

Pseudowire QoS

In This Chapter

This section describes pseudowire QoS configurations.

Topics in this section include:

- [Applicability on page 1654](#)
- [Overview on page 1655](#)
- [Configuration on page 1660](#)
- [Conclusion on page 1677](#)

Applicability

This example is applicable to the 7950 XRS-16c/20, 7750 SR-7/12, 7750 SR-c4/12 and 7450 ESS-6/6v/7/12 platforms when all network IP interfaces are on IOM3-XP/IMM (FP2 and above) hardware.

The configuration was tested on release 11.0R4. There are no other specific pre-requisites for this configuration.

Overview

A pseudowire (PW) provides a virtual connection across an IP or MPLS network between services configured on provider edge (PE) devices. From SR OS R10. OR 1, it is possible to provide specific QoS to either a single pseudowire or a multiple pseudowires. This is supported for the following:

- SDP
 - MPLS
 - GRE
- Epipe
 - Including vc-switching and dynamic MS-PW
 - PBB-epipe
 - BGP-VPWS (from 11.0R1)
- VPLS
 - Mesh and spoke SDP
 - LDP signaled pseudowires
 - BGP-AD signaled pseudowires
 - I-VPLS, B-VPLS
 - BGP-VPLS
- Spoke termination on IES/VPRN (both Epipe and Ipipe)
- Apipe (from R10.0R4)
- Cpipe (from R10.0R4)
- Fpipe (from R10.0R4)
- Ipipe (from R10.0R4)

It is supported at ingress on both Ethernet and POS/TDM ports on an IOM3-XP/IMM and only on Ethernet ports at egress.

Bandwidth control is achieved using queue-groups which are implemented per FP (flexpath) at the ingress and per port at the egress (these being relative to the data path through the system), as shown in [Figure 253](#) and [Figure 254](#), respectively.

