

Resource Guide

Implementing QoS for WX/WXC Application Acceleration Platforms



Juniper Networks, Inc.
1194 North Mathilda Avenue
Sunnyvale, CA 94089 USA
408 745 2000 or 888 JUNIPER
www.juniper.net

Table of Contents

WX/WXC Outbound QoS Overview	3
When Should You Use Congestion Control?.....	3
When Should You Configure QoS?	3
Congestion Control (Bandwidth Detection) vs. Manual QoS Configuration	4
Congestion Control	4
The Basics of Juniper WX/WXC Outbound QoS.....	5
Dedicated Mode vs. Oversubscribed Mode.....	5
What to Remember if Using Dedicated Mode:	6
What to Remember if Using Oversubscribed Mode:.....	6
Get the Bandwidth Settings Right	7
Understand the Link Overhead (or, What's This 2% Thing?).....	8
ATM Link Overhead (a Whopping 15%)	8
Why Only 80% of Bandwidth Can be Allocated.....	9
Traffic Classes	9
QoS Templates.....	10
The Difference Between Weighted Fair Queuing and Weighted Strict Priority Queuing	10
Who Gets What with WFQ	10
• Dedicated Mode Calculation:.....	10
• Oversubscribed Mode Calculation:	11
Guaranteed Bandwidth in Oversubscribed Mode (It's All Relative)	12
Using Priority	12
Always Apply a Template to "Other Traffic"	12
Appendix A: Oversubscribed QoS Calculation Example	13

WX/WXC Outbound QoS Overview

The purpose of this guide is to provide an understanding of how the WX/WXC application acceleration platforms implement quality of service (QoS). The information contained in this guide is based on the QoS implementation in WXOS 5.2.3.x software. Configuration details are covered in the Operators Guide and are not discussed here.

The WX/WXC platforms provide a rich set of QoS tools for managing bandwidth. It's quite easy to spend a lot of time fine-tuning your QoS configuration, carefully identifying your bandwidth usage, creating many traffic classes and meticulously prioritizing applications. Fortunately, most networks don't require that much effort; more often than not, only a few traffic classes are required and only a handful of applications need to be prioritized in order to run effectively.

The most important requirement for effectively setting up QoS is a clear understanding of the bandwidth that is available. Quite often customers are unsure of the exact bandwidth they have between sites. Bandwidth oversubscription or a lack of knowledge about one's own network is common; large enterprises may have hundreds of sites, and details regarding each and every WAN circuit are not always readily available. This can make effectively implementing QoS quite a challenge.

The WX/WXC platforms offer a feature called congestion control to address this issue. Congestion control will dynamically detect the available bandwidth to a given destination, allowing customers to quickly roll out new sites or increase or decrease the bandwidth without having to worry about constantly tuning the WAN optimizer. Congestion control and QoS can work together or independently; by enabling congestion control, many customers can avoid detailed QoS configurations altogether. For customers who need specific control over their traffic, they can either configure QoS alone to manage the traffic or a combination of QoS and congestion control.

When Should You Use Congestion Control?

You should use congestion control when:

- There is variable bandwidth
- Bandwidth is unknown
- You want to keep things simple

When Should You Configure QoS?

You should configure QoS if:

- Greater traffic control or bandwidth guarantees for managing traffic are needed
- There is loss or out-of-order data in the network.
- *Note:* Enabling the Forward Error Correction (FEC) feature may also help in this case; see the Operators Guide for details)
- You need to squeeze every last bit of performance out of your available bandwidth

Congestion Control (Bandwidth Detection) vs. Manual QoS Configuration

While these are often discussed as completely different concepts, QoS and congestion control are not mutually exclusive. In fact, you can have both QoS and congestion control configured for the same endpoint, or you can have some endpoints configured only for QoS and others configured only for congestion control. There are many cases where both technologies are desirable; because of this flexibility, it is sometimes difficult to determine the best configuration option for a given situation. This guide will discuss the available configuration options and describe when they are generally applicable.

