

QoS Scheduler Policies

In This Section

This section provides information to configure QoS scheduler and port scheduler policies using the command line interface.

Topics in this section include:

- [Overview on page 470](#)
- [Basic Configurations on page 493](#)
- [Service Management Tasks on page 502](#)

Overview

Scheduler Policies

Virtual schedulers are created within the context of a scheduler policy that is used to define the hierarchy and parameters for each scheduler. A scheduler is defined in the context of a tier which is used to place the scheduler within the hierarchy. Three tiers of virtual schedulers are supported. Root schedulers are defined without a parent scheduler meaning it is not subject to obtaining bandwidth from a higher tier scheduler. A scheduler has the option of enforcing a maximum rate of operation for all child queues and schedulers associated with it.

Because a scheduler is designed to arbitrate bandwidth between many inputs, a metric must be assigned to each child queue or scheduler vying for transmit bandwidth. This metric indicates whether the child is to be scheduled in a strict or weighted fashion and the level or weight the child has to other children.

Egress Port-Based Schedulers

In previous releases, HQoS root (top tier) schedulers always assumed that the configured rate was available, regardless of egress port level oversubscription and congestion. This resulted in the possibility that the aggregate bandwidth assigned to queues was not actually available at the port level. When the HQoS algorithm configures queues with more bandwidth than available on an egress port, actual bandwidth distribution to queues on the port will be solely based on the action of the hardware scheduler. This can result in a forwarding rate at each queue that is very different than the desired rate.

The port-based scheduler feature was introduced to allow HQoS bandwidth allocation based on available bandwidth at the egress port level. The port-based scheduler works at the egress line rate of the port to which it is attached. Port-based scheduling bandwidth allocation automatically includes the Inter-Frame Gap (IFG) and preamble for packets forwarded on queues servicing egress Ethernet ports. However, on PoS and SDH based ports, the HDLC encapsulation overhead and other framing overhead per packet is not known by the system. Instead of automatically determining the encapsulation overhead for SDH or SONET queues, the system provides a configurable frame encapsulation efficiency parameter that allows the user to select the average encapsulation efficiency for all packets forwarded out the egress queue.

A special port scheduler policy can be configured to define the virtual scheduling behavior for an egress port. The port scheduler is a software-based state machine managing a bandwidth allocation algorithm that represents the scheduling hierarchy shown in [Figure 16 on page 473](#).

The first tier of the scheduling hierarchy manages the total frame based bandwidth that the port scheduler will allocate to the eight priority levels.

The second tier receives bandwidth from the first tier in two priorities, a “within-cir” loop and an “above-cir” loop. The second tier “within-cir” loop provides bandwidth to the third tier “within-cir” loops, one for each of the eight priority levels. The second tier “above-cir” loop provides bandwidth to the third tier “above-cir” loops for each of the eight priority levels.

The “within-cir” loop for each priority level on the third tier supports an optional rate limiter used to restrict the maximum amount of “within-cir” bandwidth the priority level can receive. A maximum priority level rate limit is also supported that restricts the total amount of bandwidth the level can receive for both “within-cir” and “above-cir”. The amount of bandwidth consumed by each priority level for “within-cir” and “above-cir” is predicated on the rate limits described and the ability for each child queue or scheduler attached to the priority level to use the bandwidth.

The priority 1 “above-cir” scheduling loop has a special two tier strict distribution function. The high priority level 1 “above-cir” distribution is weighted between all queues and schedulers attached to level 1 for “above-cir” bandwidth. The low priority distribution for level 1 “above-cir” is reserved for all orphaned queues and schedulers on the egress port. Orphans are queues and schedulers that are not explicitly or indirectly attached to the port scheduler through normal parenting conventions. By default, all orphans receive bandwidth after all parented queues and schedulers and are allowed to consume whatever bandwidth is remaining. This default behavior for orphans can be overridden on each port scheduler policy by defining explicit orphan port parent association parameters.

Ultimately, any bandwidth allocated by the port scheduler is given to a child queue. The bandwidth allocated to the queue is converted to a value for the queue’s PIR (maximum rate) setting. This way, the hardware schedulers operating at the egress port level will only schedule bandwidth for all queues on the port up to the limits prescribed by the virtual scheduling algorithm.

The following lists the bandwidth allocation sequence for the port virtual scheduler:

1. Priority level 8 offered load up to priority CIR
2. Priority level 7 offered load up to priority CIR
3. Priority level 6 offered load up to priority CIR
4. Priority level 5 offered load up to priority CIR
5. Priority level 4 offered load up to priority CIR
6. Priority level 3 offered load up to priority CIR
7. Priority level 2 offered load up to priority CIR
8. Priority level 1 offered load up to priority CIR
9. Priority level 8 remaining offered load up to remaining priority rate limit
10. Priority level 7 remaining offered load up to remaining priority rate limit

11. Priority level 6 remaining offered load up to remaining priority rate limit
12. Priority level 5 remaining offered load up to remaining priority rate limit
13. Priority level 4 remaining offered load up to remaining priority rate limit
14. Priority level 3 remaining offered load up to remaining priority rate limit
15. Priority level 2 remaining offered load up to remaining priority rate limit
16. Priority level 1 remaining offered load up to remaining priority rate limit
17. Priority level 1 remaining orphan offered load up to remaining priority rate limit (default orphan behavior unless orphan behavior has been overridden in the scheduler policy)

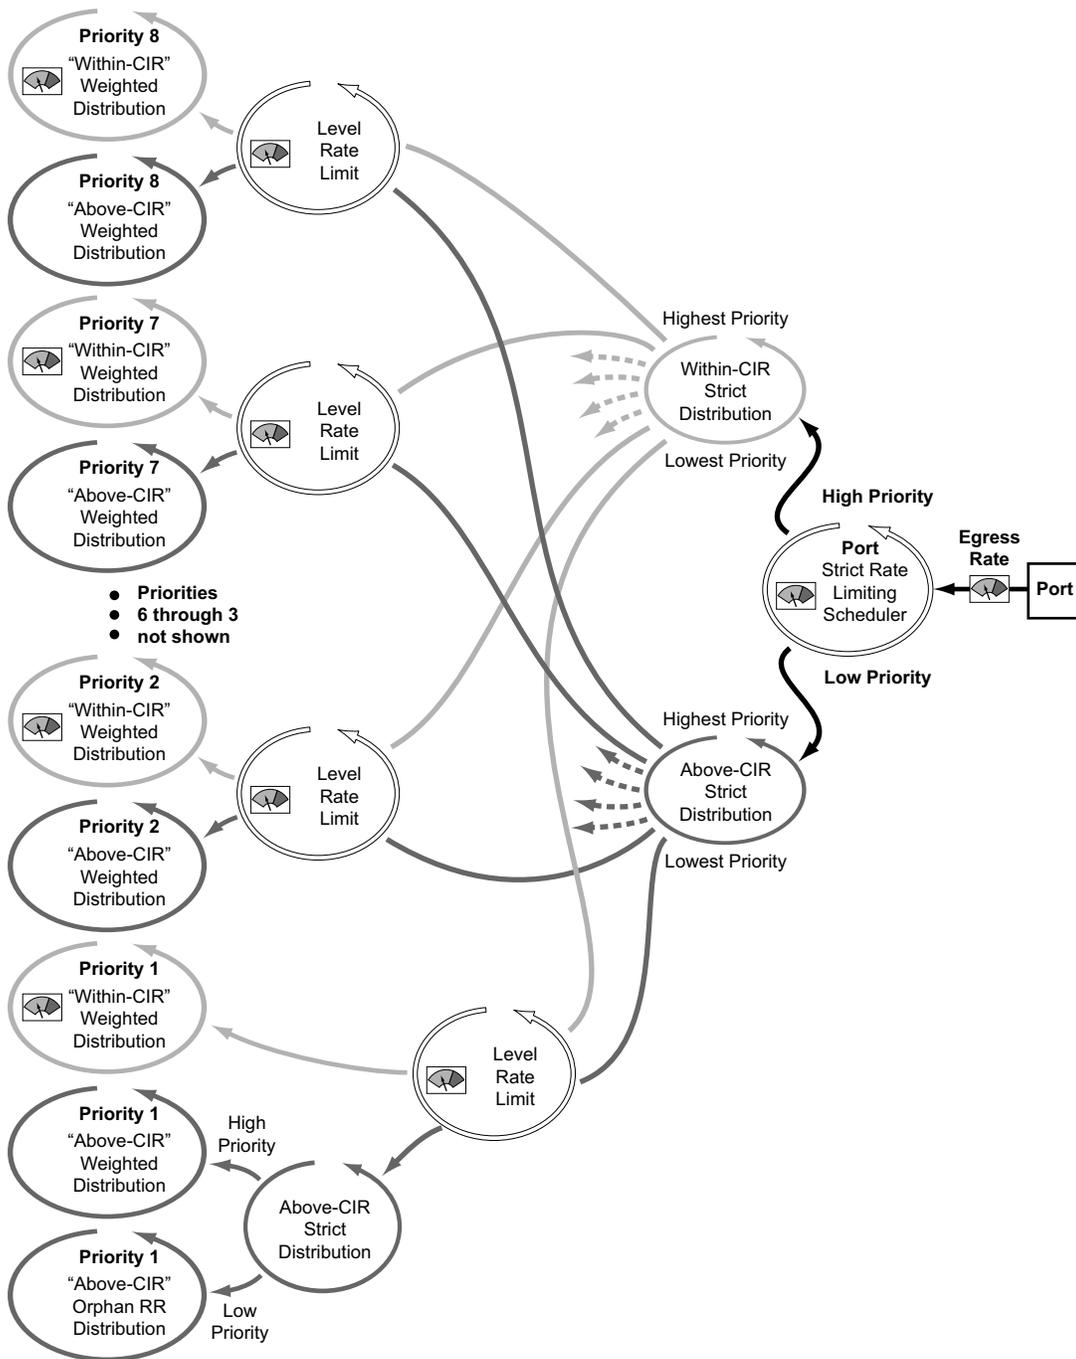
When a queue is inactive or has a limited offered load that is below its fair share (fair share is based on the bandwidth allocation a queue would receive if it was registering adequate activity), its operational PIR must be set to some value to handle what would happen if the queues offered load increased prior to the next iteration of the port virtual scheduling algorithm. If an inactive queue's PIR was set to zero (or near zero), the queue would throttle its traffic until the next algorithm iteration. If the operational PIR was set to its configured rate, the result could overrun the expected aggregate rate of the port scheduler.

