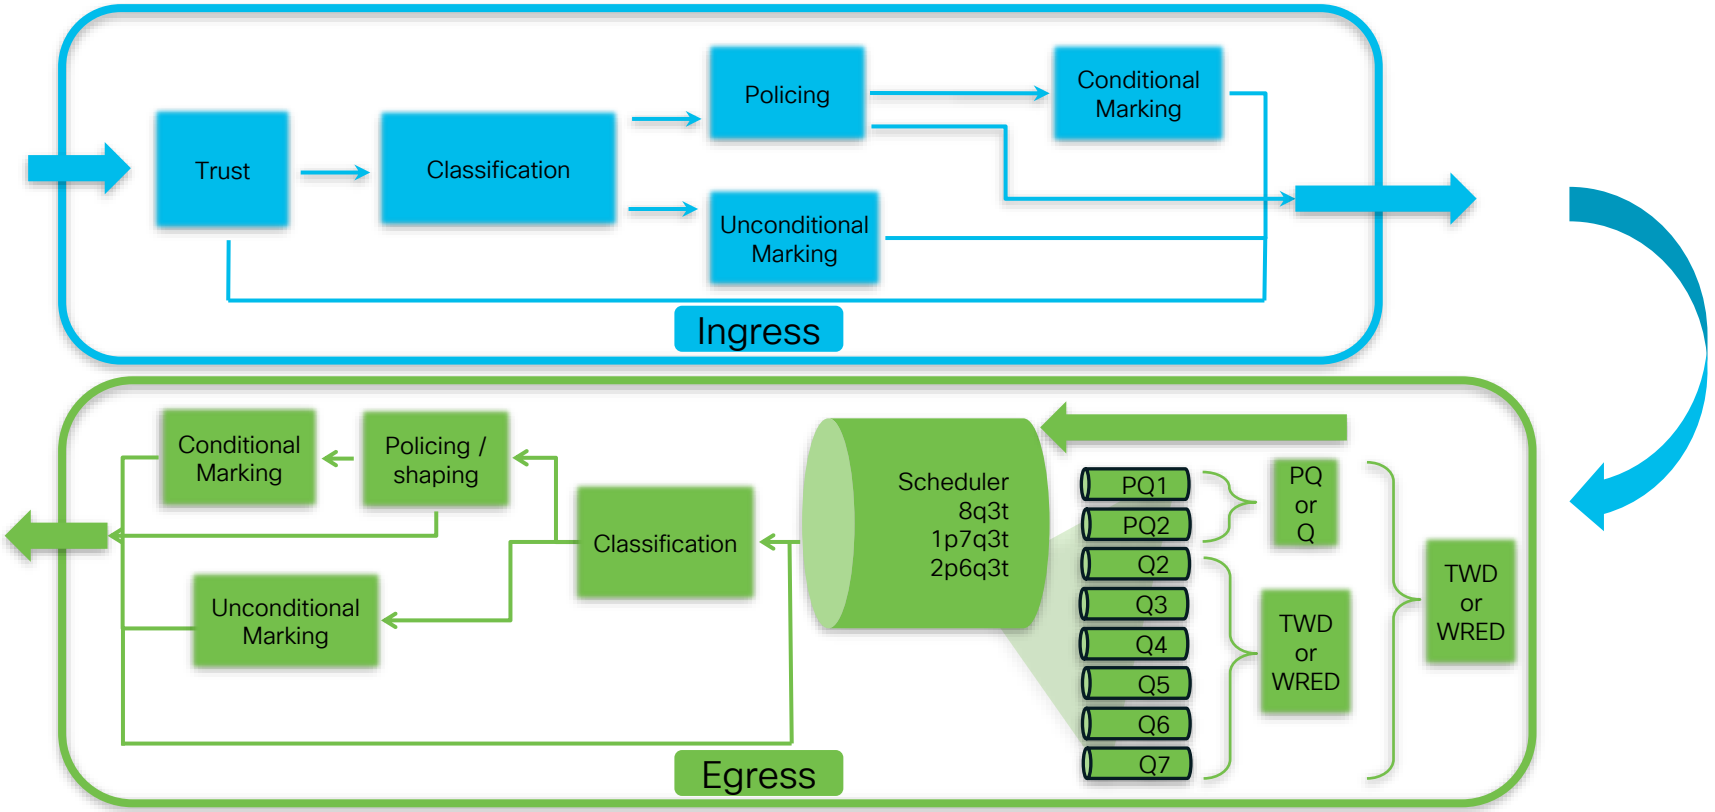

WRED: up to 4 queues with UADP 2.0x; up to 8 queues with UADP 3.0x

UADP QoS forwarding
ingress and egress

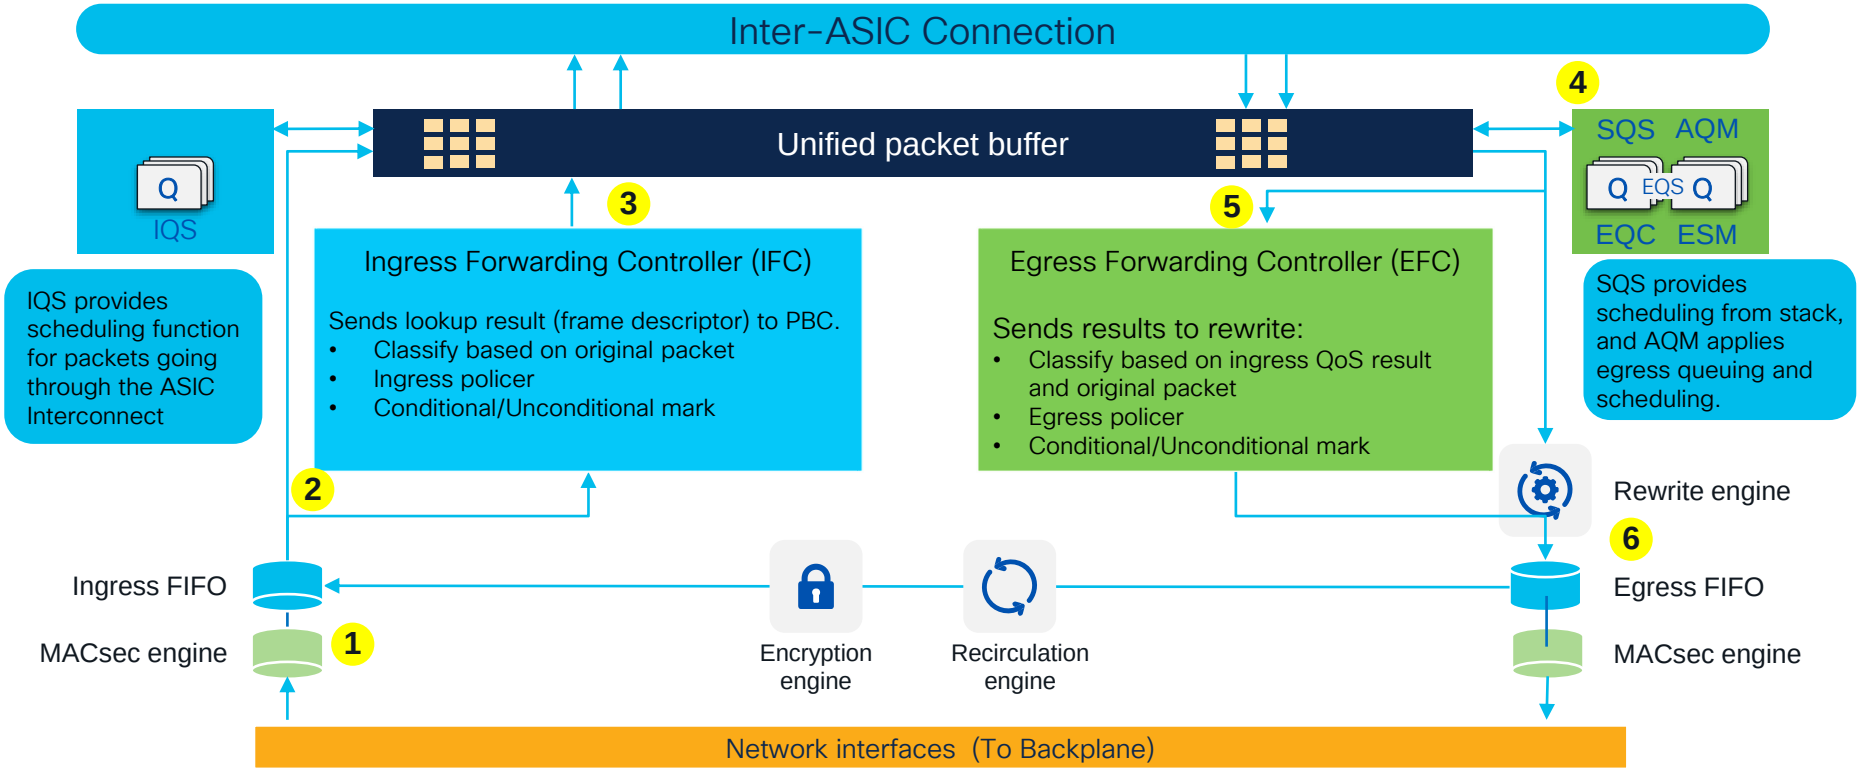

UADP forwarding

1. Packet arrives at ingress port, PHY converts the signal and serializes the bits, and then it sends to network interface ports. Network interface passes packet to ingress MACsec engine. MACsec engine decrypts the packet if needed and passes unencrypted packet to ingress FIFO.
2. IFC snoops the packet between FIFO and PBC. IFC returns lookup result (frame descriptor) to PBC.
3. PBC uses the frame descriptor to determine the egress port. Egress on same ASIC, so result to moved to EQS.
4. EQS schedule the packet for egress process. EQS – replication, scheduling, and queue management. PBC sends packet with new frame descriptor and enqueues the frame.
5. EFC snoops the packet between PBC and rewrite engine. EFC performs egress lookup functions to learn SRC MAC, egress SPAN, etc. and sends results to rewrite engine.
6. Rewrite engine rewrites packets and sends through the egress FIFO. MACsec engine encrypts packet prior to placing it on NIF.

UADP QoS Default

- Catalyst 9000 Switches with UADP ASICs
  - QoS enabled
  - All ports trust at layer2 and layer3
  - Two queues (neither set as priority)

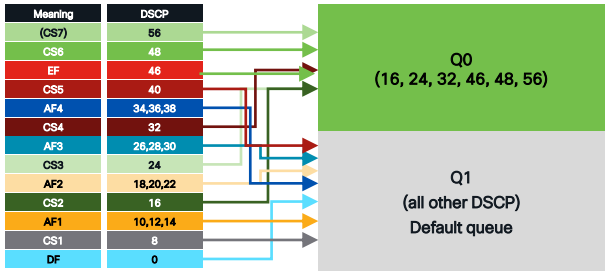

Classification, Marking and Policing

Classification and Marking

- Identify traffic
  - Access Control Lists (ACLs)
  - DSCP
  - IP precedence
  - CoS
  - QoS Group (local with the switch)
  - EXP (MPLS)
  - Network-Based Application Recognition (NBAR) protocols *
  - VLANs
- Marking
  - Conditional or unconditional
  - Table map (default-class)
  - QoS group (local within switch)

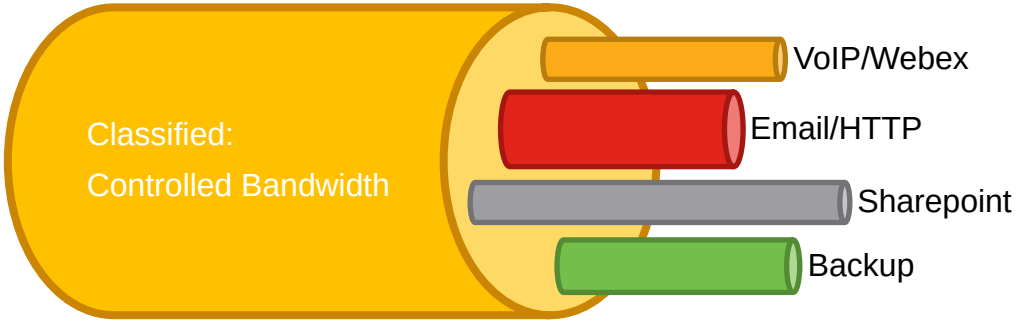

* Access platforms

Policing – Limit the traffic

1 rate 2 color

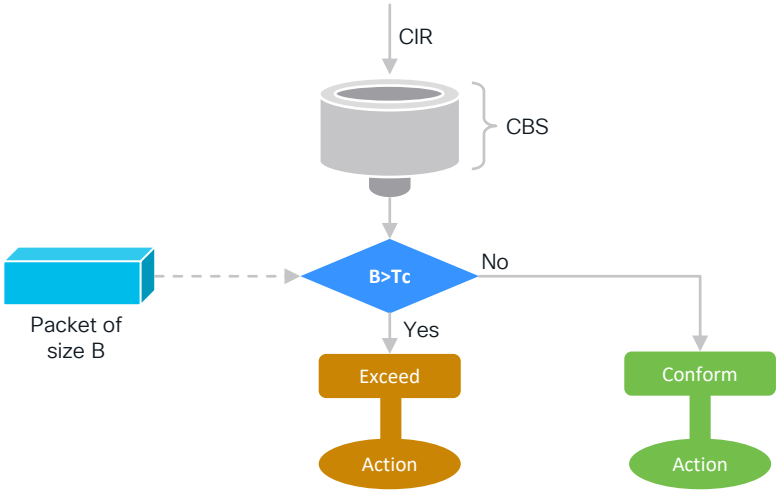


```
police cir 1g bc 3125000
  conform-action set-dscp-transmit af41
  exceed-action drop
```

2 rate 3 color

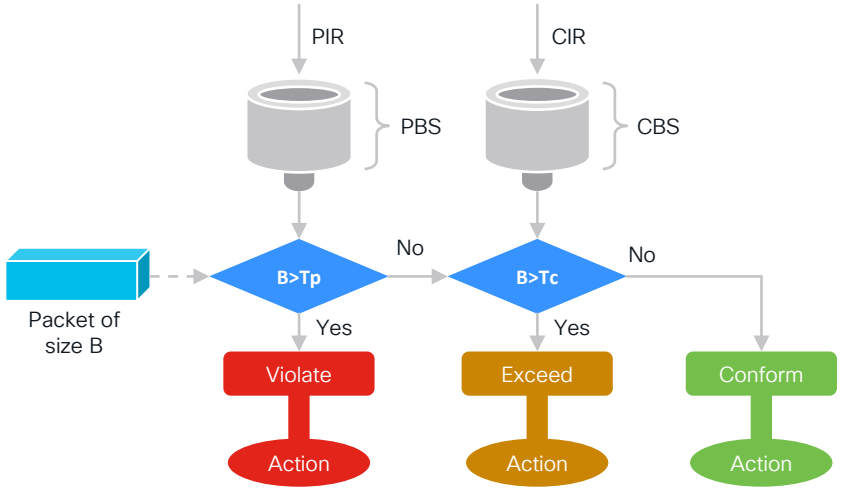


```
police cir percent 10 pir percent 50
  conform-action transmit
  exceed-action set-dscp-transmit dscp table MARKDOWN
  violate-action drop
```

CIR – Committed Information Rate
PIR – Peak Information Rate

PBS- Peak Burst Size
CBS – Committed Burst Size

Queueing, Scheduling and Shaping

Queueing

- Separate the traffic into the queues
- Traffic in different queue can be treated differently
- Up-to 8 queues per interface, 2 of which can be priority-queues (PQ).
- Both priority-queues are strict priority queues.
- Policer or a shaper on the priority queue will limit the traffic to the configured value regardless of the traffic level on other queues.