Congestion Control

The simplest way to install the WX/WXC products is to use congestion control to manage bandwidth. This approach works well and minimizes the amount of additional work required by the user. Congestion control automatically senses the bandwidth available on the WAN and adjusts throughput accordingly. Even with congestion control, you can still prioritize traffic, enable DSCP marking, guarantee bandwidth, etc. If you have an environment where you may be increasing or decreasing bandwidth to various locations over time and don't want to adjust the configured bandwidth to each site every time it changes, then congestion control will automatically learn the available bandwidth when it changes.

Bandwidth allocation is dynamic when congestion control is enabled. If you have oversubscribed WAN bandwidth or variable speed links like MPLS/VPN or Frame Relay in your environment, then you will want to enable congestion control. However, if your environment has large amounts of loss or large amounts of out-of-order traffic, then congestion control alone may not be the best choice.

Congestion control is similar to TCP in that it will respond to loss and out-of-order traffic by reducing the amount of traffic it sends, essentially adjusting the size of the WAN pipe on the fly. If you have applications that always require a specific amount of bandwidth, then congestion control's dynamic bandwidth adjustment may not be suitable for your environment.

Congestion control also allows you to set minimum speeds on a per-endpoint basis, guaranteeing a certain amount of bandwidth is always available for those endpoints. For example, if you have a 2MB circuit with a 1MB CIR that occasionally bursts to 2MB, you can set the maximum bandwidth level to 2MB in QoS and a minimum speed of 1MB in congestion control.

In a network with mixed technologies, you might do the following.

- Configure dedicated links directly with QoS.
- Configure congestion control with a minimum speed for oversubscribed MPLS networks where you need to guarantee minimum bandwidth levels for each site but still need to send at maximum data rates.
- Configure congestion control with no minimum for Internet VPNs with no available traffic rate guarantees.

Note: When configuring minimum bandwidth in congestion control, the WX/WXC devices will not reduce bandwidth below the configured minimum even if there is loss or congestion in the network.

TIP: In the case of loss, you can configure congestion control to be more aggressive (i.e. backs off less) by changing the mode of congestion control. This can be changed by doing the following:

- Go to Admin -> tools-> command line interface
 - Enter the following lines in the window:


```
configure qos outbound set congestion-control-action-on-loss scps-slow-backoff
commit
```
 - Press "Submit"
- Remember to save your configuration.

The Basics of Juniper WX/WXC Outbound QoS

Dedicated Mode vs. Oversubscribed Mode

When configuring QoS, you first need to pick a model: Dedicated Mode or Oversubscribed Mode.

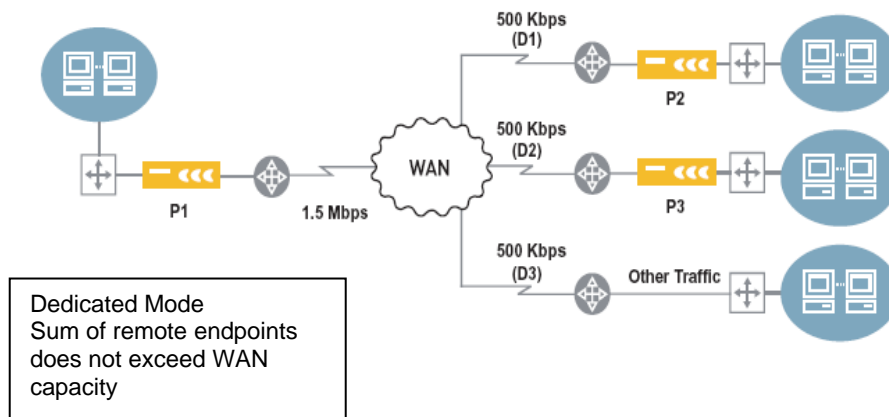


Figure 1: Dedicated Mode

In dedicated mode, the sum of the remote WAN speeds is equal to or less than the local WAN router's WAN speed. Configuring QoS, either manually or via congestion control, is relatively simple in this example.




Figure 2: Oversubscribed Mode

In oversubscribed mode, the sum of the remote WAN speeds exceeds the local WAN router's capacity. In this type of environment, congestion control is typically a good choice, but you can still manually configure QoS and it will work properly. In this example, direct QoS configuration is not as straightforward as it might seem. See the Oversubscribed Mode discussion later in this document for more details.

