



>BUSINESS MADE **SIMPLE**

**NORTEL**

Ethernet Routing Switch

8600

Engineering Lab note

> Traffic Shaping Application on the ERS 8600  
Connected to TLS

Core Systems Engineering

Document Date: March 2, 2007

Document Number: NN46205-203

Document Version: 01.02

Nortel is a recognized leader in delivering communications capabilities that enhance the human experience, ignite and power global commerce, and secure and protect the world's most critical information. Serving both service provider and enterprise customers, Nortel delivers innovative technology solutions encompassing end-to-end broadband, Voice over IP, multimedia services and applications, and wireless broadband designed to help people solve the world's greatest challenges. Nortel does business in more than 150 countries. For more information, visit Nortel on the Web at [nortel.com](http://nortel.com).

**NORTEL NETWORKS CONFIDENTIAL:** This document contains material considered to be proprietary to Nortel. No part of it shall be disclosed to a third party for any reason except after receiving express written permission from Nortel and only after securing agreement from the third party not to disclose any part of this document. Receipt of this document does not confer any type of license to make, sell or use any device based upon the teachings of the document. Receipt of the document does not constitute a publication of any part hereof and Nortel explicitly retains exclusive ownership rights to all proprietary material contained herein. This restriction does not limit the right to use information contained herein if it is obtained from any other source without restriction.

Nortel Business Made Simple, Nortel, the Nortel logo, and the Globemark are trademarks of Nortel Networks.

All other trademarks are the property of their owners.

Copyright © 2007 Nortel Networks. All rights reserved. Information in this document is subject to change without notice. Nortel assumes no responsibility for any errors that may appear in this document.

#### **Disclaimer**

This engineering document contains the best information available at the time of publication in terms of supporting the application and engineering of Nortel products in the customer environment. They are solely for use by Nortel customers and meant as a guide for network engineers and planners from a network engineering perspective. All information is subject to interpretation based on internal Nortel test methodologies which were used to derive the various capacity and equipment performance criteria and should be reviewed with Nortel engineering primes prior to implementation in a live environment.



## Abstract

With the growth of Ethernet-based services for providing connectivity between multiple sites, an increasing number of enterprises use the ERS 8600 to communicate with a metropolitan area network/wide area network (MAN/WAN) Transparent local area network (LAN) Service (TLS).

## References

[1] “Configuring QoS and Filtering for the Ethernet Routing Switch 8600 R Modules”  
Available at <http://www.nortel.com/support>

[2] “Technical Configuration Guide for Ethernet Routing Switch 8600 RSeries Modules Filters and QoS”  
Available at <http://www.nortel.com/support>



# Table of Contents

<b>REFERENCES .....</b>	<b>4</b>
<b>LIST OF FIGURES .....</b>	<b>6</b>
<b>LIST OF TABLES.....</b>	<b>6</b>
<b>1. INTRODUCTION .....</b>	<b>7</b>
<b>2. NETWORK OVERVIEW.....</b>	<b>8</b>
<b>3. CONFIGURING THE ERS 8600 .....</b>	<b>9</b>
3.1 Planning.....	9
3.2 Applying an Egress Queue Set .....	9
<b>4. OPERATION .....</b>	<b>13</b>
4.1 Operating initially without congestion .....	13
4.2 Operating with congestion .....	15
4.3 Sending traffic using QoS Level 5 .....	16
4.4 Using Default Queue Set with Port Level Shaping.....	16
<b>5. DESIGN CONSIDERATIONS .....</b>	<b>18</b>



## List of Figures

Figure 1 Network to test shaping on the ERS 8600 .....	8
Figure 2 QOSEqs (QoS Egress Queues) dialog box .....	10
Figure 3 QOSEqs dialog box showing new queue set .....	11
Figure 4 QoS Egress Queue Set Queues dialog box.....	11
Figure 5 QoS Egress Queue Set Queues dialog box.....	13
Figure 6 Ixia statistics dialog box.....	15
Figure 7 QoS Egress Queue Set Queues dialog box.....	17