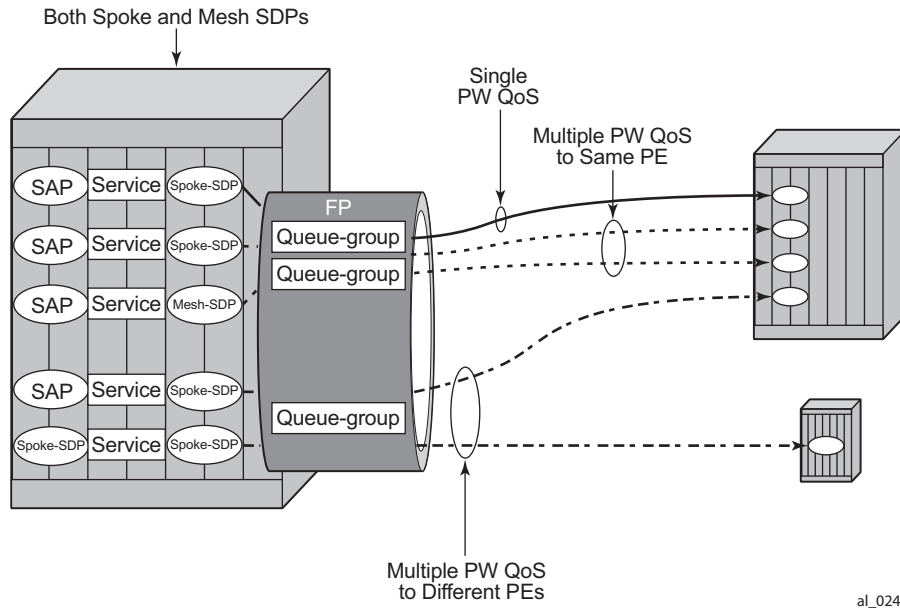


Figure 253: Ingress PW QoS

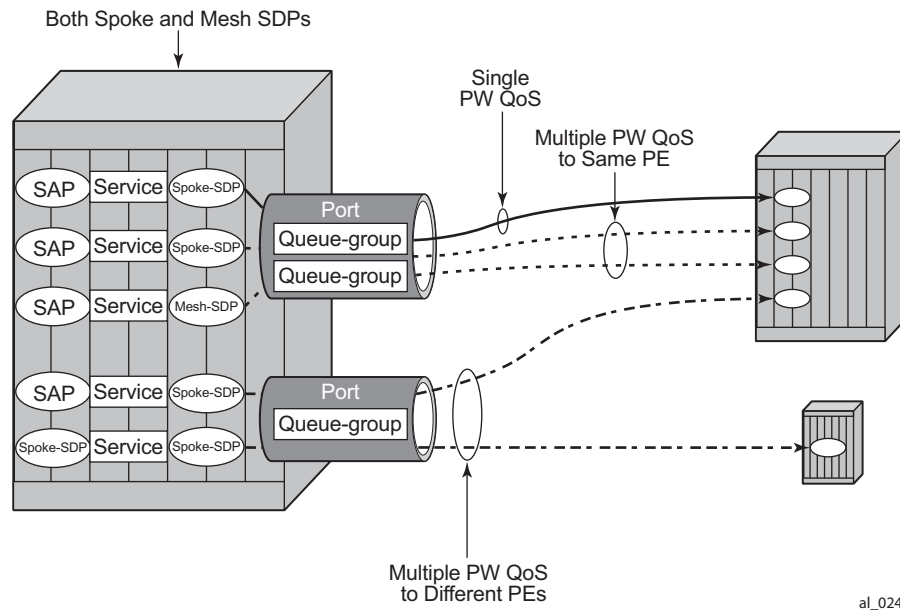


Figure 254: Egress PW QoS

Bandwidth control is applied independently for ingress and egress, and can be set up for a single pseudowire or for multiple pseudowires where the remote services are located on a single PE or on multiple PEs.

It is possible to benefit from Hierarchical QoS which can be configured under the queue-groups, but this is beyond the scope of this example.

The ingress and egress classification and egress marking is configured by applying a network QoS policy to each pseudowire.

Ingress QoS

Ingress QoS is achieved using a queue group which is applied to an ingress FP on a card. Queue groups applied to an FP can only contain policers, not queues. The network QoS policy applied to the pseudowire redirects forwarding classes (FCs) to the individual queue group (unicast or multipoint) policers. The actual queue group to be used is defined separately to the network QoS policy, thereby allowing the network QoS policies to be independent from the queue groups used and therefore both are reusable.

Ingress bandwidth control does not take into account the outer Ethernet header, the MPLS labels/control word or GRE headers, or the FCS of the incoming frame. The configuration allows an offset to be added or subtracted from the received frame size in order to change the actual length used for the bandwidth control.

For example: if the same ingress rate is configured on a pseudowire (without a control word) and a dot1q SAP, what packet-byte-offset needs to be used on the pseudowire in order to achieve the same throughput as on the SAP?


- SAP — The following shows the bytes in the frame that are used by default on a policer for the rate at a SAP ingress.

6B	6B	4B	2B	xxxxB	4B
Source MAC	Dest. MAC	802.1Q	Ether Type	Payload	CRC/FCS

al_0247


- VPLS Pseudowire — For a tagged (**vc-type vlan**) pseudowire, it would be necessary to add 4 bytes using the packet-byte-offset applied to the ingress policer in order to achieve the same throughput as on the SAP. This compensates for the omission of the FCS that is included on the SAP and so needs to be added.

Egress QoS



al_0248

- VPRN Pseudowire — For an Ipipe (**vc-type** ipipe) pseudowire, it would be necessary to add 22 bytes using the packet-byte-offset to the ingress policer to achieve the same throughput as on the SAP. This compensates for the omission of the source and destination MAC addresses (12 bytes), Ether type (2 bytes), VLAN tag (4 bytes) and the FCS (4 bytes) that are included on the SAP and so needs to be added.



al_0249

The ingress classification is configured in the ingress section of the network QoS policy and is based on the outer encapsulation header only, the outer Ethernet header (dot1p/DE), MPLS labels (EXP) or GRE headers (DSCP). At an egress LER, the `ler-use-dscp` is applicable only to IES and VPRN pseudowires.


Egress QoS

Egress QoS is achieved using a queue group which is applied to an egress port. Queue groups applied to a port can contain both policers and queues. The network QoS policy applied to the pseudowire redirects forwarding classes (FCs) to the individual queue group policers/queues. The actual queue group to be used is defined separately to the network QoS policy, thereby allowing the network QoS policies to be independent from the queue groups used and therefore both are reusable.

Egress bandwidth control does takes into account the outer Ethernet header, MPLS labels/control word or GRE headers, and the FCS of the outgoing frame. The configuration allows an offset to be added or subtracted from the sent frame size in order to affect the actual length used for the bandwidth control.


For example, if the same egress rate is configured on a pseudowire (without a control word) and a dot1q SAP, what packet-byte-offset needs to be used on the pseudowire in order to achieve the same throughput as on the SAP?

- SAP — The following shows the bytes in the frame that are used by default on a policer/queue at a SAP egress.




al_0250

- VPLS Pseudowire — For a tagged (**vc-type vlan**) pseudowire, it would be necessary to subtract 22 bytes using the packet-byte-offset applied to the egress policer/queue applied to achieve the same throughput as on the SAP. This compensates for the MPLS header (source and destination MAC addresses (12 bytes), Ether type (2 bytes), two labels (8 bytes)) that is not included on the SAP and needs to be subtracted.



al_0251

- VPRN Pseudowire — For an Ipipe (**vc-type ipipe**) pseudowire, it would be necessary to subtract 4 bytes using the packet-byte-offset applied to the egress policer/queue applied to achieve the same throughput as on the SAP. This compensates for the MPLS header (source and destination MAC addresses (12 bytes), Ether type (2 bytes), two labels (8 bytes)) that is not included on the SAP so is subtracted, and the source and destination MAC addresses (12 bytes), dot1q header (4 bytes) and Ether type (2 bytes) of the SAP frame which needs to be added. This results in subtracting 4 bytes.



al_0252

The egress classification and marking is configured in the egress section of the network QoS policy. DSCP/prec egress reclassification is supported from release R10.0R4 for IES and VPRN spoke SDPs. The egress marking affects the outer encapsulation header, the outer Ethernet header (dot1p/DE), MPLS labels (EXP) or GRE headers (DSCP).

