



Quality of Service Simulation

Quality of Service (QoS) is a means of ensuring high-quality performance for critical applications. The concept is that since requirements of some users and services are more critical than others, some traffic requires preferential treatment.

Using WAE Design QoS features, you can ensure these mission-critical service levels are met without reactively expanding or over-provisioning the network. These QoS features are available for undifferentiated traffic, for service classes alone, for interface queues alone, and for service classes that are mapped to queues.

Undifferentiated traffic—Aggregate traffic on an interface.

Service class—A user-defined classification of traffic that is not discovered by WAE Collector. Examples could include voice, video, and data. Service classes apply to the entire network unless you map them to specific queues.

Queue—In live networks, traffic waits in conceptual lines (queues) and then is forwarded over an interface on a per-queue basis according to QoS parameters. Similarly in WAE Design, each queue has a set of user-defined QoS parameters (interface queue properties) that specify how these queues are prioritized and what percentage of traffic they carry. An interface may contain zero or more queues that are discoverable by WAE Collector. You can also manually create and configure them. The traffic per queue is also discovered.

QoS Parameters

In WAE Design, QoS requirements are defined by policies and interface queue properties ([Figure 10-1](#)).

Policy—Maximum percentage of traffic capacity that can be utilized by either a service class or by undifferentiated traffic. There are two policies: one for normal operation and one for worst-case scenarios. Policies set on service classes do not affect QoS requirements of any other service class. Nor would this parameter have any effect on live network behavior.

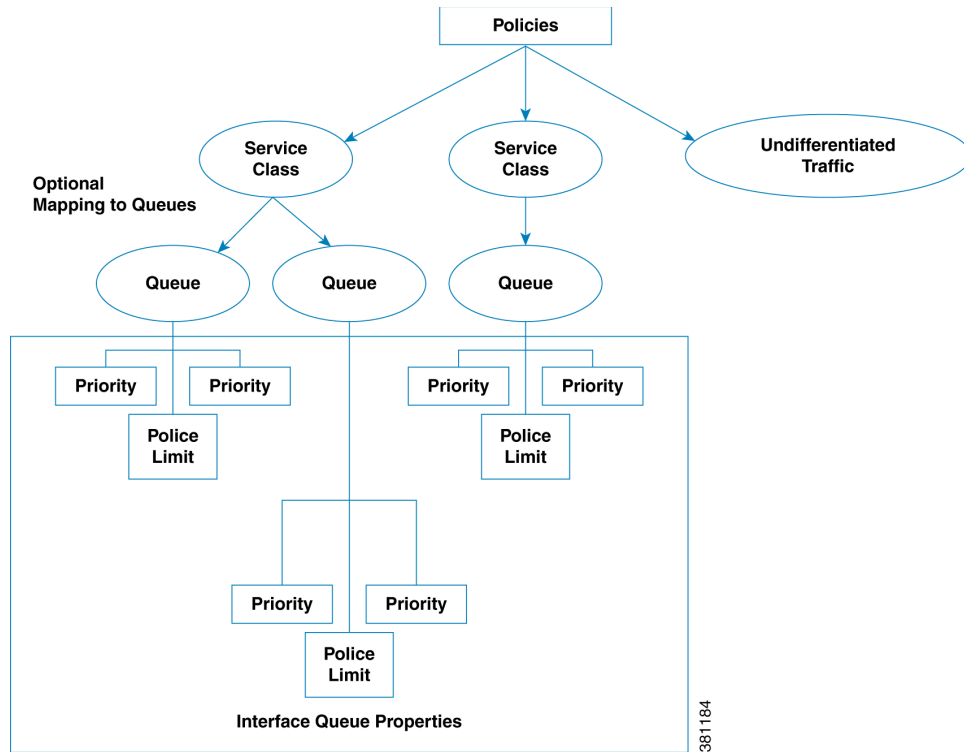
In addition to setting policies for the network, you can further refine the policy to a group of interfaces (called *interface policy group*). For instance, you might need to model behavior for service classes offered only on interfaces of a specific capacity, or you might need to observe service class traffic on only one area of the network.

Interface queue properties—Configured parameters that would affect routing behavior in a live network. In WAE Design, the interface queue properties are priority, weight, and police limit.