## List of Tables

Table 1 Bandwidth and QoS requirements for traffic types.....	9
Table 2 Summary of traffic classes, bandwidth requirements and mappings .....	12

# 1. Introduction

With the growth of Ethernet-based services for providing connectivity between multiple sites, an increasing number of enterprises use the ERS 8600 to communicate with a MAN/WAN TLS. The service provider offers a Service Level Agreement (SLA) that specifies the various attributes of the service including the amount of bandwidth provided. Typically, the bandwidth provided is less than the physical port speed. For example, an enterprise can subscribe for 50 Mb/s of bandwidth but the physical Ethernet connection to the ERS 8600 is a fiber Gigabit Ethernet port.

In these scenarios, Traffic Shaping applies to the ERS 8600 interfacing port to control the amount of traffic leaving the ERS 8600 to the service provider network. In doing so, the ERS 8600 determines which packets to drop when the subscription rate is exceeded based on the priority set in each packet. Without Traffic Shaping, the service provider edge device, which rate limits traffic entering its port, discards traffic regardless of the priority. This latter scenario is disastrous for enterprises that run voice, video, and other applications that require prioritization over Best Effort traffic.

The Route Switch Processor (RSP), a programmable network processor on the ERS 8600 R modules, provides traffic shaping for each Class of Service (CoS). This lab note provides a real-life application of this capability. The following sections describe configuration, operation, and design considerations.

## 2. Network overview

In the CSE lab, the following network was constructed to test shaping on the ERS 8600.