Configuration

The configuration of pseudowire QoS is described using an Epipe pseudowire. The topology is shown in Figure 255.

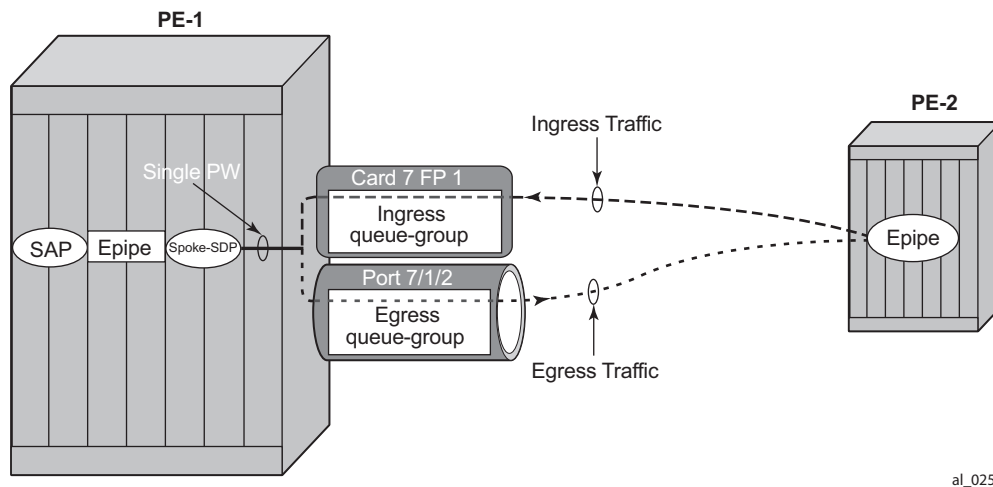


Figure 255: Example Epipe Pseudowire Topology

The following pre-requisite configuration is assumed to be in place:

- Hardware provisioning
- IP address and routing
- MPLS protocols
- SDP
- Epipe service, including the SAP
- SAP QoS policies

Traffic is sent across a virtual leased line between PE-1 and PE-2 using Epipes with a pseudowire configured as a spoke SDP on each PE. The QoS is applied to the pseudowire at the ingress and egress of PE-1.

The following configuration is required for applying pseudowire QoS:

- Create the ingress and egress queue groups.
These contain the ingress policer and egress policer/queue definitions.

- Create an instance of the ingress queue group on the ingress FP and instance of the egress queue group on the port that will be used for the pseudowire traffic.
- Create a network QoS policy to redirect the traffic to the ingress and egress queue groups, and to perform the ingress classification and egress marking.
- Apply the network QoS policy, together with the reference to the ingress and egress queue group instances, to the spoke SDP representing the pseudowire.

The traffic consists of two bidirectional flows, one in FC BE and one in FC EF. At the ingress of the pseudowire, each FC is assigned to its own policer, whereas at the egress of the pseudowire, FC BE is assigned to a queue and FC EF is assigned to a policer.

Although this example makes use of both ingress and egress queue groups, the focus is pseudowire QoS, so the full details of queue group configuration are not covered.

Create the Ingress and Egress Queue Groups

Queue groups are created using templates, which are separate for ingress and egress. The following shows the queue group templates configured.

```
configure qos
  queue-group-templates
    ingress
      queue-group "ingress-queue-group" create
      policer 1 create
        rate 6000
        packet-byte-offset add 4
      exit
      policer 2 create
        rate 4000
        packet-byte-offset add 4
      exit
    exit
  exit
  egress
    queue-group "egress-queue-group" create
    queue 1 best-effort create
      rate 6000
      xp-specific
        packet-byte-offset subtract 22
    exit
    policer 1 create
      rate 4000
      packet-byte-offset subtract 22
    exit
  exit
exit
```

Create the Ingress FP and Egress Port Queue Group Instances

The ingress queue group has two policers associated with it; policer 1 will be used for the FC BE traffic and policer 2 will be used for the FC EF traffic. The configuration of policers in an ingress queue group is the same as that in a sap-ingress QoS policy, with the exception that the percent-rate is not supported within the queue group.

In order to achieve the same ingress throughput as that when applying the same rates to policers on a dot1q tagged SAP, the packet-byte-offset adds 4 bytes to the packet length for both policers.