To accommodate inactive queues, the system calculates a Minimum Information Rate (MIR) for each queue. To calculate each queue's MIR, the system determines what that queue's Fair Information Rate (FIR) would be if that queue had actually been active during the latest iteration of the virtual scheduling algorithm. For example, if three queues are active (1, 2, and 3) and two queues are inactive (4 and 5), the system first calculates the FIR for each active queue. Then it recalculates the FIR for queue 4 assuming queue 4 was active with queues 1, 2, and 3 and uses the result as the queue's MIR. The same is done for queue 5 using queues 1, 2, 3, and 5. The MIR for each inactive queue is used as the operational PIR for each queue.

Service/Subscriber Egress Port Bandwidth Allocation

The port-based egress scheduler can be used to allocate bandwidth to each service or subscriber associated with the port. While egress queues on the service can have a child association with a scheduler policy on the SAP or multi-service site, all queues must vie for bandwidth from an egress port. Two methods are supported to allocate bandwidth to each service or subscriber queue:

1. Service or subscriber queue association with a scheduler on the SAP or multi-service site which is itself associated with a port-level scheduler.
2. Service or subscriber queue association directly with a port-level scheduler.

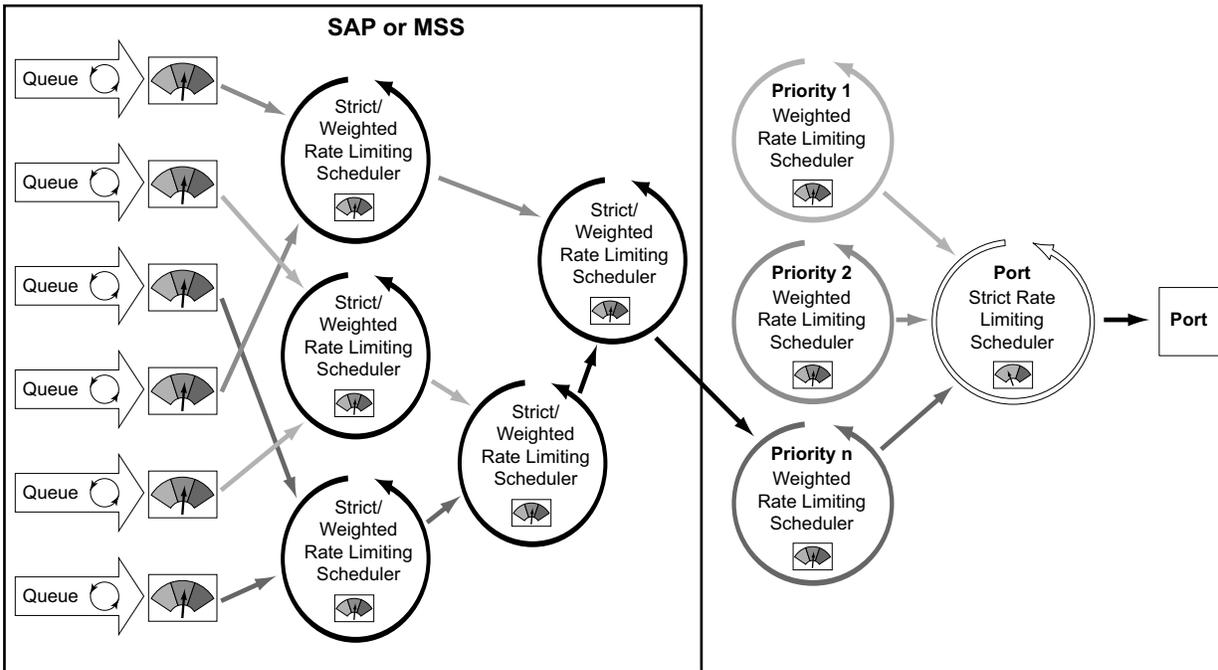


OSSG130

Figure 16: Port Level Virtual Scheduler Bandwidth Allocation Based on Priority and CIR

Service or Subscriber Scheduler Child to Port Scheduler Parent

The service or subscriber scheduler to port scheduler association model allows for multiple services or subscriber to have independent scheduler policy definitions while the independent schedulers receive bandwidth from the scheduler at the port level. By using two scheduler policies, available egress port bandwidth can be allocated fairly or unfairly depending on the desired behavior. Figure 17 graphically demonstrates this model.



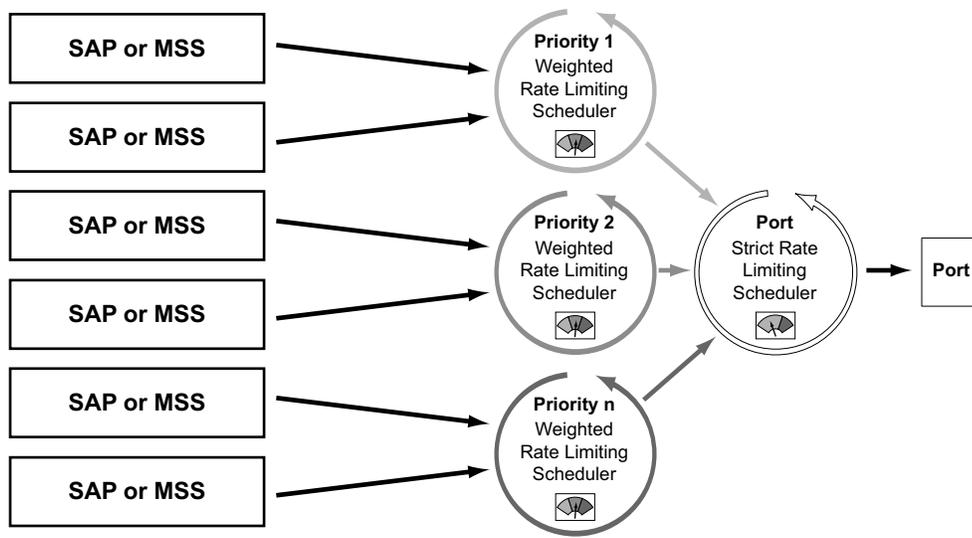
OSSG131

Figure 17: Two Scheduler Policy Model for Access Ports

Once a two scheduler policy model is defined, the bandwidth distribution hierarchy allocates the available port bandwidth to the port schedulers based on priority, weights, and rate limits. The service or subscriber level schedulers and the queues they service become an extension of this hierarchy.

Due to the nature of the two scheduler policy, bandwidth is allocated on a per-service or per subscriber basis as opposed to a per-class basis. A common use of the two policy model is for a carrier-of-carriers mode of business. In essence, the goal of a carrier is to provide segments of

bandwidth to providers who purchase that bandwidth as services. While the carrier does not concern itself with the interior services of the provider, it does however care how congestion affects the bandwidth allocation to each provider's service. As an added benefit, the two policy approach provides the carrier with the ability to preferentially allocate bandwidth within a service or subscriber context through the service or subscriber level policy without affecting the overall bandwidth allocation to each service or subscriber. Figure 18 shows a per-service bandwidth allocation using the two scheduler policy model. While the figure shows services grouped by scheduling priority, it is expected that many service models will place the services in a common port priority and use weights to provide a weighted distribution between the service instances. Higher weights provide for relatively higher amounts of bandwidth.