- The **priority** identifies the precedence of the queue. For example, traffic in a priority 1 queue is routed before traffic in a priority 2 queue. Queues with the same priority evenly share the capacity based on weighted-round robin (WRR) calculations. You can change this behavior using the weight and police limit parameters. There are an unlimited number of priorities, though most networks only use no more than three. By default, queues do not have priorities.
- The **weight** is the percentage of preference given to queues of an equal priority level, which enables the network to fairly distribute the load amongst available resources. For example, if 10 Gbps were passing through a 10GbE interface on two priority 1 queues, by default 5 Gbps would pass through each queue. However, if you set the weight of one queue to 75% and the other to 25%, the distribution would be 7 Gbps and 2.5 Gbps, respectively. By default, all queues have a weight of 100%.
- The **police limit** is the maximum percentage of available capacity permitted through a queue of a given priority level, thereby preventing traffic from higher priority queues from starving lower priority queues. For instance, if the interface is a 20GbE and a priority 1 queue has a police limit of 40%, then only 8 Gbps of interface traffic can go through this queue. By default, all queues have a police limit of 100%. To see examples of this *starvation*, refer to the examples in the [Policies and QoS Bound Calculations](#) section where you can see that lower priority queues received zero traffic due to priority settings.

| For More Information... | See... |
|-------------------------------------|---|
| Create service classes | Create Service Classes |
| Create and editing policies | Create and Edit Service Class Policies |
| Create interface policy groups | Create or Edit Policy Groups for Interfaces |
| Create Interface queues | Create Queues |
| Assign queues to interfaces | Assign Queues to Interfaces |
| Edit interface queue properties | Edit Interface Queues Properties (QoS Requirements) |
| Map service classes to queues | Map Service Classes to Queues |
| Import QoS model into the plan file | <i>WAE Design Integration and Development Guide</i> |

Figure 10-1 Policies and Interface Queue Parameters



QoS Bound and QoS Violations

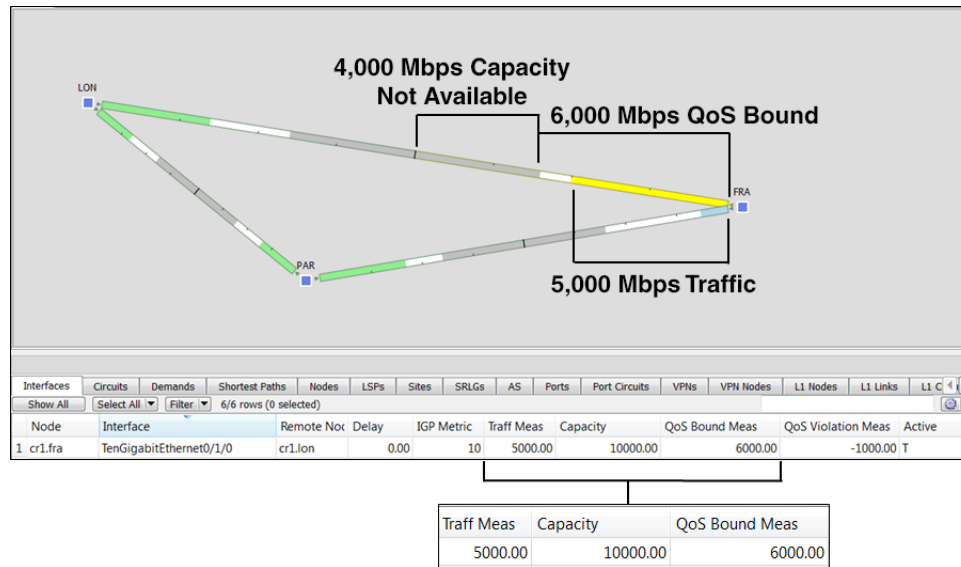
WAE Design uses the concepts of *QoS bound* and *QoS violation* as a way of identifying whether QoS parameters are being met or surpassed, thus better enabling you to plan for service requirements across the network. Policies and queue properties determine the QoS bound calculation. In turn, this calculation determines whether there is a violation.