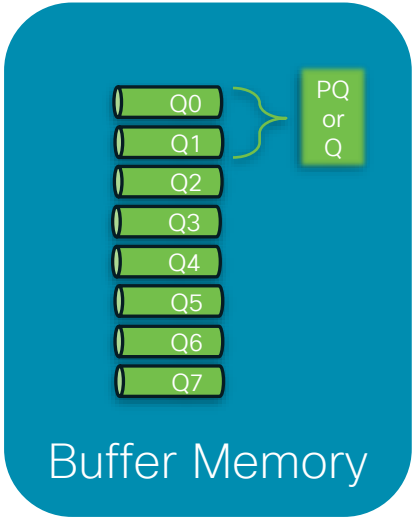

Note: Queue classification - DSCP/COS/IP Prec/QoS Group

Queueing

Priority queue

```
policy-map 2P6Q3T
  class VOICE-PQ1
    priority level 1
```

- Priority queue.
- As long as there is traffic in priority queue all other queues will not be serviced.

```
policy-map 2P6Q3T
  class VOICE-PQ1
    priority level 1
      police rate percent 10
```

- Priority queue.
- **Limit (police)** traffic to 10 percent of link speed regardless the utilization of other queues

```
policy-map 2P6Q3T
   class VOICE-PQ1
     priority level 1 percent 10
```

- Enable strict priority queue.
- **Shape** traffic to 10 percent of link speed regardless the utilization of other queues

Scheduling – UADP

- Scheduling defines the order of transmission of traffic out of the queues
- Different type of queues are served differently
  - Strict priority queues
    - Always serviced first
    - With 2 PQs, level1 over level 2
  - Normal queues
    - Served only after priority queues are empty
    - Use Weighted Round Robin (WRR) for scheduling
- WRR servers normal queue based on the weight and packet size
- Egress Queue System (EQS) is the component on the UADP ASIC responsible for the scheduling

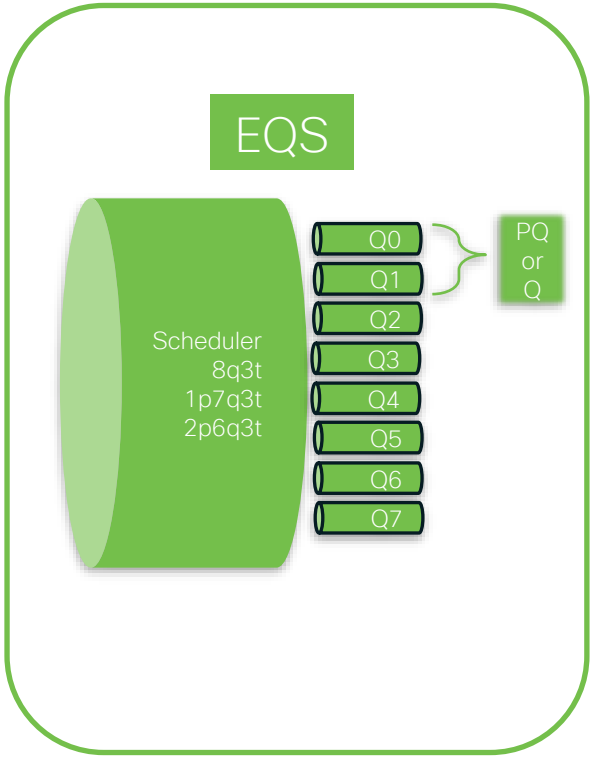

Scheduling - Example

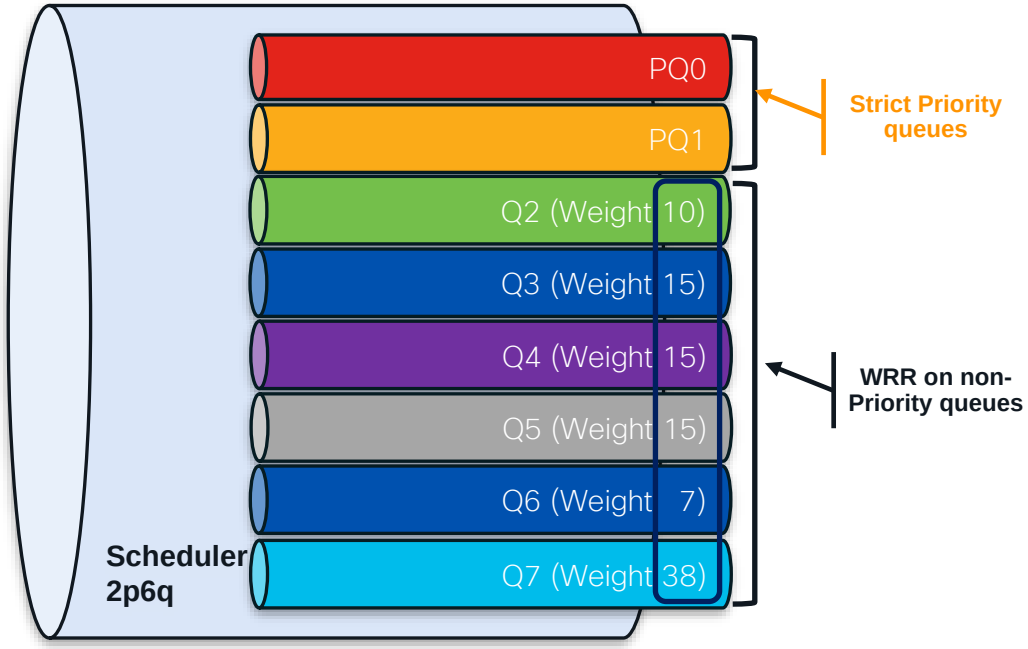

Policy Map Configuration

```
policy-map 2P6Q3T
  class VOICE-PQ1
    priority level 1
  class VIDEO-PQ2
    priority level 2
  class CONTROL-MGMT-QUEUE
    bandwidth remaining percent 10
  class MULTIMEDIA-CONFERENCING-QUEUE
    bandwidth remaining percent 15
  class MULTIMEDIA-STREAMING-QUEUE
    bandwidth remaining percent 15
  class TRANSACTIONAL-DATA-QUEUE
    bandwidth remaining percent 15
  class SCAVENGER-BULK-DATA-QUEUE
    bandwidth remaining percent 7
  class class-default
    bandwidth remaining percent 38
```

Shaping

- Smooth out traffic peaks, microburst, with preserving traffic
- Control traffic rate to the desired value with buffering.
- Usually in the egress direction

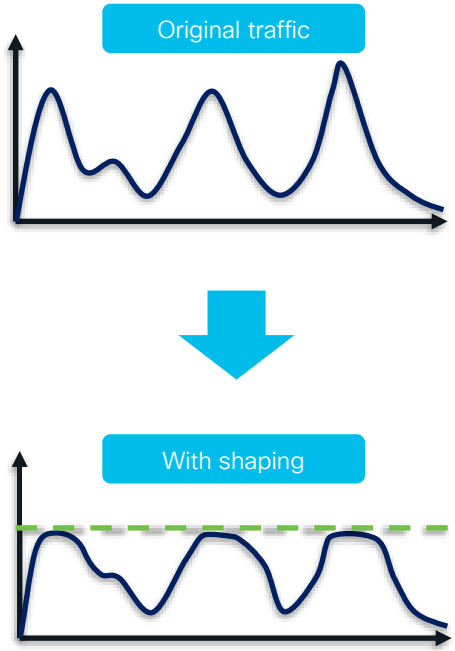

Shaping Example

```
policy-map Shaper
    class Transactions
    shape average percent 30
```

UADP Congestion Management

Where do we need congestion management?

2P6Q3T Example

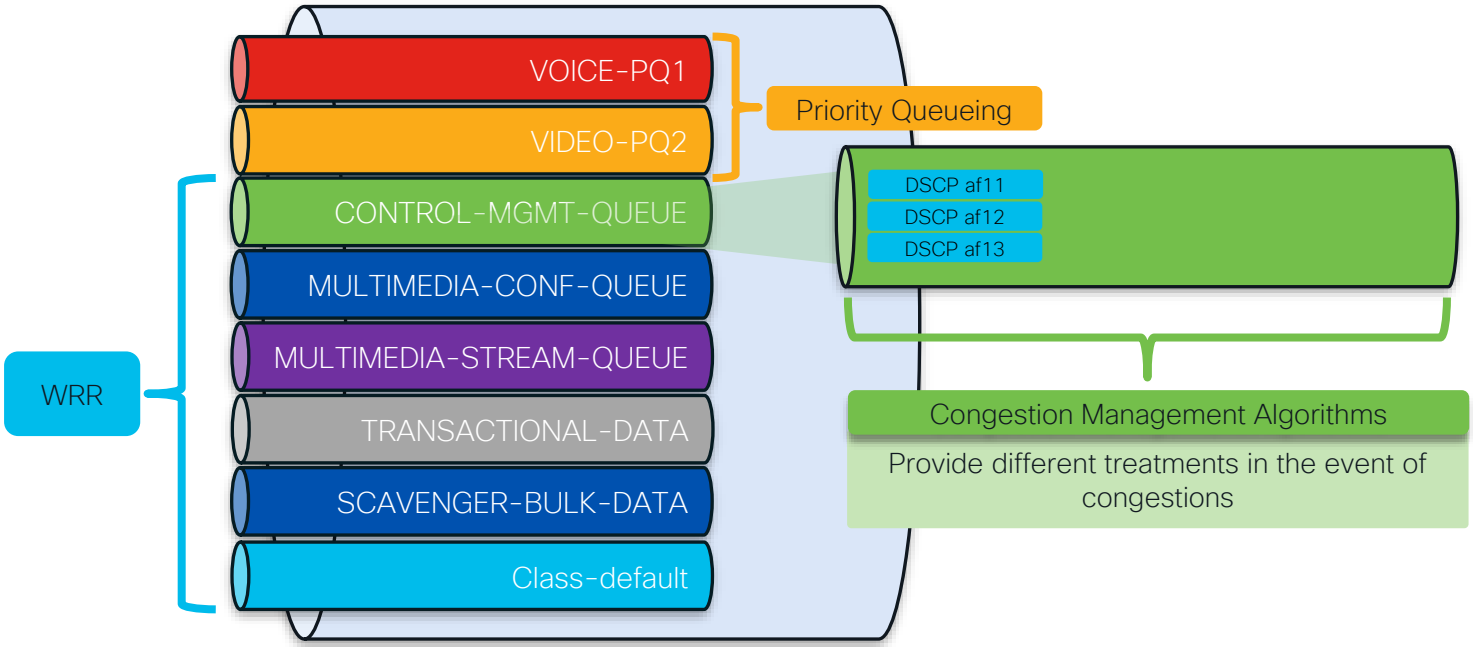

UADP - Congestion Management

Weighted Tail Drop (WTD)

- Default
- For non-priority queues
- Up to 3 thresholds per queue, one threshold per QoS tag
- Each queue need to use same QoS tag type

Weighted Random Early Detection (WRED)

- For non-priority queues
- Up to 4 queues with UADP 2.0X and up to 8 queues with UADP 3.0X
- Up to 3 threshold pairs per queue
- Each queue need to same QoS tag type

Weighted -> Multiple pair of thresholds

Congestion Management Tools

- Tail Drop (TD)
  - Drop packets at tail of the queue
  - Single threshold per queue

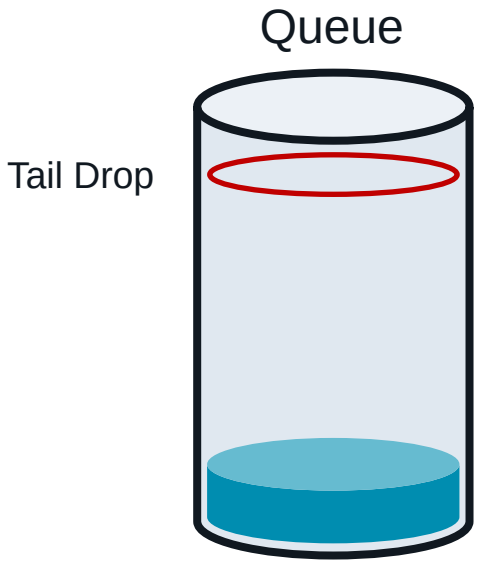

Congestion Management Tools

- Tail Drop (TD)
  - Drop packets at tail of the queue
  - Single threshold per queue

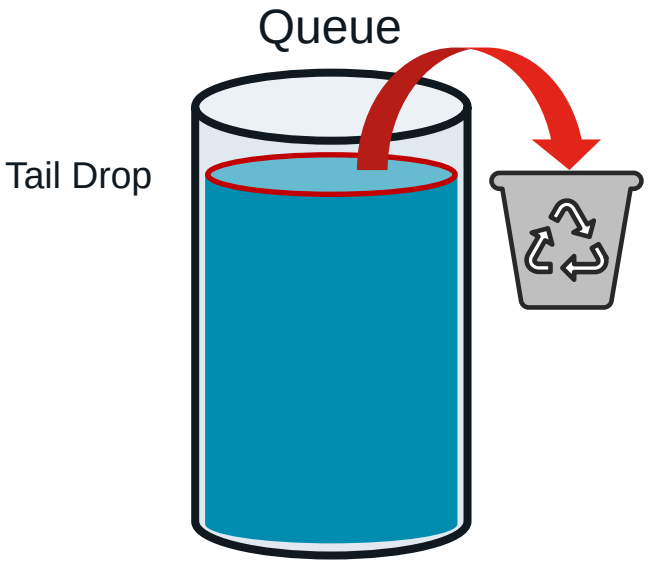

WTD – UADP Example

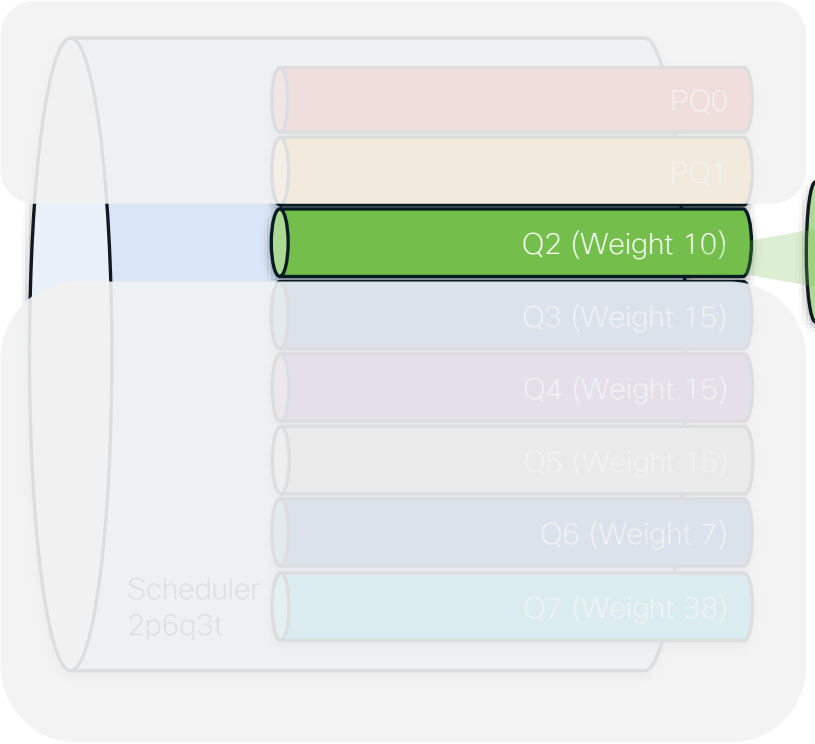

Three thresholds to conditionally drop specific traffic in the event of congestion

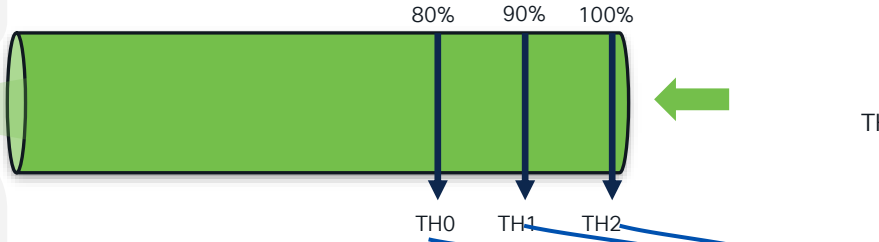

TH: threshold

Policy Map Configuration

```
policy-map 2P6Q3T
 class DATA-QUEUE
  queue-limit dscp values af13 percent 80
  queue-limit dscp values af12 percent 90
  queue-limit dscp values af11 percent 100
```

Congestion Management Tools

- Tail Drop (TD)
  - Drop packets at tail of the queue
  - Single threshold per queue

- Weighted Random Early Drop (WRED)
  - One or more thresholds per queue
  - Threshold associated with priority
  - Buffer usage below threshold no affect