The egress queue group has one queue (queue 1) that will be used for the FC BE traffic and one policer (policer 1) that will be used for the FC EF traffic. The configuration of policers in an egress queue group is the same as that in a sap-egress QoS policy, with the exception that the percent-rate is not supported within the queue group. The configuration of queues in an egress queue group is the same as in a sap-egress QoS policy, with the exception that the avg-frame-overhead is not supported within the queue group.

In order to achieve the same egress throughput as that when applying the same rates to policers/queues on a dot1q tagged SAP, the packet-byte-offset subtracts 22 bytes from the packet length for both the policer and queue.

Rates have been configured such that the ingress and egress capacity of the BE traffic is 6Mb/s and 4Mb/s for the EF traffic.

Create the Ingress FP and Egress Port Queue Group Instances

The queue group templates are then applied as individual instances to the ingress FP and egress port; using instances allows the reuse of the same template.

Below is the ingress FP configuration. From a QoS perspective, it is also possible to configure a policer-control-policy under the ingress queue group in order to perform hierarchical policing. From R11.0R4, the configuration supports overrides for both the policer-control-policy parameters and some of the queue group policer parameters.

```
configure
  card 7
    card-type imm5-10gb-xfp
    mda 1
      no shutdown
    exit
    fp 1
      ingress
        network
          queue-group "ingress-queue-group" instance 1 create
        exit
      exit
    exit
  exit
  no shutdown
exit
```

Below is the egress port configuration. From a QoS perspective, it is also possible to configure under the egress queue group a policer-control-policy in order to perform hierarchical policing, a scheduler-policy in order to perform hierarchical shaping and overrides for some of the queue group queue parameters.

```
configure
  port 7/1/2
    ethernet
      network
        egress
          queue-group "egress-queue-group" instance 1 create
          exit
        exit
      exit
    exit
  exit
  no shutdown
exit
```

If there are redundant network interfaces over which the pseudowire traffic can enter or exit the system, it is necessary to configure any ingress FP and egress port queue groups consistently across all possible interfaces to be used by the pseudowire to ensure the QoS is always applied. If a queue group configuration was omitted, the pseudowire would not be subject to the QoS defined in that queue group.

If a LAG is used, the system only allows the egress port queue group to be added or removed from the LAG primary port, thereby keeping the LAG configuration consistent. However, this is not possible at the ingress as the queue-group is applied at the FP, so it is necessary to ensure that the ingress queue group is applied consistently on all FP's corresponding to the configured LAG.

Create the Network QoS Policy

A network QoS policy is created to redirect ingress and egress traffic to the respective queue groups, and perform ingress classification (in this example).

The redirection to the queue group policer/queue is performed per FC.

At ingress, traffic can be redirected to policers (being the same or different policers) based on the traffic type. Unicast traffic is redirected to a policer specified by the policer command and will use the ingress shared policer-output-queues to access the switch fabric. All multipoint traffic is redirected to the policer specified by the multicast-policer command (for example with a pseudowire configured in a VPLS service, all broadcast, unknown and multicast traffic will use this policer). The multipoint traffic accesses the switch fabric using the Ingress Multicast Path Management queues. It is possible to individually redirect one traffic type (unicast or multipoint) within an FC to a queue group policer while allowing the other traffic type to use default network queues.

At egress, traffic can be redirected to a queue or to a policer. The policed traffic will exit the egress port using one of the default network queues (with the queue chosen by FC assignment) or optionally can use a queue in the egress queue group if configured in the port-redirect-group command following the policer parameter.

Any FC not redirected to a queue-group, will continue to use the regular default network ingress and egress queues.

The syntax for the FC redirection is as follows.

```
config# qos
  network <network-policy-id> [create]
    ingress
      fc <fc-name>
        fp-redirect-group multicast-policer <policer-id>
        fp-redirect-group policer <policer-id>
    egress
      fc <fc-name>
        port-redirect-group {queue <queue-id>|
          policer <policer-id> [queue <queue-id>]}
```

The required commands are shown below.

```
configure qos
  network 10 create
    ingress
      lsp-exp 5 fc ef profile in
      fc be
        fp-redirect-group policer 1
      exit
      fc ef
        fp-redirect-group policer 2
```

```

        exit
    exit
    egress
        fc be
            port-redirect-group queue 1
        exit
        fc ef
            port-redirect-group policer 1
        exit
    exit
exit

```

At ingress, the FC BE and FC EF traffic are redirected to the two policers in the queue-group applied to the FP. At egress, the two FCs are redirected to the queue and policer in the queue group applied to the egress port.

The ingress classification required here is for the traffic which is received with exp=5 to be in FC EF.