QoS Bound—Maximum interface capacity available without violating these QoS requirements. A separate QoS bound is calculated for both policy and interface queue properties.

| QoS Bound For... | Calculation Based On... |
|--------------------------------|--|
| Interface Queues | Combination of interface queue properties, or in live networks, it is the capacity percentage that is discovered. |
| Service class | Policy |
| Service class mapped to queues | The lower of these two calculations is used. <ul style="list-style-type: none"> • Policy for service class • Queue properties for queues |
| Undifferentiated traffic | Policy |

In the plot, the QoS bound on an interface is a combination of the color and white (Figure 10-2). The interface capacity that is not available because it exceeds the QoS bound is in gray. Columns that convey the QoS bound information are QoS Bound Meas, QoS Bound Meas (%), QoS Bound Sim, and QoS Bound Sim (%).

Figure 10-2 QoS Bound and Available Interface Capacity

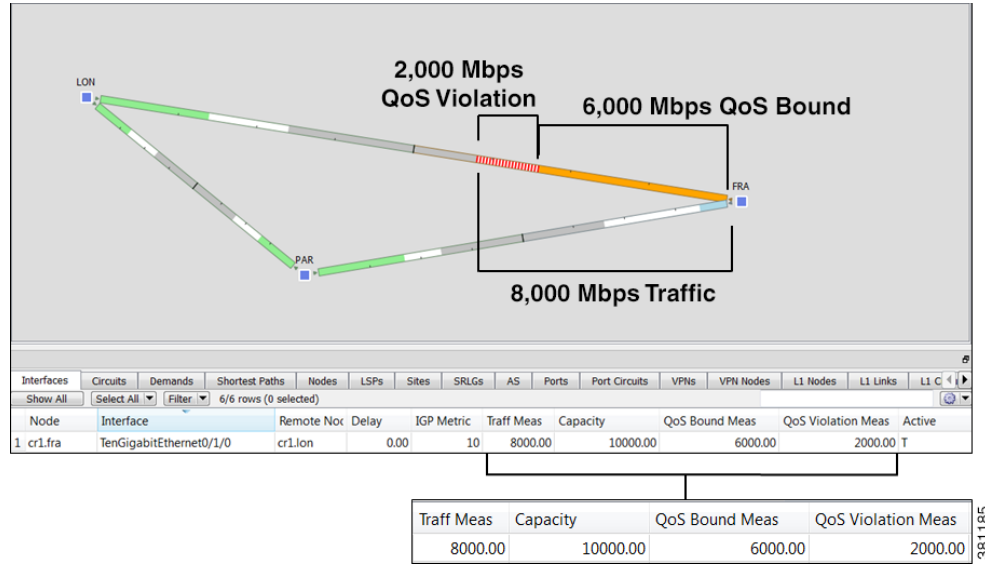


QoS bound calculations are a set of decisions being made to determine how to raise traffic on the queues until that traffic cannot be raised any further. This capacity, or the reason the traffic cannot be raised further, is defined both by the QoS parameters and the amount of traffic. For example, when traffic arrives at Queue X, WAE Design fixes the traffic on all other queues and then determines how it can raise the traffic on Queue X until some other traffic blocks it.

For those queues that do not reach full capacity, unused queue capacity is made available for other queues.

QoS Violation—Total traffic minus the capacity permitted for the queue (QoS bound). A violation occurs if the maximum QoS capacity allotted through policies and interface queue properties is exceeded. If the number appearing in the QoS Violation column is positive, then the allotted capacity has been surpassed. If the number is negative, the allotted capacity has not been reached. In the plot, traffic exceeding the QoS bound appears in red and white stripes on the interface in violation (Figure 10-3).

Figure 10-3 Example QoS Violation



Policies and QoS Bound Calculations

If no other QoS parameters are set via the interface queue properties of priority, weight, and police limit, the QoS bound is equivalent to the policy set.

Table 10-1 Example Policies and QoS Bound Calculations

| Example Configuration | QoS Bound | QoS Violation (Positive # = Violation) |
|--|--|--|
| Interface capacity = 10,000 Mbps Undifferentiated traffic = 5,000 Mbps Normal operation policy = 60% | 6,000 Mbps (60%) | -1,000 Mbps (-10%) Since this number is negative, there is no capacity violation (Figure 10-2). |
| Interface capacity = 10,000 Mbps Undifferentiated traffic = 8,000 Mbps Normal operation policy = 60% | 6,000 Mbps (60%) | 2,000 Mbps (20%) Since this number is positive, there is a capacity violation (Figure 10-3). |
| Interface capacity = 10,000 Mbps Voice traffic = 6,000 Mbps Video traffic = 2,000 Mbps Voice normal operation policy = 90% Video normal operation policy = 60% | Voice = 9,000 Mbps (90%) Video = 6,000 Mbps (60%) | Voice = -3,000 Mbps (30%) Video = -4000 Mbps (40%) |

Interface Queue Properties and QoS Bound Calculations

WAE Design simultaneously calculates QoS bound for each queue in the interface. In so doing, WAE Design uses the interface queue parameters (priority, weight, and police limit) and the traffic measured or simulated for all queues in the interface. Priority is always considered first. If there are queues of equal priority, then weight is applied next.

