



Ethernet Routing Switch

45XX, 48XX, 55xx, 56xx, 59xx

Virtual Services Platform

7000

**Engineering**

> Ethernet Routing Switch Simplified  
QoS Configuration Using Traffic  
Profile Filter Sets Technical  
Configuration Guide

**Avaya Networking**

**Document Date: September 2015**

**Document Number: NN48500-624**

**Document Version: 4.0**

© 2015 Avaya Inc.  
All Rights Reserved.

#### Notices

While reasonable efforts have been made to ensure that the information in this document is complete and accurate at the time of printing, Avaya assumes no liability for any errors. Avaya reserves the right to make changes and corrections to the information in this document without the obligation to notify any person or organization of such changes.

#### Documentation disclaimer

Avaya shall not be responsible for any modifications, additions, or deletions to the original published version of this documentation unless such modifications, additions, or deletions were performed by Avaya. End User agree to indemnify and hold harmless Avaya, Avaya's agents, servants and employees against all claims, lawsuits, demands and judgments arising out of, or in connection with, subsequent modifications, additions or deletions to this documentation, to the extent made by End User.

#### Link disclaimer

Avaya is not responsible for the contents or reliability of any linked Web sites referenced within this site or documentation(s) provided by Avaya. Avaya is not responsible for the accuracy of any information, statement or content provided on these sites and does not necessarily endorse the products, services, or information described or offered within them. Avaya does not guarantee that these links will work all the time and has no control over the availability of the linked pages.

#### Warranty

Avaya provides a limited warranty on this product. Refer to your sales agreement to establish the terms of the limited warranty. In addition, Avaya's standard warranty language, as well as information regarding support for this product, while under warranty, is available to Avaya customers and other parties through the Avaya Support Web site: <http://www.avaya.com/support>

Please note that if you acquired the product from an authorized reseller, the warranty is provided to you by said reseller and not by Avaya.

#### Licenses

THE SOFTWARE LICENSE TERMS AVAILABLE ON THE AVAYA WEBSITE, [HTTP://SUPPORT.AVAYA.COM/LICENSEINFO/](http://support.avaya.com/licenseinfo/) ARE APPLICABLE TO ANYONE WHO DOWNLOADS, USES AND/OR INSTALLS AVAYA SOFTWARE, PURCHASED FROM AVAYA INC., ANY AVAYA AFFILIATE, OR AN AUTHORIZED AVAYA RESELLER (AS APPLICABLE) UNDER A COMMERCIAL AGREEMENT WITH AVAYA OR AN AUTHORIZED AVAYA RESELLER. UNLESS OTHERWISE AGREED TO BY AVAYA IN WRITING, AVAYA DOES NOT EXTEND THIS LICENSE IF THE SOFTWARE WAS OBTAINED FROM ANYONE OTHER THAN AVAYA, AN AVAYA AFFILIATE OR AN AVAYA AUTHORIZED RESELLER, AND AVAYA RESERVES THE RIGHT TO TAKE LEGAL ACTION AGAINST YOU AND ANYONE ELSE USING OR SELLING THE SOFTWARE WITHOUT A LICENSE. BY INSTALLING, DOWNLOADING OR USING THE SOFTWARE, OR AUTHORIZING OTHERS TO DO SO, YOU, ON BEHALF OF YOURSELF AND THE ENTITY FOR WHOM YOU ARE INSTALLING, DOWNLOADING OR USING THE SOFTWARE (HEREINAFTER REFERRED TO INTERCHANGEABLY AS "YOU" AND "END USER"), AGREE TO THESE TERMS AND CONDITIONS AND CREATE A BINDING CONTRACT BETWEEN YOU AND AVAYA INC. OR THE APPLICABLE AVAYA AFFILIATE ("AVAYA").

#### Copyright

Except where expressly stated otherwise, no use should be made of the Documentation(s) and Product(s) provided by Avaya. All content in this documentation(s) and the product(s) provided by Avaya including the selection, arrangement and design of the content is owned either by Avaya or its licensors and is protected by copyright and other intellectual property laws including the sui generis rights relating to the protection of databases. You may not modify, copy, reproduce, republish, upload, post, transmit or distribute in any way any content, in whole or in part, including any code and software. Unauthorized reproduction, transmission, dissemination, storage, and or use without the express written consent of Avaya can be a criminal, as well as a civil offense under the applicable law.

#### Third Party Components

Certain software programs or portions thereof included in the Product may contain software distributed under third party agreements ("Third Party Components"), which may contain terms that expand or limit rights to use certain portions of the Product ("Third Party Terms"). Information regarding distributed Linux OS source code (for those Products that have distributed the Linux OS source code), and identifying the copyright holders of the Third Party Components and the Third Party Terms that apply to them is available on the Avaya Support Web site: <http://support.avaya.com/Copyright>.

#### Trademarks

The trademarks, logos and service marks ("Marks") displayed in this site, the documentation(s) and product(s) provided by Avaya are the registered or unregistered Marks of Avaya, its affiliates, or other third parties. Users are not permitted to use such Marks without prior written consent from Avaya or such third party which may own the Mark. Nothing contained in this site, the documentation(s) and product(s) should be construed as granting, by implication, estoppel, or otherwise, any license or right in and to the Marks without the express written permission of Avaya or the applicable third party. Avaya is a registered trademark of Avaya Inc. All non-Avaya trademarks are the property of their respective owners.

#### Downloading documents

For the most current versions of documentation, see the Avaya Support. Web site: <http://www.avaya.com/support>

#### Contact Avaya Support

Avaya provides a telephone number for you to use to report problems or to ask questions about your product. The support telephone number is 1-800-242-2121 in the United States. For additional support telephone numbers, see the Avaya Web site: <http://www.avaya.com/support>.

## Acronym Key

Throughout this guide the following acronyms will be used:

ACL	Access Control List
AF	Assured Forwarding PHB
CLI	Command Line Interface
CS	Class Selector Code Point
DSCP	Differentiated Services Code Point
EF	Expedited Forwarding PHB
ERS	Ethernet Routing Switch
HTTP	Hypertext Transfer Protocol
ICMP	Internet Control Message Protocol
IETF	Internet Engineering Task Force
MIB	Management Information Base
NMS	Network Management Station
NSNA	Nortel Secure Network Access
PHB	Per Hop Behavior
PDU	Protocol Data Unit
QoS	Quality of Service

## Revision Control

No	Date	Version	Revised By	Remarks
1	May 2011	1.0	PRMGT	Update to Avaya format
2	December 2011	2.0	John Vant Erve	Update to include traffic meters and shaping
3	August 2014	3.0	John Vant Erve	Added VSP 7000 and Cisco IOS ACL equivalents
4	September 2015	4.0	John Vant Erve	Added ERS 5900 and updated QoS precedence levels supported on ERS 4800

# Table of Contents

Figures .....	7
Tables.....	8
1. Introduction .....	10
2. Background .....	10
2.1 Benefits .....	10
3. Filter Set Anatomy.....	11
3.1 QoS Primer.....	11
3.2 Filter Set Essentials .....	12
4. Traffic-Profile Filter Set Usage .....	16
4.1 Traffic-Profiles Filter Set Configuration .....	16
4.2 Displaying Traffic-Profile Filter Set Data .....	20
4.3 Traffic-Profile Filter Set Usage Example.....	21
4.4 Troubleshooting .....	28
4.4.1 Resource Exhaustion .....	28
4.4.2 Classifier Block Formation.....	28
4.4.3 Interface Capabilities.....	29
5. QoS Egress Queue Shaping.....	30
5.1 Configuring QoS Interface Shaper .....	30
5.2 Configuring QoS interface queue shaper.....	31
5.3 Configuration Examples using Meters and Shaper.....	32
5.3.1 ERS 4000 switch: UDP/TCP Port Range to remark traffic and dst-IP to Meter traffic .....	32
5.3.2 ERS 4000 switch: UDP/TCP Port Range using meter to remark traffic.....	33
5.3.3 Port Shaper .....	34
5.3.4 Per Queue Shaper .....	35
6. Appendix .....	37
6.1 Command Output: 'show qos traffic-profile classifier' .....	37
6.2 Legacy CLI Command Comparison.....	42
7. Default QoS Settings.....	43
7.1 Default Settings: ERS 4000 .....	43
7.2 Default Settings: ERS 5000 .....	47
7.3 Default Settings: ERS 5900 .....	52
7.4 Default Settings: VSP 7000.....	55
7.5 QoS Precedence Levels – Masks used by various applications .....	57
8. Cisco to Avaya Comparison.....	58

---

9. Reference Documentation .....	59
----------------------------------	----

Figures

Figure 1: QoS Policy Components..... 11

# Tables

Table 1: Filter Set Element Evaluation Order and Action Selection .....	13
Table 2: Filter Content and Block Compatibility Examples .....	13
Table 3: Traffic-Profile Classifier Command Classifier Options.....	17
Table 4: Traffic-Profile Classifier Command Action Operation .....	18
Table 5: Traffic-Profile Classifier Command Element Options .....	18
Table 6: Traffic-Profile Set Command Metering Options .....	19
Table 7: Traffic-Profile Set Command Action Options .....	19



## Conventions

This section describes the text, image, and command conventions used in this document.

### Symbols



Tip – Highlights a configuration or technical tip.



Note – Highlights important information to the reader.



Warning – Highlights important information about an action that may result in equipment damage, configuration or data loss.

### Text

**Bold** text indicates emphasis.

*Italic* text in a Courier New font indicates text the user must enter or select in a menu item, button or command:

```
ERS5520-48T# show running-config
```

Output examples from Avaya devices are displayed in a Lucida Console font:

```
ERS5520-48T# show sys-info
```

```

Operation Mode:      Switch
MAC Address:        00-12-83-93-B0-00
PoE Module FW:      6370.4
Reset Count:        83
Last Reset Type:     Management Factory Reset
Power Status:        Primary Power
Autotopology:        Enabled
Pluggable Port 45:   None
Pluggable Port 46:   None
Pluggable Port 47:   None
Pluggable Port 48:   None
Base Unit Selection: Non-base unit using rear-panel switch
sysDescr:            Ethernet Routing Switch 5520-48T-PWR
HW:02               FW:6.0.0.10   SW:v6.2.0.009
Mfg Date:12042004    HW Dev:H/W rev.02
    
```

# 1. Introduction

The defining characteristic of the Quality of Service (QoS) support on the Ethernet Routing Switch (ERS) stackable platforms and Virtual Services Platform 7000 has always been flexibility. Hardware capabilities have been fully exposed and exploited to give the user a plethora of knobs with which to manipulate the system's QoS capabilities. Over time, this has resulted in a sometimes dizzying array of options being made available. This, in turn, has steepened the learning curve, making it more time-consuming to initiate even the most basic of QoS operations.

While there will always be a demand for comprehensive and advanced configuration options, a streamlined and easy-to-learn QoS configuration mechanism has become a necessity. The support must provide a reasonable amount of filtering and action options while requiring no more than a few intuitive configuration steps to achieve the desired result. The Traffic-Profile Filter Set support has been introduced to satisfy this need.

# 2. Background

The process of defining filters, meters, actions and policies using the standard set of Command Line Interface (CLI) commands and advanced web pages is tedious at best. Developed nearly a decade ago in accordance with the emerging IETF standards and closely aligned to the QoS Management Information Base (MIB) data model, the command set is quite flexible but bulky and onerous to use. Access Control List (ACL) support was introduced to provide some process improvement in this area. Unfortunately, the ACL support by definition lacks the flexibility necessary for it to assume the role of a general-purpose configuration mechanism.

Filter set templates were introduced to support Nortel Secure Network Access (NSNA) and User Based Policy (UBP) applications. Based on the templates, filter set data (i.e., a named collection of filters, actions, meters and policies) is associated with a port under the direction of a non-QoS application (e.g., NSNA, UBP). Allowing the customer to leverage this functionality directly, in a fashion similar to ACL usage, could provide the flexibility and ease-of-use necessary of the configuration mechanism being sought.