What to Remember if Using Dedicated Mode:

- Accurately configure the bandwidth for each endpoint
- DO NOT apply a policy to Other Traffic (by default, no policy exists)
- There should be no or very little contention for bandwidth if using dedicated mode.

What to Remember if Using Oversubscribed Mode:

- Accurately configure the outbound speed
- Accurately identify a bandwidth setting for Other Traffic
- Configure a policy for Other Traffic (mandatory)
- Accurately configure bandwidth for each endpoint
- A good alternative might be to just use congestion control as described in the configurations section in this document.

TIP: On the QoS overview screen, you can change the "Guaranteed Bandwidth" display from percentage to Kbps by selecting "Kbps" from the pull-down menu next to "Show bandwidth as."

If variable or unknown bandwidth conditions exist, it is better to configure devices for congestion control

The biggest difference between QoS and congestion control is the way "Other Traffic" is handled. Other Traffic is traffic that is either passed through because it does not match a reduction policy on the WX/WXC platform or is explicitly excluded from reduction. If you are in Oversubscribed Mode, you must configure a policy and bandwidth value for Other

Traffic to ensure it gets the proper QoS treatment. If you do not configure this properly, it will severely impact traffic that is passed through. Since you are managing pass-through traffic in this mode, it is critical to properly configure the outbound speed for all pass-through traffic and to configure the endpoint speeds correctly as well.

The downside of performing QoS on Other Traffic is that in high-bandwidth environments where compression, QoS and acceleration are applied to tunneled traffic, it can impact overall performance of the WX/WXC device.

If configuring QoS in oversubscribed mode, always use the wizard; it will make all appropriate calculations when configuring policies for each endpoint. If you do not use the wizard, be prepared to do some serious math.

In dedicated mode, the WX/WXC platform will pass through Other Traffic by default since there should be no contention for this bandwidth. If a policy is applied to Other Traffic in dedicated mode, it will act like oversubscribed mode and perform QoS on the Other Traffic. Therefore, unless there is a specific reason for doing so, do not apply a policy to Other Traffic when in Dedicated Mode.

When in Dedicated Mode, you need to make sure you know how much bandwidth each remote site has so you can configure the bandwidth accordingly in the QoS endpoints section. If you configure the value too high, you can overdrive the WAN, causing traffic to drop and performance to suffer.

Get the Bandwidth Settings Right

If you manually configure QoS (using the QoS wizard), it is critical to get the correct values for all WAN circuit speeds involved. If the bandwidth is configured too high, it can overdrive the middle of the network, causing dropped traffic and degraded performance. If the bandwidth is configured too low, it will not be able to fully utilize the available capacity, resulting in sub-optimal performance.

It is very important to clearly understand the differences between the outbound speed of the local device and the bandwidth to each remote site.

The outbound speed setting in the QoS wizard sets the speed of the outgoing WAN interface for the attached router. This setting is only used if configuring QoS using Oversubscribed Mode, since the sum of the attached WAN devices typically exceeds the bandwidth of the local circuit speed.

For example, assume a router is attached to a frame relay network with six sites, each site connected via a 1500Kbps WAN circuit. Since there is no way you can send 9000Kbps of traffic on a 1500Kbps circuit, this would be configured as oversubscribed.

If a network falls into the oversubscribed category like this, it is generally best to configure everything with a large WAN bandwidth value and enable congestion control to manage bandwidth.

The best environment for configuring outbound QoS is a network that fits the Dedicated Mode model. In a dedicated environment, the most important thing to know is the bandwidth to the remote sites. If this value is set too high, it overdrive the middle of the network and start dropping packets; if it is set too low, bandwidth is not being fully utilized.

Understand the Link Overhead (or, What's This 2% Thing?)