Apply Network QoS Policy with Queue Group Instances to the Spoke SDP

To apply the QoS to the pseudowire, the following commands can be used, dependent on the service type.

```

config# service {apipe|cpipe|epipe|fpipe|ipipe} <service-id>
    spoke-sdp <sdp-id:vc-id>
        ingress
            qos <network-policy-id> fp-redirect-group <queue-group-name>
                                                    instance <instance-id>
        egress
            qos <network-policy-id> port-redirect-group <queue-group-name> instance <instance-
            id>

config# service {ies|vprn} <service-id>
    interface <ip-int-name>
        spoke-sdp <sdp-id:vc-id>
            ingress
                qos <network-policy-id> fp-redirect-group <queue-group-name> instance <instance-id>
            egress
                qos <network-policy-id> port-redirect-group <queue-group-name>
                    instance <instance-id>

config# service vpls <service-id>
    {spoke-sdp|mesh-sdp} <sdp-id:vc-id>
        ingress
            qos <network-policy-id> fp-redirect-group <queue-group-name> instance <instance-id>
        egress
            qos <network-policy-id> port-redirect-group <queue-group-name>
                instance <instance-id>

```

Apply Network QoS Policy with Queue Group Instances to the Spoke SDP

For services using BGP auto-discovery to signal the pseudowire, the QoS configuration is included in the pseudowire template.

```
config# service pw-template <policy-id>
  ingress
    qos <network-policy-id> fp-redirect-group <queue-group-name> instance <instance-id>
  egress
    qos <network-policy-id> port-redirect-group <queue-group-name>
                                     instance <instance-id>
```

To propagate changes in a pw-template to existing BGP-AD pseudowires, it is necessary to use the following command:

```
tools perform service eval-pw-template policy-id
```

Note that the allow-service-impact parameter is not required for changing the ingress or egress QoS definition as these do not affect the operational state of the pseudowire.

QoS applied directly to a pseudowire, using the above commands, takes precedence over any QoS applied to the network interface (using a network QoS policy with or without queue group redirection).

Note that each time a pseudowire uses a network egress port queue group an FP resource is allocated. This only requires that the pseudowire egress QoS is configured with a port-redirect-group, and will occur even if there are no FCs redirected using a port-redirect-group within the configured network QoS policy. The resources used can be seen using the **tools dump system-resources** command and is listed under Egr Network Queue Group Mappings which is part of the total for the “Dynamic Service Entries “.

As an Epipe is used in this example, QoS is configured directly under a spoke SDP.

```
configure service
  epipe 1 customer 1 create
    spoke-sdp 1:1 vc-type vlan create
      ingress
        qos 10 fp-redirect-group "ingress-queue-group" instance 1
      exit
      egress
        qos 10 port-redirect-group "egress-queue-group" instance 1
      exit
      no shutdown
    exit
  no shutdown
exit
```

The created network QoS policy is applied at both ingress and egress, with the ingress referencing the ingress queue group instance applied to the FP and the egress referencing the egress queue group instance applied to the port.

Show Output

The configured ingress queue group can be shown, including the details of the configured policers and where it is applied.

```
*A:PE-1# show qos queue-group "ingress-queue-group" ingress detail
=====
QoS Queue-Group Ingress
=====
-----
QoS Queue Group
-----
Group-Name      : ingress-queue-group
Description     : (Not Specified)
-----
...
-----
Queue Group FP Maps
=====
Card Num      Fp Num      Instance      Type
-----
7             1           1             Network
-----
Entries found: 1
-----
-----
Queue Group Policer
=====
Policer Id     : 1
Description    : (Not Specified)
PIR Adptn     : closest          CIR Adptn      : closest
Parent        : none             Level          : 1
Weight        : 1             Adv. Cfg Plcy : none
Admin PIR     : 6000          Admin CIR     : 0
CBS           : def           MBS           : def
Hi Prio Only  : def           Pkt Offset    : 4
Profile Capped : Disabled
StatMode      : minimal
-----
Policer Id     : 2
Description    : (Not Specified)
PIR Adptn     : closest          CIR Adptn      : closest
Parent        : none             Level          : 1
Weight        : 1             Adv. Cfg Plcy : none
Admin PIR     : 4000          Admin CIR     : 0
CBS           : def           MBS           : def
Hi Prio Only  : def           Pkt Offset    : 4
Profile Capped : Disabled
StatMode      : minimal
```