OSSG132

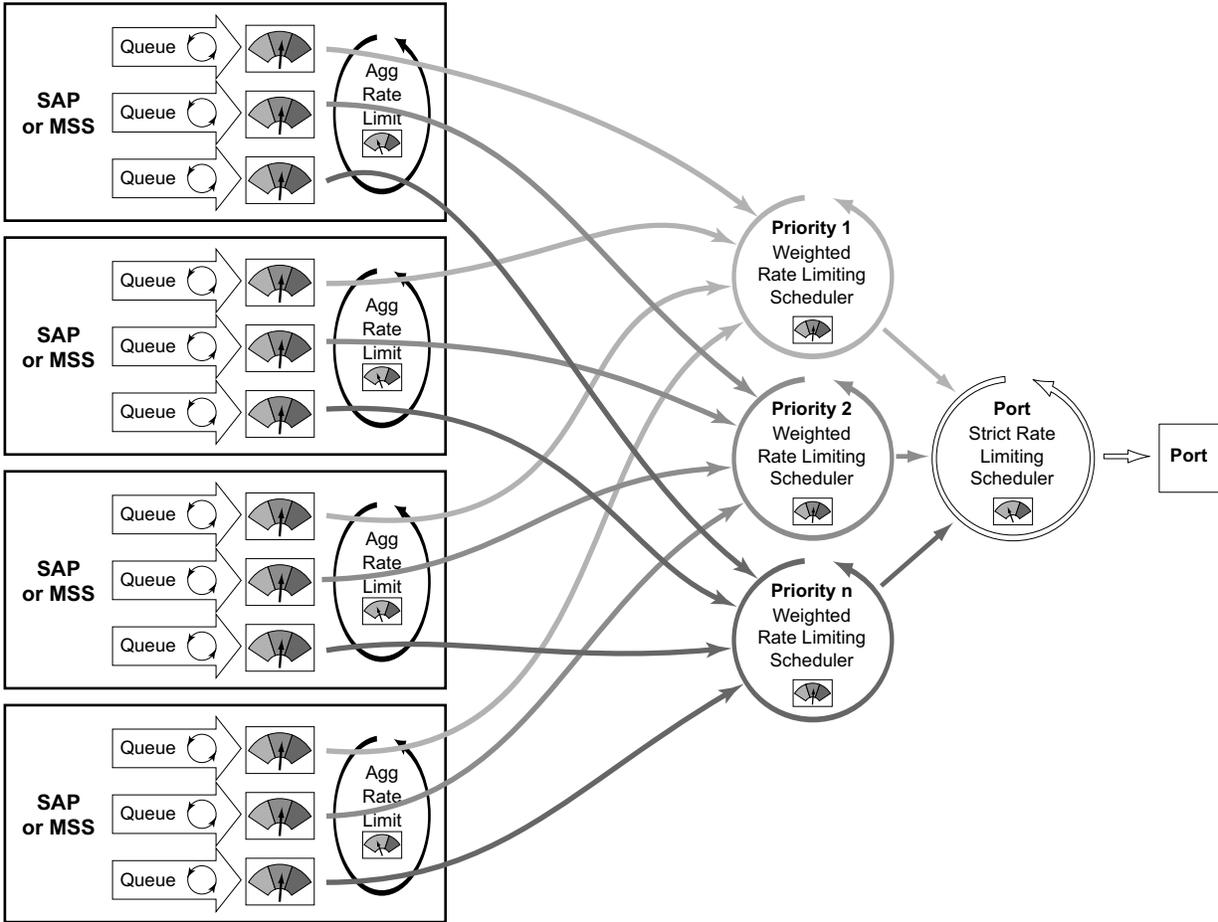
Figure 18: Schedulers on SAP or Multi-Service Site Receive Bandwidth From Port Priority Levels

Direct Service or Subscriber Queue Association to Port Scheduler Parents

The second model of bandwidth allocation on an egress access port is to directly associate a service or subscriber queue to a port-level scheduler. This model allows the port scheduler hierarchy to allocate bandwidth on a per class or priority basis to each service or subscriber queue. This allows the provider to manage the available egress port bandwidth on a service tier basis ensuring that during egress port congestion, a deterministic behavior is possible from an aggregate perspective. While this provides an aggregate bandwidth allocation model, it does not inhibit per service or per subscriber queuing. [Figure 19](#) demonstrates the single, port scheduler policy model.

[Figure 19](#) also demonstrates the optional aggregate rate limiter at the SAP, multi-service site or subscriber level. The aggregate rate limiter is used to define a maximum aggregate bandwidth at which the child queues can operate. While the port-level scheduler is allocating bandwidth to each child queue, the current sum of the bandwidth for the service or subscriber is monitored. Once the aggregate rate limit is reached, no more bandwidth is allocated to the children associated with the SAP, multi-service site, or subscriber. Aggregate rate limiting is restricted to the single scheduler policy model and is mutually exclusive to defining SAP, multi-service site, or subscriber scheduling policies.

The benefit of the single scheduler policy model is that the bandwidth is allocated per priority for all queues associated with the egress port. This allows a provider to preferentially allocate bandwidth to higher priority classes of service independent of service or subscriber instance. In many cases, a subscriber can purchase multiple services from a single site (VoIP, HSI, Video) and each service can have a higher premium value relative to other service types. If a subscriber has purchased a premium service class, that service class should get bandwidth before another subscriber's best effort service class. When combined with the aggregate rate limit feature, the single port-level scheduler policy model provides a per-service instance or per-subscriber instance aggregate SLA and a class based port bandwidth allocation function.



OSSG133

Figure 19: Direct Service or Subscriber Association to Port Scheduler Model

Frame and Packet-Based Bandwidth Allocation

A port-based bandwidth allocation mechanism must consider the effect that line encapsulation overhead plays relative to the bandwidth allocated per service or subscriber. The service or subscriber level bandwidth definition (at the queue level) operates on a packet accounting basis. For Ethernet, this includes the DLC header, the payload and the trailing CRC. This does not include the IFG or the preamble. This means that an Ethernet packet will consume 20 bytes more bandwidth on the wire than what the queue accounted for. When considering HDLC encoded PoS or SDH ports, the overhead is variable based on '7e' insertions (and other TDM framing issues). The HDLC and SONET/SDH frame overhead is not included for queues forwarding on PoS and SDH links.

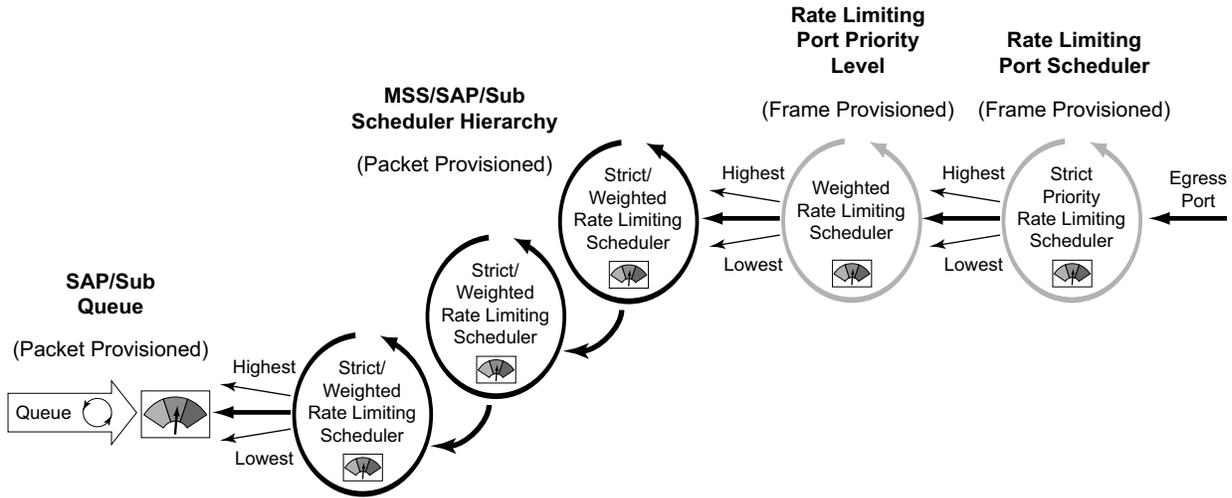
The port-based scheduler hierarchy must translate the frame based accounting (on-the-wire bandwidth allocation) it performs to the packet based accounting in the queues. When the port scheduler considers the maximum amount of bandwidth a queue should get, it must first determine how much bandwidth the queue can use. This is based on the offered load the queue is currently experiencing (how many octets are being offered the queue). The offered load is compared to the queues configured CIR and PIR. The CIR value determines how much of the offered load should be considered in the "within-cir" bandwidth allocation pass. The PIR value determines how much of the remaining offered load (after "within-cir") should be considered for the "above-cir" bandwidth allocation pass.