- Queues with priority 1 share all available interface capacity. Their weight and police limits further refine how much each priority 1 queue can use (their QoS bound). Each priority 1 queue can borrow available capacity from other priority 1 queues up to the limit of their QoS bound.
- The available capacity for priority 2 queues is the total interface capacity less all capacity consumed by priority 1 queues. The process then begins again for all priority 2 queues. Their weight and police limits determine their QoS bound, and priority 2 queues can borrow capacity from each other up to the limits set by the QoS bound.
- This process continues for each successive priority level. Traffic that is outside any QoS bound is dropped to the lowest priority of all traffic on the interface.

For discovered networks with measured traffic, if no WAE Design QoS parameters are set, the QoS bound is based on whatever capacity percentages the live network has for each queue.

Priority

Provided policies are not set that further affect the QoS bound, a queue's QoS bound is calculated as follows.

- Priority 1 QoS bound = 100% of the interface capacity
- Priority 2 QoS bound = Total interface capacity - amount of traffic consumed by priority 1 queues
- Priority 3 QoS bound = Total interface capacity - amount of traffic consumed by (priority 1 + priority 2 queues)
- QoS bound for each succeeding priority follows this same pattern where the traffic consumed by all higher priority queues is subtracted from the total interface capacity.

Table 10-2 Examples of Priority QoS Bound Calculations

| Example Configuration | QoS Bound | QoS Violation (Positive # = violation) | QoS Bound Calculations |
|---|--------------------------------------|--|--|
| Interface capacity = 20,000 Mbps EF traffic = 6,000 Mbps; priority = 1 BE traffic = 3,000 Mbps; no priority set | EF = 20,000 Mbps BE = 14,000 Mbps | EF = -14,000 Mbps BE = -11,000 Mbps | EF = Total interface capacity because it is the only priority 1 queue BE = 20,000 (interface capacity) - 6,000 (consumed by higher priority queues) |
| Interface capacity = 10,000 Mbps EF Traffic = 6,000 Mbps; priority = 1 BE Traffic = 5,000 Mbps; priority = 2 | EF = 10,000 Mbps BE = 4,000 Mbps | EF = -4,000 Mbps BE = 1,000 Mbps | EF = Total interface capacity because it is the only priority 1 queue BE = 10,000 (interface capacity) - 6,000 (consumed by higher priority queues) |

Weight

The weight identifies the forwarding precedence for queues of equal priority. If weights for queues of the same priority do not add up to 100%, weights are converted proportionally so they do add up to 100%.

Table 10-3 Examples of Weight QoS Bound Calculations

| Example Configuration | QoS Bound | QoS Violation (Positive # = Violation) | QoS Bound Calculations |
|--|--------------------------------------|--|---|
| Interface capacity = 10,000 Mbps AF1 traffic = 3,000 Mbps; priority = 1; weight = 100% AF2 traffic = 6,000 Mbps; priority = 1; weight = 100% | AF1 = 5,000 Mbps AF2 = 7,000 Mbps | AF1 = -2,000 Mbps AF2 = -1,000 Mbps | AF1 = Half of capacity for priority 1 queues since both queues have equal weights AF2 = 5,000 (half of capacity) + 2,000 (unused AF1 capacity) |
| Interface capacity = 10,000 Mbps AF1 = 5,000 Mbps; priority = 1; weight = 60% AF2 traffic = 6,000 Mbps; priority = 1; weight = 40% | AF1 = 6,000 Mbps AF2 = 5,000 Mbps | AF1 = -1,000 Mbps AF2 = 1,000 Mbps | AF1 = 60% of capacity for all priority 1 queues AF2 = 10,000 (interface capacity) - 5,000 (consumed by AF1 queue) |

Police Limits

Priority 1 queues have 100% of the interface traffic, and thus starve out the remaining queues. To prevent this queue starvation, use police limits to configure how much of the maximum percentage should be available for a given priority level.