Apply Network QoS Policy with Queue Group Instances to the Spoke SDP

Similar information can be shown for the egress queue group, including the details of the configured queue and policer and again where it is applied.

```
*A:PE-1# show qos queue-group "egress-queue-group" egress detail
=====
QoS Queue-Group Egress
=====
-----
QoS Queue Group
-----
Group-Name      : egress-queue-group
Description     : (Not Specified)
-----
-----
Q  CIR Admin PIR Admin CBS          HiPrio PIR Lvl/Wt   Parent   BurstLimit(B)
   CIR Rule  PIR Rule  MBS          CIR Lvl/Wt   Wred-Queue Slope
   Named-Buffer Pool          Adv Config Policy Name
-----
1  0          6000    def          def         1/1        None      default
   closest   closest   def          0/1        disabled   default
   (not-assigned)          (not-assigned)
...
-----
Queue Group Ports (network)
=====
Port  Sched Pol  Policer-Ctrl-Pol  Acctg Pol  Stats  Description  QGrp-Instance
-----
7/1/2                                No          1
...
-----
Queue Group Policer
=====
Policer Id      : 1
Description     : (Not Specified)
PIR Adptn      : closest          CIR Adptn      : closest
Parent         : none          Level          : 1
Weight        : 1          Adv. Cfg Plcy: none
Admin PIR     : 4000         Admin CIR      : 0
CBS           : def          MBS           : def
Hi Prio Only  : def          Pkt Offset    : -22
Profile Capped : Disabled
StatMode      : minimal
...

```


The following command shows where the ingress queue group has been applied.

```
*A:PE-1# show qos queue-group ingress association
=====
QoS Queue-Group Ingress
=====
...
-----
QoS Queue Group
-----
Group-Name      : ingress-queue-group
Description     : (Not Specified)
...
=====
Queue Group FP Maps
=====
Card Num      Fp Num      Instance      Type
-----
7             1             1             Network
-----
Entries found: 1
...
=====
```

A similar command shows where the egress queue group has been applied.

```
*A:PE-1# show qos queue-group egress association
=====
QoS Queue-Group Egress
=====
-----
QoS Queue Group
-----
Group-Name      : egress-queue-group
Description     : (Not Specified)
...
=====
Queue Group Ports (network)
=====
Port  Sched Pol  Policer-Ctrl-Pol  Acctg Pol  Stats  Description  QGrp-Instance
-----
7/1/2                No                1
-----
...
=====
```

Apply Network QoS Policy with Queue Group Instances to the Spoke SDP

The ingress queue group applied to the FP on card 7 can be shown.

```
*A:PE-1# show card 7 fp 1 ingress queue-group "ingress-queue-group" instance 1
mode network
=====
Card:7 Net.QGrp: ingress-queue-group Instance: 1
=====
Group Name      : ingress-queue-group
Description     : (Not Specified)
Pol Ctl Pol    : None                      Acct Pol      : None
Collect Stats  : disabled
```

In order to show the details of the policers in the ingress FP queue group, the following command can be used.

```
*A:PE-1# show qos policer card 7 fp 1 queue-group "ingress-queue-group" instance 1 network
detail
=====
Policer Info (Net-FPQG-1-ingress-queue-group:1->1), Slot 7
=====
Policer Name      : Net-FPQG-1-ingress-queue-group:1->1
Direction         : Ingress                      Fwding Plane   : 1
Depth PIR         : 0 Bytes                      Depth CIR      : 0 Bytes
Depth FIR         : 0 Bytes
MBS               : 7680 B                       CBS            : 0 KB
Hi Prio Only     : 768 B                       Pkt Byte Offset : 4
Admin PIR        : 6000 Kbps                    Admin CIR      : 0 Kbps
Oper PIR         : 6000 Kbps                    Oper CIR       : 0 Kbps
Oper FIR         : 6000 Kbps
Stat Mode        : minimal
PIR Adaption     : closest                      CIR Adaption   : closest
Adv.Cfg Plcy    : None                        Profile Capped  : disabled
Parent Arbiter Name: (Not Specified)
-----
Arbiter Member Information
-----
Offered Rate     : 0 Kbps
Level            : 0                            Weight         : 0
Parent PIR      : 0 Kbps                      Parent FIR     : 0 Kbps
Consumed        : 0 Kbps
-----
=====
Policer Info (Net-FPQG-1-ingress-queue-group:1->2), Slot 7
=====
Policer Name      : Net-FPQG-1-ingress-queue-group:1->2
Direction         : Ingress                      Fwding Plane   : 1
Depth PIR         : 0 Bytes                      Depth CIR      : 0 Bytes
Depth FIR         : 0 Bytes
MBS               : 5 KB                       CBS            : 0 KB
Hi Prio Only     : 512 B                       Pkt Byte Offset : 4
Admin PIR        : 4000 Kbps                    Admin CIR      : 0 Kbps
Oper PIR         : 4000 Kbps                    Oper CIR       : 0 Kbps
Oper FIR         : 4000 Kbps
Stat Mode        : minimal
PIR Adaption     : closest                      CIR Adaption   : closest
Adv.Cfg Plcy    : None                        Profile Capped  : disabled
```

```

Parent Arbiter Name: (Not Specified)
-----
Arbiter Member Information
-----
Offered Rate      : 0 Kbps
Level             : 0
Weight            : 0
Parent PIR        : 0 Kbps
Parent FIR        : 0 Kbps
Consumed          : 0 Kbps
-----
Network Interface Association
-----
No Association Found.
-----
SDP Association
-----
Policer Info (1->1:1->10), Slot 7
-----
SDP Association Count : 1
-----