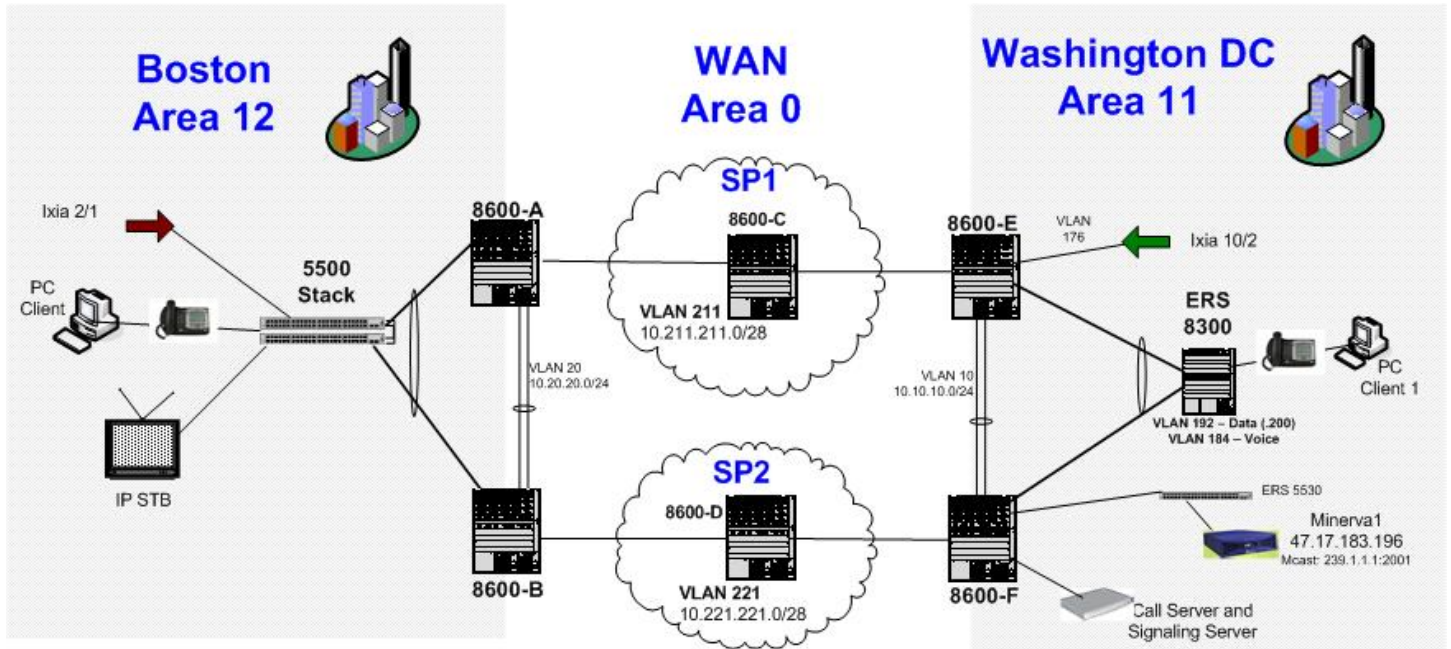


Figure 1 Network to test shaping on the ERS 8600

The network simulates two sites within an enterprise network—in this case, Boston and Washington. The two sites interconnect using two different service providers, referred to as SP1 and SP2. Each service provider provides 50 Mb/s of bandwidth to each site. Each service provider presents a Gigabit Ethernet (GbE) fiber to connect to the WAN.

The WAN uses open shortest path first (OSPF) Area 0, which touches all sites. Non-zero OSPF areas are configured at each enterprise site. At each site, a point-to-point virtual local area network (VLAN) exists between the two ERS 8600s in Area 0 for continuity.

All ERS 8600s ran release 4.1 software with R modules. Note that shaping has been supported since release 4.0.0 on all R modules.

Shaping was configured on the egress port of 8600-E in Washington for communication with SP1. Various types of traffic with different Classes of Service were generated so that all flows from Washington to Boston used 8600-E. The three types of traffic were as follows:

**Voice:** A phone call was placed between Boston and Washington. The media path for the voice traffic was marked with Expedited Forwarding (EF) DiffServ Code Point marking (DSCP) which corresponds to Premium Class of Service as defined in the Nortel Service Class architecture.



**Video:** A Minerva video encoder in Washington continuously streamed video to an IP Set Top Box (STB) in Boston. A filter policy, configured on the ERS 5530 in Washington, marked this traffic as Gold Class of Service. (In the lab network, using Gold for video was not an issue because only three streams of traffic represented each Class of Service. In a real deployment, consider using Platinum-level Class of Service for video.)

**Best Effort:** An Ixia traffic generator sent a 1-Gb/s stream of 512-byte packets from Washington to Boston with Best Effort Class of Service.

## 3. Configuring the ERS 8600

### 3.1 Planning

Before you configure the ERS 8600, it is important that you have a global Quality of Service (QoS) policy in place. This means that the enterprise customer understands the types of traffic flows in the network and defines policies (or filters) at the edge of the network that prioritize these various flows. As described above, in this lab configuration three types of traffic exist with Premium (VoIP), Gold (video), and Best Effort (Ixia traffic generator) classifications.

In addition to classifying flows into various Classes of Service, you must determine the amount of bandwidth allocated to each Class of Service. For example, in this scenario, the enterprise shapes traffic to the service provider at 50 Mb/s. How much bandwidth should be allocated for VoIP, video, and background data? This is discussed in further detail below; for now, assume 1 to 10 Mb/s for VoIP and video, and 20 to 50 Mb/s for Best Effort data. The following table indicates the bandwidth and QoS requirements for the three traffic types.

**Table 1 Bandwidth and QoS requirements for traffic types**

Traffic	Bandwidth requirements	QoS Level (CoS)
Voice	1–10 Mb/s	6
Gold - Video/Mcast	1–10 Mb/s	4
Standard/Best Effort	20–50 Mb/s	1

### 3.2 Applying an Egress Queue Set

On each R module port, an assigned Egress Queue Set defines the following:

- the amount of bandwidth that should be reserved for a CoS
- the rate at which traffic in a CoS can burst to
- the type of queue that each CoS uses (for example, strict, weighted queue, or low priority queue)
- the number of buffers allocated to the queue

By default, an Egress Queue Set applies to 10/100/1000BaseT ports on the 8648GTR module and a second Egress Queue Set applies to GbE and 10 GbE interfaces (see [1] in References for more information). However, you can create, customize, and apply additional queue sets to individual ports.