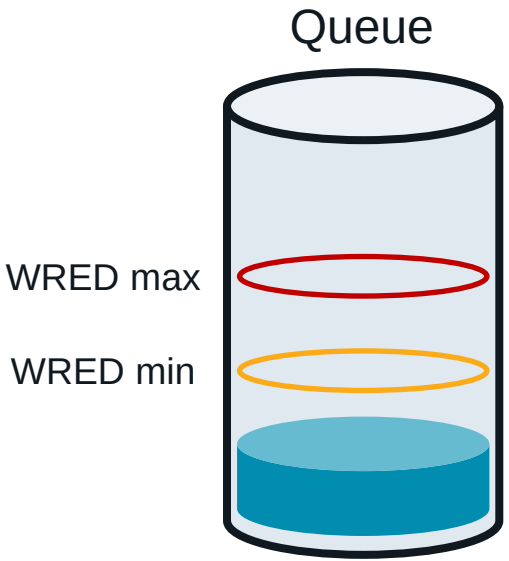

Congestion Management Tools

- Tail Drop (TD)
  - Drop packets at tail of the queue
  - Single threshold per queue

- Weighted Random Early Drop (WRED)
  - One or more thresholds per queue
  - Threshold associated with priority
  - Buffer usage below threshold no affect
  - Buffer usage over min threshold = random drops

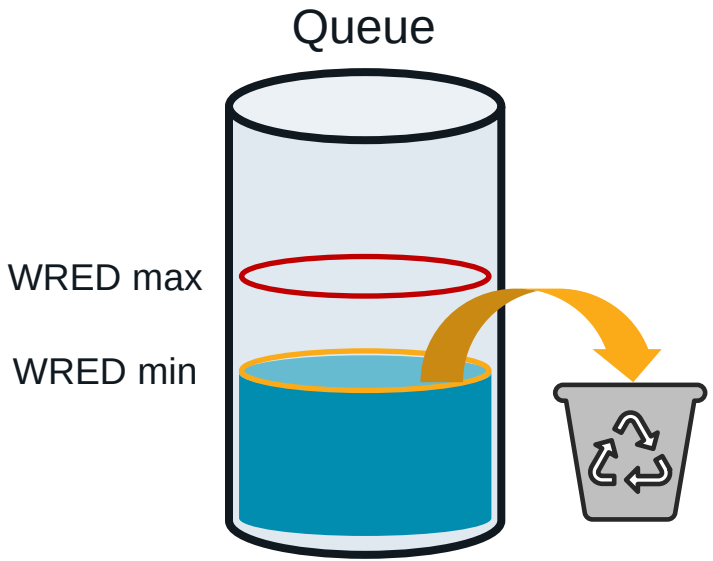

Congestion Management Tools

- Tail Drop (TD)
  - Drop packets at tail of the queue
  - Single threshold per queue

- Weighted Random Early Drop (WRED)
  - One or more thresholds per queue
  - Threshold associated with priority
  - Buffer usage below threshold no affect
  - Buffer usage over min threshold = random drops
  - Buffer usage over max threshold = all traffic drop

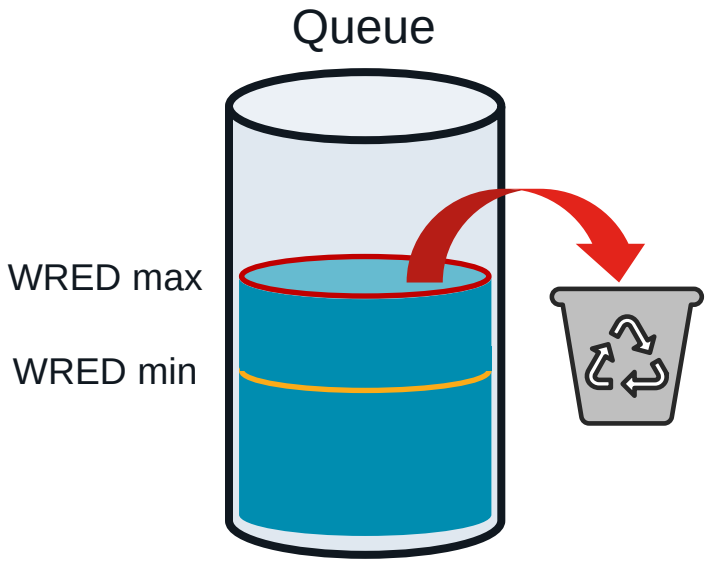

WRED – UADP Example

- Shown two pairs of WRED thresholds
- UADP supports up to 3 pairs of thresholds

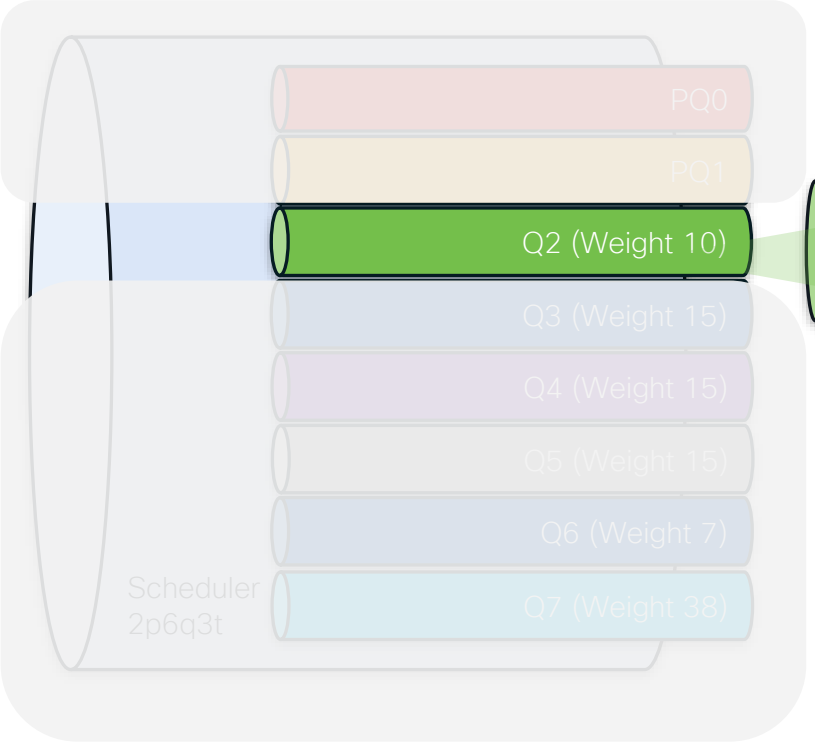

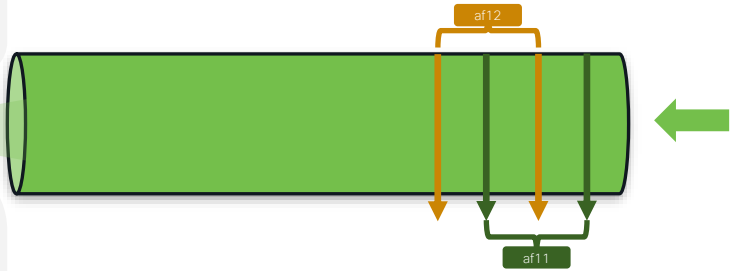

Policy Map Configuration

```
policy-map 2P6Q3T
 class DATA-QUEUE
   bandwidth remaining percent <number>
   queue-buffers ratio <number>
   random-detect dscp-based
   random-detect dscp af11 percent 70 90
   random-detect dscp af12 percent 60 80
```

Buffers

Buffers

- Allocation
  - Dedicated and shared: use dedicated first then shared
  - Dynamic Threshold Scale (DTS): Algorithm to managed the shared buffer

Dedicated

- Allocated to each port on boot.
- Cannot be dynamically changed/edited

Shared

- Dynamically assigned to ports for burst absorption.
- Returned to common pool when not in use.

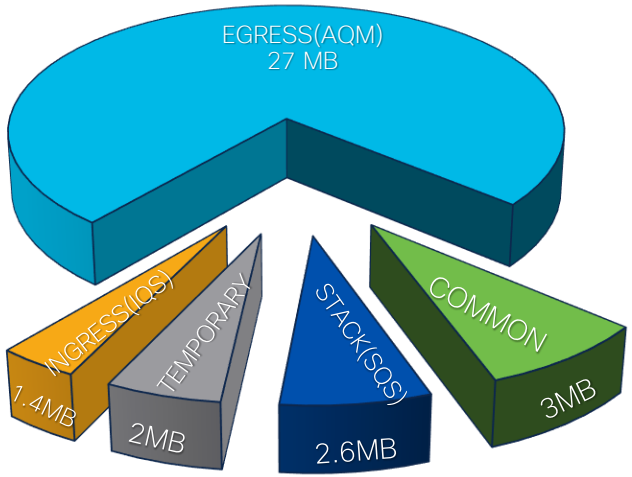

UADP 3.0

- UADP 3.0 specific
  - Buffer can be shared across two cores
  - “qos share-buffer” to enable the unified buffer

Dynamic Threshold Scale (DTS)

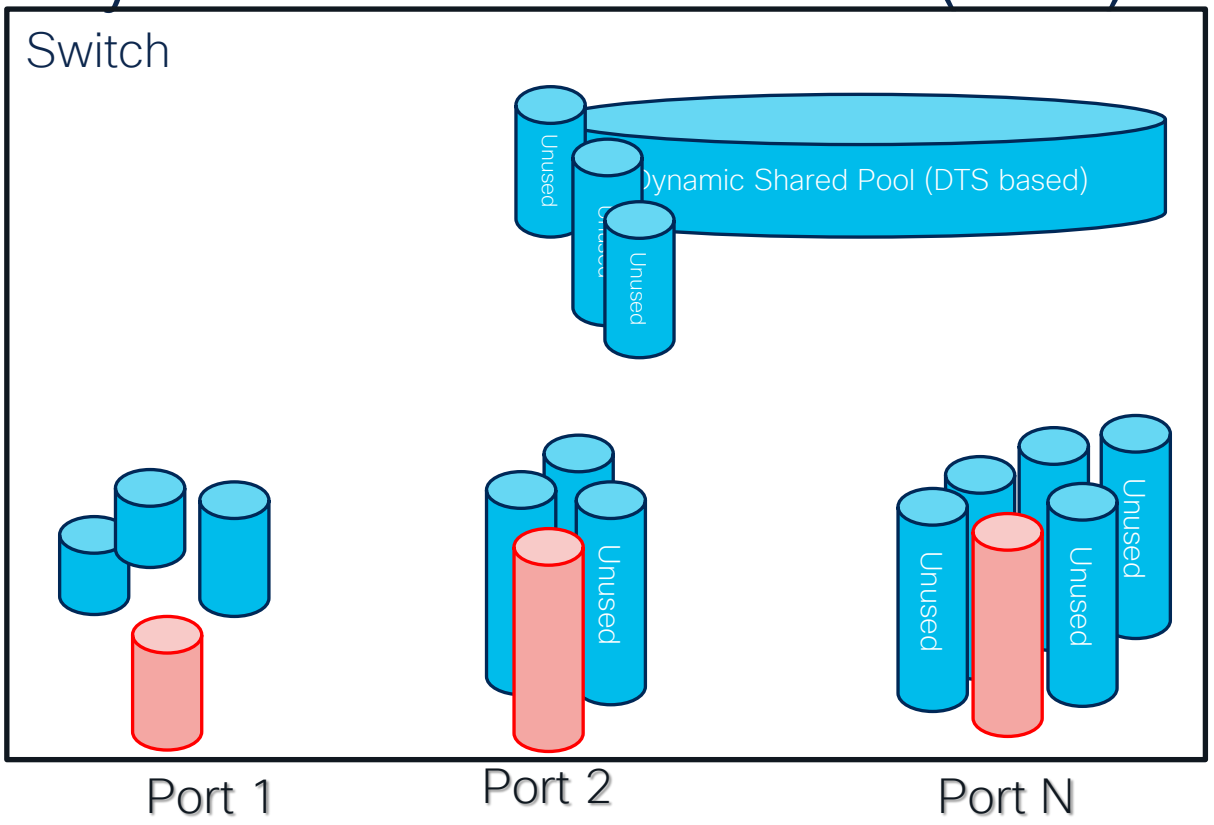

- Shared buffer is good for burst absorption.
- Dedicated buffer is good for predicated performance for each port.
- Buffer management is flexible: Dedicated plus shared.
- Configurable dedicated threshold per port/queue
- Configurable global maximum shared threshold
- Automatically adjusted depends on the available shared pool

The famous softmax multiplier

What does it do?

- `qos queue-softmax-multiplier ...`

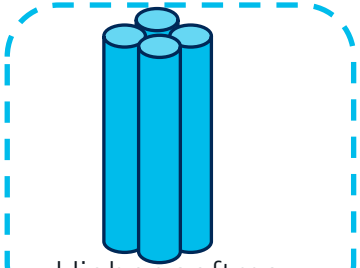

Higher softmax

- Higher burst absorption by individual ports
- Lower concurrent burst spread across ports

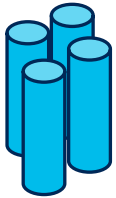

Lower softmax

- Lower burst absorption by individual ports
- Higher concurrent burst spread across ports

UADP 2.0	Max value of 1200
UADP 3.0	Max 1200 by default and up to 4800 with shared buffer

Unused buffers are returned to the shared pool. No impact from maxing out the softmax multiplier

Dynamic Threshold Scale (DTS)
buffer allocation graph

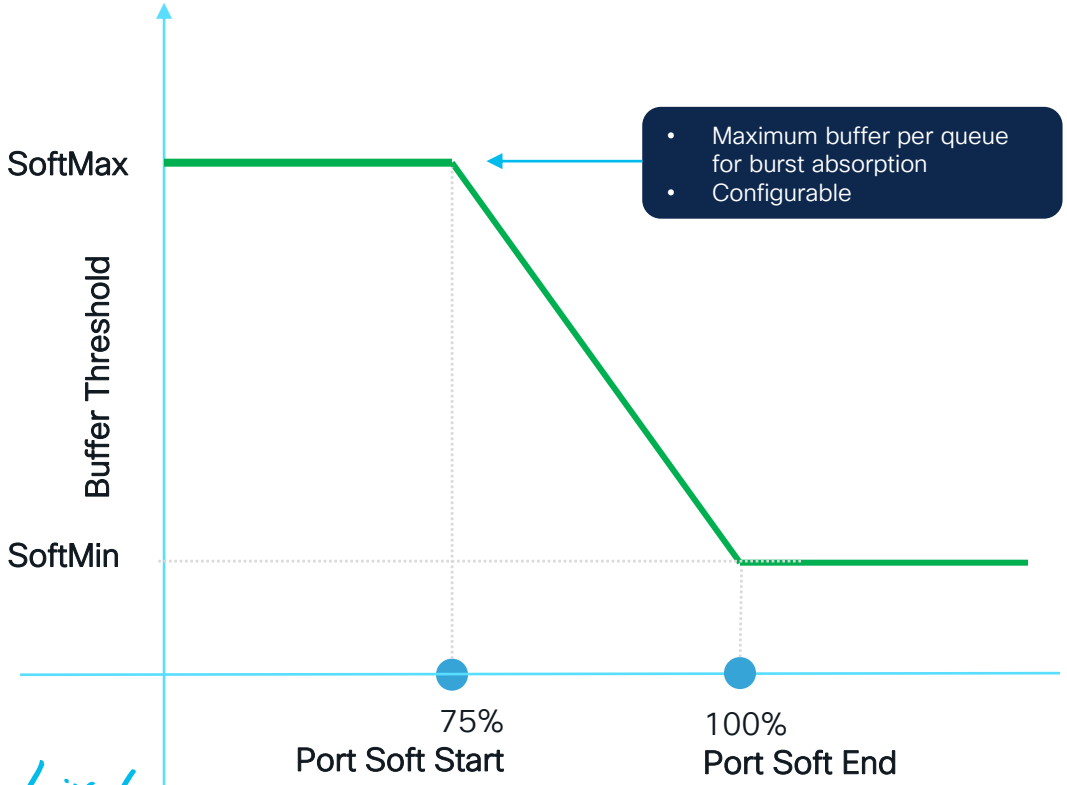

- Maximum buffer per queue for burst absorption
- Configurable

- **SoftMin** – Minimum Shared buffer given to the port
- **SoftMax** – Maximum Shared buffer the port can consume from shared Pool
- **Port Soft Start** – Moment when the Softmax start to reduce with step base
- **Port Soft End** – Moment when SoftMin and SoftMax are equal