## 2.1 Benefits

Traffic-profile filter set support offers several advantages over the standard QoS CLI support, as well as the deployed ACL functionality:

- Streamlined command set: filter set definition and installation can be completed in as few as 2 commands, replacing potentially 7 standard CLI QoS commands.
- Combined IP and L2 options: deployed ACL support forces the user to define IP or L2 ACLs. Filter set classifier options include both IP and L2 data.
- Meter availability: a filter set may be associated with metering criteria that is applied to all filter set policies (individual filters or blocks of filters). Metering is currently not supported with ACLs.
- No implicit drop: an ACL is terminated by an implicit 'drop all' prohibiting ACL layering on a port. This limitation is eliminated with filter sets.
- Addition/deletion support: filter set classifier elements (i.e., filters/actions) may be added/deleted while the filter set is in-use (i.e., associated with a port). This type of manipulation is not supported with ACLs.
- Additional filtering options: latest IP/L2 filters options are available in conjunction with filter sets.

## 3. Filter Set Anatomy

Before delving into the specifics of the traffic-profile filter set support, an understanding of the components that comprise a filter set is needed. As a starting point, a quick review of QoS essentials is warranted.

### 3.1 QoS Primer

The most basic QoS operation is that of identifying traffic and manipulating it in some way. To accomplish this task, classifiers, actions and policies must be defined. Classifiers (a.k.a. filters) are used to identify traffic by specifying which protocol fields need to be examined and what these fields must contain for a filter match to occur. A single classifier can target one or more fields to be matched. If multiple fields are targeted, all of the fields being examined must contain the specified data for the packet being inspected to match the filter.

Multiple classifiers may be combined into a block (i.e., a classifier block) so that they may be referenced as a unit. Only classifiers that are compatible (e.g., match the same or compatible protocol fields) may be combined into a block.

Actions specify the treatment packets matching a classifier are to receive. Typically this involves dropping the packets or remarking certain fields that are used to dictate traffic treatment both internally (as the packet egresses the switch) and downstream as the packet is processed by other switching equipment.

A policy acts as a coupling that ties together the various QoS components. A policy identifies a classifier (or classifier block) for traffic identification and an action to be initiated when packets matching the classification criteria are encountered. Ports on which to install the filter data are included in the policy specification. A policy also includes a precedence value which dictates the order in which policies (i.e., classifiers and actions) are applied when multiple policies target the same port. A limited number of precedence levels (a.k.a. masks) are available per port. This limit dictates the number of policies that can be applied to a given port at the same time.

A policy specification may also include metering criteria, to control ingress traffic, and statistics support to keep track of the traffic processed by the policy.

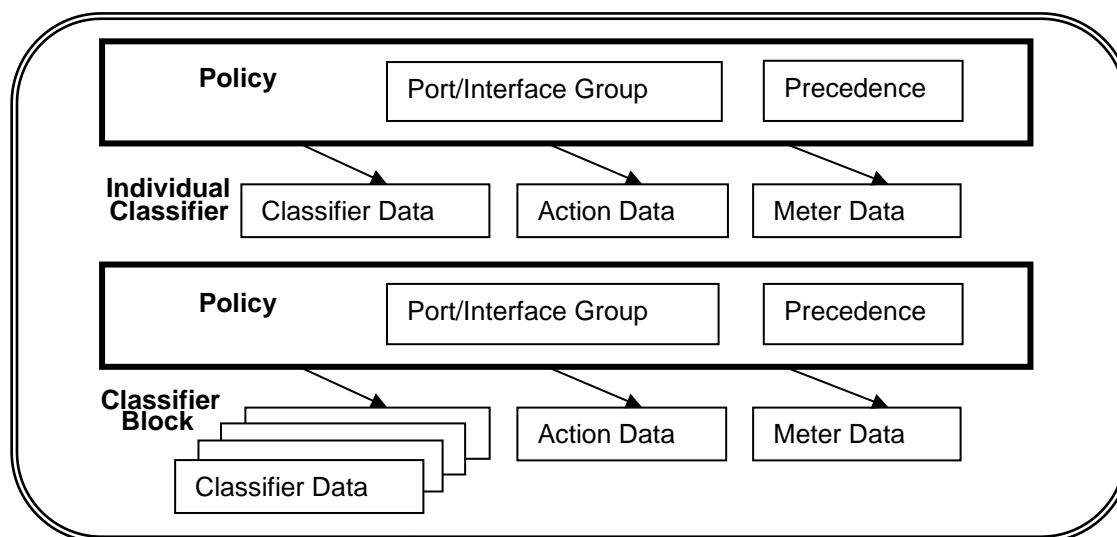


Figure 1: QoS Policy Components

## 3.2 Filter Set Essentials

At the most basic level, a filter set is a collection of policies that are identified as a single, named unit. Each policy references classifier and action criteria for identifying and processing traffic. The policies within a set are ordered, meaning that they are applied to ingress traffic in a specific order. This evaluation order is important because, once a policy “hit” occurs (i.e., traffic matching the policy’s filter criteria is identified) and actions are initiated, policies with a higher evaluation order may not be enforced even if they match the traffic being processed.

A filter set is associated with a port. The same filter set can be applied to multiple ports but can only be applied once to any given port. When a filter set is installed, limited resources are consumed (e.g., filter/meter/counter resources and precedence levels/masks). If the necessary resources are not available, the filter set installation will fail. Partial installation is not supported. Active (a.k.a. in-use) filter sets can be modified. Classifier elements can be added and deleted, pending resource availability. In-use filter set modification may impact traffic processing as the installed data is being updated.

Let us consider each of these items in more detail.

**Filter Set Classifier Elements:** a filter set classifier element identifies the protocol fields and their content that are to be used for traffic identification. Actions to be initiated for traffic matching the specified filter data are included as well. Elements are “named” such that all classifier elements that comprise a filter set share the same name (i.e., the filter set name). Elements have a specific type as well (e.g., NSNA, UBP, traffic-profile) though this is set automatically during configuration and, for the most part, is transparent to the user. Each element also contains an evaluation order value and filter block information, if appropriate.

**Filter Set Element Evaluation Order:** each filter set classifier element contains an evaluation order. This value dictates the order in which classifier elements associated with the same filter set (i.e., elements with the same filter set name) are applied. Elements with a low evaluation order are applied before elements with a higher evaluation order. The evaluation order can determine the actions initiated when a packet matches multiple classifier elements in a filter set (see Table 1).

An evaluation order must be unique within a filter set. The evaluation order for a classifier block is determined by the lowest evaluation order of the elements that are members of the block. Members of a classifier block may be assigned the same evaluation order, though this will impact the evaluation order of block member classifier data on platforms that honor classifier block member precedence.

Element Evaluation Order	Action Criteria	Action Initiated
1	Update DSCP: 20	
2	Drop Packet	Drop Packet
Reason: no action conflicts – both performed (drop overrides remarking)		
1	Update DSCP: 46	Update DSCP: 46
2	Update 802.1p: 6	Update 802.1p: 6
Reason: no action conflicts – both performed (complimentary actions)		
1	Don't Drop Packet	Don't Drop Packet
2	Drop Packet	

Reason: conflicting actions – resolve conflict using evaluation order		
1	Update DSCP: 46	Update DSCP: 46
2	Update DSCP: 0	
Reason: conflicting actions – resolve conflict using evaluation order		

**Table 1: Filter Set Element Evaluation Order and Action Selection**

**Filter Set Element Block Definition:** filter set classifier elements can be combined into a block to make better use of scarce resources. When installed, a single filter set classifier element consumes one precedence level (i.e., resource usage equivalent to a policy referencing a single filter). A block containing any number of filter set classifier elements still only consumes one precedence level (i.e., resource usage equivalent to a policy referencing a filter block) when installed. Combining compatible filter set classifier elements into blocks can substantially impact resource usage.

Which filter set elements can be combined into a block? As a general rule, classifiers matching the exact same protocol fields or the exact same portions thereof (e.g., same IP address bits) are compatible. On the ERS stackable platforms (i.e., 4XXX/56XX/59XX) and VSP7000 these restrictions have been eased quite a bit. Classifiers matching the same protocol fields, regardless of the whether full or partial field matches are specified, are compatible. Furthermore, filters matching different but compatible fields (i.e., fields from a hardware-defined set of compatible fields) can be combined into a block. Table 2 contains a few examples of classifier content that is compatible and non-compatible for blocking purposes.

Filter Content		Blocking Compatible	
Filter A	Filter B	ERS55XX	ERS56XX/ERS4XXX/ERS59XX VSP7000
Source IP 1.2.3.4/32	Source IP 5.6.7.8/32	Yes	Yes
DSCP 23, Protocol TCP	DSCP 38, Protocol UDP	Yes	Yes
Src MAC 11:22:33:FF:FF:FF	Src MAC 44:55:66:FF:FF:FF	Yes	Yes
Dst L4 Port Min/Max 0/31	Dst L4 Port Min/Max 0/31	Yes	Yes
Destination IP 1.2.3.4/32	Destination IP 5.6.0.0/16	No	Yes
DSCP	TCP_Ctrl	No	Yes
Src L4 Port Min/Max 123/161	Src L4 Port Min/Max 20/30	No	Yes
Dst MAC 08:81:FF:FF	Src MAC 01:00:81:FF:FF:FF	No	Yes
Destination MAC	Source IP	No	No

**Table 2: Filter Content and Block Compatibility Examples**

**Filter Set Resource Usage:** as previously noted, a filter set is essentially a named collection of policies that reference classifier and action criteria for identifying and processing traffic. Policies, filters and related QoS components (e.g., meters, counters) consume hardware resources when applied to a port. A filter set therefore requires hardware resources equivalent to the sum of the resources required to support the individual policies, filters, etc. that comprise the filter set.

When installed, a single filter set classifier element or a classifier block consumes one precedence mask (i.e., resource usage equivalent to a policy referencing a single filter or a filter block). Thus the number of precedence levels/masks required for filter set installation equals the number of classifiers and unique classifier blocks included in the filter set. To ensure deterministic behavior, the required number of precedence levels must be available and consecutive for a filter set installation to be successful. Consider the following filter set:

```
qos traffic-profile classifier name fivePrecs dst-ip 1.2.3.4/32 drop-action enable
eval-order 1

qos traffic-profile classifier name fivePrecs dst-ip 1.2.3.5/32 drop-action enable
eval-order 2

qos traffic-profile classifier name fivePrecs dst-ip 1.2.3.6/32 drop-action enable
eval-order 3

qos traffic-profile classifier name fivePrecs ds-field 23 drop-action enable block
dscpBlk eval-order 10

qos traffic-profile classifier name fivePrecs ds-field 24 drop-action enable block
dscpBlk eval-order 11

qos traffic-profile classifier name fivePrecs protocol 111 drop-action enable block
protoBlk eval-order 20

qos traffic-profile classifier name fivePrecs protocol 112 drop-action enable block
protoBlk eval-order 21

qos traffic-profile set port 1-2 name fivePrecs track-statistics aggregate
```

Filter set “fivePrecs” contains 3 classifiers and 2 unique classifier blocks requiring that 5 consecutive precedence levels be available on ports to which the filter set is applied. There is no requirement regarding which precedence levels are available. The highest suitable starting precedence will be selected automatically when the filter set is installed. For example, under a default configuration scenario, “fivePrecs” would consume precedence masks 9 – 13 if installed on port 1 and 2 of a 5698:

5650TD(config)#*show qos diag*

Unit/Port	Mask Precedence Usage															
	16	15	14	13	12	11	10	9	8	7	6	5	4	3	2	1
1/1																
1/2																

AR=ARP DH=DHCP Q=QoS

On an ERS 4500 switch, we could not apply the above traffic profile as by default, only precedence levels 3 – 6 are available. However, we could disable DHCP-Relay freeing up precedence level 7 or we could ether assign an interface group of trusted or unrestricted to free up both precedence levels 1 and 2. For example, assuming we disable DHPC-Relay, “fivePrecs” would consume precedence masks 3 - 7 if installed on port 1 and 2 of a 4526:

```
no ip dhcp-relay

qos traffic-profile classifier name fivePrecs dst-ip 1.2.3.4/32 drop-action enable
eval-order 1
```

```

qos traffic-profile classifier name fivePrecs dst-ip 1.2.3.5/32 drop-action enable
eval-order 2

qos traffic-profile classifier name fivePrecs dst-ip 1.2.3.6/32 drop-action enable
eval-order 3

qos traffic-profile classifier name fivePrecs ds-field 23 drop-action enable block
dscpBlk eval-order 10

qos traffic-profile classifier name fivePrecs ds-field 24 drop-action enable block
dscpBlk eval-order 11

qos traffic-profile classifier name fivePrecs protocol 111 drop-action enable block
protoBlk eval-order 20

qos traffic-profile classifier name fivePrecs protocol 112 drop-action enable block
protoBlk eval-order 21

qos traffic-profile set port 1-2 name fivePrecs track-statistics aggregate
4526GTX-60pwr(config)#show qos diag

```

Unit/Port		Mask	Precedence	Usage															
		8	7	6	5	4	3	2	1										
1/1	AR	Q	Q	Q	Q	Q	Q	Q	Q										
1/2	AR	Q	Q	Q	Q	Q	Q	Q	Q										

Please note that with the ERS 4800 using release 5.8 or higher increases the number of QoS filter precedences from 8 to 16. Starting at release 5.8, support for the ERS4500 series is no longer supported.