In this case, we create an Egress Queue Set to apply to the 8600-E GbE port that connects to SP1.

1. In the Device Manager, select Egress Queue Set from the QoS menu.

A dialog box appears listing the two default Egress Queue Sets and the ports that apply (if you run release 4.0.x software, the default queue sets are not shown).

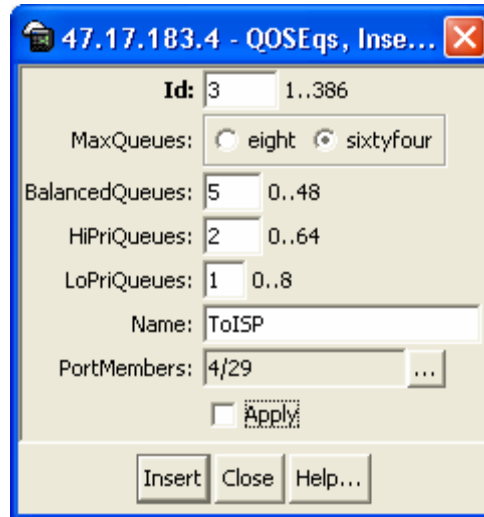


Figure 2 QOSEqs (QoS Egress Queues) dialog box

2. Click **Insert** to create a new Egress Queue Set:
3. Set the ID to 3 or higher as IDs 1 and 2 are already in use by the default Egress Queue Sets.
4. Because the queue set applies to a GbE port, select **sixtyfour** for **MaxQueues**.

A MaxQueues of eight applies only to 8648GTR ports.

Even though up to 64 queues can be defined, only 8 queues are defined in this case to simplify the configuration.

5. Set the number of **BalancedQueues** to 5, **HiPriQueues** to 2, and **LoPriQueues** to 1.

The High Priority queues are strictly dequeued, Balanced queues use Weighted Fair Queueing (WFQ), and the Low Priority queue is dequeued last.

6. Give the Egress Queue Set a name, and set the port to which the queue set should apply. You can specify the port later after the Queue Set queue configuration is complete.
7. When you finish, click **Insert**.

The new queue set appears in the QOSEqs dialog box, as shown in the following figure.

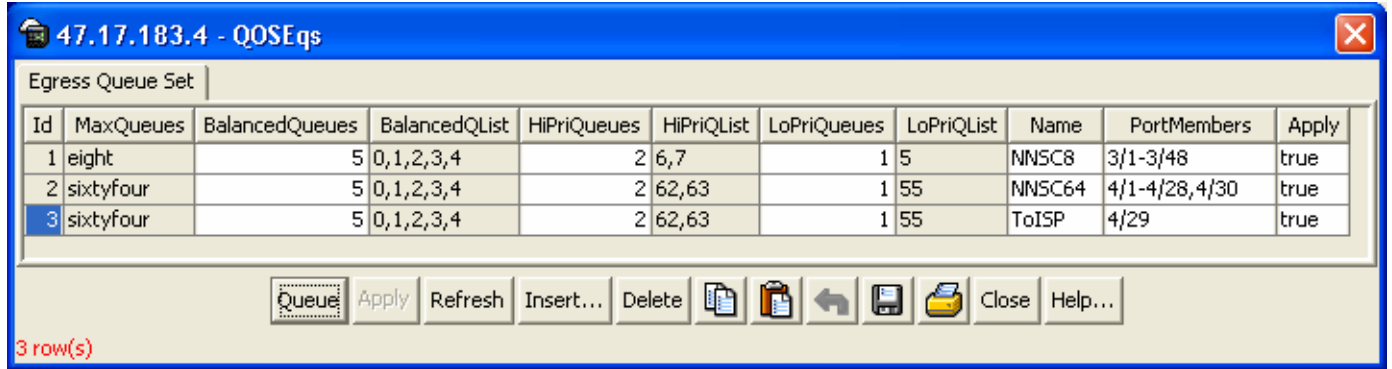


Figure 3 QOSEqs dialog box showing new queue set

8. After you create the queue set, select Queue Set Id 3, and click **Queue**.

The QoS Egress Queue Set Queues dialog box appears showing the default settings.