HQoS

UADP Hierarchical QoS (HQoS)

HQoS (two-level hierarchy) allows you to perform the following functions:

- Classification
- Policing
- Shaping

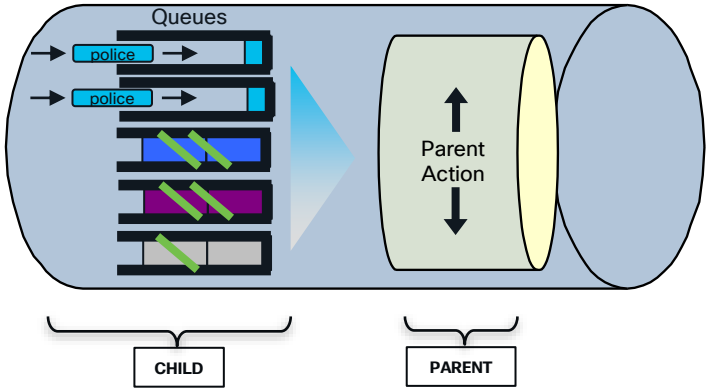

Child Policy	Parent Policy
Classification + Policing	Shaping
	Marking
Classification + Marking	Policing
	Shaping

Silicon One Q200 QoS (VoQ)

Silicon One QoS Terminology

Term	Explanation
VoQ	Virtual Output Queues between Ingress and Egress
Packet Color	Used for congestion management to prioritize packets to be dropped
Traffic-Class	Internal tag used by the Silicon One ASIC to differentiate packet priority
Traffic Manager	The block in Q200 responsible for scheduling
Traffic/Transmit Scheduler	When the OQ can send traffic out to the wire
Credit Scheduler	When the VoQ can send traffic to the Output queue
SMS	Shared Memory Sub-system - Primary Buffering system
HBM	High Bandwidth Memory - Secondary deep Buffering system used during congestion

VoQ and Head of Line Blocking (HoL)

Buffer types – Silicon One vs UADP

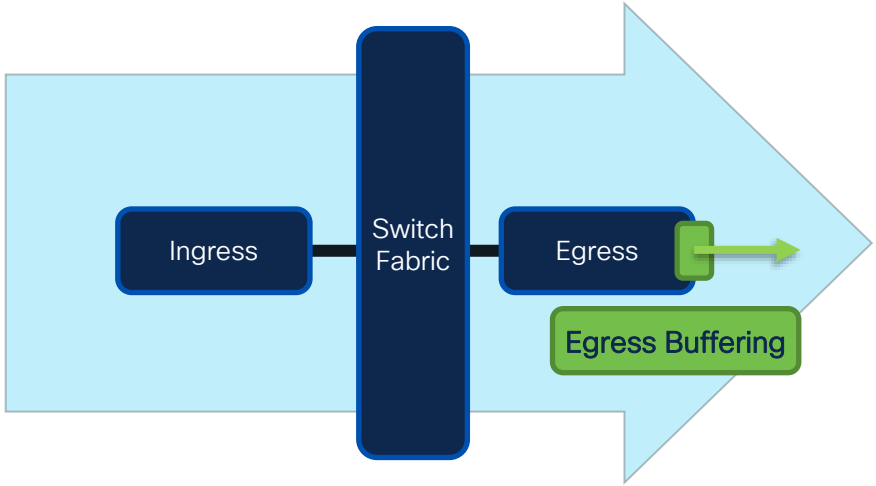

UADP

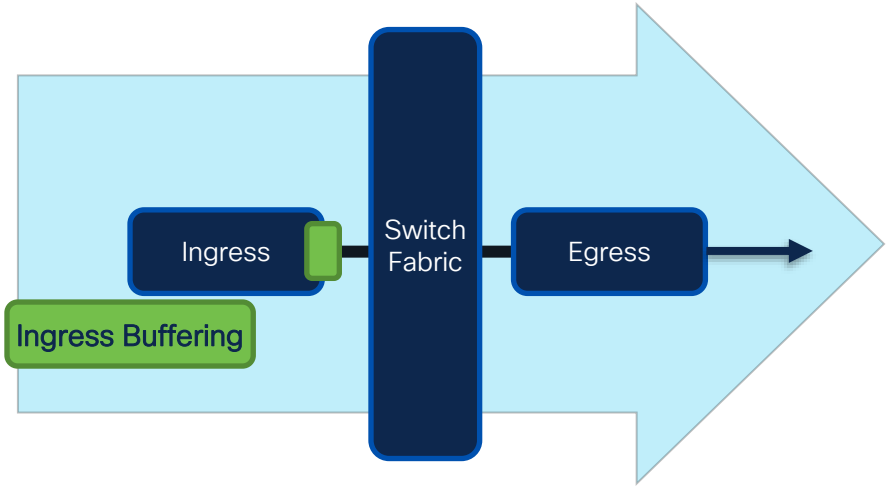

Silicon One Q200

Ingress buffering – Head of Line Blocking

What is the Problem?