For Ethernet queues (queues associated with an egress Ethernet port), the packet to frame conversion is relatively easy. The system multiplies the number of offered packets by 20 bytes and adds the result to the offered octets ($\text{offeredPackets} \times 20 + \text{offeredOctets} = \text{frameOfferedLoad}$). This frame-offered-load value represents the amount of line rate bandwidth the queue is requesting. The system computes the ratio of increase between the offered-load and frame-offered-load and calculates the current frame based CIR and PIR. The frame-CIR and frame-PIR values are used as the limiting values in the "within-cir" and "above-cir" port bandwidth distribution passes.

For PoS or SDH queues, the packet to frame conversion is more difficult to dynamically calculate due to the variable nature of HDLC encoding. Wherever a '7e' bit or byte pattern appears in the data stream, the framer performing the HDLC encoding must place another '7e' within the payload. Since this added HDLC encoding is unknown to the forwarding plane, the system allows for an encapsulation overhead parameter that can be provisioned on a per queue basis. This is provided on a per queue basis to allow for differences in the encapsulation behavior between service flows in different queues. The system multiplies the offered load of the queue by the encapsulation-overhead parameter and adds the result to the offered load of the queue ($\text{offeredOctets} \times \text{configuredEncapsulationOverhead} + \text{offeredOctets} = \text{frameOfferedLoad}$). The frame-offered-load value is used by the egress PoS/SDH port scheduler in the same manner as the egress Ethernet port scheduler above.

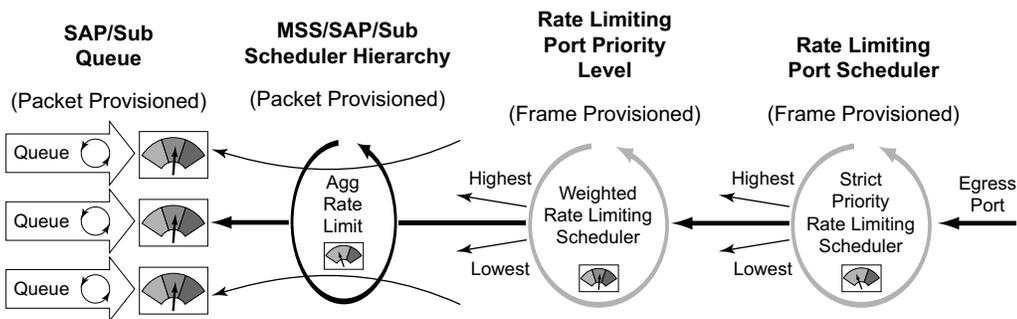
From a provisioning perspective, queues and service level (and subscriber level) scheduler policies are always provisioned with packet-based parameters. The system will convert these values to

frame-based on-the-wire values for the purpose of port bandwidth allocation. However, port-based scheduler policy scheduler maximum rates and CIR values are always interpreted as on-the-wire values and must be provisioned accordingly. Figure 20 and Figure 21 provide a logical view of bandwidth distribution from the port to the queue level and shows the packet or frame-based provisioning at each step.



OSSG134

Figure 20: Port Bandwidth Distribution for Service and Port Scheduler Hierarchies



OSSG135

Figure 21: Port Bandwidth Distribution for Direct Queue to Port Scheduler Hierarchy

Queue Parental Association Scope

A **port-parent** command in the sap-egress and network-queue QoS policy queue context defines the direct child/parent association between an egress queue and a port scheduler priority level. The **port-parent** command is mutually exclusive to the already-existing **parent** command, which associates a queue with a scheduler at the SAP, multi-service site or subscriber profile level. It is possible to mix local parented (parent to service or subscriber level scheduler) and port parented queues with schedulers on the same egress port.

The **port-parent** command only accepts a child/parent association to the eight priority levels on a port scheduler hierarchy. Similar to the local **parent** command, two associations are supported, one for “within-cir” bandwidth (cir-level) and a second one for “above-cir” bandwidth (level). The “within-cir” association is optional and can be disabled by using the default “within-cir” weight value of 0. In the event that a queue with a defined parent port is on a port without a port scheduler policy applied, that queue will be considered an orphaned queue. If a queue with a parent command is defined on a port and the named scheduler is not found due a missing scheduler policy or a missing scheduler of that name, the queue will be considered orphaned as well.

A queue can be moved from a local (on the SAP, multi-service site, or subscriber profile) parent to a port parent priority level simply by executing the **port-parent** command. Once the **port-parent** command is executed, any local parent information for the queue is lost. The queue can also be moved back to a local parent at anytime by executing the local parent command. Lastly, the local parent or port parent association can be removed at any time by using the no version of the appropriate parent command.

Service or Subscriber-Level Scheduler Parental Association Scope

The **port-parent** command in the scheduler-policy scheduler context (at all tier levels) allows a scheduler to be associated with a port scheduler priority level. The **port-parent** command is mutually exclusive to the **parent** command for schedulers at tiers 2 and 3 within the scheduler policy. The **port-parent** command is the only parent command allowed for schedulers in tier 1.

The **port-parent** command only accepts a child/parent association to the eight priority levels on a port scheduler hierarchy. Similar to the normal local parent command, two associations are supported, one for “within-cir” bandwidth (cir-level) and a second one for “above-cir” bandwidth (level). The “within-cir” association is optional and can be disabled by using the default “within-cir” weight value of 0. In the event that a scheduler with a port parent defined is on a port without a port scheduler policy applied, that scheduler will be considered an orphaned scheduler.

A scheduler in tiers 2 and 3 can be moved from a local (within the policy) parent to a port parent priority level simply by executing the **port-parent** command. Once the **port-parent** command is executed, any local parent information for the scheduler is lost. The schedulers at tiers 2 and 3 can also be moved back to a local parent at anytime by executing the local parent command. Lastly, the

local parent or port parent association can be removed at anytime by using the no version of the appropriate parent command. A scheduler in tier 1 can only be associated with a port parent and that port parent definition can be added or removed at anytime.

Network Queue Parent Scheduler

Network queues support port scheduler parent priority-level associations. Using a port scheduler policy definition and mapping network queues to a port parent priority level, HQoS functionality is supported providing eight levels of strict priority and weights within the same priority. A network queue's bandwidth is allocated using the "within-cir" and "above-cir" scheme normal for port schedulers.

Queue CIR and PIR percentages when port-based schedulers are in effect will be based on frame-offered-load calculations. Figure 22 demonstrates port-based virtual scheduling bandwidth distribution.

A network queue with a port parent association exists on a port without a scheduler policy defined will be considered to be orphaned.

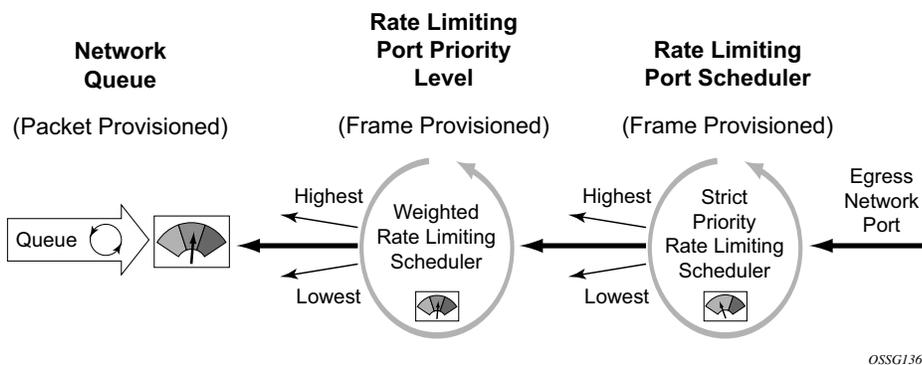


Figure 22: Bandwidth Distribution on Network Port with Port-Based Scheduling

Foster Parent Behavior for Orphaned Queues and Schedulers

All queues and schedulers on a port that has a port-based scheduler policy configured will be subject to bandwidth allocation through the port-based schedulers. All queues and schedulers that are not configured with a scheduler parent are considered to be orphaned when port-based scheduling is in effect. This includes access and network queue schedulers at the SAP, multi-service site, subscriber and port level.

By default, orphaned queues and schedulers are allocated bandwidth after all queues and schedulers in the parented hierarchy have had bandwidth allocated “within-cir” and “above-cir”. In essence, an orphaned scheduler or queue can be considered as being foster parented by the port scheduler. Orphaned queues and schedulers have an inherent port scheduler association as shown below:

- Within-CIR priority = 1
- Within-CIR weight = 0
- Above-CIR priority = 1
- Above-CIR weight = 0

The above-CIR weight = 0 value is only used for orphaned queues and schedulers on port scheduler enabled egress ports. The system interprets weight=0 as priority level 0 and will only distribute bandwidth to level 0 once all other properly parented queues and schedulers have received bandwidth. Orphaned queues and schedulers all have equal priority to the remaining port bandwidth.