If QoS or non-QoS applications are consuming resources on the target port, a different starting precedence will be selected if one is available. In the following example, precedence 10 is consumed by the IPSG application and 5 consecutive precedence levels are only available starting at precedence 9:

```

5650TD(config)#show qos diag
Unit/Port      Mask Precedence Usage
              16  15  14  13  12  11  10  9  8  7  6  5  4  3  2  1
-----
1/1            AR  DH                IS  Q  Q  Q  Q  Q                Q  Q
1/2            AR  DH

```

AR=ARP DH=DHCP Q=QoS IS=IPSG

Aside from precedence masks, filter set installation consumes filter (one per classifier element) and potentially meter and counter resources. The availability of these resources, tracked on a per-port (ERS55XX) or per-port cluster (ERS4XXX/ERS56XX/59XX/ERS59XX) basis, may limit whether a filter set can be applied to a port. It is also important to remember that adding classifier elements to an in-use filter set impacts resource requirements by changing the precedence mask requirements and/or the filter, meter, etc. resource requirements.

**Filter Set Port Assignment:** filter sets are associated with a port in different ways. Non-QoS applications, such as NSNA and UBP, use filter sets as a template. The filter set, when defined, is not associated with port but an instance of the filter set is created and installed based on application operation. Traffic-profile filter sets are associated with a port when defined by the user. A valid, QoS-enabled port must be provided by the administrator for the filter set installation to succeed.

**Filter Set Metering:** the ability to meter traffic processed by a filter set is currently limited. Metering data may be defined for the filter set overall but not for the individual classifiers. When metering data is specified, a meter is associated with each classifier and unique classifier block. Each meter has the same characteristics (e.g., committed rate, burst size) and is independently evaluated (i.e., there is no relationship between meters associated with different filter set classifiers/blocks). Conforming (i.e., in-profile) traffic is processed based on the action criteria associated with the individual classifier or block



member with the lowest evaluation order. Non-conforming (i.e., out-of-profile traffic) traffic is processed based on the out-of-profile action criteria associated with the overall filter set.

**Filter Set Statistics Tracking:** by default, a counter is associated with each filter set classifier element and unique classifier block. For a classifier element, the counter tracks matches of the specified classification criteria. For classifier blocks, the counter tracks the total filter data hits across all the classifiers in the block. Different statistics tracking modes are available. The default mode is referred to as 'aggregate' mode because of the way counts are totaled, or aggregated, across all the filters referenced by a policy. Should the user wish to track statistics for each classifier block member, the 'individual' mode should be specified. Note that counter resource availability determines whether either 'aggregate' or 'individual' statistics tracking modes can be used. If statistics tracking is not necessary or a filter set can only be installed with statistics tracking disabled (i.e., counter resources are exhausted), the statistics tracking mode 'disabled' can be specified.

## 4. Traffic-Profile Filter Set Usage

With the filter set essentials covered, it is time to specifically examine traffic-profile filter set usage. Configuration and display command syntax will be presented first. Next a series of usage examples will be presented, followed by a discussion of the areas in which configuration issues more commonly arise.

### 4.1 Traffic-Profiles Filter Set Configuration

Traffic-profile filter set configuration begins with specifying the classifier criteria used to identify the traffic being targeted and with specifying the actions to be taken when matching traffic is identified. The 'qos traffic-profile classifier' CLI command is used for this purpose:

- Command: **qos traffic-profile classifier name [filterSetName] {classifierData} {actionData} {elemData}**
  - name [filterSetName] – mandatory command variable that associates the element with a label. Name labels are used to identify elements belonging to the same traffic-profile filter set (i.e., members of a given filter set share the same name label). A 16 character (maximum) name string is supported.
  - classifierData – optional command variables used for traffic identification purposes.
  - actionData – optional command variables used to specify actions to be performed when traffic matching the classifier data is identified.
  - elemData – optional command variables that facilitate filter set element ordering and efficient resource utilization.

A large number of classifier options are available to facilitate traffic identification. Zero (i.e., a fully defaulted element), one or more classifier options can be specified in a single filter set classifier element. Certain combinations of classifier options may be prohibited based on the ERS stackable platform and the underlying QoS hardware.



Classifier Option	Description
addr-type <addrtype>	IP address type (IPv4/IPv6) classifier criteria
src-ip <src-ip-info>	Source IP address (or portion thereof) classifier criteria
dst-ip <dst-ip-info>	Destination IP address (or portion thereof) classifier criteria
ds-field <dscp>	DiffServ code point classifier criteria
protocol <protocoltype>	IPv4 protocol classifier criteria
next-header <headertype>	IPv6 next-header classifier criteria
src-port-min <port> src-port-max <port>	Minimum and maximum Layer 4 source port classifier criteria. Both values must be specified.
dst-port-min <port> dst-port-max <port>	Minimum and maximum Layer 4 destination port classifier criteria. Both values must be specified
flow-id <flowid>	IPv6 flow identifier classifier criteria
ip-flag <ip-flags>	IPv4 fragment flag classifier criteria
ipv4-options <no-opt   with-opt>	IPv4 option status classifier criteria
tcp-control <tcp-flags>	TCP control flag classifier criteria – <Urg   Ack   Psh   Rst   Syn   Fin>
src-mac <src-mac>	Source MAC address (or portion thereof) classifier criteria
src-mac-mask <src-mac-mask>	Source MAC address mask
dst-mac <dst-mac>	Destination MAC address (or portion thereof) classifier criteria
dst-mac-mask <dst-mac-mask>	Destination MAC address mask
pkt-type <etherII   llc   snap>	Layer 2 packet format classifier criteria
vlan-min <vid-min> vlan-max <vid-max>	Minimum and maximum (outer) VLAN ID classifier criteria. Both values must be specified
ivlan-min <vid-min> ivlan-max <vid-max>	Minimum and maximum inner VLAN ID classifier criteria. Both values must be specified
vlan-tag <vtag>	VLAN tag status classifier criteria
ethertype <etype>	EtherType value classifier criteria
priority <ieee1p-seq>	802.1p user priority value classifier criteria

**Table 3: Traffic-Profile Classifier Command Classifier Options**

Zero (i.e., leave traffic unchanged), one or more action options can be specified in a single filter set classifier element. Certain options may not be allowed if they are inappropriate for the traffic being targeted. For example, updating the DSCP value in a non-IP packet is not allowed for obvious reasons.

Action Option	Description
drop-action <drop   pass>	Drop or pass (do not drop) targeted traffic
update-dscp <0-63>	Update DSCP value in IP packets in targeted traffic
update-1p <0-7>	Update 802.1p user priority value in targeted traffic
set-drop-prec <high-drop   low-drop>	Set the drop precedence value associated with targeted traffic. Potentially impacts egress queuing

**Table 4: Traffic-Profile Classifier Command Action Operation**

Optional command variables that facilitate filter set element ordering and efficient resource utilization may be specified as well. The evaluation order value dictates the order in which filter set elements are applied, potentially impacting action selection. A unique evaluation order value is required when more than one element is included in the filter set. Multiple filter set elements may be combined into a block by labeling block members with the same block name. As previously noted, filter set elements must contain compatible classifier criteria (determined by the ERS stackable platform and underlying QoS hardware) to be members of the same block. A 16 character (maximum) block name string is supported.

Element Option	Description
block <block-name>	Label used to identify block members in the context of the filter set
eval-order <1-255>	Element evaluation order in the context of the filter set

**Table 5: Traffic-Profile Classifier Command Element Options**

Traffic-profile filter set classifier criteria can be deleted. The 'no qos traffic-profile classifier' CLI command is used for this purpose:

- Command: **no qos traffic-profile classifier name [filterSetName] {elemData}**
  - name [filterSetName] – mandatory command variable that identifies the filter set being targeted by the deletion operation.
  - elemData – optional command variables that facilitate filter set element identification.

All of the classifier elements associated with a filter set may be deleted at once or individual filter set classifier elements may be deleted. All filter set classifier elements may be deleted by providing the target filter set name with no additional options. If the filter set is in-active (i.e., not currently applied to any ports in the system), all related classifier elements will be deleted. The operation is rejected if the filter set is currently active. The optional command variable 'eval-order' (Table 5) can be used to target an individual filter set classifier element for deletion. Note that empty filter sets are not supported meaning that the last classifier element in an active filter set cannot be deleted.

Once the classifier data has been defined, the traffic-profile filter set needs to be instantiated and applied to a port. Metering options and additional action options, applied to the filter set as a whole, can be specified as well. The 'qos traffic-profile set' CLI command is used for this purpose:

- Command: **qos traffic-profile set port [port] name [filterSetName] {meterData} {actionData}**

- port [port] – mandatory command variable the identifies the port to which the filter set will be applied. A valid, QoS-enabled port must be provided.
- name [filterSetName] – mandatory command variable that associates the filter set with a label. Name labels are used to identify classifier elements belonging to the same traffic-profile filter set (i.e., members of a given filter set share the same name label). A 16 character (maximum) name string is supported.
- meterData – optional command variables used to specify metering criteria.
- actionData – optional command variables used to specify additional actions to be performed.

As previously discussed, the ability to meter traffic processed by a traffic-profile filter set is currently limited. Metering data may be defined for the filter set overall using the options presented in Table 6. When metering data is specified, a meter is associated with each classifier and unique classifier block. Each meter has the same characteristics (determined by the command options) and is independently evaluated (i.e., there is no relationship between meters associated with different filter set classifiers/blocks). Conforming (i.e., in-profile) traffic is processed based on the action criteria associated with the individual classifier or block member with the lowest evaluation order. Non-conforming (i.e., out-of-profile traffic) traffic is processed according to the out-of-profile action options.

Meter Option	Description
committed-rate <64-10230000> Kbps	Committed rate metering criteria
committed-burst-size <burst-size-options>	Committed burst size metering criteria
max-burst-rate <64-4294967295> Kbps	Maximum burst rate metering criteria
max-burst-duration <1-4294967295> ms	Maximum burst duration metering criteria
drop-out-action <enable   disable>	Drop (enable) or pass (disable) out-of-profile packets
update-dscp-out-action <0-63>	Update DSCP value in out-of-profile IP packets

**Table 6: Traffic-Profile Set Command Metering Options**

Optional command variables can be used to initiate actions that are applied to the overall filter set. The ability to specify the statistics tracking mode is one such action. Statistics tracking is enabled by default for all traffic-profile filter sets using the 'aggregate' mode. Statistics may be tracked at a finer level of granularity (i.e., per classifier) using the 'individual' statistics tracking mode. Statistics tracking can be bypassed and no counter resources consumed using the 'disable' mode.