Table 10-4 Examples of Police Limit QoS Bound Calculations

| Example Configuration | QoS Bound | QoS Violation (Positive # = Violation) | QoS Bound Calculations |
|---|---|---|--|
| Interface capacity = 10,000 Mbps EF traffic = 1,000 Mbps; priority = 1; police limit = 50% BE traffic = 2,000 Mbps; priority = 2 | EF = 5,000 Mbps BE = 9,000 Mbps | EF = -4,000 Mbps BE = -7,000 Mbps | EF = 50% of total interface capacity BE = 10,000 (interface capacity) - 1,000 (capacity consumed by EF) |
| Interface capacity = 10,000 Mbps EF traffic = 1,000 Mbps; priority = 1; police limit = 5% BE traffic = 2,000 Mbps; priority = 2 | F = 500 Mbps BE = 9,500 Mbps | EF = 500 Mbps BE = -7500 Mbps | EF = 5% of total interface capacity BE = 10,000 (interface capacity) - 500 (capacity consumed by EF) |
| Interface capacity = 10,000 Mbps EF = 3,000 Mbps; priority = 1; police limit = 20% AF1 traffic = 4,000 Mbps; priority = 2; police limit = 75% AF2 traffic = 2,500 Mbps; priority = 2; police limit 25% | EF = 2,000 Mbps AF1 = 6,000 Mbps AF2 = 4,000 Mbps | EF = 1,000 Mbps AF1 = -2,000 Mbps AF2 = -1,500 Mbps | EF = 20% of total interface capacity AF1 = 75% of (10,000 [interface capacity] - 2,000 [capacity consumed by EF]) AF2 = 10,000 (interface capacity) - 2,000 (capacity consumed by EF) - 4,000 (capacity consumed by AF1) |

Interface QoS Bound Calculations Using Multiple QoS Parameters

WAE Design calculates a QoS bound for interface queues based on all three parameters if they are all configured: priority, weight, and police limits.

Table 10-5 Examples of Interface QoS Bound Calculations Using Multiple QoS Parameters

| Example Configuration | QoS Bound | QoS Violation (Positive #= Violation) | QoS Bound Calculation |
|---|---|---|---|
| Interface capacity = 10,000 Mbps EF = 3,000 Mbps; priority = 1; police limit = 20% AF1 traffic = 4,000 Mbps; priority = 2; weight = 75% AF2 traffic = 2,500 Mbps; priority = 2; weight = 25% | EF = 2,000 Mbps AF1 = 6,000 Mbps AF2 = 4,000 Mbps | EF = 1,000 Mbps AF1 = -2,000 Mbps AF2 = -1,500 Mbps | EF = 20% of total interface capacity AF1 = Maximum of these two values. <ul style="list-style-type: none"> 75% of (10,000 [interface capacity] - 2,000 [capacity consumed by EF]) 8,000 (available capacity) - 2,500 (AF2 traffic) AF2 = Maximum of these two values. <ul style="list-style-type: none"> 25% of (10,000 [interface capacity] - 2,000 [capacity consumed by EF]) 8,000 (available capacity) - 4,000 (AF1 traffic) |

Service Class QoS Bound Calculations Using Multiple QoS Parameters

If service classes have policies and they are mapped to queues, WAE Design calculates a QoS bound for both. WAE Design then uses the lowest value of the two so as to enforce restrictions in the strictest possible manner.

Example:

Interface capacity = 10,000 Mbps

QoS bound for service class = 50%, or 5,000 Mbps based on policy

QoS bound for EF queue = 7,500 Mbps based on combined parameters of priority, weight, and police limit

The QoS bound for this service class is 5,000 Mbps because the policy QoS bound calculation is lower.