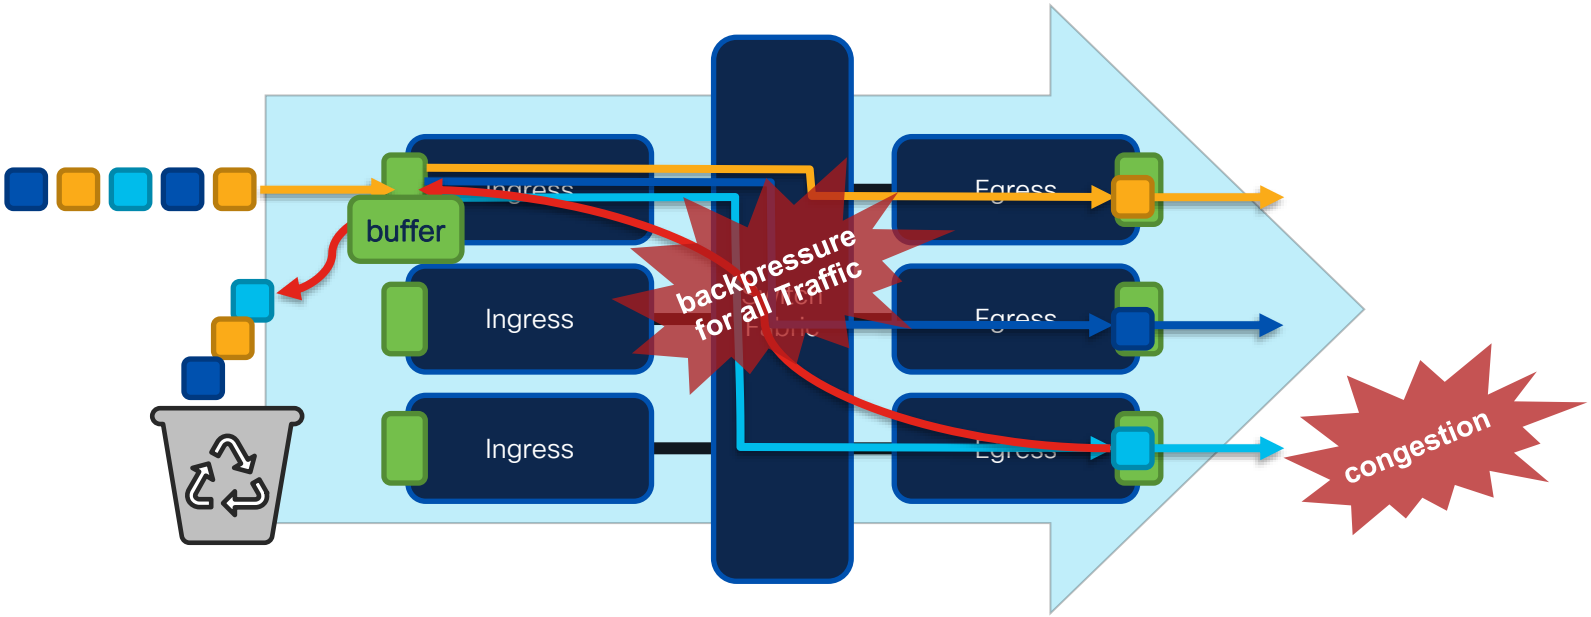

Egress buffering - UADP

Avoid Head of Line Blocking

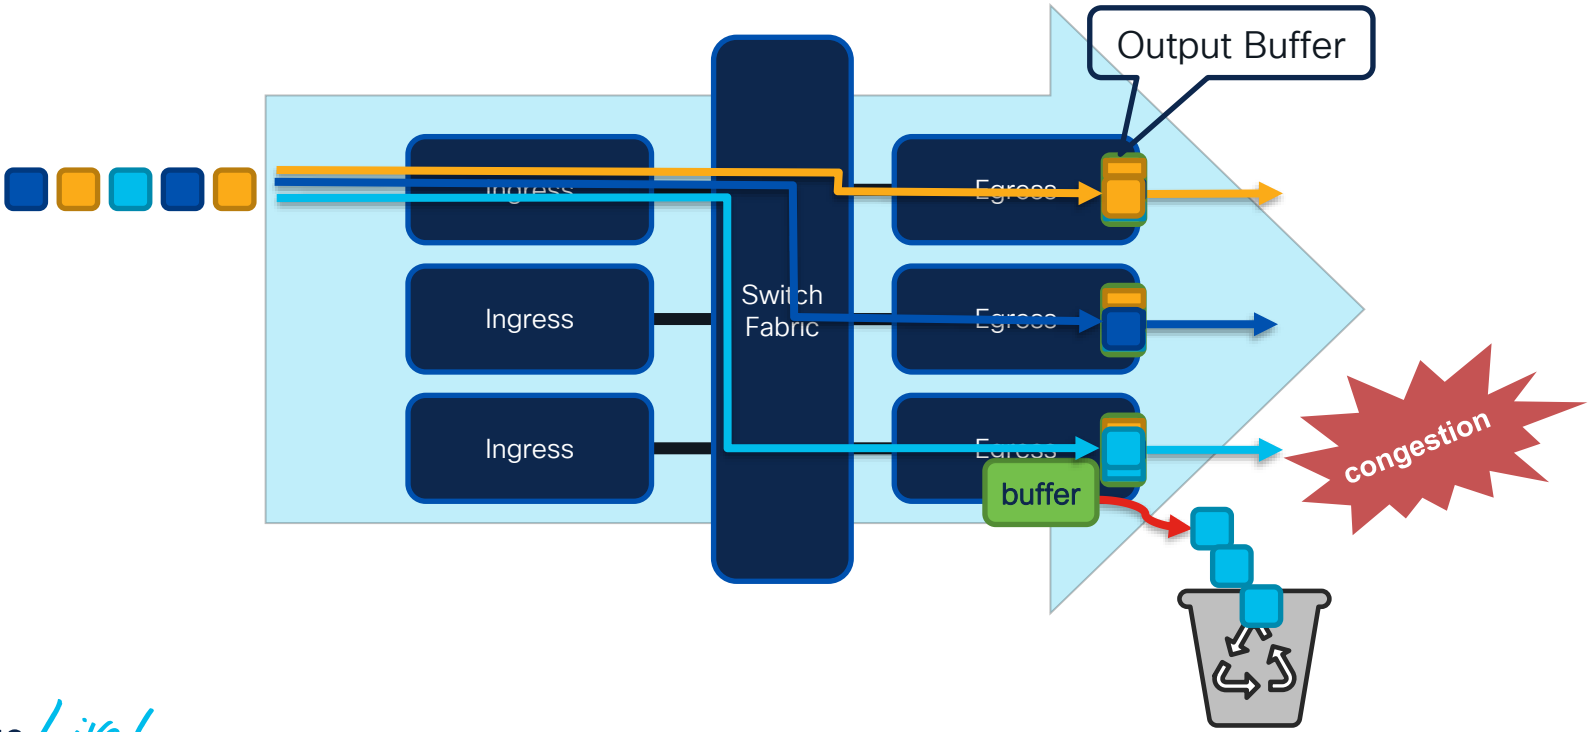

Virtual Output Queuing – Silicon One ASIC

Avoid Head of Line block

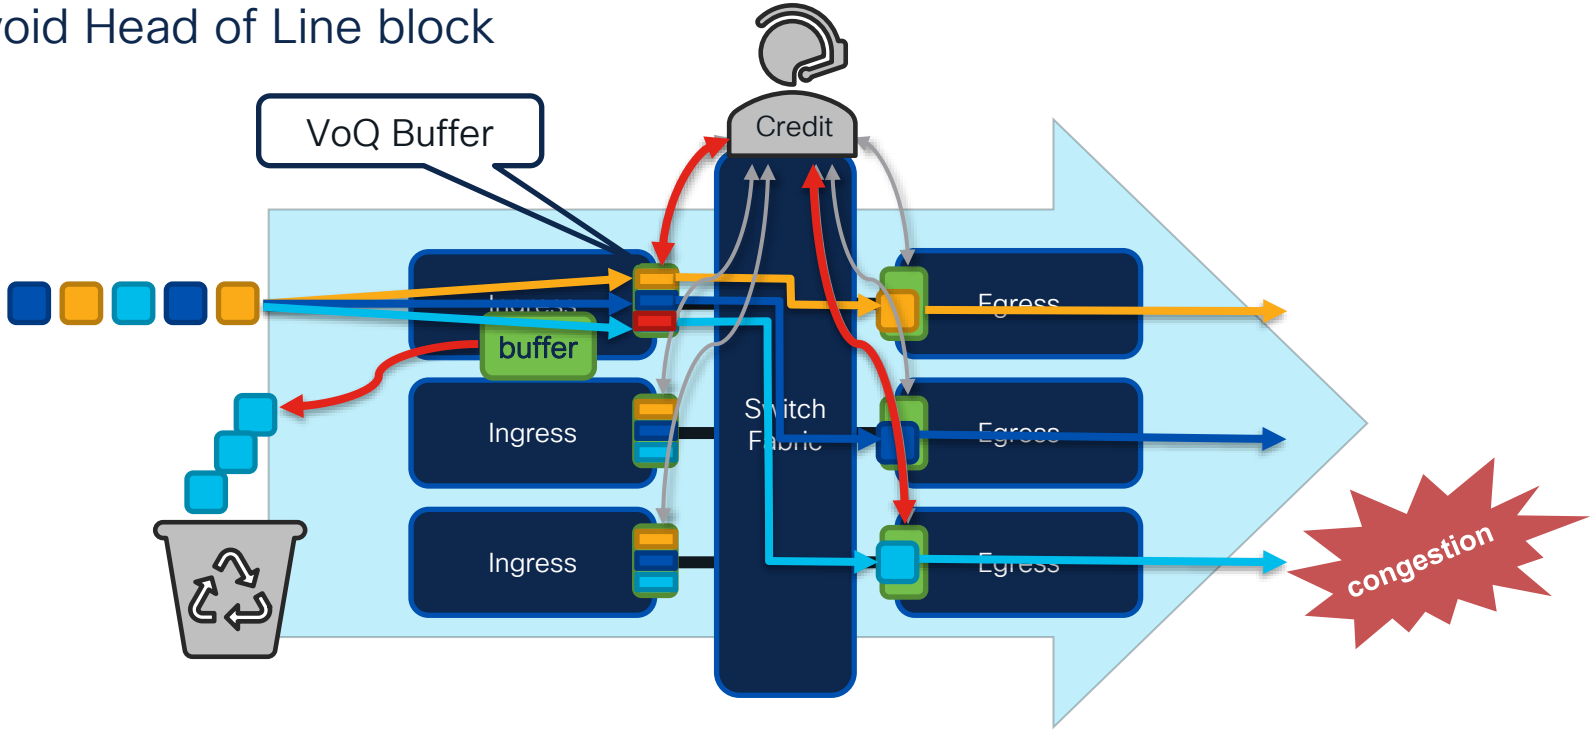

Silicon One QoS Overview

Silicon One

Features Mapping

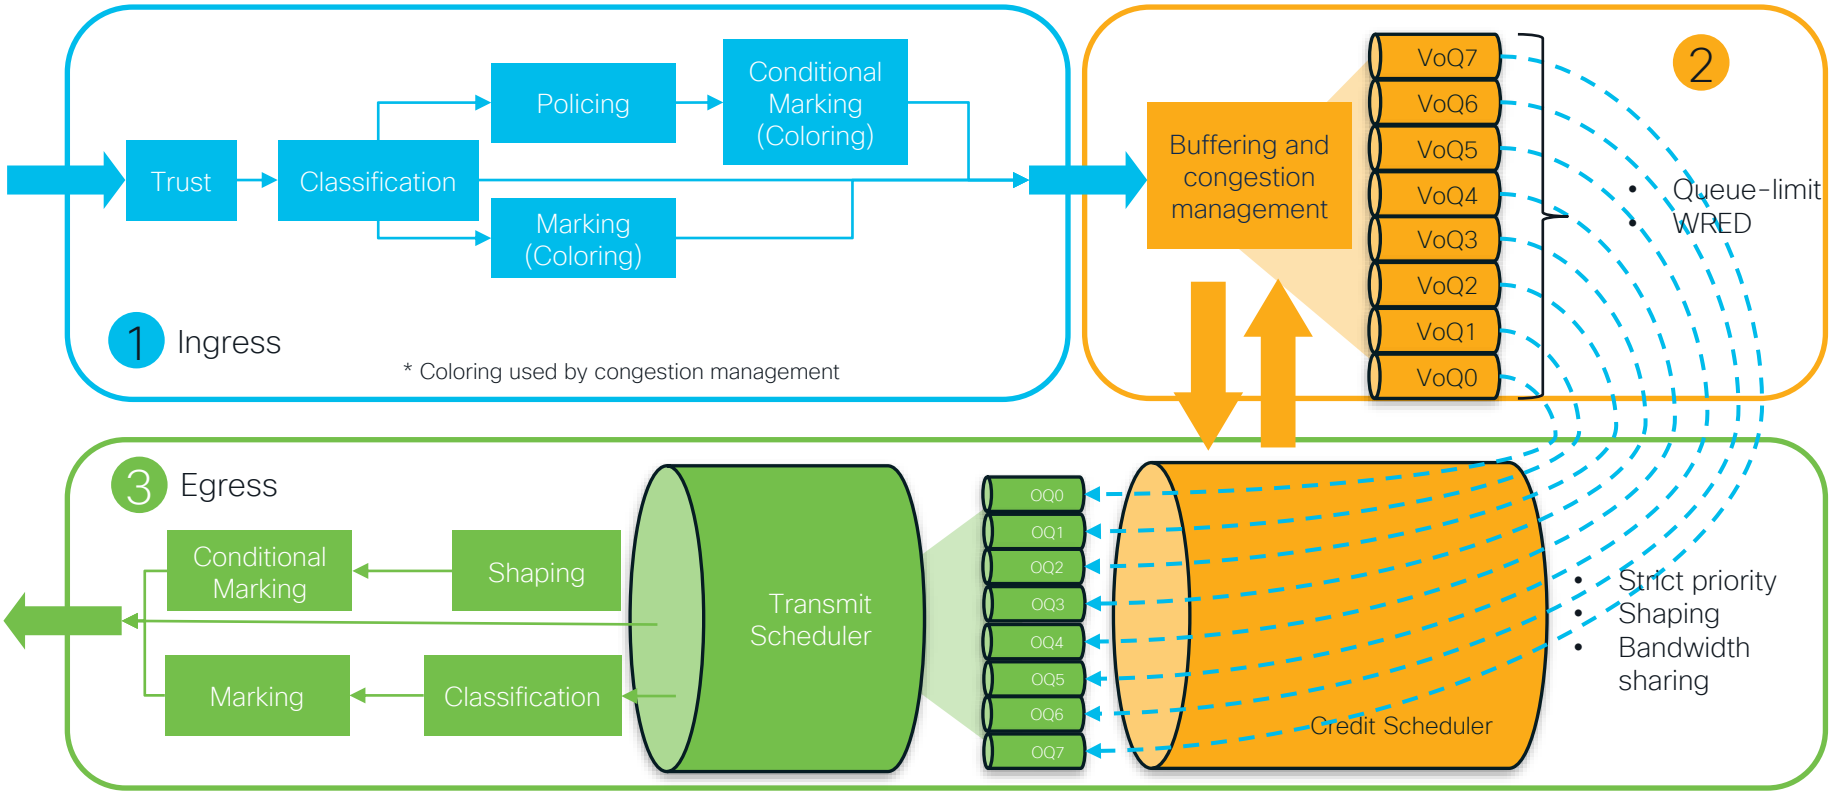

Silicon One ASIC mapping

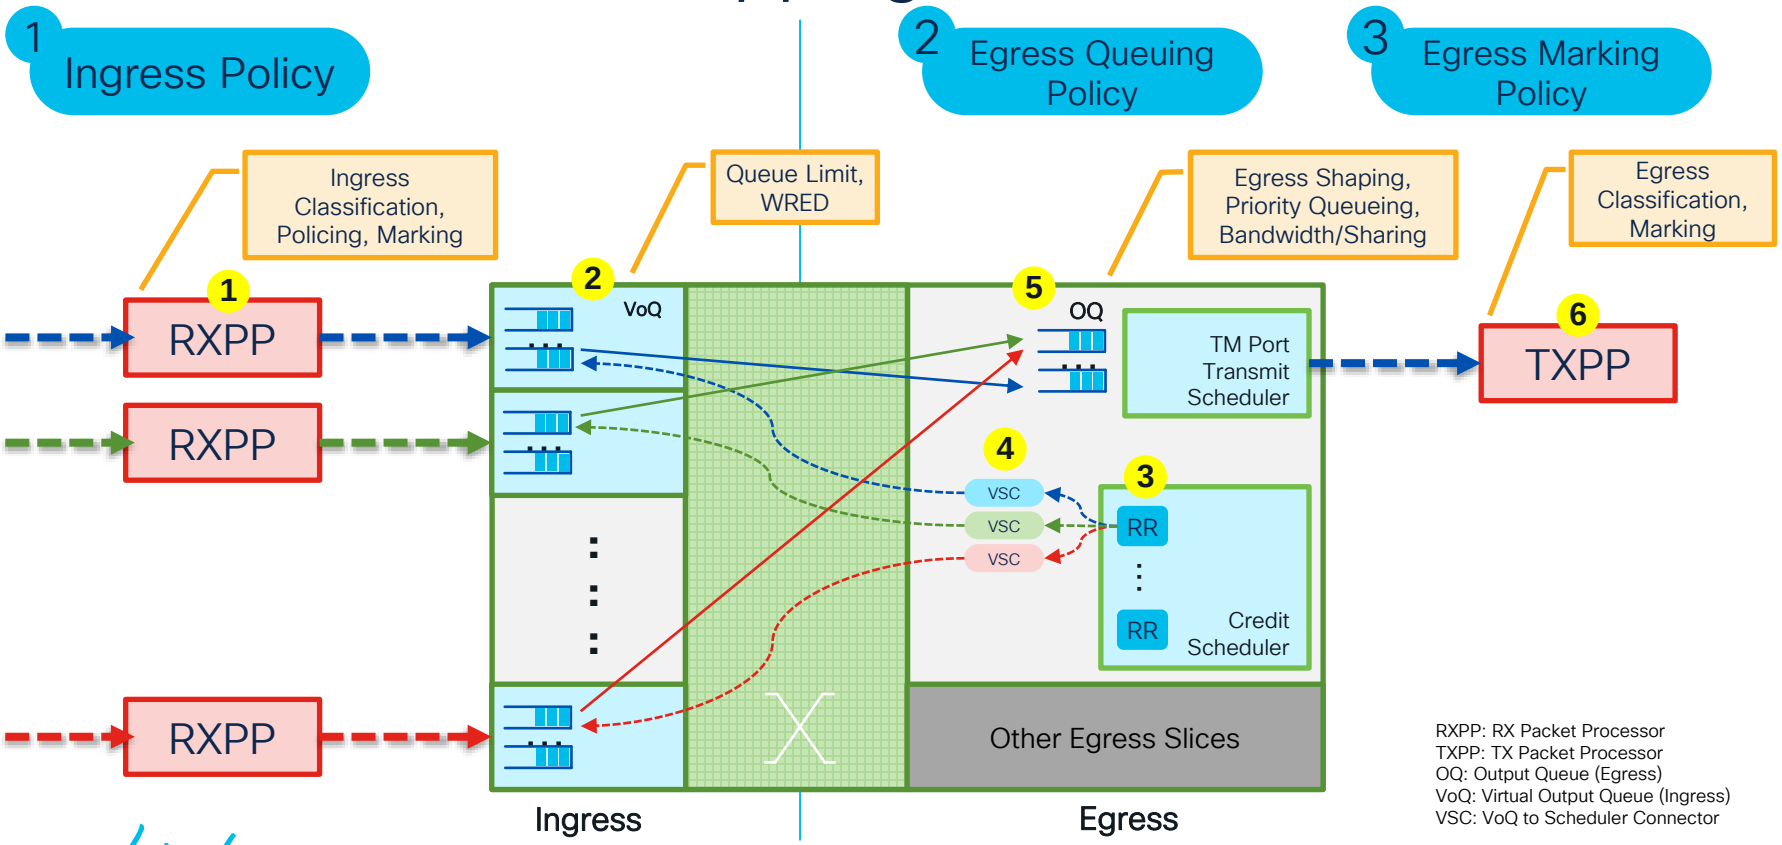

RXPP: RX Packet Processor
TXPP: TX Packet Processor
OQ: Output Queue (Egress)
VoQ: Virtual Output Queue (Ingress)
VSC: VoQ to Scheduler Connector

Silicon One ASIC mapping Hidden slide

1. Packet from the ingress interface hits the Receive Interface Group (Rx IFG). RxPP (Receive packet processor) consists of both the corresponding Interface Group as well as the Network Processing Unit (RxNPU) for the corresponding slice.

2. From the RxIFG, the packet descriptor is sent to corresponding VoQ (Virtual Output Queue) where it is queued and forwarded once it receives a credit.

3. Each VoQ connects to the scheduler via a VSC (Virtual queue scheduler) The scheduler uses a round robin (RR) algorithm to provide credits each corresponding VSC.

4. The credit is sent from the VSC to the corresponding VoQ. Once the VoQ receives the credit, it can forward the traffic to the egress.

5. The packet is sent from VoQ to corresponding OQ.

6. Once at the OQ, any marking operations, if any, are performed and the packet egresses out the TxNPU and TxIFG (collectively called TxPP) out of the switch.

Traffic class

- S1 ASIC uses traffic classes to map traffic to different queues. ”**traffic-class**” is local significant to the switch only
- 3-bit field => 8 values, traffic-class <0 - 7>
- Traffic-class 0 - lowest priority (maps to class-default); traffic-class 7 - highest priority (traffic-class 1 to 6 can be non-priority)
- **Ingress policies** classify packets to specific traffic classes
- Class-maps in **egress queuing policy** can only match traffic-class

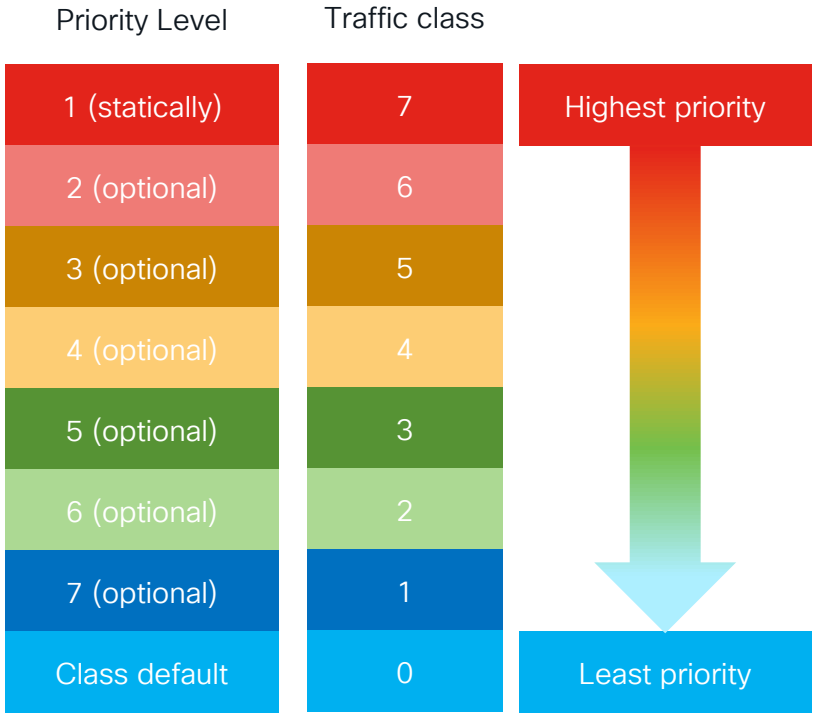

Traffic class to VoQ mapping

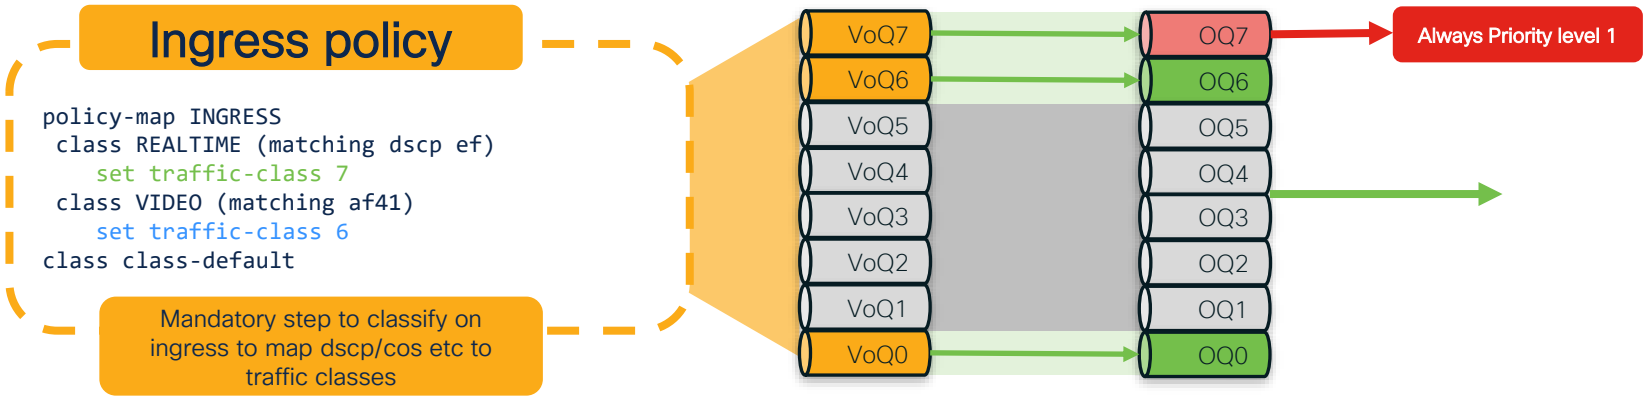

- Ingress policy determines how many VoQ are enabled.
- Each VoQ has a 1:1 mapping with corresponding OQ.