On any WAN link, some traffic overhead is always required for things like router updates, hello packets, etc. Juniper suggests allocating two percent of the total bandwidth to accommodate this traffic. Take, for example, a T-1 link with 1536Kbps of usable bandwidth; to determine the proper two percent overhead, multiply the usable bandwidth by .98. In this case, the answer is about 1505Kbps; for the sake of simplicity, just configure 1500Kbps in the wizard.

ATM Link Overhead (a Whopping 15%)

If a customer has asynchronous transfer mode (ATM) links, it is important that you know exactly what the bandwidth really is in order to determine the proper overhead. In most cases when a customer tells you the circuit size, they are actually providing the raw circuit size; if you take two percent off that number, you will still be way over the actual bandwidth, since ATM uses 53-byte cells with 5 bytes of overhead in each cell. This translates to about 9.5 percent total overhead (for simplicity's sake, we'll just round up to 10). This is commonly referred to as the ATM Cell Tax.

While this will give you the effective speed of the link, you still need to factor in the same two percent overhead as before, which brings total to 12 percent.

But that's not all. Since not everything fits neatly into 48 byte chunks (53-byte cells minus the five byte overhead), there will naturally be some cell padding. Add to that overhead related to the aal5 encapsulation type (which requires 16 bytes) and aal5Mux (which requires eight bytes).

The actual overhead is related to the amount of mismatch. The table below shows the amount of additional ATM overhead required for various packet sizes. Since the WX/WXC tries to put as much data as possible into the packet, they will typically be on the larger side. Generally speaking, you can use three percent to calculate the additional overhead, but you may need to increase this if you have much smaller packets. The end result is a 15 percent overhead for most ATM networks. This has bitten many people many times, so remember it well. A general rule of thumb for ATM links: take the link speed and multiply by .85 to get the total overhead required.

ATM Link Overhead Calculation:

Cell overhead = 10%

Link protocol overhead = 2%

Wasted cell padding = 3%

Total Overhead 15%

To determine the ATM overhead for the same T-1 circuit used in the earlier example, the formula would look like this:

1536Kbps x .85 = 1305Kbps overhead

Packet Size	ATM Cell Tax	Protocol Overhead	Number of Cells Required	aal5snap =16 Leftover bytes	aal5 Overhead Percentage	Total Overhead Percentage
1500	10	2	31.25	36	2.4	14.400
1350	10	2	28.125	42	3.1	15.111
1200	10	2	25	53	4.4	16.417
500	10	2	10.417	27.84	5.6	17.568
250	10	2	5.208	38.4	15.4	27.360

Table 1: Total ATM overhead for different packet sizes.

Why Only 80% of Bandwidth Can be Allocated

When configuring QoS, the basic goal is to guarantee a certain bandwidth percentage to each traffic class. The WX/WXC platforms allow up to 80 percent of the outbound speed to be guaranteed to traffic classes. Any remaining bandwidth that is not guaranteed is called Excess Bandwidth.

If the guaranteed bandwidth is not being fully utilized, then any portion that is not in use becomes Excess Bandwidth. For example, if 60 percent of the outbound speed is guaranteed to traffic classes but only 45 percent is being consumed at a given time, the total Excess Bandwidth at that time is 55 percent.

The reason for providing Excess Bandwidth is to make sure that there is always bandwidth available for local system resources such as SNMP updates and management traffic.

The way in which Excess Bandwidth is allocated depends on the scheduling algorithm used: Weighted Fair Queuing (WFQ) or Weighted Strict Priority (WSP). Both methods are described in the following section.

Traffic Classes

Traffic classes are used by the WX/WXC platforms to group traffic types using application definitions to categorize them. When configuring the WX/WXC platforms using the wizard, five different traffic classes are created by default: Business Critical, Business Standard, Low Latency, Prohibited and Default. The WX/WXC platforms allow up to 16 traffic classes to be defined, but for most environments, the five default categories are plenty.

Priorities and bandwidth allocations are specified by traffic class, and each class can have one or more applications. The WX/WXC platforms can guarantee up to 80 percent of the total bandwidth across all classes.

Traffic is dropped when the maximum bandwidth is exceeded or when the guaranteed bandwidth is exceeded while the circuit is fully utilized (i.e. by a burst of high-priority traffic). A minimum of 20 percent of the total bandwidth (un-guaranteed Excess Bandwidth) is available for priority sharing.