View Queue and Service Class Information

Table 10-6 View Queue and Service Class Information

| To View | Show or Select |
|---|--|
| Queue information | Show the Interface Queue table. Select the queue from the QoS drop-down menu. Both the plot and the Traff Meas and Traff Sim columns display traffic data specific to the queue type selected. |
| Per-queue traffic in the Interfaces table | Select the queue from the QoS drop-down menu. Both the plot and the Traff Meas and Traff Sim columns display data specific to the queue type selected. |

Table 10-6 View Queue and Service Class Information

| To View | Show or Select |
|---------------------------|---|
| Per-service-class traffic | Select the service class from the QoS drop-down menu. Both the plot and the Traff Meas and Traff Sim columns in the Interfaces table display data specific to the service class selected. |
| Service class demands | Show the Service Class column in the Demands table. |

View QoS Bounds and QoS Violations

As [Figure 10-2](#) and [Figure 10-3](#) demonstrate, the QoS bounds and QoS violations appear in the plot view. [Table 10-7](#) lists the available column options to display numeric values of the QoS bound calculations. For information on QoS values as they relate to VPNs, see the [VPN Simulation](#) chapter.

Table 10-7 View QoS Bounds and QoS Violations

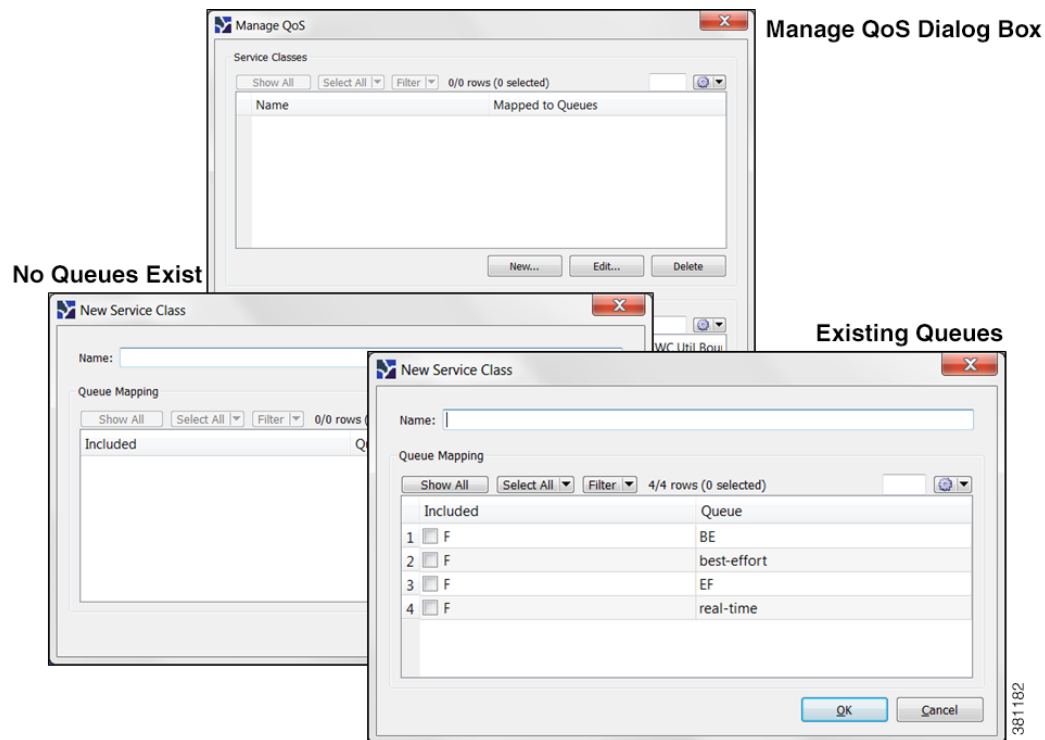
| To View | Show this Column in Interfaces, Circuits, or Interface Queues Table |
|---|---|
| Measured Traffic | |
| Maximum capacity before a QoS bound is violated under normal operations | QoS Bound Meas |
| QoS bound as a percentage of total interface capacity | QoS Bound Meas (%) |
| QoS violations under normal operations; if the number is positive, there is a violation | QoS Violation Meas |
| QoS violation as a percent of the total interface capacity | QoS Violation Meas (%) |
| Simulated Traffic | |
| Maximum capacity before a QoS bound is violated under normal operations | QoS Bound Sim |
| QoS bound as a percentage of total interface capacity | QoS Bound Sim (%) |
| QoS violations under normal operations; if the number is positive, there is a violation | QoS Violation Sim |
| QoS violation as a percent of the total interface capacity | QoS Violation Sim (%) |
| Worst-Case Traffic | |
| Maximum capacity before a QoS bound is violated under worst-case operations | WC QoS Bound |
| WC QoS bound as a percentage of total interface capacity | WC QoS Bound (%) |
| QoS violations under worst-case operations; if the number is positive, there is a violation | WC QoS Violation |
| WC QoS violation as a percent of the total interface capacity | WC QoS Violation (%) |
| Service class causing the worst-case utilization | WC Service Class |

Create Service Classes

- Step 1** Access the Manage QoS dialog box in one of two ways.
- Select Edit->Manage QoS.
 - Select Manage QoS from the QoS drop-down menu in the toolbar.