The default orphan behavior can be overridden for each port scheduler policy by using the orphan override command. The orphan override command accepts the same parameters as the port parent command. When the orphan override command is executed, all orphan queues and schedulers are treated in a similar fashion as other properly parented queues and schedulers based on the override parenting parameters.

It is expected that an orphan condition is not the desired state for a queue or scheduler and is the result of a temporary configuration change or configuration error.

Frame-Based Accounting

The standard accounting mechanism uses ‘packet based’ rules that account for the DLC header, any existing tags, Ethernet payload and the 4 byte CRC. The Ethernet framing overhead which includes the Inter-Frame Gap (IFG) and preamble (20 bytes total) are not included in packet based accounting. When frame based accounting is enabled, the 20 byte framing overhead is included in the queue CIR, PIR and scheduling operations allowing the operations to take into consideration on-wire bandwidth consumed by each Ethernet packet.

Since the native queue accounting functions (stats, CIR and PIR) are based on packet sizes and do not include Ethernet frame encapsulation overhead, the system must manage the conversion between packet based and frame based accounting. To accomplish this, the system requires that a queue operates in frame based accounting mode, and must be managed by a virtual scheduler policy or by a port virtual scheduler policy. Egress queues can use either port or service schedulers to accomplish frame based accounting, but ingress queues are limited to service based scheduling policies.

Turning on frame based accounting for a queue is accomplished through a frame based accounting command defined on the scheduling policy level associated with the queue or through a queue frame based accounting parameter on the aggregate rate limit command associated with the queues SAP, multi-service site or subscriber context.

Operational Modifications

To add frame overhead to the existing QoS Ethernet packet handling functions, the system uses the already existing virtual scheduling capability of the system. The system currently monitors each queue included in a virtual scheduler to determine its offered load. This offered load value is interpreted based on the queues defined CIR and PIR threshold rates to determine bandwidth offerings from the queues virtual scheduler. When egress port based virtual scheduling was added, frame based usage on the wire was added to allow for the port bandwidth to be accurately allocated to each child queue on the port.

Existing Egress Port Based Virtual Scheduling

The port based virtual scheduling mechanism takes the native packet based accounting results from the queue and adds 20 bytes to each packet to derive the queue's frame based offered load. The ratio between the frame based offered load and the packet based offered load is then used to determine the effective frame based CIR and frame based PIR thresholds for the queue. Once the port virtual scheduler computes the amount of bandwidth allowed to the queue (in a frame based fashion), the bandwidth is converted back to a packet based value and used as the queue's operational PIR. The queue's native packet based mechanisms continue to function, but the maximum operational rate is governed by frame based decisions.

Queue Behavior Modifications for Frame Based Accounting

The frame based accounting feature extends this capability to allow the queue CIR and PIR thresholds to be defined as frame based values as opposed to packet based values. The queue continues to internally use its packet based mechanisms, but the provisioned frame based CIR and PIR values are continuously revalued based on the ratio between the calculated frame based offered load and actual packet based offered load. As a result, the queue's operational packet based CIR and PIR are accurately modified during each iteration of the virtual scheduler to represent the provisioned frame based CIR and PIR.

Virtual Scheduler Rate and Queue Rate Parameter Interpretation

Normally, a scheduler policy contains rates that indicate packet based accounting values. When the children queues associated with the policy are operating in frame based accounting mode, the parent schedulers must also be governed by frame based rates. Since either port based or service based virtual scheduling is required for queue frame based operation, enabling frame based operation is configured at either the scheduling policy or aggregate rate limit command level. All queues associated with the policy or the aggregate rate limit command will inherit the frame based accounting setting from the scheduling context.

When frame based accounting is enabled, the queues CIR and PIR settings are automatically interpreted as frame based values. If a SAP ingress QoS policy is applied with a queue PIR set to 100Mbps on two different SAPs, one associated with a policy with frame based accounting enabled and the other without frame based accounting enabled, the 100Mbps rate will be interpreted differently for each queue. The frame based accounting queue will add 20 bytes to each packet received by the queue and limit the rate based on the extra overhead. The packet based accounting queue will not add the 20 bytes per packet and thus allow more packets through per second.

Similarly, the rates defined in the scheduling policy with frame based accounting enabled will automatically be interpreted as frame based rates.

The port based scheduler aggregate rate limit command always interprets its configured rate limit value as a frame based rate. Setting the frame based accounting parameter on the aggregate rate limit command only affects the queues managed by the aggregate rate limit and converts them from packet based to frame based accounting mode.

Configuring Port Scheduler Policies

Port Scheduler Structure

Every port scheduler supports eight strict priority levels with a two pass bandwidth allocation mechanism for each priority level. Priority levels 8 through 1 (level 8 is the highest priority) are available for port-parent association for child queues and schedulers. Each priority level supports a maximum rate limit parameter that limits the amount of bandwidth that may be allocated to that level. A CIR parameter is also supported that limits the amount of bandwidth allocated to the priority level for the child queue's offered load, within their defined CIR. An overall maximum rate parameter defines the total bandwidth that will be allocated to all priority levels.

Special Orphan Queue and Scheduler Behavior

When a port scheduler is present on an egress port or channel, the system ensures that all queues and schedulers receive bandwidth from that scheduler to prevent free-running queues which can cause the aggregate operational PIR of the port or channel to oversubscribe the bandwidth available. When the aggregate maximum rate for the queues on a port or channel operate above the available line rate, the forwarding ratio between the queues will be affected by the hardware schedulers on the port and may not reflect the scheduling defined on the port or intermediate schedulers. Queues and schedulers that are either explicitly attached to the port scheduler using the port-parent command or are attached to an intermediate scheduler hierarchy that is ultimately attached to the port scheduler are managed through the normal eight priority levels. Queues and schedulers that are not attached directly to the port scheduler and are not attached to an intermediate scheduler that itself is attached to the port scheduler are considered orphaned queues and, by default, are tied to priority 1 with a weight of 0. All weight 0 queues and schedulers at priority level 1 are allocated bandwidth after all other children and each weight 0 child is given an equal share of the remaining bandwidth. This default orphan behavior may be overridden at the port scheduler policy by using the orphan-override command. The orphan-override command accepts the same parameters as the port-parent command. When the orphan-override command is executed, the parameters will be used as the port parent parameters for all orphans associated with a port using the port scheduler policy.

Packet to Frame Bandwidth Conversion

Another difference between the service level scheduler-policy and the port level port-scheduler-policy is in bandwidth allocation behavior. The port scheduler is designed to offer on-the-wire bandwidth. For Ethernet ports, this includes the IFG and the preamble for each frame and represents 20 bytes total per frame. The queues and intermediate service level schedulers (a

service level scheduler is a scheduler instance at the SAP, multi-service site or subscriber profile level) operate based on packet overhead which does not include the IFG or preamble on Ethernet packets. In order for the port based virtual scheduling algorithm to function, it must convert the queue and service scheduler packet based required bandwidth and bandwidth limiters (CIR and rate PIR) to frame based values. This is accomplished by adding 20 bytes to each Ethernet frame offered at the queue level to calculate a frame based offered load. Then the algorithm calculates the ratio increase between the packet based offered load and the frame based offered load and uses this ratio to adapt the CIR and rate PIR values for the queue to frame-CIR and frame-PIR values. When a service level scheduler hierarchy is between the queues and the port based schedulers, the ratio between the average frame-offered-load and the average packet-offered-load is used to adapt the scheduler's packet based CIR and rate PIR to frame based values. The frame based values are then used to distribute the port based bandwidth down to the queue level.

Packet over SONET (PoS) and SDH queues also operate based on packet sizes and do not include on-the-wire frame overhead. Unfortunately, the port based virtual scheduler algorithm does not have access to all the frame encapsulation overhead occurring at the framer level. Instead of automatically calculating the difference between packet-offered-load and frame-offered-load, the system relies on a provisioned value at the queue level. This avg-frame-overhead parameter is used to calculate the difference between the packet-offered-load and the frame-offered-load. This difference is added to the packet-offered-load to derive the frame-offered-load. Proper setting of this percentage value is required for proper bandwidth allocation between queues and service schedulers. If this value is not attainable, another approach is to artificially lower the maximum rate of the port scheduler to represent the average port framing overhead. This, in conjunction with a zero or low value for avg-frame-overhead, will ensure that the allocated queue bandwidth will control forwarding behavior instead of the low level hardware schedulers.