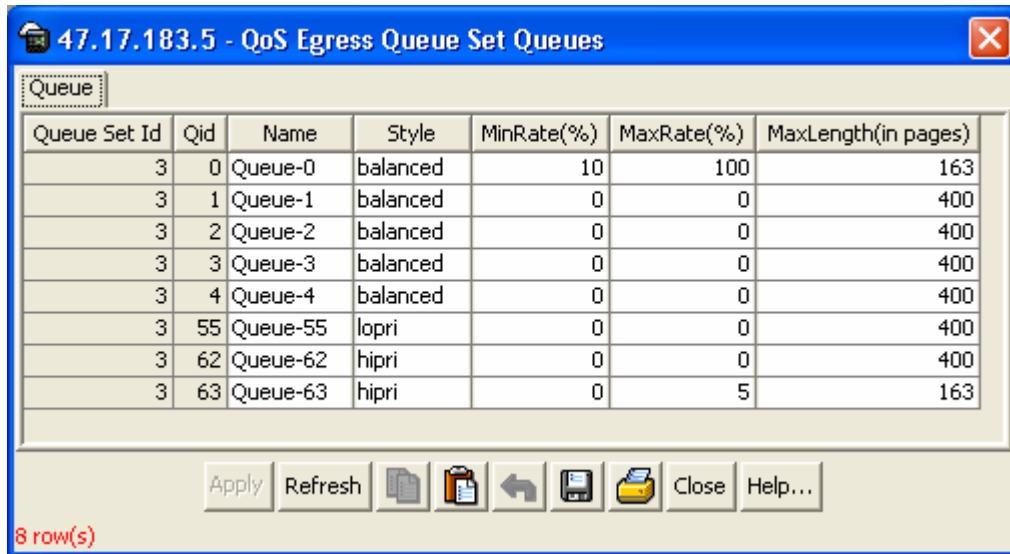


Figure 4 QoS Egress Queue Set Queues dialog box

With respect to queue numbering, the High Priority Queues numbers start at 63. By default, Queue 63 is for Network Control (Critical) traffic. Typically, queue 62 is for Premium traffic such as VoIP. The Balanced Queues start at 0. In this case, there are five queues, so they are numbered 0 to 4 with 0 being the highest Balanced queue. Finally, the Low Queue is number 55, by default, and represents the lowest Class of Service.

The Min Rate and Max Rate parameters represent the percentage of bandwidth guaranteed for each Class of Service as well as the Maximum Rate at which a given traffic class can burst to. Because this value is a

percentage, the actual bandwidth value depends on whether the port is a GbE or 10 GbE port. The smallest value, or granularity, is 1 percent.

For High Priority Queues, the Min Rate is always 0 and the Max Rate represents the guaranteed bandwidth. For the rest of the queues, the Min Rate represents the guaranteed bandwidth and the Max Rate represents the maximum bandwidth. For Balanced Queues, the ratio of the Min Rates determines the weight of each queue used by the WFQ algorithm. The sum of all guaranteed bandwidths must be less than or equal to 100 percent. Queues share bandwidth. For example, queue 63 defaults to 5 percent guaranteed bandwidth but in reality, it uses a small fraction of that, thus a low-priority queue uses the extra bandwidth.

The Max Length represents the number of buffers allocated to the Queue. In general, high priority queues require fewer buffers than low priority queues that tend to back up when congestion occurs. You can modify this value for each Queue.

As mentioned in the previous subsection, the enterprise must be aware of the various flows in the network, their Class of Service, and the amount of bandwidth required for each Class of Service. In this case, the Egress Queue set applies to a GbE interface so 1 percent of bandwidth represents a granularity of 10 Mb/s. For voice and video, 10 Mb/s is allocated to each. For all other Best Effort traffic 20 Mb/s is guaranteed with burst up to 50 Mb/s corresponding to the SLA. The following table summarizes the traffic classes, bandwidth requirements, the mapping to the individual Egress Queue Set queues.