The amount of bandwidth allocated to each traffic class, and each class's priority, is managed by applying a QoS template to the endpoint.

QoS Templates

Priorities and bandwidth defined for each traffic class constitute a template. Multiple QoS templates can be defined. QoS templates can be applied on a per-endpoint basis; in most environments, you can define a single QoS template that can be applied to all endpoints.

For a dedicated WAN, if the same bandwidths and priorities are applied to each endpoint, the setup wizard produces the same results as if entering the QoS settings individually.

For an oversubscribed WAN, the wizard adjusts the template percentages so that the guaranteed portion of the outbound speed is distributed fairly across all classes and endpoints. If this is attempted manually, a lot of complex math will be required to get the right guarantees allocated properly to the endpoints.

The Difference Between Weighted Fair Queuing and Weighted Strict Priority Queuing

- **Weighted Strict Priority (WSP):** With WSP, queues are created for each priority. Excess Bandwidth is allocated by based only on priority; that is, the class with the highest priority gets all the Excess Bandwidth it needs before any Excess Bandwidth is allocated to the class with the next highest priority.
- **Weighted Fair Queuing (WFQ):** With WFQ, queues are created for each traffic class, and Excess Bandwidth is allocated as described in the following section. The allocation depends on whether the WAN is dedicated or oversubscribed.

Who Gets What with WFQ

If bandwidth guarantees have been defined, this bandwidth is allocated first. Once all guarantees have been met, the Excess Bandwidth is divided up between the various traffic classes using the following methods.

- **Dedicated Mode Calculation:**

To calculate the percentage of Excess Bandwidth allocated to a traffic class for a specific remote endpoint (since priorities start with zero, they must be incremented by one for this calculation), use the following formula:

$$\text{(Class Priority + 1)} / \text{(Sum of active class priorities + 1 for each class)}$$

For example, for the five standard classes where four classes have priority zero and one class has priority 7, the priority 7 class receives the following minimum percentage of Excess Bandwidth:

$$\text{Excess\%} = 8/12 = 66\%$$

Note that if only one class has traffic, then that class receives 100% of the Excess Bandwidth.

To calculate the minimum Excess Bandwidth for a class in Kbps, use the following formula:

$$\text{(Excess\%)} \text{(Remote WAN speed - Total class guarantee in Kbps)}$$

For example, if the Excess% is 66 percent, the remote WAN speed is 500 Kbps, and the guaranteed bandwidth for all classes is 80 percent, the minimum Excess Bandwidth calculation would be $(.66)(500 - 500 \times .8) = 66$ Kbps

- **Oversubscribed Mode Calculation:**

The Excess Bandwidth percentage for a class on a specific endpoint is calculated in the same manner as a dedicated WAN, except that the priorities must be totaled across all remote endpoints. For example, if three endpoints are using the same classes and priorities as in the dedicated example above, the minimum Excess Bandwidth for the priority 7 class is:

Excess% = Class Priority +1/(sum of all endpoints active class priorities +1 for each class)

Excess% = $8/(12 + 12 + 12) = 22\%$

To calculate the minimum Excess Bandwidth for a class in Kbps, use the following formula:

(Excess%) x (Outbound speed - All endpoint class guarantees in Kbps)

In this case, the remaining bandwidth was 900Kbps, so we get the following result: $.22 \times 900 = 198$ Kbps

Note: You must calculate the sum of the guaranteed bandwidths for each class on each remote endpoint. To calculate the sum of the guaranteed bandwidth in Kbps, do the following:

For each endpoint, add up the total percentage of all guarantees combined. For example, if we use the same template for all endpoints – high 10 percent; medium 20 percent; and low 10 percent – the total would be 40 percent. Now use the following formula to determine the dedicated bandwidth in Kbps for each endpoint:

Endpoint% x (Outbound Speed/ sum of all endpoint speeds)

If we use the oversubscribed example in Figure 2, we have an Outbound Speed of 1500Kbps and three endpoints with speeds of 1500Kbps, 500Kbps and 60Kbps. Therefore, our calculation would look like this.