```

The details of the queue and policer in the egress queue group applied to port 7/1/2 can also be shown.

```

*A:PE-1# show port 7/1/2 queue-group egress "egress-queue-group" network instance 1
=====
Ethernet port 7/1/2 Network Egress queue-group
=====
Group Name      : egress-queue-group Instance-Id   : 1
Description     : (Not Specified)
Sched Policy    : None
Collect Stats   : disabled
Acct Pol        : None
Agg. Limit      : -1

Queues
-----
Queue-Group     : egress-queue-group Instance-Id   : 1      Queue-Id    : 1
Description     : (Not Specified)
Admin PIR       : 6000*
PIR Rule        : closest*
CBS             : def*
Hi Prio         : def*
Admin CIR       : 0*
CIR Rule        : closest*
MBS             : def*

Policers
-----
Queue-Group     : egress-queue-group Instance-Id   : 1      Policer-Id   : 1
Description     : (Not Specified)
Admin PIR       : 4000*
PIR Rule        : closest*
CBS             : def*
Hi Prio         : def*
Admin CIR       : 0*
CIR Rule        : closest*
MBS             : def*

* means the value is inherited

```

Apply Network QoS Policy with Queue Group Instances to the Spoke SDP

The network QoS policy can be shown with the details of the configured FC redirection and ingress classification used on the pseudowire.

```
*A:PE-1# show qos network 10 detail
=====
QoS Network Policy
=====
-----
Network Policy (10)
-----
Policy-id          : 10                Remark          : False
Forward Class     : be                Profile         : Out
LER Use DSCP      : False
Description       : (Not Specified)
...
-----
LSP EXP Bit Map          Forwarding Class          Profile
-----
5                        ef                        In
...
-----
Egress Forwarding Class Mapping
-----
FC Value           : 0                FC Name         : be
- DSCP Mapping
Out-of-Profile    : be                In-Profile      : be
...
DE Mark           : None
Redirect Grp Q    : 1                Redirect Grp Plcr: None
...
FC Value           : 5                FC Name         : ef
...
DE Mark           : None
Redirect Grp Q    : None             Redirect Grp Plcr: 1
-----
Ingress Forwarding Class Mapping
-----
FC Value           : 0                FC Name         : be
Redirect UniCast Plcr : 1            Redirect MultiCast Plcr : None
...
FC Value           : 5                FC Name         : ef
Redirect UniCast Plcr : 2            Redirect MultiCast Plcr : None
...
-----
```

The details of the configuration of the pseudowire QoS can be seen when showing the details of the SDP within the Epipe service.

```
*A:PE-1# show service id 1 sdp 1:1 detail
=====
Service Destination Point (Sdp Id : 1:1) Details
=====
-----
Sdp Id 1:1  -(192.0.2.2)
-----
Description      : (Not Specified)
SDP Id           : 1:1                               Type           : Spoke
Spoke Descr     : (Not Specified)
VC Type         : VLAN                               VC Tag         : 0
Admin Path MTU  : 0                                 Oper Path MTU   : 9190
Delivery        : MPLS
Far End         : 192.0.2.2
Tunnel Far End  : 192.0.2.2                         LSP Types      : LDP
Hash Label      : Disabled                           Hash Lbl Sig Cap : Disabled
Oper Hash Label : Disabled
Admin State     : Up                                 Oper State      : Up
...
Ingress Qos Policy : 10                             Egress Qos Policy : 10
Ingress FP QGrp   : ingress-queue-group             Egress Port QGrp  : egress-queue*
Ing FP QGrp Inst  : 1                               Egr Port QGrp Inst: 1
```

The usage of the “Egr Network Queue Group Mappings” out of the total number of “Dynamic Service Entries” can be seen with the following command. Only one egress QoS pseudowire redirection has been configured.

```
*A:PE-1# tools dump system-resources
Resource Manager info at 005 m 07/31/13 13:11:03.355:

Hardware Resource Usage for Slot #7, CardType imm5-10gb-xfp, Cmplx #0:
-----+-----+-----+-----
...
      Dynamic Service Entries |      65535|      1|      65534
          Subscriber Hosts |           |      0|
          Encap Group Members |           |      0|
Egr Network Queue Group Mappings |           |      1|
```

It is possible to show the statistics on the ingress FP queue group used by the pseudowire.

```
*A:PE-1# show card 7 fp 1 ingress queue-group "ingress-queue-group" instance 1 mode net-
work statistics

=====
Card:7 Net.QGrp: ingress-queue-group Instance: 1
=====
Group Name      : ingress-queue-group
Description     : (Not Specified)
Pol Ctl Pol    : None                               Acct Pol       : None
```