Aggregate Rate Limits for Directly Attached Queues

When all queues for a SAP, multi-service site or subscriber instance are attached directly to the port scheduler (using the port-parent command), it is possible to configure an agg-rate-limit for the queues. This is beneficial since the port scheduler does not provide a mechanism to enforce an aggregate SLA for a service or subscriber and the agg-rate-limit provides this ability. Queues may be provisioned directly on the port scheduler when it is desirable to manage the congestion at the egress port based on class priority instead of on a per service object basis.

The agg-rate-limit is not supported when one or more queues on the object are attached to an intermediate service scheduler. In this event, it is expected that the intermediate scheduler hierarchy will be used to enforce the aggregate SLA. Attaching an agg-rate-limit is mutually exclusive to attaching an egress scheduler policy at the SAP or subscriber profile level. Once an aggregate rate limit is in effect, a scheduler policy cannot be assigned. Once a scheduler policy is assigned on the egress side of a SAP or subscriber profile, an agg-rate-limit cannot be assigned.

Since the sap-egress policy defines a queue's parent association before the policy is associated with a service SAP or subscriber profile, it is possible for the policy to either not define a port-parent association or define an intermediate scheduler parenting that does not exist. As stated above, queues in this state are considered to be orphaned and automatically attached to port scheduler priority 1. Orphaned queues are included in the aggregate rate limiting behavior on the SAP or subscriber instance they are created within.

SAP Egress QoS Policy Queue Parenting

A sap-egress QoS policy queue may be associated with either a port parent or an intermediate scheduler parent. The validity parent definition cannot be checked at the time it is provisioned since the application of the QoS policy is not known until it is applied to an egress SAP or subscriber profile. It is allowed to have port or intermediate parenting decided on a queue by queue basis, some queues tied directly to the port scheduler priorities while other queues are attached to intermediate schedulers.

Network Queue QoS Policy Queue Parenting

A network-queue policy only supports direct port parent priority association. Intermediate schedulers are not supported on network ports or channels.

Egress Port Scheduler Overrides

Once a port scheduler has been associated with an egress port, it is possible to override the following parameters:

- The max-rate allowed for the scheduler.
- The maximum rate for each priority level 8 through 1.
- The CIR associated with each priority level 8 through 1.

The orphan priority level (level 1) has no configuration parameters and cannot be overridden.

Applying a Port Scheduler Policy to a Virtual Port

In order to represent a downstream network aggregation node in the local node scheduling hierarchy, a new scheduling node, referred to as virtual port, and vport in CLI have been introduced. The vport operates exactly like a port scheduler except multiple vport objects can be configured on the egress context of an Ethernet port.

[Figure 23](#) illustrates the use of the vport on an Ethernet port of a Broadband Network Gateway (BNG). In this case, the vport represents a specific downstream DSLAM.

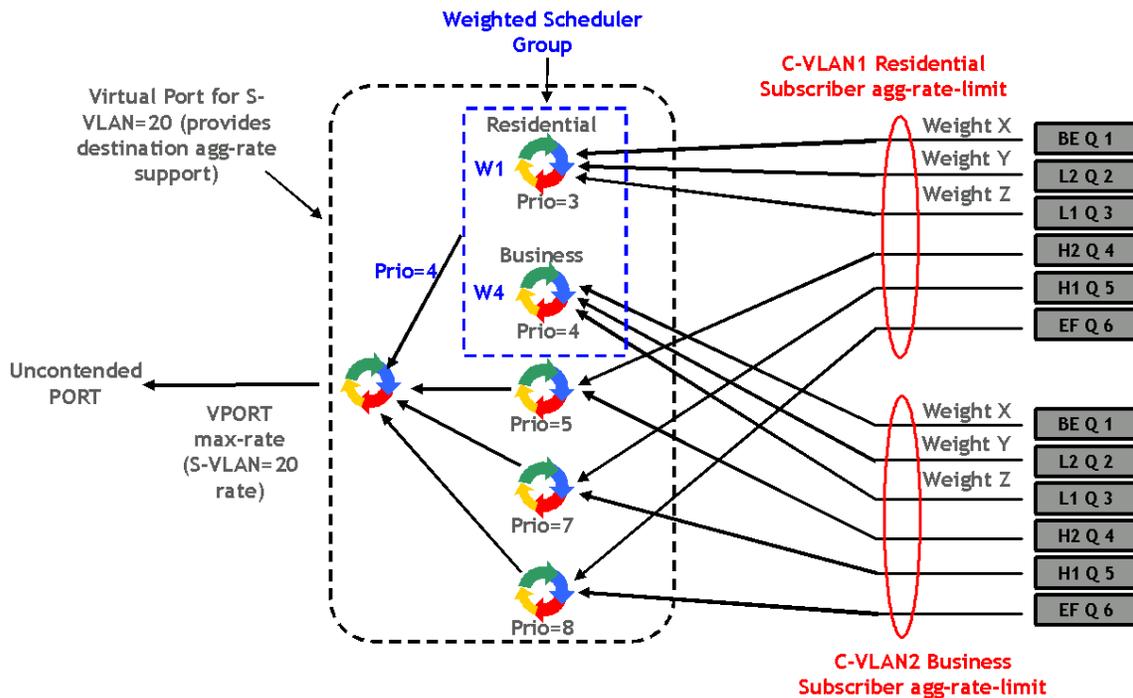


Figure 23: Applying a Port Scheduler Policy to a Vport

The user adds a vport to an Ethernet port using the following command:

CLI Syntax: `configure>port>ethernet>access>egress>vport vport-name create`

The vport is always configured at the port level even when a port is a member of a LAG. The vport name is local to the port it is applied to but must be the same for all member ports of a LAG. It however does not need to be unique globally on a chassis.

The user applies a port scheduler policy to a vport using the following command:

CLI Syntax: `configure>port>ethernet>access>egress>vport>port-scheduler-policy port-scheduler-policy-name`

A vport cannot be parented to the port scheduler when it is using a port scheduler policy itself. It is thus important the user ensures that the sum of the **max-rate** parameter value in the port scheduler policies of all vport instances on a given egress Ethernet port does not oversubscribe the port's rate. If it does, the scheduling behavior degenerates to that of the H/W scheduler on that port. A vport which uses an **agg-rate-limit** can be parented to a port scheduler. This is explained in Section Applying Aggregate Rate Limit to a VPORT.

Each subscriber host queue is port parented to the vport which corresponds to the destination DSLAM using the existing port-parent command:

CLI Syntax: `configure>qos>sap-egress>queue>port-parent [weight weight] [level level] [cir-weight cir-weight] [cir-level cir-level]`

This command can parent the queue to either a port or to a vport. These operations are mutually exclusive in CLI as explained above. When parenting to a vport, the parent vport for a subscriber host queue is not explicitly indicated in the above command. It is determined indirectly. The determination of the parent vport for a given subscriber host queue is described in the 7750 SR OS Triple Play Guide.

Currently, only subscriber host queues can be parented to a vport.

Applying Aggregate Rate Limit to a VPORT

The user can apply an aggregate rate limit to the vport and apply a port scheduler policy to the port.

This model allows the user to over subscribe the Ethernet port. The application of the **agg-rate-limit** option is mutually exclusive with the application of a port scheduler policy to a vport.

When using this model, a subscriber host queue with the **port-parent** option enabled is scheduled within the context of the port's port scheduler policy. More details are provided in the 7750 SR OS Triple Play Guide.

Weighted Scheduler Group in a Port Scheduler Policy

The existing port scheduler policy defines a set of eight priority levels with no ability of grouping levels within a single priority. In order to allow for the application of a scheduling weight to groups of queues competing at the same priority level of the port scheduler policy applied to the vport, or to the Ethernet port, a new group object is defined under the port scheduler policy:

CLI Syntax: `configure>qos>port-scheduler-policy>group group-name rate pir-rate[cir cir-rate]`

Up to eight groups can be defined within each port scheduler policy. One or more levels can map to the same group. A group has a rate and optionally a cir-rate and inherits the highest scheduling priority of its member levels. For example, the scheduler group shown in the vport in [Figure 23](#) consists of level priority 3 and level priority 4. It thus inherits priority 4 when competing for bandwidth with the standalone priority levels 8, 7, and 5.

In essence, a group receives bandwidth from the port or from the vport and distributes it within the member levels of the group according to the weight of each level within the group. Each priority

level will compete for bandwidth within the group based on its weight under congestion situation. If there is no congestion, a priority level can achieve up to its rate (cir-rate) worth of bandwidth.

The mapping of a level to a group is performed as follows:

CLI Syntax: `configure>qos>port-scheduler-policy>level priority-level rate
pir-rate [cir cir-rate] group group-name [weight weight-in-group]`

Note that CLI will enforce that mapping of levels to a group are contiguous. In other words, a user would not be able to add priority level to group unless the resulting set of priority levels is contiguous.