Traffic color (marking) – discard-class

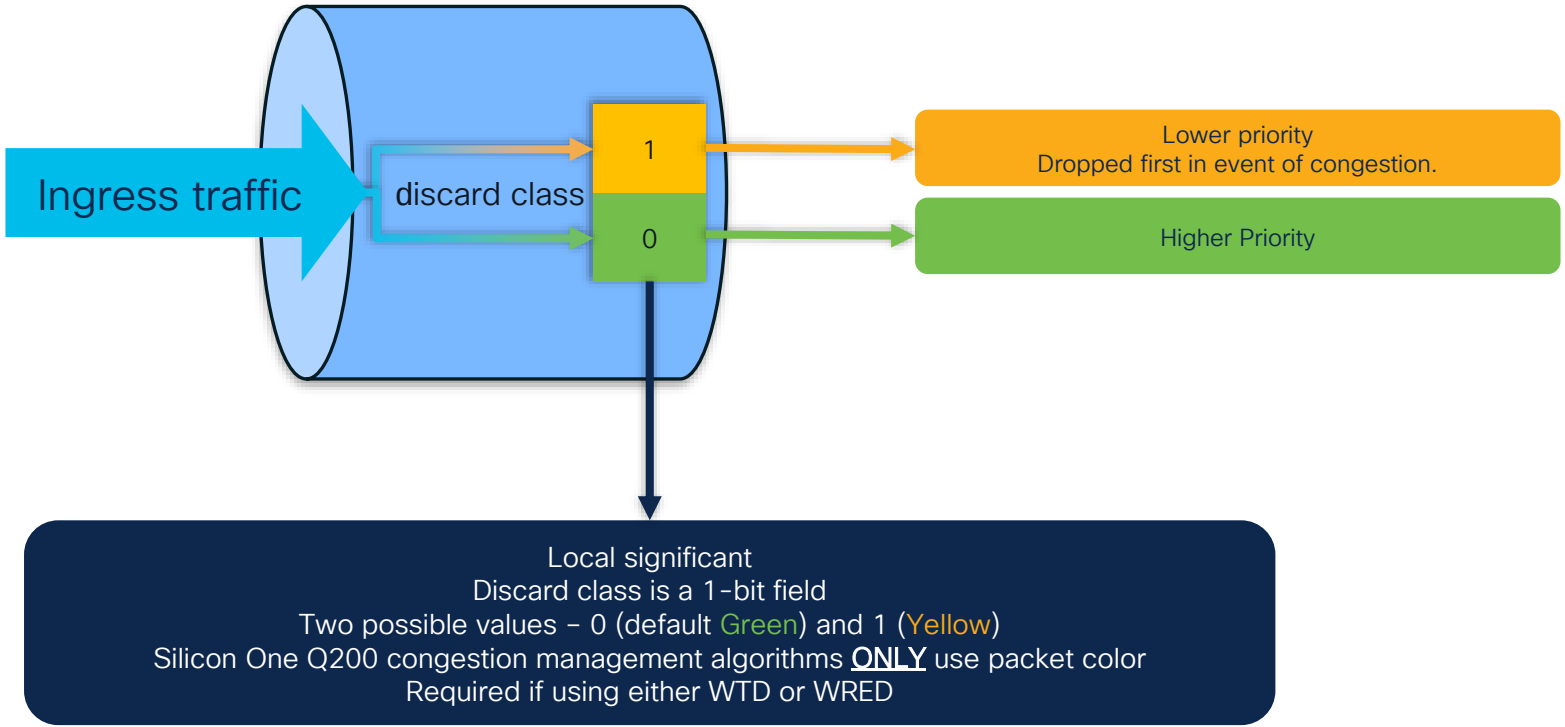

Silicon One Q200 QoS Default

- QoS enabled
- All ports trust at layer2 and layer3
- Two queues (traffic-class 7 and traffic-class 0, traffic-class7 is priority level 1)

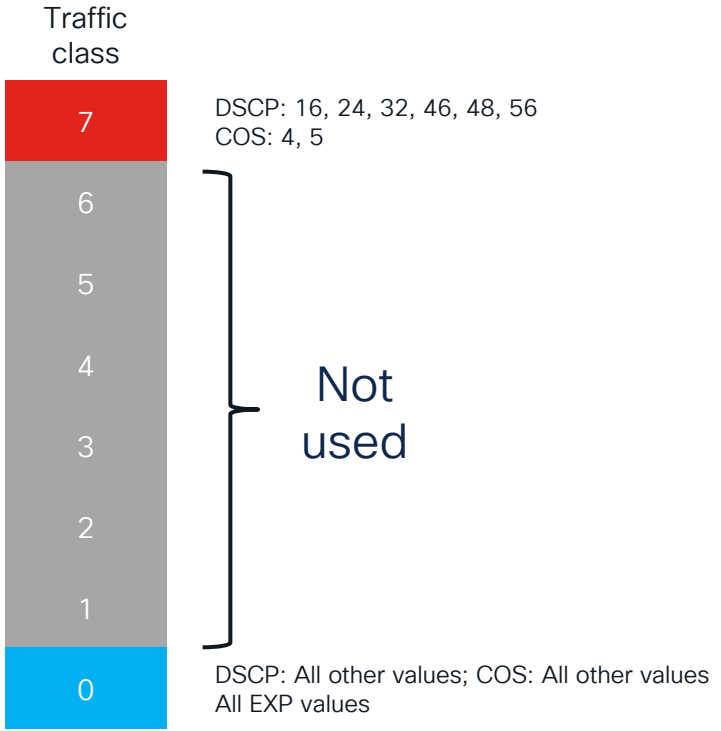

Classification, Marking and Policing

Classification and Marking

- Identify traffic
  - Access Control Lists (ACLs)
  - DSCP
  - IP precedence
  - CoS
  - QoS Group (local with the switch)
  - EXP (MPLS)
  - VLANs
- Marking (coloring)
  - Conditional or unconditional
  - Table map *
  - QoS group (local within switch)
  - Traffic-class (local within switch)
  - Traffic-color (local within switch)

* for policer markdown

Policing – Limit the traffic

1 rate 2 color

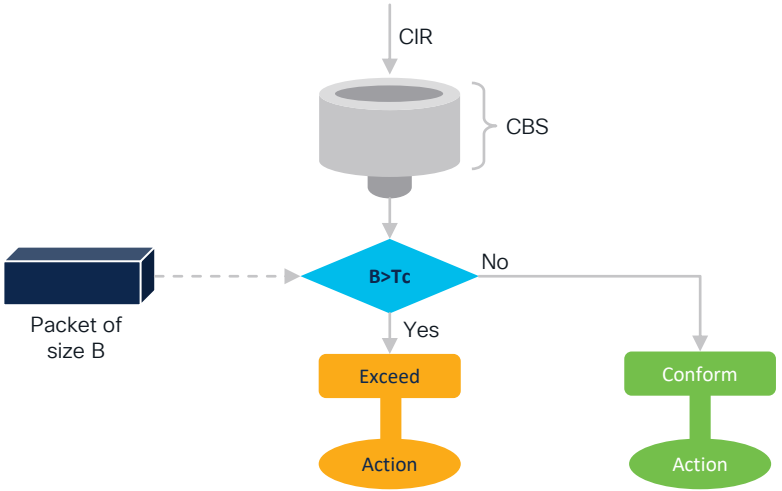


```
police cir 1g
  conform-action transmit
  exceed-action drop
```

2 rate 3 color

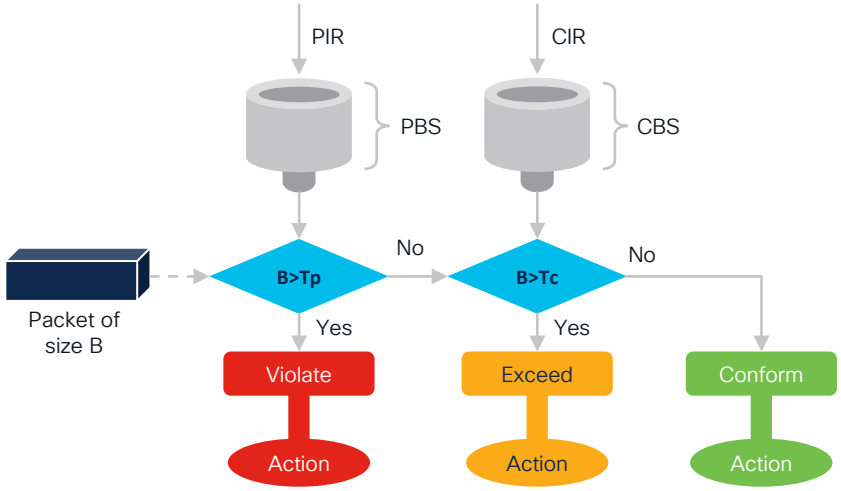


```
police cir percent 10 pir percent 50
  conform-action transmit
  exceed-action set-discard-class-transmit 1
  violate-action drop
```

Traffic color

CIR – Committed Information Rate
PIR – Peak Information Rate

PBS- Peak Burst Size
CBS – Committed Burst Size

Policing and marking/coloring example

Unconditional Traffic Marking/Coloring

```
policy-map ingress-policy
 class class-5-green
  set traffic-class 5
 class class-5-yellow
  set traffic-class 5
  set discard-class 1
```

Different class-map
Same traffic-class

1R2C Policing:

```
policy-map test-police-1R2C
 class dscp1
  set traffic-class 3
  police rate 10g bps
     conform-action transmit
     exceed-action drop
  !
```

Conditional Traffic Marking/Coloring

```
policy-map ingress-policy
 class class-5
   set traffic-class 5
   police rate 5g bps peak-rate 10g bps
    exceed-action set-discard-class-transmit 1
```

2R3C Policing:

```
policy-map test-police-2R3C
 class dscp1
  set traffic-class 3
  police rate 10g bps peak-rate 20g bps
   conform-action transmit
   exceed-action set-discard-class-transmit 1
   violate-action drop
```

Egress Toolset:

Queueing, Shaping and Scheduling

Queueing

Virtual output Queue (VoQ)
- 8 VoQ on each ingress slices for each interface
- Each traffic-class maps to a VoQ (multiple traffic-classes can map to same VoQ)

- VoQ maps to output Queue.

Output Queue
- 8 output queues (egress) for each interface
- Up to 7 strict priorities (level 1 – highest)
- Traffic-class 7 is always priority level 1, priority level is optional for other traffic-classes

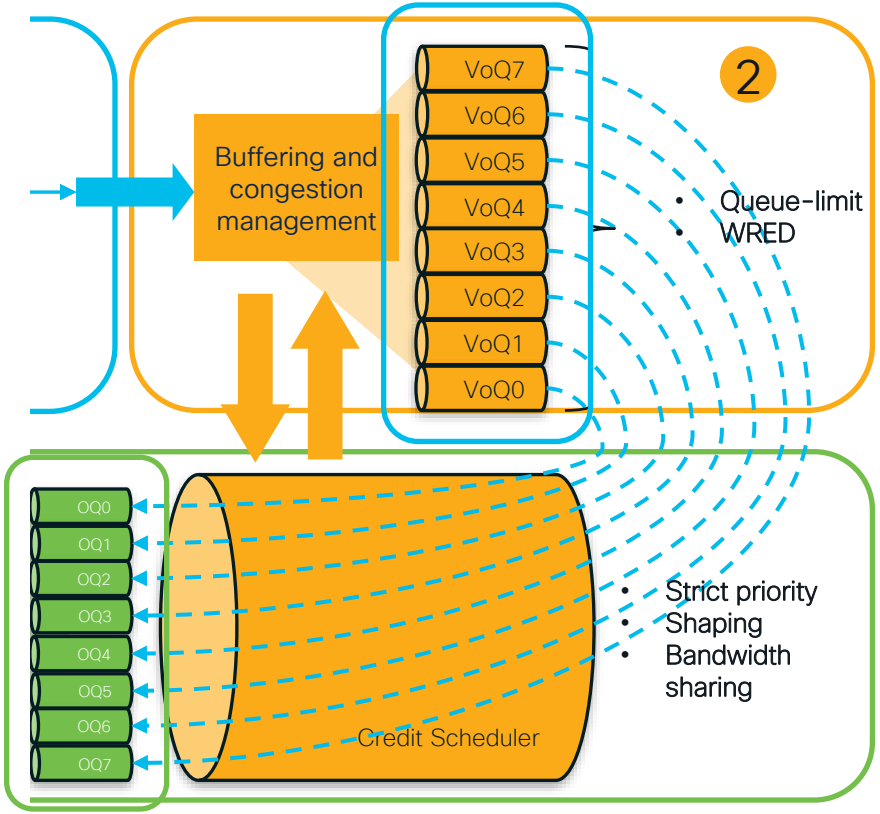

Scheduling

- Packet schedule from VoQ to OQ based on a credit scheduling system
- Packets are buffered at ingress (VoQ)
- Different type of queues are served differently
  - Strict priority queues
    - Always serviced first
    - Up to 7 PQs
  - Normal queues (without priority configured)
    - Served only after priority queues are empty
    - Use Weighted Round Robin (WRR) for scheduling

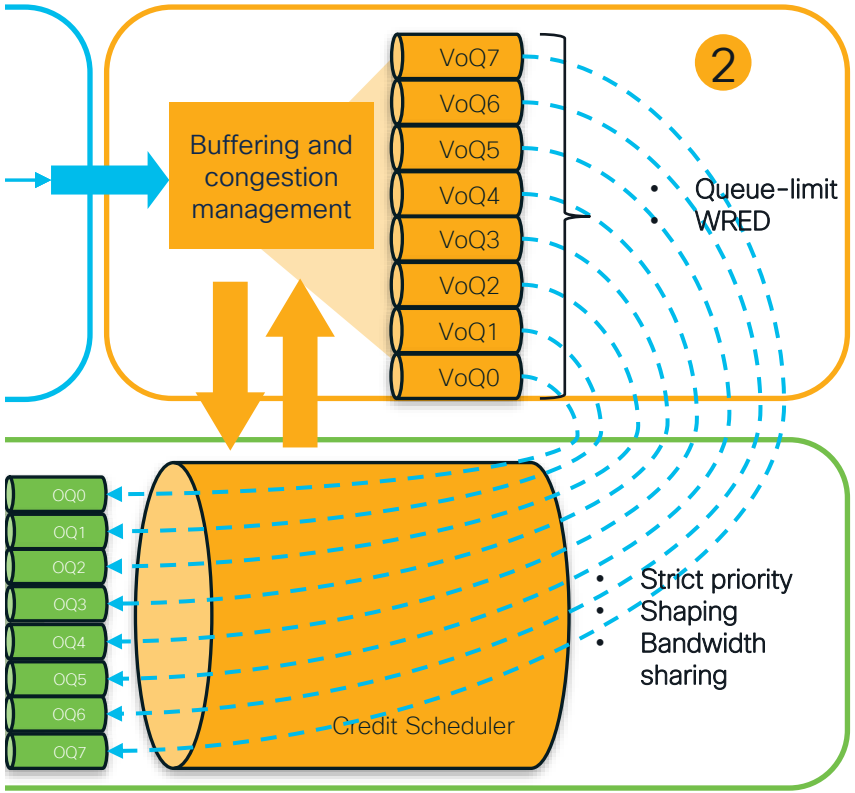

Scheduling - Example

```
class-map match-any tc-7
 match traffic-class 7
class-map match-any tc-6
 match traffic-class 6
…
class-map match-any tc-1
 match traffic-class 1
```

Map traffic to the queues

```
policy-map egress-policy
 class tc-7
  priority-level 1
 class tc-6
  Priority-level 2
 class tc-5
  bandwidth remaining ratio 1
 class t-4
  bandwidth remaining ratio 1
 class tc-3
  bandwidth remaining ratio 1
 class tc-2
  bandwidth remaining ratio 1
 class tc-1
  bandwidth remaining ratio 1
 class class-default
  bandwidth remaining ratio 4
```

- Two priority queues here
- Level 1 has the absolutely priority over level 2

- Use “bandwidth remaining ratio” to assign weight
- This example gives a higher weight to class-default and same weight to rest of the queues
- Served weighted round robin around 6 queues if there isn’t any traffic on the two PQs

Shaping

- Smooth out traffic peaks, microburst, with preserving traffic
- Control traffic rate to the desired value with buffering.
- Usually in the egress direction
- Can be applied on all classes, regardless of priority level.