**Table 2 Summary of traffic classes, bandwidth requirements and mappings**

Traffic	BW Requirements	Min BW %	Max BW %	QoS Level	Queue ID	Queuing Method
Control Traffic	50	NA	5 (default)	7	63–Critical/Network	Strict–Hi Priority
Voice	10	NA	1	6	62–Premium	Strict–Hi Priority
Platinum	0	0	0	5	0–Platinum	WFQ
Gold–Video/Mcast	10	1	2	4	1–Gold	WFQ
Silver	0	0	0	3	2–Silver	WFQ
Bronze	0	0	0	2	3–Bronze	WFQ
Standard/Best Effort	20	2	5	1	4–Standard (Default)	WFQ
Custom	0	0	0	0	55–Custom	Lo Priority

With this information, you can customize the Egress Queue Set as shown in the following figure. The values for MaxLength were copied from Queue Set 2, which was the guide (or template) for Queue Set 3.

Queue Set Id	Qid	Name	Style	MinRate(%)	MaxRate(%)	MaxLength(in pages)
3	0	Platinum	balanced	0	0	163
3	1	Gold-Video	balanced	1	2	163
3	2	Silver	balanced	0	0	327
3	3	Bronze	balanced	0	0	400
3	4	Standard-BE	balanced	2	5	980
3	55	Custom2	lopri	0	0	980
3	62	Premium-Voice	hipri	0	1	163
3	63	NetworkControl	hipri	0	5	163

Figure 5 QoS Egress Queue Set Queues dialog box

**IMPORTANT NOTE:** To apply changes to the Egress Queue set, first click Apply in the QoS Egress Queue Set Queues dialog box (as shown in the preceding figure), and then go to the previous QOSEqs dialog box and set the Apply Parameter (the last column in the table) to True and click Apply.

To prevent losing your changes, do not click Refresh in the QoS Egress Queue Set Queues dialog box.

- Verify that the appropriate port is a member of the Egress Queue Set. You can see this in the QOSEqs dialog box.

## 4. Operation

Before generating and testing any traffic streams, the connection between the ERS 8600 in Washington and SP2 is manually disabled. Thus, all traffic from Washington to Boston is forced from the 8600-E port connected to SP1 where the Egress Queue Set that was just configured is applied.

### 4.1 Operating initially without congestion

After you apply an Egress Queue Set to a port, you can monitor the individual Queue Statistics. Do this in the Device Manager by selecting a single port, selecting the Graph function, and then selecting the Egress Queue Set tab. Monitoring through the CLI provides a compact display.

Initially, a VoIP call was placed between Washington and Boston, and the Video stream started. As shown in the following example, the queue statistics from port 4/29 on the Washington router communicating with SP1 indicate that traffic exists in queue 62 (VoIP—Premium) and queue 1 (Gold—Video). Some control traffic transmits from

queue 63 (most likely OSPF and PIM Hellos) and some background traffic continuously pings between workstations in queue 4 (Standard—Best Effort).

```
8600-E:5# clear port stats
8600-E:5# show qos stats egress-queue-set port 4/29
```

```
=====
                        R-Module QOS Shapers Stats Table
=====
```

Port	Qid	Total pages (512 bytes per page)	Dropped pages (512 bytes per page)	Utilization %
4/29	0	0	0	0
4/29	1	9804	0	93
4/29	2	0	0	0
4/29	3	0	0	0
4/29	4	60	0	0
4/29	55	0	0	0
4/29	62	657	0	6
4/29	63	8	0	0

Next, the Ixia was configured to transmit approximately 20 Mb/s of traffic from Washington to Boston. The increase in traffic/use in queue 4 is shown in the example below:

```
8600-E:5# show qos stats egress-queue-set port 4/29
```

```
=====
                        R-Module QOS Shapers Stats Table
=====
```