- Step 2** Click New. The New Service Class dialog box that appears contains queues if they have already been discovered or if they have been manually created. Otherwise, the dialog box is empty. For instructions on how to create queues, see [Create Queues](#).
- Enter a unique name.
 - Optional: If queues exist and if you want to map this new service class to one or more queues, select them from the list. If queues do not exist, but you want them, you must manually create the queues and then return to this dialog box to select them. See [Map Service Classes to Queues](#).
 - Click OK.
- Step 3** Click OK to exit the Manage QoS dialog box. Note, if you close without clicking OK again, the new service classes are not saved.

Figure 10-4 Create Service Classes



Create Queues

WAE Design discovers queues. However, you can manually add them. Once discovered or created, queues appear in the Interface Queues table.

- Step 1** By default, queues are assigned to all interfaces. If you want this new queue to be assigned to specific interfaces, you must first select one or more interfaces.
- Step 2** Open a New Interface Queues Properties dialog box in one of two ways.
- Select the Insert->Interface Queues menu.

- Right click in the plot area then select New->Interface Queues.
- Step 3** Enter the queue name.
- Step 4** Optional: Enter the queue properties of priority, weight, and police limit. For information on how these queue properties behave, see the [Interface Queue Properties and QoS Bound Calculations](#) section.
- Step 5** Click OK. The new queue appears as an option in the QoS drop-down menu in the toolbar, as well as in the Manage QoS dialog box.
- Step 6** Optional: Map a service class to the queue. For further instructions, see the [Map Service Classes to Queues](#) section.
- Another way to create queues is to edit the name of an existing one. For instructions, see the [Assign Queues to Interfaces](#) section.
-

Assign Queues to Interfaces

The easiest way to assign queues to interfaces is to select the interface before creating the queue (see [Create Queues](#)). However, you can follow these steps to re-assign the queue to different interfaces.

- Step 1** From the Interface Queues table, select one or more queues.
- Step 2** Double click; a Properties dialog box appears.
- Step 3** From the Node list, select the node associated with the interface.
- Step 4** From the Interface list, select the interface to which you are assigning the queue.
- Step 5** Optional: You can change the queue name or add a new one.
- Step 6** Optional: Specify QoS requirements in the Priority, Weight, and Police Limit fields. By default, queues do not have a priority and both their weight and police limits are 100%. For information on how these properties behave, see [Interface Queue Properties and QoS Bound Calculations](#).
- Step 7** Click OK.
-

Map Service Classes to Queues

To map service classes to queues, those queues must first exist either because they were discovered by WAE Collector or because you manually added them. See [Assign Queues to Interfaces](#).

- Step 1** Access the Manage QoS dialog box in one of two ways.
- Select Edit->Manage QoS.
 - Select Manage QoS from the QoS drop-down menu in the toolbar.
- Step 2** From the Service Classes list, select one service class.
- Step 3** Click Edit. A Service Class dialog box appears.
- a. Select one or more queues.
 - b. Click OK.

- Step 4** Repeat #2 and #3 for each service class to which you are mapping queues.
- Step 5** Click OK to exit the Manage QoS dialog box. Note, if you close without clicking OK again, the above changes are not saved.
-

Create or Edit Policy Groups for Interfaces

Creating a policy group for interfaces enables you to set policies for the group in the Manage QoS dialog box.