When a level is not explicitly mapped to any group, it maps directly to the root of the port scheduler at its own priority like in existing behavior.

Basic Configurations

A basic QoS scheduler policy must conform to the following:

- Each QoS scheduler policy must have a unique policy ID.
- A tier level 1 parent scheduler name cannot be configured.

A basic QoS port scheduler policy must conform to the following:

- Each QoS port scheduler policy must have a unique policy name.
-

Create a QoS Scheduler Policy

Configuring and applying QoS policies is optional. If no QoS policy is explicitly applied to a SAP or IP interface, a default QoS policy is applied.

To create a scheduler policy, define the following:

- A scheduler policy ID value. The system will not dynamically assign a value.
- Include a description. The description provides a brief overview of policy features.
- Specify the tier level. A tier identifies the level of hierarchy that a group of schedulers are associated with.
- Specify a scheduler name. A scheduler defines bandwidth controls that limit each child (other schedulers and queues) associated with the scheduler.
- Specify a parent scheduler name to be associated with a level 2 or 3 tier.
- You can modify the bandwidth that the scheduler can offer its child queues or schedulers. Otherwise, the scheduler will be allowed to consume bandwidth without a scheduler-defined limit.

The following displays a scheduler policy configuration:

```
A:ALA-12>config>qos# info
#-----
echo "QoS Policy Configuration"
#-----
...
scheduler-policy "SLA1" create
  description "NetworkControl(3), Voice(2) and NonVoice(1) have strict priorities"
  tier 1
    scheduler "All_traffic" create
      description "All traffic goes to this scheduler eventually"
      rate 11000
    exit
  exit
  tier 2
```

```

        scheduler "NetworkControl" create
            description "network control traffic within the VPN"
            parent All_traffic level 3 cir-level 3
            rate 100
        exit
        scheduler "NonVoice" create
            description "NonVoice of VPN and Internet traffic will be serviced by
this scheduler"
            parent All_traffic cir-level 1
            rate 11000
        exit
        scheduler "Voice" create
            description "Any voice traffic from VPN and Internet use this scheduler"
            parent All_traffic level 2 cir-level 2
            rate 5500
        exit
    exit
    tier 3
        scheduler "Internet_be" create
            parent NonVoice cir-level 1
        exit
        scheduler "Internet_priority" create
            parent NonVoice level 2 cir-level 2
        exit
        scheduler "Internet_voice" create
            parent Voice
        exit
        scheduler "VPN_be" create
            parent NonVoice cir-level 1
        exit
        scheduler "VPN_nc" create
            parent NetworkControl
            rate 100 cir 36
        exit
        scheduler "VPN_priority" create
            parent NonVoice level 2 cir-level 2
        exit
        scheduler "VPN_reserved" create
            parent NonVoice level 3 cir-level 3
        exit
        scheduler "VPN_video" create
            parent NonVoice level 5 cir-level 5
            rate 1500 cir 1500
        exit
        scheduler "VPN_voice" create
            parent Voice
            rate 2500 cir 2500
        exit
    exit
    exit
    sap-ingress 100 create
        description "Used on VPN sap"
    ...
-----
A:ALA-12>config>qos#

```

Applying Scheduler Policies

Apply scheduler policies to the following entities:

- Customer
 - Epipe
 - IES
 - VPLS
 - VPRN
-

Customer

Use the following CLI syntax to associate a scheduler policy to a customer's multiservice site:

CLI Syntax: `config>customer customer-id
multiservice-site customer-site-name
egress
scheduler-policy scheduler-policy-name
ingress
scheduler-policy scheduler-policy-name`

Epipe

Use the following CLI syntax to apply QoS policies to ingress and/or egress Epipe SAPs:

CLI Syntax: `config>service# epipe service-id [customer customer-id]
sap sap-id
egress
scheduler-policy scheduler-policy-name
ingress
scheduler-policy scheduler-policy-name`

CLI Syntax: `config>service# epipe service-id [customer customer-id]
sap sap-id
egress
qos sap-egress-policy-id
ingress
qos sap-ingress-policy-id`

The following output displays an Epipe service configuration with SAP scheduler policy SLA2 applied to the SAP ingress and egress.

```
A:SR>config>service# info
-----
epipe 6 customer 6 vpn 6 create
  description "Distributed Epipe service to west coast"
  sap 1/1/10:0 create
    ingress
      scheduler-policy "SLA2"
      qos 100
    exit
  egress
    scheduler-policy "SLA2"
    qos 1010
  exit
exit
...
-----
A:SR>config>service#
```

IES

Use the following CLI syntax to apply scheduler policies to ingress and/or egress IES SAPs:

CLI Syntax: config>service# ies *service-id* [*customer customer-id*]
interface *ip-int-name*
 sap *sap-id*
 egress
 scheduler-policy *scheduler-policy-name*
 ingress
 scheduler-policy *scheduler-policy-name*

The following output displays an IES service configuration with scheduler policy SLA2 applied to the SAP ingress and egress.

```
A:SR>config>service# info
-----
ies 88 customer 8 vpn 88 create
  interface "Sector A" create
    sap 1/1/1.2.2 create
      ingress
        scheduler-policy "SLA2"
        qos 101
      exit
    egress
      scheduler-policy "SLA2"
      qos 1020
    exit
  exit
exit
no shutdown
exit
```

```
-----  
A:SR>config>service#
```

VPLS

Use the following CLI syntax to apply scheduler policies to ingress and/or egress VPLS SAPs:

CLI Syntax: config>service# vpls *service-id* [*customer customer-id*]
sap *sap-id*
egress
scheduler-policy *scheduler-policy-name*
ingress
scheduler-policy *scheduler-policy-name*

The following output displays an VPLS service configuration with scheduler policy SLA2 applied to the SAP ingress and egress.

```
A:SR>config>service# info  
-----  
...  
vpls 700 customer 7 vpn 700 create  
description "test"  
stp  
shutdown  
exit  
sap 1/1/9:0 create  
ingress  
scheduler-policy "SLA2"  
qos 100  
exit  
egress  
scheduler-policy "SLA2"  
exit  
exit  
spoke-sdp 2:222 create  
exit  
mesh-sdp 2:700 create  
exit  
no shutdown  
exit  
...  
-----  
A:SR>config>service#
```

VPRN

Use the following CLI syntax to apply scheduler policies to ingress and/or egress VPRN SAPs:

```
CLI Syntax: config>service# vprn service-id [customer customer-id]  
                interface ip-int-name  
                    sap sap-id  
                        egress  
                            scheduler-policy scheduler-policy-name  
                    ingress  
                        scheduler-policy scheduler-policy-name
```

The following output displays a VPRN service configuration with the scheduler policy SLA2 applied to the SAP ingress and egress.

```
A:SR7>config>service# info  
-----  
...  
    vprn 1 customer 1 create  
        ecmp 8  
        autonomous-system 10000  
        route-distinguisher 10001:1  
        auto-bind ldp  
        vrf-target target:10001:1  
        interface "to-cel" create  
            address 11.1.0.1/24  
            sap 1/1/10:1 create  
                ingress  
                    scheduler-policy "SLA2"  
                exit  
                egress  
                    scheduler-policy "SLA2"  
                exit  
            exit  
        exit  
        no shutdown  
    exit  
    epipe 6 customer 6 vpn 6 create  
-----  
A:SR7>config>service#
```

Creating a QoS Port Scheduler Policy

Configuring and applying QoS port scheduler policies is optional. If no QoS port scheduler policy is explicitly applied to a SAP or IP interface, a default QoS policy is applied.

To create a port scheduler policy, define the following:

- A port scheduler policy name.
- Include a description. The description provides a brief overview of policy features.

Use the following CLI syntax to create a QoS port scheduler policy.

Note that the **create** keyword is included in the command syntax upon creation of a policy.

CLI Syntax:

```
config>qos
  port-scheduler-policy scheduler-policy-name [create]
    description description-string
    level priority-level rate pir-rate [cir cir-rate]
    max-rate rate
    orphan-override [level priority-level] [weight percent]
                   [cir-level priority-level] [cir-weight cir-weight]
```

The following displays a scheduler policy configuration example:

```
*A:ALA-48>config>qos>port-sched-plcy# info
-----
      description "Test Port Scheduler Policy"
      orphan-override weight 50 cir-level 4 cir-weight 50
-----
*A:ALA-48>config>qos>port-sched-plcy#
```

Configuring Port Parent Parameters

The **port-parent** command defines a child/parent association between an egress queue and a port based scheduler or between an intermediate service scheduler and a port based scheduler. The command may be issued in three distinct contexts; **sap-egress>queue** *queue-id*, and **network-queue> queue** *queue-id* and **scheduler-policy>scheduler** *scheduler-name* the **network-queue>queue** *queue-id* context. The **port-parent** command allows for a set of within-cir and above-cir parameters that define the port priority levels and weights for the queue or scheduler. If the port-parent command is executed without any parameters, the default parameters are assumed.