Port	Qid	Total pages (512 bytes per page)	Dropped pages (512 bytes per page)	Utilization %
4/29	0	0	0	0
4/29	1	149496	0	33
4/29	2	0	0	0
4/29	3	0	0	0
4/29	4	289591	0	64
4/29	55	0	0	0
4/29	62	10035	0	2
4/29	63	111	0	0

## 4.2 Operating with congestion

Next, the Ixia was configured to transmit 1 Gb/s of traffic from Washington to Boston. This causes congestion on the shaped port communicating with SP1. The queue statistics show that packets are dropped (Dropped Pages).

8600-E:5# show qos stats egress-queue-set port 4/29

```

=====
                        R-Module QOS Shapers Stats Table
=====
Port      Qid      Total pages      Dropped pages      Utilization
          (512 bytes per page) (512 bytes per page) %
-----
4/29      0         0                 0                   0
4/29      1      1158344           0                   3
4/29      2         0                 0                   0
4/29      3         0                 0                   0
4/29      4     34405121         23771454            96
4/29     55         0                 0                   0
4/29     62      77762            0                   0
4/29     63         880              0                   0
=====
    
```

While this traffic stream ran, voice and video quality remained unchanged. The Ixia statistics dialog box (see the following figure) shows that from port 10/2 in Washington 234960 packets per second transmitted and only 12399 packets per second were received at Ixia port 2/1 in Boston. Because the stream consisted of 512-byte packets the approximate receive rate in Boston is:

$$512 \text{ bytes/packet} * 8 \text{ bits/byte} * 12399 \text{ packets/second} = 50.786 \text{ Mb/s}$$

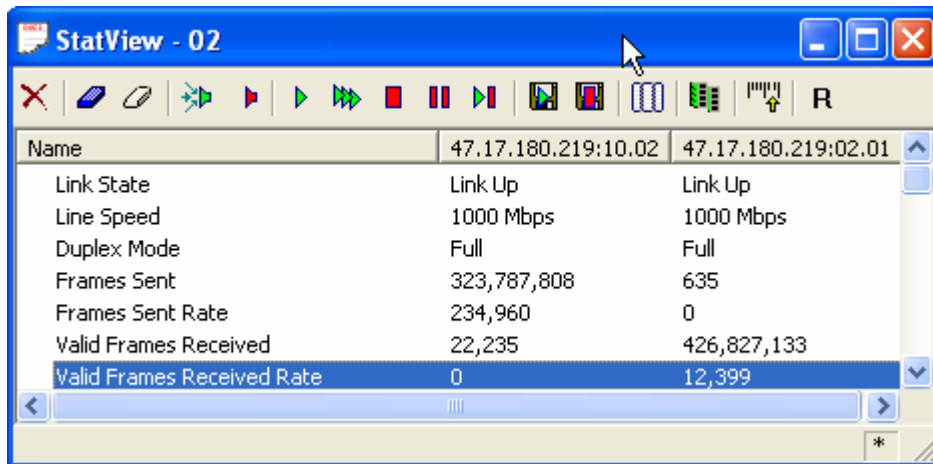


Figure 6 Ixia statistics dialog box

## 4.3 Sending traffic using QoS Level 5

To test the Egress Queue, the Ixia was configured to transmit a packet stream with QoS level 5. This corresponds to the Platinum CoS or Queue 0 configured with 0 percent Min and Max bandwidth. The queue statistics on port 4/29 on 8600-E show the following.

```
8600-E:5# show qos stats egress-queue-set port 4/29
```

```
=====
                          R-Module QOS Shapers Stats Table
=====
```

Port	Qid	Total pages (512 bytes per page)	Dropped pages (512 bytes per page)	Utilization %
4/29	0	2320544	2320544	99
4/29	1	4920	0	0
4/29	2	0	0	0
4/29	3	0	0	0
4/29	4	5	0	0
4/29	55	0	0	0
4/29	62	365	0	0
4/29	63	6	0	0

As expected all received packets dropped.