Action Option	Description
track-statistics <disable   aggregate   individual>	Statistics tracking mode (default mode: aggregate)

**Table 7: Traffic-Profile Set Command Action Options**

Active traffic-profile filter sets can be deleted based on name and/or the associated port. The 'no qos traffic-profile set' CLI command is used for this purpose:

- Command: **no qos traffic-profile set [name [filterSetName] | port [port] name [filterSetName]]**

- name [filterSetName] – optional command variable that identifies the filter set being targeted for the deletion operation.
- port [port] - optional command variable that identifies the port being targeted for the filter set deletion operation. A filter set name must also be provided if a port value is specified.

Filter set deletion requires a filter set name be provided to identify the target filter set. A port value may be provided as well. If a port value is specified, the instance of the filter set associated with the port will be deleted. If only a filter set name is provided (i.e., no port value is specified), all instances of the filter set will be deleted regardless of port.

An active traffic-profile filter set may be modified. New classifier elements may be defined, pending determination of resource availability, and existing elements may be deleted. The aforementioned traffic-profile classifier commands (i.e., '**[no] qos traffic-profile classifier [...]**') are used for this purpose. No other commands are required.

## 4.2 Displaying Traffic-Profile Filter Set Data

Traffic-profile filter set data can be displayed in a number of different ways. Traffic-profile filter set classifier elements can be displayed based on a filter set name or all defined classifier elements can be displayed:

- Command: **show qos traffic-profile classifier {name [filterSetName]}**
  - name [filterSetName] – optional command variable that identifies the filter set being targeted for the display operation. Only classifier elements associated with the identified filter set will be displayed.

Active (i.e., in-use) filter set data can be displayed based on filter set name (i.e., display all instances of a filter set) or based on port and name data (i.e., display all instances of a filter set associated with the specified port):

- Command: **show qos traffic-profile set name [filterSetName]**
  - name [filterSetName] – mandatory command variable that identifies the filter set being targeted for the display operation. All instances of the filter set will be displayed.
- Command: **show qos traffic-profile set port [port] name [filterSetName]**
  - port [port] – mandatory command variable the identifies the port being targeted for the display operation.
  - name [filterSetName] – mandatory command variable that identifies the filter set being targeted for the display operation. All instances of the filter set associated with the specified port will be displayed.

An abbreviated display, presenting just the traffic-profile filter set to port associations, is available as well:

- Command: **show qos traffic-profile interface**

Filter set statistics can be displayed on a per-port or on a per-port and precedence basis. The port-based display presents in-profile and out-of-profile (if metering is requested) counts for all precedence values consumed by the filter set instance. Providing a specific precedence target causes the individual counts associated with classifier elements installed with the identified precedence to be displayed. This can be used to display counts for classifier block members.

- Command: **show qos traffic-profile statistics port [port] name [filterSetName] {precedence [prec]}**
  - port [port] – mandatory command variable the identifies the port being targeted for the display operation.

- name [filterSetName] – mandatory command variable that identifies the filter set being targeted for the display operation.
- precedence [prec] – optional command variable used to identify a specific precedence target for the display operation.

To assist with filter set installation and related potential issues, filter set diagnostic information can be displayed. Precedence mask, filter, meter and counter requirements can be displayed in-total and on a per-precedence basis. Interface capabilities that are required for successfully filter set installation are presented as well:

- Command: **show qos traffic-profile diags name [filterSetName]**
  - name [filterSetName] – mandatory command variable that identifies the filter set being targeted for the display operation.

## 4.3 Traffic-Profile Filter Set Usage Example

To help fill in any blanks regarding the traffic-profile commands that have been covered, a usage example will now be presented. Consider a scenario where there is a need to monitor, and perhaps prioritize, management traffic on the network. As a first step, management traffic needs to be identified (some command abbreviations used for brevity):

```
qos traffic-profile classifier name mgmtTraf protocol 17 dst-port-min 161 dst-port-max 161 drop-action disable eval-order 10
qos traffic-profile classifier name mgmtTraf protocol 6 dst-port-min 80 dst-port-max 80 drop-action disable eval-order 20
qos traffic-profile classifier name mgmtTraf protocol 6 dst-port-min 23 dst-port-max 23 drop-action disable eval-order 21
qos traffic-profile classifier name mgmtTraf src-ip 10.30.31.15/32 drop-action disable eval-order 30
qos traffic-profile classifier name mgmtTraf src-ip 10.30.31.47/32 drop-action disable eval-order 31
qos traffic-profile classifier name mgmtTraf src-ip 10.30.31.48/32 drop-action disable eval-order 32
qos traffic-profile classifier name mgmtTraf protocol 1 drop-action disable eval-order 40
```

Traffic-profile filter set “mgmtTraf” targets SNMP and HTTP PDUs, telnet traffic, ping (ICMP Echo) traffic and all communication from specific Network Management Systems (NMSs). Prior to filter set installation, switch resource usage is as follows:

```
5650TD(config)#show qos diag
```

Unit/Port	Mask Precedence Usage																		
	16	15	14	13	12	11	10	9	8	7	6	5	4	3	2	1			
1/1	AR	DH											IS					Q	Q
1/2	AR	DH																Q	Q

AR=ARP DH=DHCP Q=QoS IS=IPSG

Unit/Port	Prec	Filter	Meter	Cntr	Non QoS		Non QoS		Filter	Meter	Cntr	RngChk
					Filter	Meter	Filter	Meter				
		Used	Used	Used	Used	Used	Used	Total	Total	Total	Used	

1	/1	-24,49	16	0	0	0	25	25	256	128	128
			15	0	0	0	50	25	256	128	128
			14	0	0	0	0	0	256	128	128
			13	0	0	0	0	0	256	128	128
			12	0	0	0	0	0	256	128	128
			11	0	0	0	0	0	256	128	128
			10	0	0	0	0	0	256	128	128
			9	0	0	0	0	0	256	128	128
			8	0	0	0	0	0	256	128	128
			7	0	0	0	0	0	256	128	128
			6	0	0	0	1	0	256	128	128

For this example, switch port 1 has been designated by the system administrator as the preferred network management traffic port:

```
qos traffic-profile set port 1 name mgmtTraf
```

After filter set installation, switch resource usage reflects the resources consumed by “mgmtTraf” (7 precedence levels/masks and 7 filters/counters):

```
5650TD(config)#show qos diag
```

Unit/Port	Mask Precedence Usage															
	16	15	14	13	12	11	10	9	8	7	6	5	4	3	2	1
-----	-----															
1/1	AR	DH	Q	Q	Q	Q	Q	Q	Q		IS				Q	Q
1/2	AR	DH													Q	Q

AR=ARP DH=DHCP Q=QoS IS=IPSG

Unit/Port	Prec	Non QoS			Non QoS			Filter Total	Meter Total	Cntr Total	RngChk Used
		Filter Used	Meter Used	Cntr Used	Filter Used	Meter Used	Cntr Used				
1 /1 -24,49	16	0	0	0	25	25		256	128	128	
	15	0	0	0	50	25		256	128	128	
	14	1	0	1	0	0		256	128	128	
	13	1	0	1	0	0		256	128	128	
	12	1	0	1	0	0		256	128	128	
	11	1	0	1	0	0		256	128	128	
	10	1	0	1	0	0		256	128	128	
	9	1	0	1	0	0		256	128	128	
	8	1	0	1	0	0		256	128	128	
	7	0	0	0	0	0		256	128	128	
	6	0	0	0	1	0		256	128	128	

Suppose the administrator decides that port 1 is to be used exclusively for network management traffic. All non-management traffic should be dropped. This is easily accomplished by adding a catch-all filter at the end of the filter set to match and drop all non-network management traffic. Note that the evaluation order dictates where in the filter set the new classifier element will be placed (i.e., last to be evaluated). All filter set classifier elements with an evaluation order lower than the “drop-all” filter need to ensure that the identified network management traffic is ‘passed’ (action: drop disable) so that the new classifier doesn’t accidentally impact the wrong flows:

```
qos traffic-profile classifier name mgmtTraf drop-action enable eval-order 100
```

A new NMS has been purchased and it requires access to the network:

```
qos traffic-profile classifier name mgmtTraf src-ip 10.30.31.49/32 drop-action
disable eval-order 33
```

**// ERROR// % Traffic Prof filter set elem count exceeds available resources**

An issue was encountered while trying to add an element to the filter set. The error message indicates a resource availability issue. A check of the system resource usage levels exposes the issue:

```
5650TD(config)#show qos diag
```

Unit/Port	Mask Precedence Usage															
	16	15	14	13	12	11	10	9	8	7	6	5	4	3	2	1
1/1	AR	DH	Q	Q	Q	Q	Q	Q	Q	Q	IS				Q	Q
1/2	AR	DH													Q	Q

AR=ARP DH=DHCP Q=QoS IS=IPSG

Unit/Port	Prec	Filter	Meter	Cntr	Non QoS		Filter	Meter	Filter	Meter	Cntr	RngChk
					Used	Used						
1 /1 -24,49	16	0	0	0	25	25	256	128	128			
	15	0	0	0	50	25	256	128	128			
	14	1	0	1	0	0	256	128	128			
	13	1	0	1	0	0	256	128	128			
	12	1	0	1	0	0	256	128	128			
	11	1	0	1	0	0	256	128	128			
	10	1	0	1	0	0	256	128	128			
	9	1	0	1	0	0	256	128	128			
	8	1	0	1	0	0	256	128	128			
	7	1	0	1	0	0	256	128	128			
	6	0	0	0	1	0	256	128	128			

Adequate precedence mask resources are not available to support the filter set expansion. A maximum of 8 consecutive precedence levels can currently be consumed by the filter set on port 1. Traffic-profile filter set “mgmtTraf” already consumes 8 precedence levels. Filter set optimization is necessary. Several of the classifier elements can be combined into blocks to improve resource usage. First the elements being replaced are deleted. Next the elements are redefined with a ‘block’ component:

```
no qos traffic-profile classifier name mgmtTraf eval-order 20
no qos traffic-profile classifier name mgmtTraf eval-order 21

qos traffic-profile classifier name mgmtTraf protocol 6 dst-port-min 80 dst-port-
max 80 drop-action disable block tcp eval-order 20

qos traffic-profile classifier name mgmtTraf protocol 6 dst-port-min 23 dst-port-
max 23 drop-action disable block tcp eval-order 21

no qos traffic-profile classifier name mgmtTraf eval-order 30
no qos traffic-profile classifier name mgmtTraf eval-order 31
no qos traffic-profile classifier name mgmtTraf eval-order 32

qos traffic-profile classifier name mgmtTraf src-ip 10.30.31.15/32 drop-action
disable block nmsCluster eval-order 30
```

```

qos traffic-profile classifier name mgmtTraf src-ip 10.30.31.47/32 drop-action
disable block nmsCluster eval-order 31

qos traffic-profile classifier name mgmtTraf src-ip 10.30.31.48/32 drop-action
disable block nmsCluster eval-order 32

qos traffic-profile classifier name mgmtTraf src-ip 10.30.31.49/32 drop-action
disable block nmsCluster eval-order 33

```

Post filter set optimization, the resource utilization level is much improved:

```
5650TD(config)#show qos diag
```

Unit/Port	Mask Precedence Usage															
	16	15	14	13	12	11	10	9	8	7	6	5	4	3	2	1
1/1	AR	DH	Q	Q	Q	Q	Q					IS			Q	Q
1/2	AR	DH													Q	Q