- Step 1** Select one or more interfaces.
- Step 2** Double-click to open the Properties dialog box.
- Step 3** Select the Advanced tab.
- Step 4** In the QoS Policy Group field, you have three options
- To add this interface to an existing policy group, select it from the drop-down menu.
 - To change the name of an existing policy, select it and then type over it.
 - To add a new policy group, select the empty row in the drop-down menu and then type the name of the new group.
- Step 5** Click OK.
- Step 6** To assign a service class to this policy group, use the Manage QoS option in the QoS drop-down in the toolbar. For further instructions, see the [Create and Edit Service Class Policies](#) section.
-

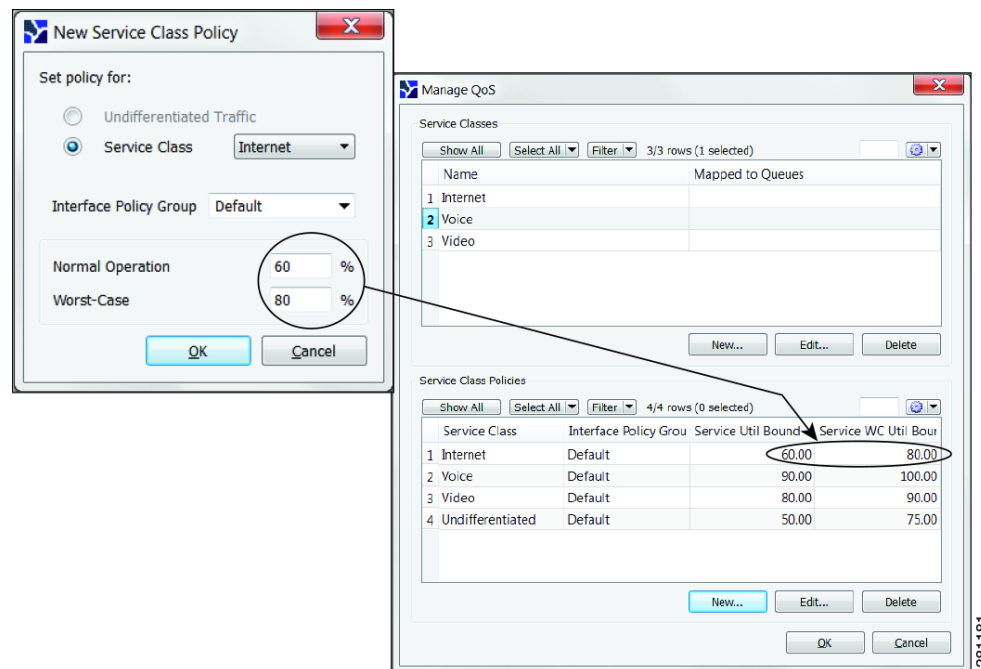
Create and Edit Service Class Policies

You can configure policies for undifferentiated traffic and for service classes. For information on how policies behave, see [Policies and QoS Bound Calculations](#).

- Step 1** Access the Manage QoS dialog box in one of two ways.
- Select Edit->Manage QoS.
 - Select Manage QoS from the QoS drop-down menu in the toolbar.
- Step 2** In the Service Class Policies section, click New, or select a service class row and click Edit. A Service Class Policy dialog box appears.
- a. If creating a policy for undifferentiated traffic, select that option. If creating a policy for an existing service class, select it from the Service Class drop-down menu.
 - b. Optional: To apply this service class mapping to a group of interfaces, select from or enter the name in the Interface Policy Group drop-down menu. You can enter a name that does not exist and then create the policy group for a set of interfaces. For further instructions, see the [Create or Edit Policy Groups for Interfaces](#) section.
 - c. In the Normal Operation field, enter the percentage of bandwidth capacity that you do not want this interface (or group of interfaces) to exceed for this traffic or service class under normal conditions.

- d. In the Worst-Case field, enter the percentage of bandwidth capacity that you do not want this interface (or group of interfaces) to exceed for this traffic or service under worst-case operating conditions.
 - e. Click OK. These values now appear in the Manage QoS dialog box where you can edit them as needed.
- Step 3** Click OK to exit the Manage QoS dialog box. Note, if you close without clicking OK again, the above changes are not saved.

Figure 10-5 Manage QoS dialog box



Edit Interface Queues Properties (QoS Requirements)

The three parameters are priority, weight, and police limit. For information on how these queue properties behave, see the [Interface Queue Properties and QoS Bound Calculations](#) section.

- Step 1** In the Interface Queues table, select one or more queues.
- Step 2** Right click. The Interface Queues Properties dialog box appears.
- Step 3** Change one or more fields to create the desired QoS requirement.
- Step 4** Click OK.