## 4.4 Using Default Queue Set with Port Level Shaping

The ERS 8600 supports port-level shaping, which you configure by locating the port and specifying the rate (kbps) at which to shape and enable the feature. In this case,

```
8600-E:5/config/ethernet/4/29# shape 50000 enable
```

This setting is independent of the Egress Queue Set configuration. Thus, when you configure the Min and Max bandwidth percentage as part of an egress queue, the value is a percentage of the physical port speed (for example, GbE) and not the port-level shaped rate.

To repeat the previous congestion test, the port-level shaping was enabled and the default queue set (ID 2) applied to the GbE port on 8600-E facing SP1.



Queue Set Id	Qid	Name	Style	MinRate(%)	MaxRate(%)	MaxLength(in pages)
2	0	Platinum	balanced	10	100	163
2	1	Gold	balanced	10	100	163
2	2	Silver	balanced	5	100	327
2	3	Bronze	balanced	15	100	327
2	4	Standard(Default)	balanced	5	100	980
2	55	Custom	lopri	0	100	980
2	62	Premium	hipri	0	50	163
2	63	Critical/Network	hipri	0	5	163

**Figure 7 QoS Egress Queue Set Queues dialog box**

When congesting the port with 1 Gb/s of Best Effort traffic generated by the Ixia, the video and voice streams remained unaffected because each is guaranteed bandwidth (50 Mb/s for Voice/Premium and 10 Mb/s for Video/Gold). The Ixia port in Boston received approximately 50 Mb/s, which means that the port-level shaping applied.

## 5. Design considerations

When the ERS 8600 interacts with an Ethernet-based WAN/MAN service, the enterprise must consider other issues:

**Egress Queue Set Bandwidth Granularity:** In this case, the bandwidth granularity is 10 Mb/s or 1 percent of a 1 GbE interface. Suppose the SLA to the service provider is 10 Mb/s; then, the bandwidth for the various Classes of Service cannot be reserved. Instead of a GbE interface between the ERS 8600 and the service provider, the enterprise can ask for a 100 Mb/s interface that they can use on an 8648GTR module. If this is not possible, then the GbE interface from the service provider can terminate on another device (for example, ES 470) which then presents a 100 Mb/s interface to the ERS 8600 R module, thereby providing bandwidth granularity that would be 1 Mb/s, or 1 percent of 100 Mb/s.

**Hub and Spoke design:** In this Lab Note, the connectivity existed between two sites. What if three or more sites exist? With Data Center consolidation, this can mean that one site receives and transmits more traffic than the other sites. Call this the Hub site. In this case, the enterprise can consider having a higher SLA at the Hub site than the other sites, as increased bandwidth is required. Even if congestion occurs, using the QoS mechanisms described in this paper ensures that high priority traffic always reaches its destination.

**QoS in service provider network:** The enterprise needs to make sure that the Service Provider does not modify the CoS markings (for example, 802.1p bits or DSCP) of the data packets. In addition, the enterprise needs to understand whether the service provider honors or ignores the CoS markings when traffic forwards through its network.

## Contact us

If you purchased a service contract for your Nortel product from a distributor or authorized reseller, contact the technical support staff for that distributor or reseller for assistance.

If you purchased a Nortel service program, contact Nortel Technical Support. To obtain contact information online, go to [www.nortel.com/contactus](http://www.nortel.com/contactus)

From the Technical Support page, you can open a Customer Service Request online or find the telephone number for the nearest Technical Solutions Center. If you are not connected to the Internet, call 1-800-4NORTEL (1-800-466-7835) to learn the telephone number for the nearest Technical Solutions Center.

An Express Routing Code (ERC) is available for many Nortel products and services. When you use an ERC, your call is routed to a technical support person who specializes in supporting that product or service. To locate an ERC for your product or service, go to [www.nortel.com/erc](http://www.nortel.com/erc)

To provide feedback or report a problem in this document, go to [www.nortel.com/documentfeedback](http://www.nortel.com/documentfeedback)