Shaping Example:

```
policy-map type queueing egress-queueing
  class tc7
    priority level 1
    shape average 1g
   class tc6
     priority level 2
     shape average 5g
...

   class class-default
     shape average 5g
```

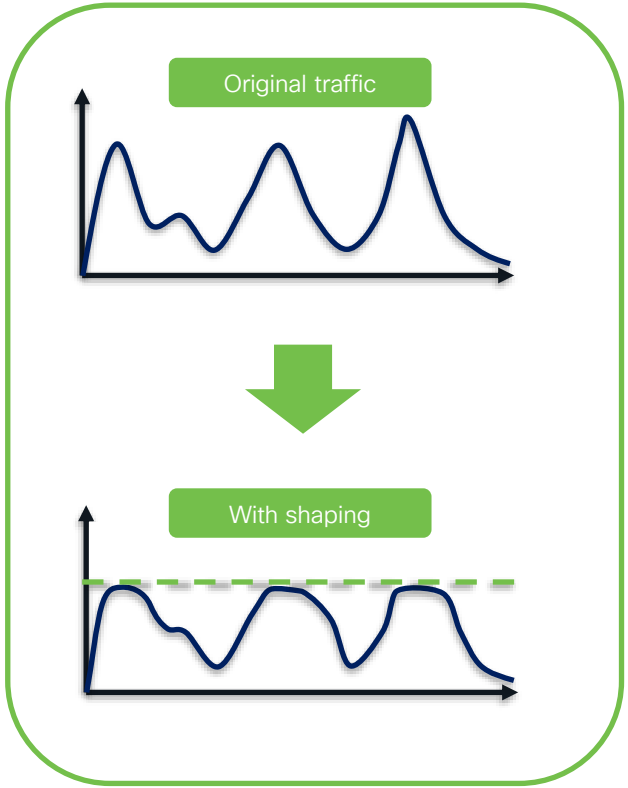

Egress Marking

Egress Marking

- Used to change packet tags of packets egressing the switch.
- A separate policy-map apart from the queueing policy-map.
- If both queueing and marking egress policy-maps are applied, marking happens after queueing actions.
- ACL matching in egress is no supported.

```
class-map match-any dscp-af41
 match dscp af41
!

policy-map egress-map
 class dscp-af41
  set dscp af31
 !

interface interface <#>
 service-policy output egress-map
```

Congestion Management

Weighted Tail Drop (WTD)

Ingress Policy-map

```
policy-map ingress-policy
      class class-5-green
        set traffic-class 5
      class class-5-yellow
        set traffic-class 5
        set discard-class 1
```

```
…
class-map match-any traffic-class-5
 match traffic-class 5
…
```

Egress Policy-map

```
policy-map type queueing queue-policy
 ...
 class traffic-class-5
  queue-limit 100000000 bytes
  queue-limit discard-class 1 50000000 bytes
```

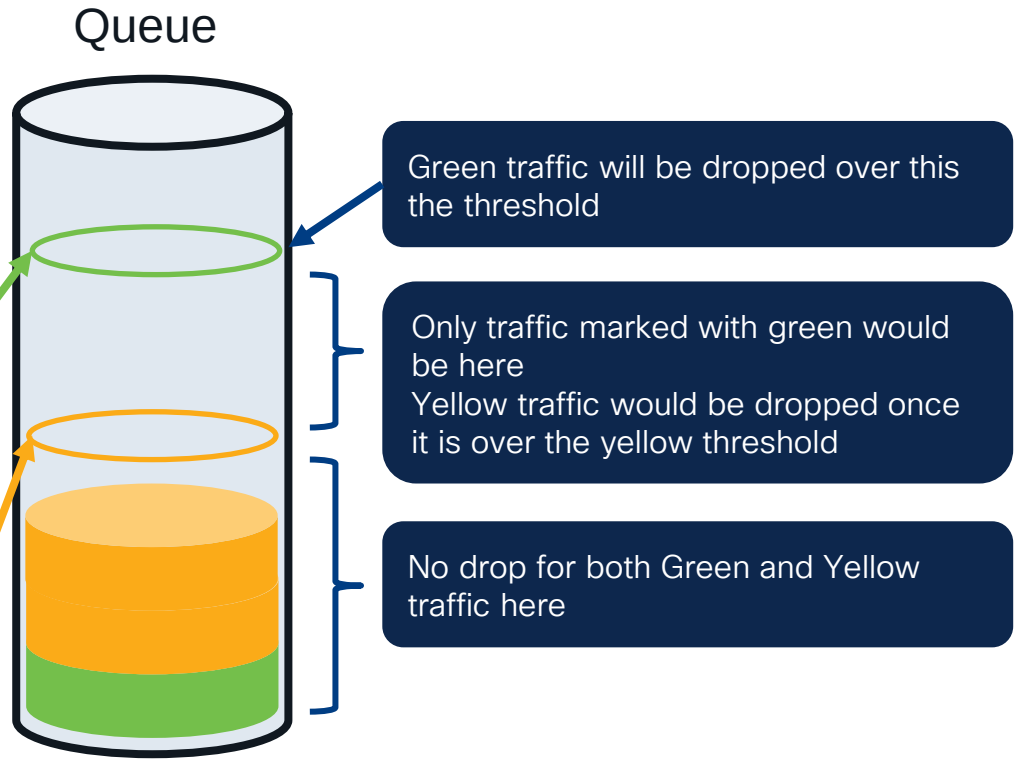

Weighted Random Early Drop (WRED)

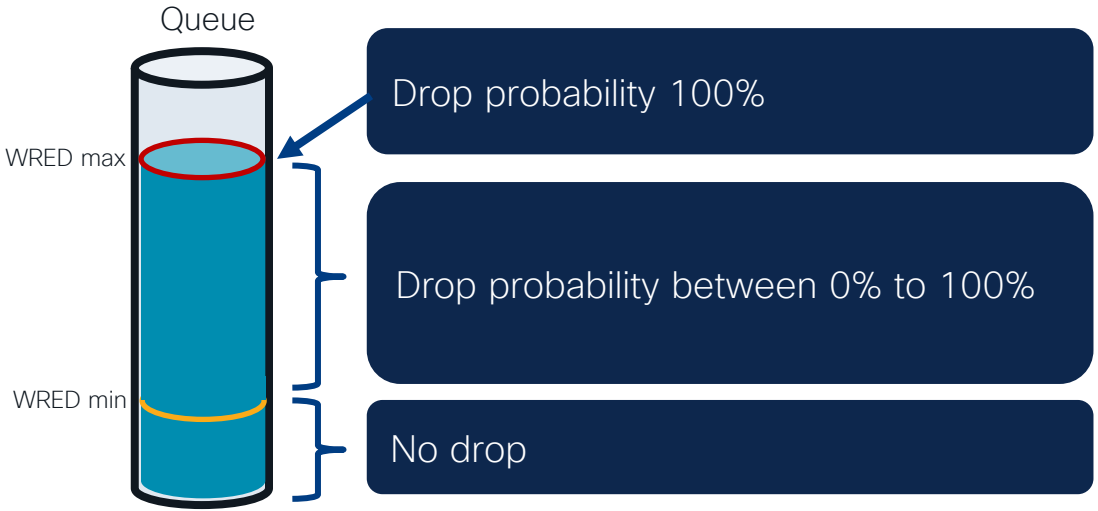

WRED – Drop probability

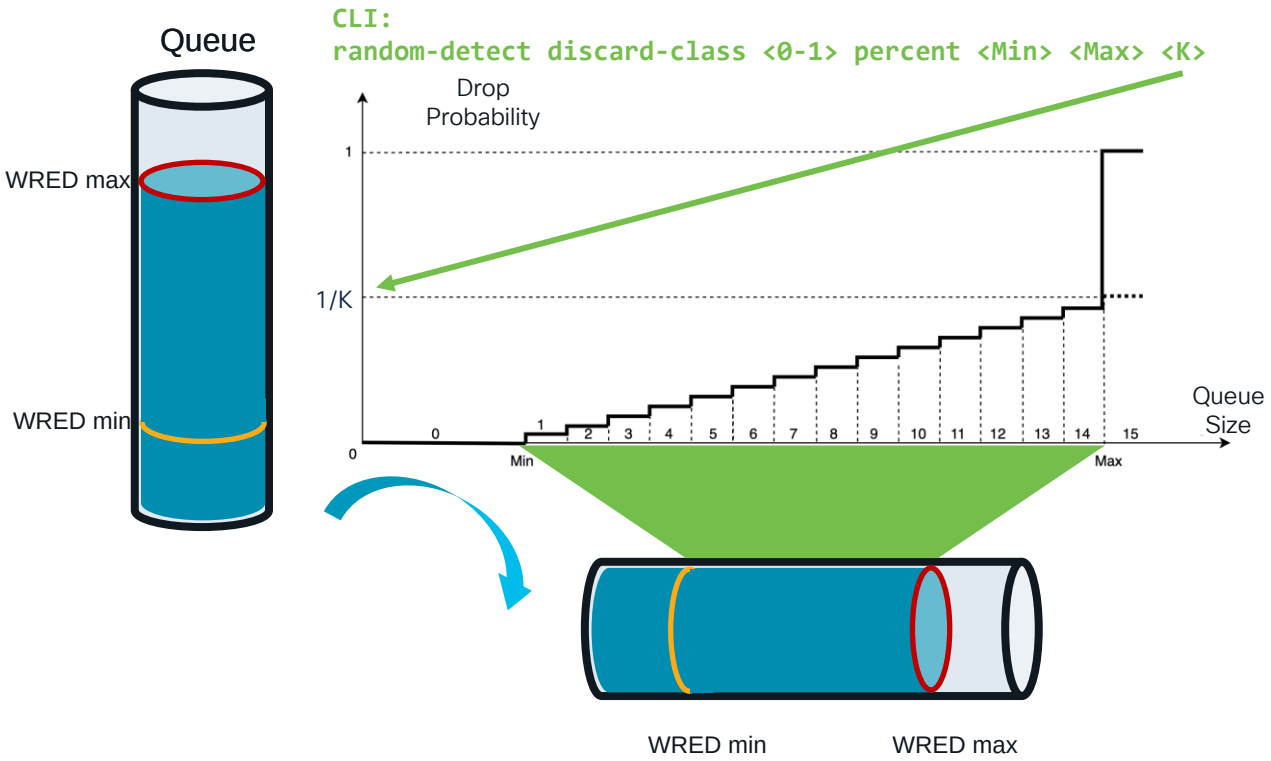

- Drop probability increases as the queue utilization increases
- Silicon One ASIC provides 16 regions (drop probabilities)
- Silicon One ASIC provides a knob to influence the drop probability

WRED – Example

Ingress Policy-map

```
policy-map ingress-policy
…
 class class-5-green
  set traffic-class 5
 class class-5-yellow
  set traffic-class 5
  set discard-class 1
…
```

Ingress policy with marking/coloring of packets

class-map

```
class tc5
  match traffic-class 5
…
```

Map the ingress class to one of the traffic-class

Egress Policy-map

```
policy-map type queueing queue-policy
...
 class tc5
  random-detect discard-class-based
  random-detect discard-class 0 percent 80 90 5
  random-detect discard-class 1 percent 40 70 2
...
```

- Green traffic has higher Min and Max threshold comparing to yellow traffic
- Green traffic also has higher forwarding probability (lower drop probability) comparing to yellow traffic

Buffers

Silicon One Buffers

2 different buffers to address 2 different requirements.

- Shared Memory Sub-system (SMS) buffers (80MB):
  - Low latency packet queueing (video/voice packets)
  - Shallow specialized pool of buffers for faster queueing.
- High Bandwidth Memory (HBM) buffers (8GB):
  - Deep pool of on-demand buffers for guaranteed delivery.
  - Reserve to absorb occasional micro-bursts
  - Absorb speed over-subscription from ingress to egress.

Local vs. HBM Buffers

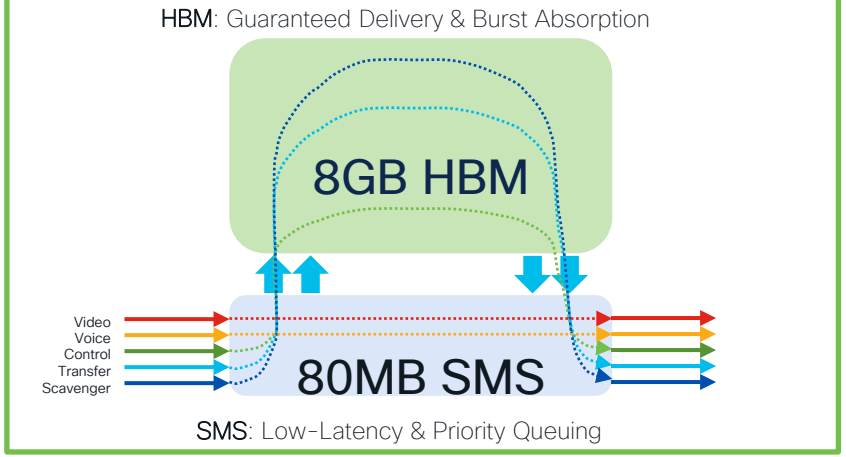

- Packet will always hit the SMS buffers first.
- SMS send the packet to HBM if additional buffers are needed.
- HBM CANNOT send the packet to the output queue, it has to be sent to the SMS again to be sent to the egress.

HQoS

Silicon One Hierarchical QoS (HQoS)

HQoS (two-level hierarchy) allows a parent and child policies on an interface for greater granularity. The Cisco Silicon One supports shaping as parent action.

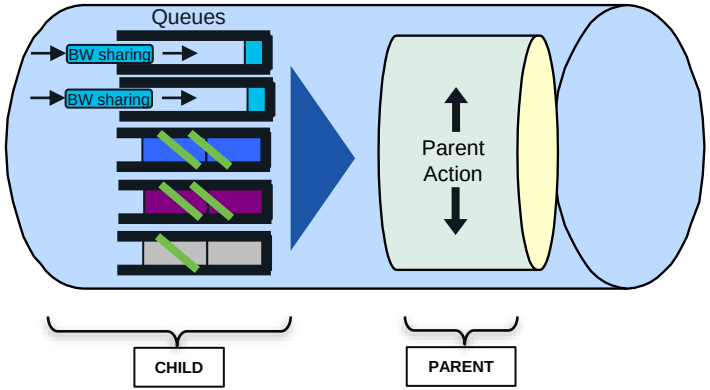

Child Action	Parent Action
Bandwidth sharing + Priority	Shaping

Silicon One vs UADP QoS

QoS tools on UADP and Silicon One ASICs 1/2

Features	UADP ASIC	Silicon One ASIC (Q200)
Trust	Trust all ports by default	Trust all ports by default
Classification	Based on Packet header and ACL for both ingress and egress	Based on packet header and ACL for ingress **Based on packet header for egress. Must classify to traffic-classes at ingress.**
Marking	Header, Table-map, QoS-Group for ingress Header and table-map for egress	Header, Table-map, QoS-group, **traffic-class, discard-class** for ingress Header and table-map for egress
Policing	Both ingress and egress	Ingress only

QoS tools on UADP and Silicon One ASICs 2/2

Features	UADP ASIC	Silicon One ASIC (Q200)
Queueing	Based on header or QoS group Bandwidth and Bandwidth remaining	Based on **traffic-class** **Bandwidth remaining**
Buffering	Dedicated and shared buffer with DTS	SMS: Low-latency & priority queueing HBM: Guaranteed Delivery & Burst Absorption
Shaping	Egress	Egress
Congestion Management	WTD: three thresholds per class WRED: three thresholds' pairs per class	WTD: **two** threshold per class WRED: **two** thresholds' pairs per class

QoS Config Migration

Config Migration steps

1. Define the problem/behavior addressed with QoS.

   Simply copy-pasting existing configs between platform families will always throw errors due to differences in syntax and supported actions between platforms.

2. How many times do you want to split your traffic – Upto 8 queues possible with our ASIC

   Its often not as much as you think you need. Broad generalized splits often are more efficient than granular splits

3. Do you want multiple strict priority classes?

   Know what strict means. All traffic coming into it will be serviced at the expense of other classes.

4. Define traffic shaping/policing or sharing between queues.

   Police/shape priority queues. Use weights to control bandwidth sharing with remaining queues

5. Do you want differential priority for packets within same class? – Use WTD or WRED

   Advanced configuration options, not required for most use cases.