AR=ARP DH=DHCP Q=QoS IS=IPSG

Unit/Port	Prec	Filter	Meter	Cntr	Filter	Meter	Non QoS		Non QoS		RngChk
							Filter	Meter	Filter	Meter	
		Used	Used	Used	Used	Used	Total	Total	Total	Used	
1 /1 -24,49	16	0	0	0	25	25	256	128	128		
	15	0	0	0	50	25	256	128	128		
	14	<b>1</b>	0	<b>1</b>	0	0	256	128	128		
	13	<b>2</b>	0	<b>1</b>	0	0	256	128	128		
	12	<b>4</b>	0	<b>1</b>	0	0	256	128	128		
	11	<b>1</b>	0	<b>1</b>	0	0	256	128	128		
	10	<b>1</b>	0	<b>1</b>	0	0	256	128	128		
	9	0	0	0	0	0	256	128	128		
	8	0	0	0	0	0	256	128	128		
	7	0	0	0	0	0	256	128	128		
	6	0	0	0	1	0	256	128	128		

It may prove useful to pay closer attention to the NMS cluster and bandwidth consumption. Which NMS is generating the most traffic? Is too much bandwidth being consumed by the NMS cluster or by any of the other network management streams? Adding metering criteria and improving statistics tracking granularity should provide some answers:

```

no qos traffic-profile set port 1 name mgmtTraf

qos traffic-profile set port 1 name mgmtTraf committed-rate 64 committed-burst
16384 drop-action enable track-statistics individual

```

Meter usage and additional counter consumption is reflected in the resource usage data:

```
5650TD(config)#show qos diag
```

Unit/Port	Mask Precedence Usage															
	16	15	14	13	12	11	10	9	8	7	6	5	4	3	2	1
1/1	AR	DH	Q	Q	Q	Q	Q					IS			Q	Q
1/2	AR	DH													Q	Q



AR=ARP DH=DHCP Q=QoS IS=IPSG

Unit/Port	Prec	Filter	Meter	Cntr	Filter	Meter	Non QoS		Cntr	RngChk
							Filter	Meter		
		Used	Used	Used	Used	Used	Total	Total	Total	Used
1 /1 -24,49	16	0	0	0	25	25	256	128	128	
	15	0	0	0	50	25	256	128	128	
	14	<b>1</b>	<b>1</b>	<b>1</b>	0	0	256	128	128	
	13	<b>2</b>	<b>1</b>	<b>2</b>	0	0	256	128	128	
	12	<b>4</b>	<b>1</b>	<b>4</b>	0	0	256	128	128	
	11	<b>1</b>	<b>1</b>	<b>1</b>	0	0	256	128	128	
	10	<b>1</b>	<b>1</b>	<b>1</b>	0	0	256	128	128	
	9	0	0	0	0	0	256	128	128	
	8	0	0	0	0	0	256	128	128	
	7	0	0	0	0	0	256	128	128	
	6	0	0	0	1	0	256	128	128	

The administrator has decided to designate another port as a secondary network management traffic interface. The “mgmtTraf” filter set thus needs to be applied to this interface (e.g., port 2). Management traffic on this port should be given less priority however. Remark network management traffic that is inappropriately prioritized (i.e., remark traffic with DSCP values CS7, CS6, EF and CS5 to AF41) to an acceptable value:

```

qos traffic-profile classifier name dscpChk ds-field 56 update-dscp 34 block
dscpRemark eval-order 1

qos traffic-profile classifier name dscpChk ds-field 48 update-dscp 34 block
dscpRemark eval-order 2

qos traffic-profile classifier name dscpChk ds-field 46 update-dscp 34 block
dscpRemark eval-order 3

qos traffic-profile classifier name dscpChk ds-field 40 update-dscp 34 block
dscpRemark eval-order 4

qos traffic-profile set port 2 name dscpChk

qos traffic-profile set port 2 name mgmtTraf committed-rate 64 committed-burst
16384 drop-action enable track-statistics individual

```

Traffic-profile filter sets are now installed on port 1 (“mgmtTraf”) and port 2 (“dscpChk”, “mgmtTraf”). Filter set details, as well as the port associations, can be queried using the ‘**show qos traffic-profile**’ commands:

```
5650TD(config)#show qos traffic-profile interface
```

	Id		Unit	Port	Filter Set Name
1	1	1			mgmtTraf
2	1	2			dscpChk
3	1	2			mgmtTraf

```
5650TD(config)#show qos traffic-profile set
```

```

Id: 1
Name: mgmtTraf
Unit/Port: 1/1

```

```
Commit Rate: 64 Kbps
Commit Burst: 16777216 Bytes
Out-Profile Drop Action: Drop
Out-Profile Remark DSCP Action: None
Non-Match Action: Defer
Storage Type: NonVolatile
```

```
Id: 2
Name: dscpChk
Unit/Port: 1/2
Non-Match Action: Defer
Storage Type: NonVolatile
```

```
Id: 3
Name: mgmtTraf
Unit/Port: 1/2
Commit Rate: 64 Kbps
Commit Burst: 16777216 Bytes
Out-Profile Drop Action: Drop
Out-Profile Remark DSCP Action: None
Non-Match Action: Defer
Storage Type: NonVolatile
```

The current filter set assignments are, of course, reflected in the resource usage data. Filter set “dscpChk” consumes 1 precedence mask (precedence 14 on port 2) and 4 filters/counters while each instance of filter set “mgmtTraf” consumes 5 precedence masks (precedence values 10 – 14 on port 1 and 9 – 13 on port 2), 5 meters and 9 filters/counters:

```
5650TD(config)#show qos diag
```

Unit/Port	Mask Precedence Usage															
	16	15	14	13	12	11	10	9	8	7	6	5	4	3	2	1
1/1	AR	DH	Q	Q	Q	Q	Q					IS			Q	Q
1/2	AR	DH	Q	Q	Q	Q	Q	Q							Q	Q

AR=ARP DH=DHCP Q=QoS IS=IPSG

Unit/Port	Prec	Non QoS			Non QoS			Filter Total	Meter Total	Cntr Total	RngChk Used
		Filter Used	Meter Used	Cntr Used	Filter Used	Meter Used	Cntr Used				
1 /1 -24,49	16	0	0	0	25	25	256	128	128		
	15	0	0	0	50	25	256	128	128		
	14	5	1	5	0	0	256	128	128		
	13	3	2	3	0	0	256	128	128		
	12	6	2	6	0	0	256	128	128		
	11	5	2	5	0	0	256	128	128		
	10	2	2	2	0	0	256	128	128		
	9	1	1	1	0	0	256	128	128		
	8	0	0	0	0	0	256	128	128		
	7	0	0	0	0	0	256	128	128		
	6	0	0	0	1	0	256	128	128		

Traffic-profile filter set statistics confirm the presence of the targeted management traffic:

```
5650TD(config)#show qos traffic-profile statistics port 1 name mgmtTraf
```

Name: mgmtTraf

IfIndex: 1

Mask Precedence	In-Profile Packets (Overflow)	Out-Profile Packets (Overflow)
-----	-----	-----
14	414	0
13	110	0
12	316	0
11	399	0
10	11	0

Statistics for individual members of a classifier block may be displayed as well:

```
5650TD(config)#show qos traffic-profile statistics port 1/1 name mgmtTraf  
precedence 12
```

Name: mgmtTraf

IfIndex: 1

Precedence: 12

Eval Order	In-Profile Packets (Overflow)	Out-Profile Packets (Overflow)
-----	-----	-----
31	219	0
32	29	0
34	68	0

Only non-zero values are displayed for efficiency purposes.

## 4.4 Troubleshooting

Special consideration should be given to a few areas where issues are more commonly encountered during traffic-profile filter set construction and installation.

### 4.4.1 Resource Exhaustion

Filter set installation consumes precedence levels/masks and filter/meter/counter resources. There is a finite amount of these resources. Furthermore, these resources are shared by QoS and non-QoS applications and may also be shared across ports in a cluster. For these reasons, resource exhaustion is one of the more common issues that are encountered. Troubleshooting resource exhaustion issues is generally straightforward. Error messages are generated when the requested resources are not available during filter set construction:

```
% Evaluation precedence conflict (Traffic Prof filter set data)
% Traffic Prof filter set rule count exceeds limit
% Traffic Prof filter set policy count exceeds limit
```

Resource-related error messages are generated during filter set installation as well:

```
% Traffic Prof filter set elem count exceeds available resources
% Traffic Prof filter set install operation failed
% Traffic Prof filter set update operation failed
```

Issues encountered during filter set construction are relatively easy to diagnose based on the provided error message. Filter set installation issues require a bit more effort to understand. The first step is always to perform an inventory of the filter set components to determine the precedence and filter/meter/counter resources that are needed. Next the resources available on the target port need to be considered. Are the required consecutive precedence levels available? Are there enough filter/meter/counter resources available to support an instance of the filter set? The '**show qos diag**' output can be used to answer these questions.

The different types of QoS hardware complicate resource issue diagnosis. With legacy ERS stackable platforms (ERS55XX), resources are allocated on a per-port basis. With the newer platforms (ERS4XXX/ERS56XX/ERS59XX/VSP7000), resources are allocated from a central pool that serves a collection of ports (referred to in this document as a "port cluster"). With a centralized resource scheme, resources consumed on one port can impact resource availability for all other ports in the same cluster. Once again, the '**show qos diag**' output can be consulted for cluster-wide resource utilization information.

### 4.4.2 Classifier Block Formation

As with resource issue diagnosis, different types of QoS hardware complicate classifier block formation. Classifier elements matching the exact same protocol fields or the exact same portions thereof (i.e., the fields are masked with the same bit-mask) are always block-compatible. On the ERS stackable platforms (i.e., ERS4XXX/56XX/59XX/VSP7000) these restrictions have been eased. Filter criteria matching the same protocol fields, regardless of the whether full or partial field matches are specified, can be combined into a block. Furthermore, classifier elements matching different but compatible fields (i.e., fields from a hardware-defined set of compatible fields) may also be combined into a block. This last class of potential block members, however, is sometimes difficult to determine.

The QoS hardware on the ERS stackable platforms (ERS4XXX/ERS56XX/ERS59XX/VSP7000) builds a 'key' which represents the classification criteria that is used to identify traffic. The same classification data can be represented several different ways (i.e., different keys can be built to match the same protocol fields). If the keys representing different classifiers are the same, the classifiers may be combined into a

block. If they aren't, they can't. If there is flexibility regarding which protocol fields can be matched to achieve the desired result, it may be beneficial to try different filter variations to see which, if any, can be combined into a block (particularly if resources are scarce). Following are some of the more common sets of fields that can be used by members of the same block:

- VLAN ID (outer), Inner VLAN ID
- Source IPv4 address, Destination IPv4 address, IP protocol, L4 source port (exact match), L4 destination port (exact match), DSCP, IP flags, TCP control flags
- Source IPv4 address, Destination IPv4 address, IP protocol, L4 source/destination range check, L4 source/destination port (exact match), DSCP, IP flags, TCP control flags
- Destination MAC address, Source MAC address, EtherType, VLAN ID (outer)
- Source MAC address, Source IPv4 address, EtherType, VLAN ID (outer)
- Destination MAC address, Destination IPv4 address, EtherType, VLAN ID (outer)

### 4.4.3 Interface Capabilities

Several different types of platforms can be combined into an ERS stack. Across the different platforms there are several variations of QoS hardware present. This translates into ports with different QoS capabilities. These differences, some of which have already been discussed, include:

- Classifier capabilities – certain protocol fields (e.g., IP flags, TCP control flags, inner VLAN ID) can only be matched on specific types of ports (deemed QoS Version 2 ports)
- Metering capabilities – supported rates, burst sizes and granularity differ across platforms/ports
- Range check support – true (random, not bit-mask constrained) range check support is not available on legacy platforms (deemed QoS Version 1 only ports)
- Classifier block flexibility – block member selection is more restrictive on legacy platforms
- Action support – IPv6 remarking is not supported on certain legacy platforms

When constructing a filter set, the target port capabilities should be kept in mind. Filter set content should be restricted to the lowest common denominator in terms of the capabilities supported by the target ports. Otherwise the filter set installation request will be rejected:

```
% Filter set requirements incompatible with target interface
```

Interface capabilities can be displayed with the 'show qos capability meter' and the 'show qos if-assign' commands.