Within-CIR Priority Level Parameters

The within-cir parameters define which port priority level the queue or scheduler should be associated with when receiving bandwidth for the queue or schedulers within-cir offered load. The within-cir offered load is the amount of bandwidth the queue or schedulers could use that is equal to or less than its defined or summed CIR value. The summed value is only valid on schedulers and is the sum of the within-cir offered loads of the children attached to the scheduler. The parameters that control within-cir bandwidth allocation are the port-parent commands **cir-level** and **cir-weight** keywords. The **cir-level** keyword defines the port priority level that the scheduler or queue uses to receive bandwidth for its within-cir offered load. The **cir-weight** is used when multiple queues or schedulers exist at the same port priority level for within-cir bandwidth. The weight value defines the relative ratio that is used to distribute bandwidth at the priority level when more within-cir offered load exists than the port priority level has bandwidth.

A **cir-weight** equal to zero (the default value) has special meaning and informs the system that the queue or scheduler does not receive bandwidth from the within-cir distribution. Instead all bandwidth for the queue or scheduler must be allocated in the port scheduler's above-cir pass.

Above-CIR Priority Level Parameters

The above-cir parameters define which port priority level the queue or scheduler should be associated with when receiving bandwidth for the queue's or scheduler's above-cir offered load. The above-cir offered load is the amount of bandwidth the queue or scheduler could use that is equal to or less than its defined PIR value (based on the queue or schedulers rate command) less any bandwidth that was given to the queue or scheduler during the above-cir scheduler pass. The parameters that control above-cir bandwidth allocation are the port-parent commands **level** and **weight** keywords. The **level** keyword defines the port priority level that the scheduler or queue uses to receive bandwidth for its above-cir offered load. The **weight** is used when multiple queues or schedulers exist at the same port priority level for above-cir bandwidth. The weight value defines the relative ratio that is used to distribute bandwidth at the priority level when more above-cir offered load exists than the port priority level has bandwidth.

CLI Syntax: config>qos# scheduler-policy *scheduler-policy-name*
tier {1 | 2 | 3}
scheduler *scheduler-name*
port-parent [level *priority-level*] [weight *priority-weight*] [cir-level *cir-priority-level*] [cir-weight *cir-priority-weight*]

CLI Syntax: config>qos#
sap-egress *sap-egress-policy-id* [create]
queue *queue-id* [{auto-expedite | best-effort | expedite}]
[priority-mode | profile-mode] [create]
port-parent [level *priority-level*] [weight *priority-weight*] [cir-level *cir-priority-level*] [cir-weight *cir-priority-weight*]

CLI Syntax: config>qos#
network-queue *network-queue-policy-name* [create]
no network-queue *network-queue-policy-name*
queue *queue-id* [multipoint] [{auto-expedite | best-effort | expedite}] [priority-mode | profile-mode] [create]
port-parent [level *priority-level*] [weight *priority-weight*] [cir-level *cir-priority-level*] [cir-weight *cir-priority-weight*]

Service Management Tasks

This section discusses the following service management tasks:

- [Deleting QoS Policies on page 502](#)
 - [Copying and Overwriting Scheduler Policies on page 505](#)
 - [Editing QoS Policies on page 507](#)
-

Deleting QoS Policies

There are no scheduler or port-scheduler policies associated with customer or service entities. Removing a scheduler or port-scheduler policy from a multi-service customer site causes the created schedulers to be removed which makes them unavailable for the ingress SAP queues associated with the customer site. Queues that lose their parent scheduler association are deemed to be orphaned and are no longer subject to a virtual scheduler. The SAPs that have ingress queues that rely on the schedulers enter into an orphaned state on one or more queues.

A QoS scheduler policy cannot be deleted until it is removed from all customer multi-service sites or service SAPs where it is applied.

```
SR7>config>qos# no scheduler-policy SLA2
MINOR: QOS #1003 The policy has references
SR7>config>qos#
```

Removing a QoS Policy from a Customer Multi-Service Site

CLI Syntax: `config>service>customer customer-id
multi-service-site customer-site-name
egress
no scheduler-policy
ingress
no scheduler-policy`

Example: `config>service>customer# multi-service-site "Test"
config>service>cust>multi-service-site# ingress
config>service>cust>multi-service-site>ingress# no
scheduler-policy`

Removing a QoS Policy from SAP(s)

CLI Syntax: config>service# {epipe|vpls} service-id [customer customer-id]
sap sap-id
egress
no scheduler policy
ingress
no scheduler policy

Example: config>service# epipe 6
config>service>epipe# sap sap 1/1/9:0
config>service>epipe>sap# egress
config>service>epipe>sap>egress# no scheduler-policy
config>service>epipe>sap>egress# exit
config>service>epipe>sap# ingress
config>service>epipe>sap>ingress#
config>service>epipe>sap>ingress# no scheduler-policy

CLI Syntax: config>service# {ies|vprn} service-id [customer customer-id]
interface ip-int-name
sap sap-id
egress
no scheduler policy
ingress
no scheduler policy

Example: config>service# vprn 1
config>service>vprn# interface "to-cel"
config>service>vprn>if# sap 1/1/10:1
config>service>vprn>if>sap# ingress
config>service>vprn>if>sap>ingress# no scheduler-policy
config>service>vprn>if>sap>ingress# exit
config>service>vprn>if>sap# egress
config>service>vprn>if>sap>egress# no scheduler-policy
config>service>vprn>if>sap>egress# exit
config>service>vprn>if>sap#

Removing a Policy from the QoS Configuration

To delete a scheduler policy, enter the following commands:

CLI Syntax: config>qos# no scheduler-policy *network-policy-id*

Example: config>qos# no scheduler-policy ***SLA1***

To delete a port scheduler policy, enter the following commands:

CLI Syntax: config>qos# no port-scheduler-policy *network-policy-id*

Example: config>qos# no port-scheduler-policy ***test1***

Copying and Overwriting Scheduler Policies

You can copy an existing QoS policy, rename it with a new QoS policy value, or overwrite an existing policy. The `overwrite` option must be specified or an error occurs if the destination policy exists.

CLI Syntax: `config>qos> copy scheduler-policy src-name dst-name [overwrite]`

Example: `config>qos# copy scheduler-policy SLA1 SLA2`

```
A:SR>config>qos#
...
#-----
echo "QoS Policy Configuration"
#-----
    scheduler-policy "SLA1" create
        description "NetworkControl(3), Voice(2) and NonVoice(1) have strict priorities"
        tier 1
            scheduler "All_traffic" create
                description "All traffic goes to this scheduler eventually"
                rate 11000
            exit
        exit
        tier 2
            scheduler "NetworkControl" create
                description "network control traffic within the VPN"
                parent "All_traffic" level 3 cir-level 3
                rate 100
            exit
            scheduler "NonVoice" create
                description "NonVoice of VPN and Internet traffic will be serviced by
this scheduler"
                parent "All_traffic" cir-level 1
                rate 11000
            exit
            scheduler "Voice" create
                description "Any voice traffic from VPN and Internet use this scheduler"
                parent "All_traffic" level 2 cir-level 2
                rate 5500
            exit
        tier 3
            scheduler "Internet_be" create
                parent "NonVoice" cir-level 1
            exit
            scheduler "Internet_priority" create
                parent "NonVoice" level 2 cir-level 2
            exit
        exit
...
    scheduler-policy "SLA2" create
        description "NetworkControl(3), Voice(2) and NonVoice(1) have strict priorities"
        tier 1
            scheduler "All_traffic" create
                description "All traffic goes to this scheduler eventually"
```

```

        rate 11000
    exit
exit
tier 2
    scheduler "NetworkControl" create
        description "network control traffic within the VPN"
        parent "All_traffic" level 3 cir-level 3
        rate 100
    exit
    scheduler "NonVoice" create
        description "NonVoice of VPN and Internet traffic will be serviced by
this scheduler"
        parent "All_traffic" cir-level 1
        rate 11000
    exit
    scheduler "Voice" create
        description "Any voice traffic from VPN and Internet use this scheduler"
        parent "All_traffic" level 2 cir-level 2
        rate 5500
    exit
exit
tier 3
    scheduler "Internet_be" create
        parent "NonVoice" cir-level 1
    exit
    scheduler "Internet_priority" create
        parent "NonVoice" level 2 cir-level 2
    exit
...
#-----
A:SR>config>qos#

```

Editing QoS Policies

You can change existing policies and entries in the CLI. The changes are applied immediately to all customer multi-service sites and service SAPs where the policy is applied. To prevent configuration errors use the copy command to make a duplicate of the original policy to a work area, make the edits, and then overwrite the original policy.