Config Migration from 6k to UADP
MLS configs

Catalyst 6K MLS config

```
!
mls qos
!
interface TenGigabitEthernet2/14
 platform qos queue-mode mode-dscp
 wrr-queue bandwidth 20 1 14 10 15 2 3
 priority-queue queue-limit 20
 wrr-queue dscp-map 1 1 0 1 2 3 4 5 6 7
 wrr-queue dscp-map 2 1 8 14 32 35 36 37 38
 wrr-queue dscp-map 3 1 10 22 24 30 49 50 51 52
 wrr-queue dscp-map 4 1 16 18
 wrr-queue dscp-map 5 1 26
 wrr-queue dscp-map 6 1 48 56
 wrr-queue dscp-map 7 1 34
```

1. Enable QoS Globally
2. Use DSCP mapping to different classes/queues
3. 1 priority queue
4. 7 normal queues WRR and the weights

Config Migration from 6k to UADP

MLS configs

Catalyst 6K MLS config

```
!
mls qos
!
interface TenGigabitEthernet2/14
 platform qos queue-mode mode-dscp
 wrr-queue bandwidth 20 1 14 10 15 2 3
 priority-queue queue-limit 20
 wrr-queue dscp-map 1 1 0 1 2 3 4 5 6 7
 wrr-queue dscp-map 2 1 8 14 32 35 36 37 38
 wrr-queue dscp-map 3 1 10 22 24 30 49 50 51 52
 wrr-queue dscp-map 4 1 16 18
 wrr-queue dscp-map 5 1 26
 wrr-queue dscp-map 6 1 48 56
 wrr-queue dscp-map 7 1 34
```

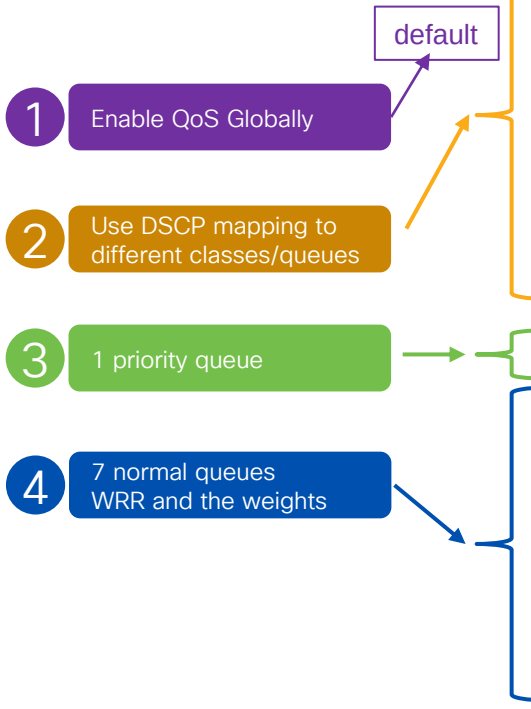

Catalyst 9K(UADP) Config

```
class-map match-any queue1
 match dscp 0 1 2 3 4 5 6 7
class-map match-any queue2
 match dscp 8 14 32 35 36 37 38
class-map match-any queue3
 match dscp 10 22 24 30 49 50 51 52
class-map match-any queue4
 match dscp 16 18
class-map match-any queue5
 match dscp 26
class-map match-any queue6
 match dscp 48 56
class-map match-any queue7
 match dscp 34
class-map match-any priority
 match dscp ef
policy-map egress-queue
 class priority
  priority level 1 percent 20
 class queue1
  bandwidth remaining percent 20
 class queue2
  bandwidth remaining percent 1
 class queue3
  bandwidth remaining percent 14
 class queue4
  bandwidth remaining percent 10
 class queue5
  bandwidth remaining percent 15
 class queue6
  bandwidth remaining percent 2
 class queue7
  bandwidth remaining percent 3
```

* Consider weight as interface speed can be much higher now

Config Migration from 6k to UADP
MQC Configs

Catalyst 6K Configuration

```
class-map type lan-queuing match-all REALTIME
 match dscp ef
class-map type lan-queuing match-all NETWORK_CONTROL
 match dscp cs6  cs7
class-map type lan-queuing match-all VIDEO
 match dscp cs3  af31  af32  af33

policy-map type lan-queuing CAMPUS_EGRESS_6800_POLICY
 class type lan-queuing REALTIME
    priority level 1
 class type lan-queuing NETWORK_CONTROL
    bandwidth remaining percent 10
 class type lan-queuing VIDEO
    bandwidth remaining percent 20
 class class-default
     random-detect dscp-based
     random-detect dscp af11 percent 80 100

Interface gig1/0/1
 service-policy type lan-queueing output
CAMPUS_EGRESS_6800_POLICY
```

1. Use DSCP mapping to different classes/queues
2. 1 priority queue
3. 3 normal queues
   WRR on non-default queue
   WRED for class-default

Config Migration from 6k to UADP
MQC Configs

Catalyst 6K Configuration

```
class-map type lan-queuing match-all REALTIME
 match dscp ef
class-map type lan-queuing match-all NETWORK_CONTROL
 match dscp cs6  cs7
class-map type lan-queuing match-all VIDEO
 match dscp cs3  af31  af32  af33

policy-map type lan-queuing CAMPUS_EGRESS_6800_POLICY
 class type lan-queuing REALTIME
    priority level 1
 class type lan-queuing NETWORK_CONTROL
    bandwidth remaining percent 10
 class type lan-queuing VIDEO
    bandwidth remaining percent 20
 class class-default
    random-detect dscp-based
    random-detect dscp af11 percent 80 100

Interface gig1/0/1
 service-policy type lan-queueing output
CAMPUS_EGRESS_6800_POLICY
```

1. Use DSCP mapping to different classes/queues
2. 1 priority queue
3. 3 normal queues
   WRR on non-default queue
   WRED for class-default

Catalyst 9K Configuration

```
class-map match-any REALTIME
 match dscp ef
class-map match-any NETWORK_CONTROL
 match dscp cs6  cs7
class-map match-any VIDEO
 match dscp cs3  af31  af32  af33

policy-map CAMPUS_EGRESS_POLICY
 class type REALTIME
    priority level 1
 class type NETWORK_CONTROL
    bandwidth remaining percent 10
 class type VIDEO
    bandwidth remaining percent 20
 class class-default
    random-detect dscp-based
    random-detect dscp af11 percent 80 100

Interface gig1/0/1
 service-policy output CAMPUS_EGRESS_POLICY
```

Migration from Catalyst 6K to Silicon One Q200

Catalyst 6K Configuration

```
class-map type lan-queuing match-all REALTIME
 match dscp ef
class-map type lan-queuing match-all NETWORK_CONTROL
 match dscp cs6  cs7
class-map type lan-queuing match-all VIDEO
 match dscp cs3  af31  af32  af33
…
policy-map type lan-queuing CAMPUS_EGRESS_6800_POLICY
 class type lan-queuing REALTIME
    priority level 1
 class type lan-queuing NETWORK_CONTROL
    bandwidth remaining percent 10
 class type lan-queuing VIDEO
    bandwidth remaining percent 20
 class class-default
     random-detect dscp-based
     random-detect dscp af11 percent 80 100
```

1. Classified Based on DSCP value
2. 4 classes (3 defined + default)

3. 4 queues
4. 1 priority queue
5. Scheduling is WRR with “bandwidth remaining”
6. Congestion management is WRED with the default class

Config Migration from Catalyst 6K to Silicon One Q200

Catalyst 6K Configuration

```
class-map type lan-queuing match-all REALTIME
 match dscp ef
class-map type lan-queuing match-all NETWORK_CONTROL
 match dscp cs6  cs7
class-map type lan-queuing match-all VIDEO
 match dscp cs3  af31  af32  af33
…
policy-map type lan-queuing CAMPUS_EGRESS_6800_POLICY
 class type lan-queuing REALTIME
    priority level 1
 class type lan-queuing NETWORK_CONTROL
    bandwidth remaining percent 10
 class type lan-queuing VIDEO
    bandwidth remaining percent 20
 class class-default
     random-detect dscp-based
     random-detect dscp af11 percent 80 100
```

1. Classified Based on DSCP value
2. 4 classes (3 defined + default)
3. 4 queues (traffic-class), traffic-7 is priority level 1

Apply policy on the ingress interface

Catalyst 9K(Q200) Configuration

```
class-map match-all REALTIME
 match dscp ef
class-map match-all NETWORK_CONTROL
 match dscp cs6  cs7
class-map match-all VIDEO
 match dscp cs3  af31  af32  af33
class-map match-all default-green
 match dscp af11


policy-map INGRESS
 class REALTIME
    set traffic-class 7
 class NETWORK_CONTROL
    set traffic-class 6
 class VIDEO
    set traffic-class 5
 class default-green
    set traffic-class 0
 class class-default
    set discard-class 1
```

Note: class-default is always assigned with traffic-class 0

Config Migration from Catalyst 6K to Silicon One Q200

Map the traffic-class marking defined on the ingress

Apply policy on the egress interface

Catalyst 6K Configuration

```
policy-map type lan-queuing CAMPUS_EGRESS_6800_POLICY
 class type lan-queuing REALTIME
    priority level 1
 class type lan-queuing NETWORK_CONTROL
    bandwidth remaining percent 10
 class type lan-queuing VIDEO
    bandwidth remaining percent 20
 class class-default
     random-detect dscp-based
     random-detect dscp af11 percent 80 100
```

Catalyst 9K(Q200) Configuration

```
class-map tc7
 match traffic-class 7
class-map tc6
 match traffic-class 6
class-map tc5
 match traffic-class 5

policy-map type queuing EGRESS
 class tc7
    priority level 1
 class tc6
    bandwidth remaining ratio 1
 class tc5
    bandwidth remaining ratio 2
 class class-default
    random-detect discard-class-based
    random-detect discard-class 0 percent 80 100
    random-detect discard-class 1 percent 40 100
```

3. 1 priority queue
4. Scheduling is WRR with “bandwidth remaining”
5. Congestion management is WRED with the default class

Note: C6K WRED default min is 40, max is 100.

Summary

Do we need QoS?

User Experience

Guaranteeing voice quality	Bandwidth Savvy Business Applications	protect network infrastructure to deal with abnormal events
Video Quality	de-prioritizing non-business applications	protecting the control planes

QoS helps define the latency priority for your traffic packets

Catalyst 9K Switching sessions @ CL23, US

Session Title	Session ID	Session Type
Catalyst 9000 Switching Family Architecture	TECARC-2446	Sun, 9:00am – 1:00pm
The Catalyst 9000 Switch Family – Access	BRKARC-2098	Wed, 1:00pm – 2:30pm
The Catalyst 9000 Switch Family – Core and Distro	BRKARC-2099	Tue, 4:00pm – 5:00pm
Catalyst 9000 SiliconOne and IOSXE Architecture & Innovations	BRKARC-2092	Wed, 2:30pm – 3:30pm
123 – Enterprise Campus Wired Design Fundamentals	BRKENS-1501	Mon, 8:00am – 9:00am
Designing Highly Available Networks using Catalyst 9000 Switches	BRKENS-2095	Tue, 4:00pm – 5:00pm
Designing High availability for your Wired and Wireless Access Networks	TECENS-2001	Sun, 9:00am – 1:00pm
Catalyst 9000 Switching QoS Deep Dive	BRKENS-2096	Mon, 1:00pm – 2:30pm
Catalyst 9000 Family Software Innovations and Solutions	TECENS-2618	Sun, 2:00pm – 6:00pm
Service Assurance with ThousandEyes on Catalyst 9000	BRKENS-1095	Thu, 11:00am – 12:00pm
Catalyst Powered Smart Buildings - Beyond PoE Connectivity	BRKENS-2091	Mon, 3:00pm – 4:30pm
Enabling Cloud Services at the Edge with App Hosting on Catalyst 9000	BRKENS-1090	Mon, 2:30pm – 3:30pm
BGP EVPN in Enterprise Campus using Catalyst 9000 Switches	BRKENS-2092	Tue, 1:00pm – 2:30pm
Building Time Sensitive Networks with Catalyst 9000 Switching Platforms	BRKENS-2098	Wed, 4:00pm – 5:00pm
The Industry's only Flood-Free mDNS Experience in the Enterprise Campus, Powered by Catalyst 9000	BRKENS-2097	Thu, 8:00am - 9:00am
Catalyst 9000 Switches and Cisco DNA Advantage	BRKENS-1093	Mon, 9:30am – 10:30am
Infrastructure as Code and the Cisco Catalyst 9000 Virtual Switch	BRKDEV-2467	Thu, 8:00am – 9:00am

Fill out your session surveys!

Attendees who fill out a minimum of four session surveys and the overall event survey will get **Cisco Live-branded socks** (while supplies last)!

Attendees will also earn 100 points in the **Cisco Live Challenge** for every survey completed.

These points help you get on the leaderboard and increase your chances of winning daily and grand prizes

Continue your education

- Visit the Cisco Showcase for related demos
- Book your one-on-one Meet the Engineer meeting
- Attend the interactive education with DevNet, Capture the Flag, and Walk-in Labs
- Visit the On-Demand Library for more sessions at www.CiscoLive.com/on-demand

The bridge to possible

Thank you

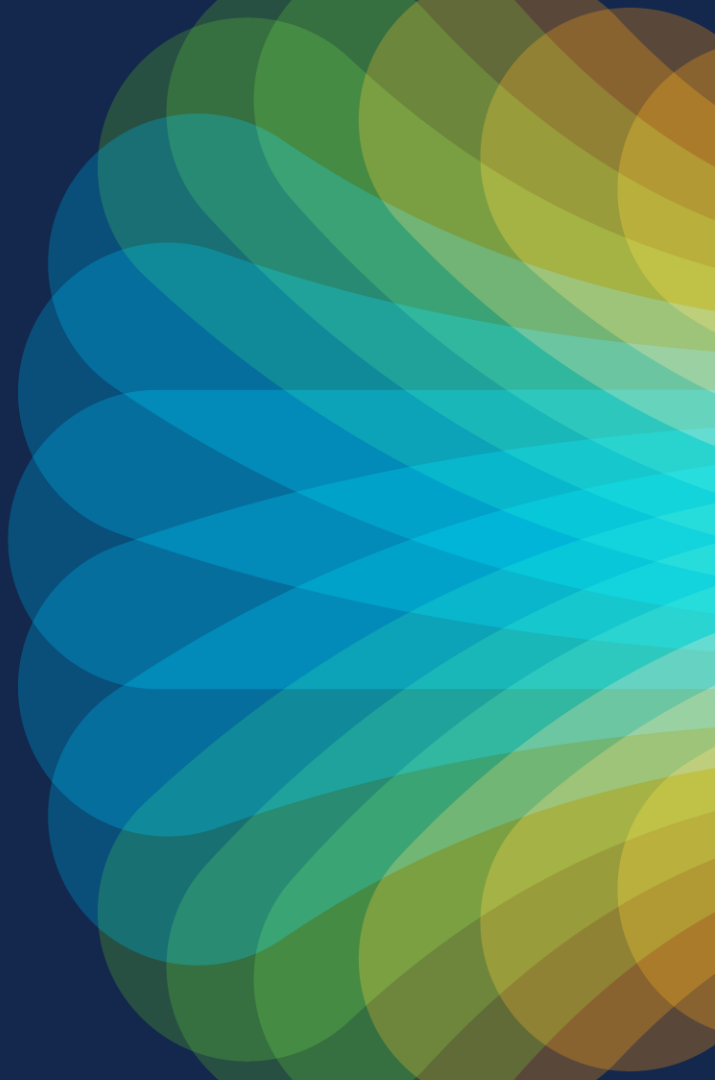

Cisco Live
Challenge

Gamify your Cisco Live experience!

Get points for attending this session!

How:

1. Open the Cisco Events App.
2. Click on 'Cisco Live Challenge' in the side menu.
3. Click on View Your Badges at the top.
4. Click the + at the bottom of the screen and scan the QR code:

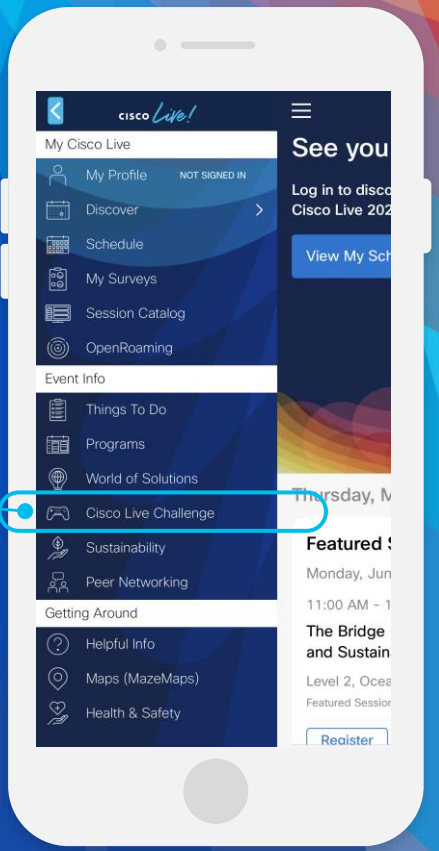

cisco Live!


cisco Live!
Let's go
#CiscoLive