Apply Network QoS Policy with Queue Group Instances to the Spoke SDP

```
Collect Stats : disabled
-----
Statistics
-----
                Packets                Octets

Ing. Policer:  1  Grp: ingress-queue-group (Stats mode: minimal)
Off. All      :                184275                23587200
Dro. All      :                36801                 4710528
For. All      :                147474                18876672

Ing. Policer:  2  Grp: ingress-queue-group (Stats mode: minimal)
Off. All      :                184274                23587072
Dro. All      :                85955                 11002240
For. All      :                98319                 12584832
```

Similar statistics can be shown for the egress port queue group used by the pseudowire.

```
*A:PE-1# show port 7/1/2 queue-group egress "egress-queue-group" network statistics
instance 1
-----
Ethernet port 7/1/2 Network Egress queue-group
-----
                Packets                Octets
Egress Queue:  1  Group: egress-queue-group  Instance-Id:  1
In Profile forwarded  :  0                0
In Profile dropped    :  0                0
Out Profile forwarded : 150989            19326592
Out Profile dropped   : 37123             4751744

Egress Policer:  1  Group: egress-queue-group  Instance-Id:  1
Stats mode: minimal
Off. All      : 188421                24117888
Dro. All      : 87894                 11250432
For. All      : 100527                12867456
```

Monitor commands are available to see the statistics (and rates on egress port queue group). As an example, the following shows the utilization on the queue and policer in the egress queue-group.

```
*A:PE-1# monitor port 7/1/2 queue-group "egress-queue-group" instance 1 egress network
egress-queue 1 repeat 1 rate
=====
Monitor Port Queue-Group Egress Network Queue Statistics
=====
-----
At time t = 0 sec (Base Statistics)
-----
                Packets                Octets

In Profile forwarded  :  0                0
In Profile dropped    :  0                0
Out Profile forwarded : 299113            38286464
Out Profile dropped   : 74155             9491840
-----
At time t = 11 sec (Mode: Rate)
-----
```

```

                Packets/sec                Octets/sec                % Port
                                                Util.
In Profile forwarded : 0                    0                    0.00
In Profile dropped   : 0                    0                    0.00
Out Profile forwarded : 5863                750436                0.06
Out Profile dropped   : 1466                187609                0.01
=====

```

```

*A:PE-1# monitor port 7/1/2 queue-group "egress-queue-group" instance 1 egress network
policer 1 repeat 1 rate
=====

```

Monitor Port Queue-Group Egress Network Policer Statistics

At time t = 0 sec (Base Statistics)

```

-----
                Packets                Octets
Off. All       : 454750                58208000
Dro. All       : 212181                27159168
For. All       : 242569                31048832
-----

```

At time t = 11 sec (Mode: Rate)

```

-----
                Packets/sec                Octets/sec                % Port
                                                Util.
Off. All       : 7326                    937716                0.07
Dro. All       : 3419                    437609                0.03
For. All       : 3907                    500108                0.04
=====

```

*A:PE-1#

As mentioned, the egress policer (FC EF) traffic exits the egress port by default using the related network queue on the port. This can be seen below.

```

*A:PE-1# show port 7/1/2 detail
=====

```

Ethernet Interface

```

-----
Description      : 10-Gig Ethernet
Interface         : 7/1/2                Oper Speed        : 10 Gbps
Link-level       : Ethernet              Config Speed      : N/A
Admin State      : up                    Oper Duplex       : full
Oper State       : up                    Config Duplex     : N/A
...
=====

```

Queue Statistics

```

-----
...
Egress Queue 6
  In Profile forwarded : 0                0
  In Profile dropped   : 0                0
-----

```

Apply Network QoS Policy with Queue Group Instances to the Spoke SDP

```

Out Profile forwarded :    102381                15357150
Out Profile dropped   :      0                    0
    
```

The throughput achieved using the above configuration can be verified in the traffic generator output. Port 202/1 is connected to PE-1 and port 204/1 is connected to PE-2.

Port /	Tx Test Packets	Rx Test Packets	Tx Test Octets	Rx Test Octets	Tx Test Throughput (Mb/s)	Rx Test Throughput (Mb/s)	Rx Packet Loss	Average Latency (us)
All Ports	29296	19531	3749888	2499968	29.999	20.000	n/a	15512.18
202/1	14648	9765	1874944	1249920	15.000	9.999	n/a	39.28
202/1->204/1, BE traffic	7324	5860	937472	750080	7.500	6.001	1464	51609.56
202/1->204/1, EF traffic	7324	3906	937472	499968	7.500	4.000	3418	39.13
204/1	14648	9766	1874944	1250048	15.000	10.000	n/a	30983.50
204/1->202/1, BE traffic	7324	5859	937472	749952	7.500	6.000	1465	39.28
204/1->202/1, EF traffic'	7324	3906	937472	499968	7.500	4.000	3418	39.27

Conclusion

This example has shown the configuration and monitoring of pseudowire QoS, providing a powerful QoS solution for pseudowire applications. QoS can be applied independently to the ingress and/or egress of a single pseudowire or multiple pseudowires.

Conclusion