## 5. QoS Egress Queue Shaping

You can use QoS Egress Queue Shaping to configure egress shaping on a port by port and queue by queue basis.

QoS shaper rate queue servicing on the Avaya ERS 4000, ERS 5600, ERS 5900, and VSP 7000 uses a weighted round robin algorithm to shape traffic. With egress queue shaping, you can specify the maximum and minimum egress shaping rates on an individual port and queue basis. You can configure shaping criteria for any or all egress queues associated with a switch port. The number of egress queues available for a port is determined by the QoS agent egress queue set value.

You can use queue shaping in conjunction with interface shaping.

Bandwidth allocation for queues is done according to Strict Priority and WRR algorithm. When shapers on queues with minimum rate are configured, the system first tries to assure minimum rate for all queues. Then the system uses the remaining bandwidth according to Strict Priority, WRR and shape maximum rate configured for each queue. In case the sum of shape minimum rates configured (queue shapers) exceeds the line rate, minimum shape rate will be assured for queue 1 and then remaining bandwidth will be distributed to the rest of the queues using WRR algorithm in order to assure the minimum rates for the rest of queues. For the same scenario ERS5600 switches may use strict priority, WRR and RR algorithms, depending on the active queue set.

### 5.1 Configuring QoS Interface Shaper

Use the following procedure to configure the interface shaping parameters for a set of ports.

To configure parameters, use the following command from Interface Configuration mode:

```
[no] qos if-shaper [name <WORD>] [shape-rate <64-10230000>] [burst-size <burst-size>] [max-burst-rate <64-4294967295>] [max-burst-duration <1-4294967295>]
```

Use the no form of this command to disable interface shaping for a set of ports.

Action Option	Description
<WORD>	Specify name for if-shaper; maximum is 16 alphanumeric characters.
shape-rate <64-10230000>	Shaping rate in kilobits/sec; range is 64-10230000 kilobits/sec.
burst-size <burst-size>	Committed burst size in Kilobytes. The value range is: 4, 8, 16, 32, 64, 128, 256, 512, 1024, 2048, 4096, 8192, 16384.
max-burst-rate <64-4294967295>	Maximum burst rate in kilobits/sec; range is 64-4294967295Kbits/sec.
max-burst-duration <1-4294967295>	Maximum burst duration in milliseconds; range is 1–4294967295 ms.

## 5.2 Configuring QoS interface queue shaper

Use the following procedure to create an egress queue shaper for one or more interfaces.

To create an egress queue shaper, use the following command from the Interface Configuration mode:

```
interface ethernet <port list or ALL>
  qos if-queue-shaper [port <portlist>] [queue <1-8>] [name <WORD>] shape-rate
    <0-10230000> shape-min-rate <0-10230000>
```




If you configure the shape rate to 0 for a specific queue or port, shaping is not performed on that queue or port.

Action Option	Description
<WORD>	Specifies an alphanumeric label used to identify the QoS interface queue shaper. Value is a character string ranging from 1–16 characters in length.
port <portlist>	Specifies the port or list of ports for which to apply egress queue shaping.
queue <1-8>	Specifies the queue for the selected interface port or ports, on which traffic is shaped. The range of available values is determined by the OoS agent default queue configuration.
shape-min-rate <0-10230000>	Specifies the minimum QoS interface queue shaping rate, in kilobits per second (Kbps). Values range from 0 to10230000 Kbps.
shape-rate <0-10230000>	Specifies the QoS interface queue shaping rate, in kilobits per second (Kbps). Values range from 0 to10230000 Kbps.

## 5.3 Configuration Examples using Meters and Shaper

### 5.3.1 ERS 4000 switch: UDP/TCP Port Range to remark traffic and dst-IP to Meter traffic



Assuming we wish to accomplish the following:

- Traffic with a destination UDP port range of 5000 to 5005, remark to Premium (DSCP = 46, p-bit = 6)
- Traffic with a destination UDP port range of 5024 to 5044, remark to Platinum (DSCP = 34, p-bit = 4)
- Traffic with a destination TCP port range of 5060 to 5061, remark to Premium (DSCP = 46, p-bit = 6)
- Traffic with a destination of 192.168.0.0/16, traffic metered to a maximum of 5 Mbps burstable to 10M Mbps for a period of 1 second
- All other traffic, pass as-is

#### 1 Traffic classifier configuration using name of *one* and add the UDP range filter to a block also named *one* to save a precedence level


```
qos traffic-profile classifier name one protocol 17 dst-port-min 5000 dst-port-max 5005 update-dscp 46 update-lp 6 block one eval-order 5
qos traffic-profile classifier name one protocol 17 dst-port-min 5024 dst-port-max 5044 update-dscp 34 update-lp 4 block one eval-order 6
qos traffic-profile classifier name one protocol 6 dst-port-min 5060 dst-port-max 5061 update-dscp 46 update-lp 6 eval-order 10
qos traffic-profile classifier name one dst-ip 192.168.0.0/16 eval-order 20
committed-rate 5000 max-burst-rate 10000 max-burst-duration 1000 drop-out-action enable
```

#### 2 Assign traffic classifier to ingress ports 1 to 11

```
qos traffic-profile set port 1-11 name one meter-mode classifier
```



## 5.3.2 ERS 4000 switch: UDP/TCP Port Range using meter to remark traffic



Assuming we wish to accomplish the following:

- Traffic with a destination UDP port range of 5000 to 5005, remark to Premium up to 10 Mbps, then remark to Silver
- Traffic with a destination UDP port range of 5024 to 5044, remark to Platinum up to 5 Mbps, then remark to Bronze
- Traffic with a destination TCP port range of 5060 to 5061, remark to Premium up to 10 Mbps, then remark to Silver
- All other traffic, pass as-is

### 1 Traffic classifier configuration using name of one

```
qos traffic-profile classifier name one protocol 17 dst-port-min 5000 dst-port-max 5005 update-dscp 46 update-lp 6 block one eval-order 5 committed-rate 10000 committed-burst-size 4 drop-out-action disable update-dscp-out-action 16

qos traffic-profile classifier name one protocol 17 dst-port-min 5024 dst-port-max 5044 update-dscp 34 update-lp 4 block one eval-order 6 committed-rate 5000 committed-burst-size 4 drop-out-action disable update-dscp-out-action 10

qos traffic-profile classifier name one protocol 6 dst-port-min 5060 dst-port-max 5061 update-dscp 46 update-lp 6 eval-order 20 committed-rate 10000 committed-burst-size 4 drop-out-action disable update-dscp-out-action 16
```

### 2 Assign traffic classifier to ingress ports 1 to 11

```
qos traffic-profile set port 1-11 name one meter-mode classifier
```

### 5.3.3 Port Shaper


Assuming we wish to accomplish the following:

- Setup up a port shaper to 40 Mbps with a burst of 80 Mbps for 10 seconds on egress port 25

#### 1 Traffic classifier configuration using name of *one*

```
interface fastEthernet 25
qos if-shaper name 40M shape-rate 40000 max-burst-rate 80000 max-burst-duration 10000
exit
```

## 5.3.4 Per Queue Shaper



Assuming we wish to accomplish the following assuming we are using egress port 25:

- Setup queue set 4 where we wish to steer certain ingress traffic to a specific queue so that we can shape the traffic
  - Traffic with destination to 10.0.0.0/24, 10.0.1.0/24, and 10.0.2.0/24 remark with p-bit 5, DSCP 32
    - Force traffic to Queue 3 with shaper set to 10M via Traffic Profile that remarks p-bit to 5 and DSCP to 32
  - All other traffic remark with p-bit 0, DSCP 0
    - Force traffic to Queue 4 with shaper set to 5M via Traffic Profile that remark p-bit to 0
- Default mapping for queue set 4
  - Queue 3: p-bit 5, DSCP 32, 34, 36, 38
  - Queue 4: p-bit 0 to 4, DSCP 0, 8, 10, 12, 14, 16, 18, 20, 22, 24, 26, 28

The following commands can be used to view the queue set details.

```
4550T-10#show qos queue-set 4
```

Set ID	Queue ID	General Discipline	Bandwidth (%)	Absolute Bandwidth (Kbps)	Bandwidth Allocation	Service Order	Size (Bytes)
4	1	Priority Queuing	100	0	Relative	1	61440
4	2	Weighted Round Robin	65	0	Relative	2	50688
4	3	Weighted Round Robin	26	0	Relative	2	45056
4	4	Weighted Round Robin	9	0	Relative	2	39424

4550T-10#show qos queue-set-assignment queue-set 4		4550T-10 (config)#show qos egressmap					
Queue Set 4		DSCP	802.1p	Priority	Drop	Precedence	New DSCP
802.1p Priority Queue							Name
0	4	0	0	High Drop	0	0	Standard Service
1	4	8	2	High Drop	8	8	Bronze Service
2	4	10	2	Low Drop	10	10	Bronze Service
3	4	16	3	High Drop	16	16	Silver Service
4	4	18	3	Low Drop	18	18	Silver Service
5	3	24	4	High Drop	24	24	Gold Service
6	1	26	4	Low Drop	26	26	Gold Service
		32	5	High Drop	32	32	Platinum Service
		34	5	Low Drop	34	34	Platinum Service
		40	6	Low Drop	40	40	Premium Service
		46	6	Low Drop	46	46	Premium Service
		48	7	Low Drop	48	48	Network Service
		56	7	Low Drop	56	56	Critical Service

## 1 Change default queue set from 2 to 4 and reboot switch

```
qos agent queue-set 4
QoS queue setting isn't effective until after reset.
boot
Reboot the unit(s) (y/n) ? y
```

## 2 Traffic classifier configuration using name of *remark* and add the dst-IP filters to a block named *one* to save a precedence level

```
interface FastEthernet ALL
qos if-queue-shaper port 25 queue 3 name 10M shape-rate 10000 shape-min-rate 10000
qos if-queue-shaper port 25 queue 4 name 5M shape-rate 5000 shape-min-rate 5000
exit
qos traffic-profile classifier name remark dst-ip 10.0.0.0/24 update-dscp 32 update-lp
5 block one eval-order 5
qos traffic-profile classifier name remark dst-ip 10.0.1.0/24 update-dscp 32 update-lp
5 block one eval-order 6
qos traffic-profile classifier name remark dst-ip 10.0.2.0/24 update-dscp 32 update-lp
5 block one eval-order 7
qos traffic-profile classifier name remark update-lp 0 eval-order 250
```

## 3 Assign traffic classifier to ingress ports 1 to 11

```
qos traffic-profile set port 1-11 name remark qos
```

## 6. Appendix

### 6.1 Command Output: 'show qos traffic-profile classifier'

The following represents the output from the 'show qos traffic-profile classifier' command for the example filter set "mgmtTraf":

```
5650TD(config)#show qos traffic-profile classifier
```

```
Id: 1
Name: mgmtTraf
Block:
Eval Order: 10
Address Type: IPv4
Destination Addr/Mask: Ignore
Source Addr/Mask: Ignore
DSCP: Ignore
IPv4 Protocol / IPv6 Next Header: UDP
Destination L4 Port Min: 161
Destination L4 Port Max: 161
Source L4 Port Min: Ignore
Source L4 Port Max: Ignore
IPv6 Flow Id: Ignore
IP Flags: Ignore
TCP Control Flags: Ignore
IPv4 Options: Ignore
Destination MAC Addr: Ignore
Destination MAC Mask: Ignore
Source MAC Addr: Ignore
Source MAC Mask: Ignore
VLAN: Ignore
VLAN Tag: Ignore
EtherType: 0x0800
802.1p Priority: All
Packet Type: Ignore
Inner VLAN: Ignore
Action Drop: No
Action Update DSCP: Ignore
Action Update 802.1p Priority: Priority 8
Action Set Drop Precedence: Low Drop
Storage Type: NonVolatile
```

```
Id: 2
Name: mgmtTraf
Block: tcp
Eval Order: 20
Address Type: IPv4
Destination Addr/Mask: Ignore
Source Addr/Mask: Ignore
DSCP: Ignore
IPv4 Protocol / IPv6 Next Header: TCP
Destination L4 Port Min: 80
Destination L4 Port Max: 80
Source L4 Port Min: Ignore
Source L4 Port Max: Ignore
IPv6 Flow Id: Ignore
IP Flags: Ignore
```