$.4 \times (1500 / (1500+500+60)) = .2912$

$D1 = 1500 \times .2912 = 436.8$ Kbps

$D2 = 500 \times .2912 = 145.6$

$D3 = 60 \times .2912 = 17.5$

Total = 599.9Kbps (rounded up to 600Kbps).

This leaves 900Kbps to allocate via WFQ.

This same formula can also be used to calculate how much any given class will have allocated to it on a per-endpoint basis:

Class% x (Outbound Speed/ sum of all endpoint speeds)

If you want to determine the bandwidth to an endpoint in Oversubscribed Mode, use this formula:

Total Endpoint% x (Outbound Speed/ sum of all endpoint speeds)

Note: This is the value used to calculate the guaranteed bandwidth. The actual data rate could be higher than this, but will not exceed the actual bandwidth of the circuit.

Guaranteed Bandwidth in Oversubscribed Mode (It's All Relative)

It is important to remember that guaranteed bandwidth in Oversubscribed Mode is a relative term — that is, it will be based on how oversubscribed the link is. You should always confirm the actual guaranteed bandwidth in Kbps from the QoS overview page; this may differ a bit from what you expect. When viewing the QoS overview, you can see the guaranteed bandwidth displayed in Kbps rather than as a percentage by selecting “Kbps” from the pulldown menu next to “Show bandwidth as;” see Appendix A: Oversubscribed QoS Calculation Example for a more detailed example of how guarantees may look in an oversubscribed network.

Using Priority

In Oversubscribed Mode, priority may play a bigger part in allocating bandwidth than guarantees. If guaranteed bandwidth in Kbps is not as high as you would like, you can increase the priority of the traffic class in the QoS template, ensuring that class a greater share of Excess Bandwidth. See Appendix A: Oversubscribed QoS Calculation Example for a more detailed look at how bandwidth is allocated by priority in a network.

Always Apply a Template to “Other Traffic”

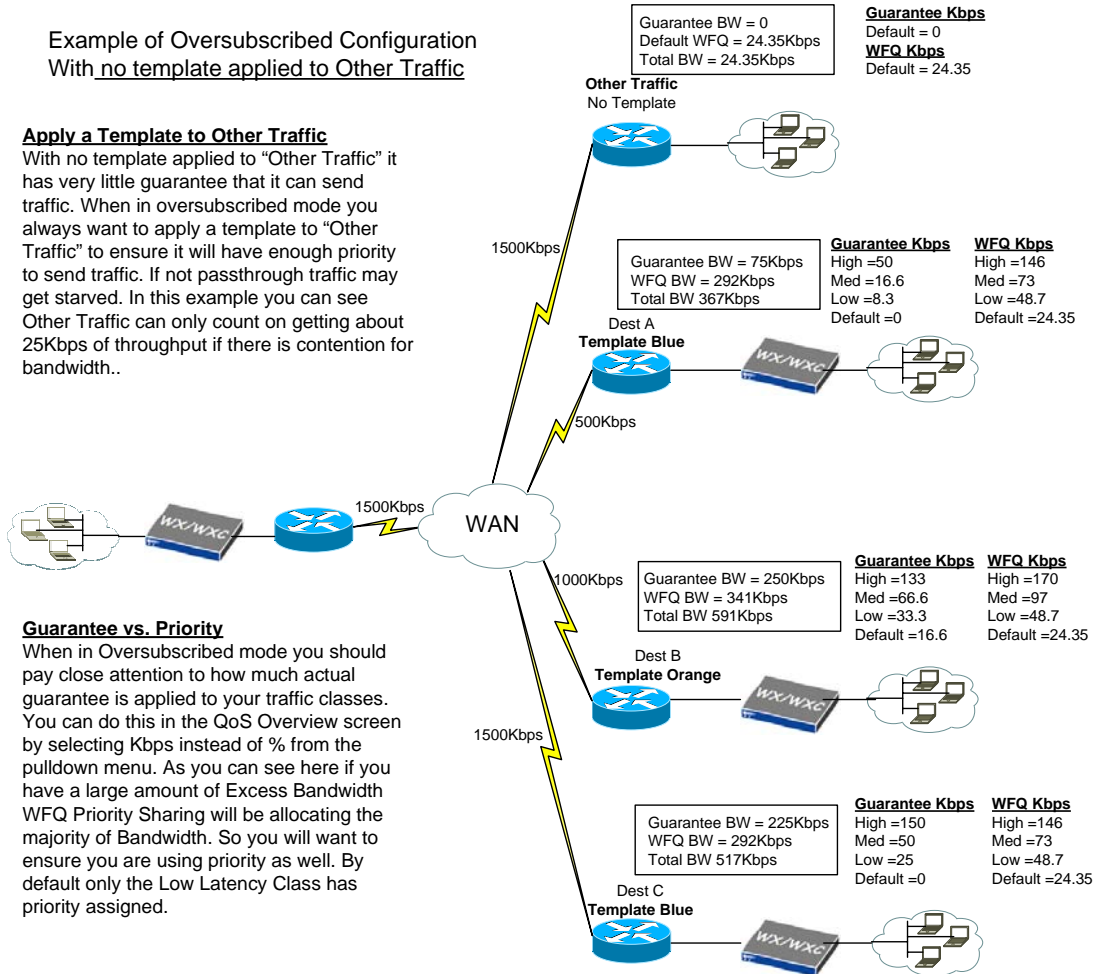
“Other Traffic” is pass-through traffic, meaning it is traffic not destined for locations where another WX/WXC device is at the other end. Most networks have some pass-through traffic, and you will want to ensure this traffic has sufficient bandwidth and priority. If you do not apply a template to Other Traffic in Oversubscribed Mode, then it will treat all traffic as default traffic and have no guarantees. This can result in pass-through traffic being starved for bandwidth, which can cause problems. Appendix A: Oversubscribed QoS Calculation Example provides additional details.

Appendix A: Oversubscribed QoS Calculation Example

Example of Oversubscribed Configuration With no template applied to Other Traffic

Apply a Template to Other Traffic
 With no template applied to "Other Traffic" it has very little guarantee that it can send traffic. When in oversubscribed mode you always want to apply a template to "Other Traffic" to ensure it will have enough priority to send traffic. If not passthrough traffic may get starved. In this example you can see Other Traffic can only count on getting about 25Kbps of throughput if there is contention for bandwidth..

Guarantee vs. Priority
 When in Oversubscribed mode you should pay close attention to how much actual guarantee is applied to your traffic classes. You can do this in the QoS Overview screen by selecting Kbps instead of % from the pulldown menu. As you can see here if you have a large amount of Excess Bandwidth WFQ Priority Sharing will be allocating the majority of Bandwidth. So you will want to ensure you are using priority as well. By default only the Low Latency Class has priority assigned.



Templates used in this example			
Template Blue		Template Orange	
High = 30%	Priority = 5	High = 40%	Priority = 6
Med = 10%	Priority = 2	Med = 20%	Priority = 3
Low = 5%	Priority = 1	Low = 10%	Priority = 1
Default = 0%	Priority = 0	Default = 5%	Priority = 0

Calculations
Class Guarantee in Kbps
 Class% x (Outbound Speed/ Sum of all Endpoint Speeds)
 Class% x Endpoint Speed = Kbps
Excess Bandwidth Distribution Per Class using WFQ
 Excess% = (Class Priority +1)/(sum of all endpoints active class priorities +1 for each class)
Excess Bandwidth Distribution Per Class in Kbps using WFQ
 Excess% x (Outbound Speed – all endpoint class Guarantees in Kbps)

Juniper Networks and the Juniper Networks logo are registered trademarks of Juniper Networks, Inc. in the United States and other countries. All other trademarks, service marks, registered trademarks, or registered service marks in this document are the property of Juniper Networks or their respective owners. All specifications are subject to change without notice. Juniper Networks assumes no responsibility for any inaccuracies in this document or for any obligation to update information in this document. Juniper Networks reserves the right to change, modify, transfer, or otherwise revise this publication without notice.