TCP Control Flags: Ignore  
 IPv4 Options: Ignore  
 Destination MAC Addr: Ignore  
 Destination MAC Mask: Ignore  
 Source MAC Addr: Ignore  
 Source MAC Mask: Ignore  
 VLAN: Ignore  
 VLAN Tag: Ignore  
 EtherType: 0x0800  
 802.1p Priority: All  
 Packet Type: Ignore  
 Inner VLAN: Ignore  
 Action Drop: No  
 Action Update DSCP: Ignore  
 Action Update 802.1p Priority: Priority 8  
 Action Set Drop Precedence: Low Drop  
 Storage Type: NonVolatile

Id: 3  
 Name: mgmtTraf  
 Block: tcp  
 Eval Order: 21  
 Address Type: IPv4  
 Destination Addr/Mask: Ignore  
 Source Addr/Mask: Ignore  
 DSCP: Ignore  
 IPv4 Protocol / IPv6 Next Header: TCP  
 Destination L4 Port Min: 23  
 Destination L4 Port Max: 23  
 Source L4 Port Min: Ignore  
 Source L4 Port Max: Ignore  
 IPv6 Flow Id: Ignore  
 IP Flags: Ignore  
 TCP Control Flags: Ignore  
 IPv4 Options: Ignore  
 Destination MAC Addr: Ignore  
 Destination MAC Mask: Ignore  
 Source MAC Addr: Ignore  
 Source MAC Mask: Ignore  
 VLAN: Ignore  
 VLAN Tag: Ignore  
 EtherType: 0x0800  
 802.1p Priority: All  
 Packet Type: Ignore  
 Inner VLAN: Ignore  
 Action Drop: No  
 Action Update DSCP: Ignore  
 Action Update 802.1p Priority: Priority 8  
 Action Set Drop Precedence: Low Drop  
 Storage Type: NonVolatile

Id: 4  
 Name: mgmtTraf  
 Block: nmsCluster  
 Eval Order: 30  
 Address Type: IPv4  
 Destination Addr/Mask: Ignore  
 Source Addr/Mask: 10.30.31.15/32  
 DSCP: Ignore  
 IPv4 Protocol / IPv6 Next Header: Ignore  
 Destination L4 Port Min: Ignore

Destination L4 Port Max: Ignore  
 Source L4 Port Min: Ignore  
 Source L4 Port Max: Ignore  
 IPv6 Flow Id: Ignore  
 IP Flags: Ignore  
 TCP Control Flags: Ignore  
 IPv4 Options: Ignore  
 Destination MAC Addr: Ignore  
 Destination MAC Mask: Ignore  
 Source MAC Addr: Ignore  
 Source MAC Mask: Ignore  
 VLAN: Ignore  
 VLAN Tag: Ignore  
 EtherType: 0x0800  
 802.1p Priority: All  
 Packet Type: Ignore  
 Inner VLAN: Ignore  
 Action Drop: No  
 Action Update DSCP: Ignore  
 Action Update 802.1p Priority: Priority 8  
 Action Set Drop Precedence: Low Drop  
 Storage Type: NonVolatile

Id: 5  
 Name: mgmtTraf  
 Block: nmsCluster  
 Eval Order: 31  
 Address Type: IPv4  
 Destination Addr/Mask: Ignore  
 Source Addr/Mask: 10.30.31.47/32  
 DSCP: Ignore  
 IPv4 Protocol / IPv6 Next Header: Ignore  
 Destination L4 Port Min: Ignore  
 Destination L4 Port Max: Ignore  
 Source L4 Port Min: Ignore  
 Source L4 Port Max: Ignore  
 IPv6 Flow Id: Ignore  
 IP Flags: Ignore  
 TCP Control Flags: Ignore  
 IPv4 Options: Ignore  
 Destination MAC Addr: Ignore  
 Destination MAC Mask: Ignore  
 Source MAC Addr: Ignore  
 Source MAC Mask: Ignore  
 VLAN: Ignore  
 VLAN Tag: Ignore  
 EtherType: 0x0800  
 802.1p Priority: All  
 Packet Type: Ignore  
 Inner VLAN: Ignore  
 Action Drop: No  
 Action Update DSCP: Ignore  
 Action Update 802.1p Priority: Priority 8  
 Action Set Drop Precedence: Low Drop  
 Storage Type: NonVolatile

Id: 6  
 Name: mgmtTraf  
 Block: nmsCluster  
 Eval Order: 32  
 Address Type: IPv4

Destination Addr/Mask: Ignore  
 Source Addr/Mask: 10.30.31.48/32  
 DSCP: Ignore  
 IPv4 Protocol / IPv6 Next Header: Ignore  
 Destination L4 Port Min: Ignore  
 Destination L4 Port Max: Ignore  
 Source L4 Port Min: Ignore  
 Source L4 Port Max: Ignore  
 IPv6 Flow Id: Ignore  
 IP Flags: Ignore  
 TCP Control Flags: Ignore  
 IPv4 Options: Ignore  
 Destination MAC Addr: Ignore  
 Destination MAC Mask: Ignore  
 Source MAC Addr: Ignore  
 Source MAC Mask: Ignore  
 VLAN: Ignore  
 VLAN Tag: Ignore  
 EtherType: 0x0800  
 802.1p Priority: All  
 Packet Type: Ignore  
 Inner VLAN: Ignore  
 Action Drop: No  
 Action Update DSCP: Ignore  
 Action Update 802.1p Priority: Priority 8  
 Action Set Drop Precedence: Low Drop  
 Storage Type: NonVolatile

Id: 7  
 Name: mgmtTraf  
 Block:  
 Eval Order: 40  
 Address Type: IPv4  
 Destination Addr/Mask: Ignore  
 Source Addr/Mask: Ignore  
 DSCP: Ignore  
 IPv4 Protocol / IPv6 Next Header: ICMP  
 Destination L4 Port Min: Ignore  
 Destination L4 Port Max: Ignore  
 Source L4 Port Min: Ignore  
 Source L4 Port Max: Ignore  
 IPv6 Flow Id: Ignore  
 IP Flags: Ignore  
 TCP Control Flags: Ignore  
 IPv4 Options: Ignore  
 Destination MAC Addr: Ignore  
 Destination MAC Mask: Ignore  
 Source MAC Addr: Ignore  
 Source MAC Mask: Ignore  
 VLAN: Ignore  
 VLAN Tag: Ignore  
 EtherType: 0x0800  
 802.1p Priority: All  
 Packet Type: Ignore  
 Inner VLAN: Ignore  
 Action Drop: No  
 Action Update DSCP: Ignore  
 Action Update 802.1p Priority: Priority 8  
 Action Set Drop Precedence: Low Drop  
 Storage Type: NonVolatile



Id: 8  
 Name: mgmtTraf  
 Block:  
 Eval Order: 100  
 Address Type: Ignore  
 Destination Addr/Mask: Ignore  
 Source Addr/Mask: Ignore  
 DSCP: Ignore  
 IPv4 Protocol / IPv6 Next Header: Ignore  
 Destination L4 Port Min: Ignore  
 Destination L4 Port Max: Ignore  
 Source L4 Port Min: Ignore  
 Source L4 Port Max: Ignore  
 IPv6 Flow Id: Ignore  
 IP Flags: Ignore  
 TCP Control Flags: Ignore  
 IPv4 Options: Ignore  
 Destination MAC Addr: Ignore  
 Destination MAC Mask: Ignore  
 Source MAC Addr: Ignore  
 Source MAC Mask: Ignore  
 VLAN: Ignore  
 VLAN Tag: Ignore  
 EtherType: Ignore  
 802.1p Priority: All  
 Packet Type: Ignore  
 Inner VLAN: Ignore  
 Action Drop: Yes  
 Action Update DSCP: Ignore  
 Action Update 802.1p Priority: Priority 8  
 Action Set Drop Precedence: Low Drop  
 Storage Type: NonVolatile

Id: 9  
 Name: mgmtTraf  
 Block: nmsCluster  
 Eval Order: 33  
 Address Type: IPv4  
 Destination Addr/Mask: Ignore  
 Source Addr/Mask: 10.30.31.49/32  
 DSCP: Ignore  
 IPv4 Protocol / IPv6 Next Header: Ignore  
 Destination L4 Port Min: Ignore  
 Destination L4 Port Max: Ignore  
 Source L4 Port Min: Ignore  
 Source L4 Port Max: Ignore  
 IPv6 Flow Id: Ignore  
 IP Flags: Ignore  
 TCP Control Flags: Ignore  
 IPv4 Options: Ignore  
 Destination MAC Addr: Ignore  
 Destination MAC Mask: Ignore  
 Source MAC Addr: Ignore  
 Source MAC Mask: Ignore  
 VLAN: Ignore  
 VLAN Tag: Ignore  
 EtherType: 0x0800  
 802.1p Priority: All  
 Packet Type: Ignore  
 Inner VLAN: Ignore  
 Action Drop: No

```
Action Update DSCP: Ignore
Action Update 802.1p Priority: Priority 8
Action Set Drop Precedence: Low Drop
Storage Type: NonVolatile
```

## 6.2 Legacy CLI Command Comparison

For those familiar with QoS configuration using the legacy CLI commands, it may be a useful exercise to compare configuration command sequences using legacy commands and the new traffic-profile filter set support. Revisiting an earlier example, suppose an administrator wishes to remark traffic that is inappropriately prioritized (i.e., remark traffic with the highest priority DSCP values CS7 and CS6 to AF41) on a specific VLAN. Using the traffic-profile filter set support, the following commands are required:

```
qos traffic classifier name dscpChk ds-field 56 vlan-min 5 vlan-max 5 update-dscp
34 block dsRem eval-order 1

qos traffic classifier name dscpChk ds-field 48 vlan-min 5 vlan-max 5 update-dscp
34 block dsRem eval-order 2

qos traffic set port 8 name dscpChk committed-rate 64 committed-burst 16384 drop
enable track-statistics individual
```

Using the legacy QoS CLI commands, a more expansive list of commands must be executed to initiate the same traffic treatment:

```
qos ip-element 1 ds-field 56
qos ip-element 2 ds-field 48
qos l2-element 1 vlan-min 5 vlan-max 5 ethertype 0x0800
qos classifier 1 set-id 1 element-type ip element-id 1
qos classifier 2 set-id 1 element-type l2 element-id 1
qos classifier 3 set-id 2 element-type ip element-id 2
qos classifier 4 set-id 2 element-type l2 element-id 1
qos classifier-block 1 block-number 1 set-id 1
qos classifier-block 2 block-number 1 set-id 2
qos action 10 update-dscp 34
qos meter 1 committed-rate 64 burst-size 16384 in-profile-action 10 out-profile-
action 1
qos policy 1 port 8 clfr-type block clfr-id 1 meter 1 precedence 14 track-
statistics individual
```

## 7. Default QoS Settings

### 7.1 Default Settings: ERS 4000

#### ***Default Settings prior to release 5.8***

- Precedence level 1, 2, 7, and 8 are used by default
  - Precedence levels 1 and 2 are used by default untrusted if-group (allQoSPolicyIfcs)
    - Can be cleared up if ports are re-assigned to QoS class of either trusted or unrestricted
  - Precedence level 7 is used by DHCP Relay
    - Can be cleared up if IP DHCP Relay is disabled – on by default
  - Precedence level 8 is used by ARP
  - Certain features if enabled can take up one or more precedence levels
    - Please see Appendix for list of precedence levels used by various features
- 128 maximum filters are available per BCM chip where each BCM chip contains 24-26 ports
  - For example, 8 ports with 16 filters would consume all 128 filters via one BCM
- Default Queue set is 2 with buffering set to large

#### ***Default Settings release 5.8 or higher – only ERS4800 series, ERS 4500 is not supported***

- Precedence level 1, 2, 14, 15 and 16 are used by default
  - Precedence levels 1 and 2 are used by default untrusted if-group (allQoSPolicyIfcs)
    - Can be cleared up if ports are re-assigned to QoS class of either trusted or unrestricted
  - Precedence level 14 is used by DHCP Relay
    - Can be cleared up if IP DHCP Relay is disabled – on by default
  - Precedence levels 15 and 16 are permanently occupied
    - Precedence level 15 is used by SPB
    - Precedence level 16 is used by ARP
  - Certain features if enabled can take up one or more precedence levels
    - Please see Appendix for list of precedence levels used by various features
  - For the ERS 4850
    - Up to 512 classifiers for each mask precedence
    - Up to 256 meters for each mask precedence for each ASIC
    - Up to 256 counters for each mask precedence for each ASIC
  - For the ERS 4826
    - Up to 256 classifiers for each mask precedence
    - Up to 128 meters for each mask precedence for each ASIC
    - Up to 128 counters for each mask precedence for each ASIC

- Default Queue set is 2 with buffering set to large

The following table describes the ports supported by each ASIC for each ERS 4000 model:

Model	ASIC Device 1	ASIC Device 2
4526FX	Port 1 - 26	
4550T	Port 1 - 24	Port 25 – 50
4550T-PWR	Port 1 - 24	Port 25 – 50
4548GT	Port 1 - 24	Port 25 – 48
4548GT-PWR	Port 1 - 24	Port 25 – 48
4526T	Port 1 – 26	
4526T-PWR	Port 1 – 26	
4526GT	Port 1 – 24	
4526GT-PWR	Port 1 – 24	
4526GT-PWR	Port 1 – 24	
4526GTX	Port 1 – 26	
4526GTX-PWR	Port 1 – 26	
4526T-PWR+	Port 1 – 26	
4550T-PWR+	Port 1 – 24	Port 25-50
4826GTS	Port 1 – 26	
4826GTS-PWR+	Port 1 – 26	
4850GTS	Port 1 – 50	
4850GTS-PWR+	Port 1 - 50	

## Verification

### 1 Verify QoS queue set and buffer:

```
4850GTS-PWR+#show qos agent
```

```
QoS Operational Mode: Enabled
```

```
QoS NVRam Commit Delay: 10 seconds
```

```
QoS Current Queue Set: 2
```

```
QoS Next Boot Queue Set: 2
```

```
QoS Current Buffering: Large
```

```
QoS Next Boot Buffering: Large
```

```
QoS UBP Support Level: Disabled
```

```
QoS Default Statistics Tracking: Aggregate
```

```
Auto QoS Mode: Disabled
```

### 2 Verify QoS interface group and assignments

```
4850GTS-PWR+#show qos if-group
```

Role Combination	Interface Class	Capabilities	Statistics Tracking	Storage Type
allQoSPolicyIfcs	Untrusted	Input 802, Input IP Aggregate		ReadOnly
\$qosDisabledIfcs	Unrestricted	Input 802, Input IP Disabled		Other

```
4850GTS-PWR+#show qos if-assign
```

Unit	Port	IfIndex	Role	Combination	Queue Set	Capability
1	1	1	allQoSPolicyIfcs	2		Version 1,2
1	2	2	allQoSPolicyIfcs	2		Version 1,2
1	49	49	allQoSPolicyIfcs	2		Version 1,2
1	50	50	allQoSPolicyIfcs	2		Version 1,2

### 3 Verify QoS precedence levels and filters / meters used

4850GTS-PWR+#*show qos diag*

Unit/Port	Mask Precedence Usage															
	16	15	14	13	12	11	10	9	8	7	6	5	4	3	2	1
1/1	AR	SB	DH												Q	Q
1/2	AR	SB	DH												Q	Q
1/3	AR	SB	DH												Q	Q
1/49	AR	SB	DH												Q	Q
1/50	AR	SB	DH												Q	Q

AR=ARP DH=DHCP EA=EAP Q=QoS SB=SPB

Unit/Port	NonQoS NonQoS										RngChk
	Prec	Filter	Meter	Cntr	Filter	Meter	Filter	Meter	Cntr		
		Used	Used	Used	Used	Used	Total	Total	Total	Used	
1 /1 -50	16	0	0	0	51	51	512	256	256		
	15	0	0	0	50	0	512	256	256		
	14	0	0	0	200	50	512	256	256		
	13	0	0	0	1	0	512	256	256		
	12	0	0	0	0	0	512	256	256		
	11	0	0	0	0	0	512	256	256		
	10	0	0	0	0	0	512	256	256		
	9	0	0	0	0	0	512	256	256		
	8	0	0	0	0	0	512	256	256		
	7	0	0	0	0	0	512	256	256		
	6	0	0	0	0	0	512	256	256		
	5	0	0	0	0	0	512	256	256		
	4	0	0	0	0	0	512	256	256		
	3	0	0	0	0	0	512	256	256		
	2	50	0	50	0	0	512	256	256		
	1	50	0	50	0	0	512	256	256		

0 /32

## 7.2 Default Settings: ERS 5000

### Default Settings

- Precedence level 1, 2, 15, and 16 are used by default
  - Precedence levels 1 and 2 are used by default untrusted if-group (allQoSPolicyIfcs)
    - Can be cleared up if ports are re-assigned to QoS class of either trusted or unrestricted
  - Precedence level 15 is used by DHCP Relay
    - Can be cleared up if IP DHCP Relay is disabled – on by default
  - Certain features if enabled can take up one or more precedence levels
    - Please see Appendix for list of precedence levels used by various features
- Up to 256 filter components per precedence per hardware device group
- Up to 128 meters per precedence per hardware device group
- Up to 128 counters per precedence per hardware device group
- Up to 16 TCP/UDP port range checkers per hardware device group
- Default Queue set is 2 with buffering set to large

When using an ERS 5500 series, it has the following limitations:

- You can configure up to 15 policies per interface (port).
- You can configure up to 63 meters per interface (port).
- You can configure up to 125 filter components per interface (port).
- When you enable tracking statistics for the policies, the switch uses one counter for each classifier for each interface (port) of the policy or a counter for each policy. You can assign up to 32 counters to an interface (port).

### Verification

#### 1 Verify QoS queue set and buffer:

```
5698-1#show qos agent
QoS Operational Mode: Enabled
QoS NVRam Commit Delay: 10 seconds
QoS Queue Set: 2
QoS Buffering: Large
QoS UBP Support Level: Disabled
QoS Default Statistics Tracking: Aggregate
QoS DoS Attack Prevention: Disabled
    Minimum TCP Header Length: 20
    Maximum IPv4 ICMP Length: 512
    Maximum IPv6 ICMP Length: 512
Auto QoS Mode: Disabled
```

QoS Trusted Processing Mode: Partial

## 2 Verify QoS interface group and assignments

5698-1#**show qos if-group**

Role Combination	Interface Class	Capabilities	Statistics Tracking	Storage Type
allQoSPolicyIfcs	Untrusted	Input 802, Input IP Aggregate		ReadOnly
\$qosDisabledIfcs	Unrestricted	Input 802, Input IP Disabled		Other
\$NsnaIfcs	Unrestricted	Input 802, Input IP Disabled		Other

5698-1#**show qos if-assign**

Unit	Port	IfIndex	Role	Combination	Queue	Set	Capability	DAPP Support
1	1	1	allQoSPolicyIfcs		42		Version 1,2	Yes
1	2	2	allQoSPolicyIfcs		42		Version 1,2	Yes
1	97	97	allQoSPolicyIfcs		50		Version 1,2	Yes
1	98	98	allQoSPolicyIfcs		50		Version 1,2	Yes



### 3 Verify QoS precedence levels and filters / meters used

5698-1# *show qos diag*

Unit/Port	Mask Precedence Usage															
	16	15	14	13	12	11	10	9	8	7	6	5	4	3	2	1
1/1	AR	DH													Q	Q
1/2	AR	DH													Q	Q
1/97	AR	DH													Q	Q
1/98	AR	DH													Q	Q

AR=ARP DH=DHCP AD=ADAC Q=QoS

Non QoS Non QoS											
Unit/Port	Prec	Filter	Meter	Cntr	Filter	Meter	Filter	Meter	Cntr	RngChk	
		Used	Used	Used	Used	Used	Total	Total	Total	Used	
1 /1 -24,97	16	0	0	0	25	25	256	128	128		
	15	0	0	0	50	25	256	128	128		
	14	0	0	0	75	0	256	128	128		
	13	0	0	0	1	0	256	128	128		
	12	0	0	0	0	0	256	128	128		
	11	0	0	0	0	0	256	128	128		
	10	0	0	0	0	0	256	128	128		
	9	0	0	0	0	0	256	128	128		
	8	0	0	0	0	0	256	128	128		
	7	0	0	0	0	0	256	128	128		
	6	0	0	0	0	0	256	128	128		
	5	0	0	0	0	0	256	128	128		
	4	0	0	0	0	0	256	128	128		
	3	0	0	0	0	0	256	128	128		
	2	25	0	25	0	0	256	128	128		
	1	25	0	25	0	0	256	128	128		
										0	
1 /25-48,98	16	0	0	0	25	25	256	128	128		

	15	0	0	0	50	25	256	128	128
	14	0	0	0	75	0	256	128	128
	13	0	0	0	0	0	256	128	128
	12	0	0	0	0	0	256	128	128
	11	0	0	0	0	0	256	128	128
	10	0	0	0	0	0	256	128	128
	9	0	0	0	0	0	256	128	128
	8	0	0	0	0	0	256	128	128
	7	0	0	0	0	0	256	128	128
	6	0	0	0	0	0	256	128	128
	5	0	0	0	0	0	256	128	128
	4	0	0	0	0	0	256	128	128
	3	0	0	0	0	0	256	128	128
	2	25	0	25	0	0	256	128	128
	1	25	0	25	0	0	256	128	128
									0
1 / 49-72	16	0	0	0	24	24	256	128	128
	15	0	0	0	48	24	256	128	128
	14	0	0	0	72	0	256	128	128
	13	1	0	1	0	0	256	128	128
	12	1	0	1	0	0	256	128	128
	11	1	0	1	0	0	256	128	128
	10	1	0	1	0	0	256	128	128
	9	0	0	0	0	0	256	128	128
	8	0	0	0	0	0	256	128	128
	7	0	0	0	0	0	256	128	128
	6	0	0	0	0	0	256	128	128
	5	0	0	0	0	0	256	128	128
	4	0	0	0	0	0	256	128	128
	3	0	0	0	0	0	256	128	128
	2	24	0	24	0	0	256	128	128
	1	24	0	24	0	0	256	128	128
									0
1 / 73-96	16	0	0	0	24	24	256	128	128
	15	0	0	0	48	24	256	128	128
	14	0	0	0	72	0	256	128	128
	13	0	0	0	0	0	256	128	128
	12	0	0	0	0	0	256	128	128
	11	0	0	0	0	0	256	128	128

10	0	0	0	0	0	256	128	128
9	0	0	0	0	0	256	128	128
8	0	0	0	0	0	256	128	128
7	0	0	0	0	0	256	128	128
6	0	0	0	0	0	256	128	128
5	0	0	0	0	0	256	128	128
4	0	0	0	0	0	256	128	128
3	0	0	0	0	0	256	128	128
2	24	0	24	0	0	256	128	128
1	24	0	24	0	0	256	128	128

0

## 7.3 Default Settings: ERS 5900

### Default Settings

- Precedence level 1, 2, 14, 15 and 16 are used by default
  - Precedence levels 1 and 2 are used by default untrusted if-group (allQoSPolicyIfcs)
    - Can be cleared up if ports are re-assigned to QoS class of either trusted or unrestricted
  - Precedence level 14 is used by DHCP Relay
    - Can be cleared up if IP DHCP Relay is disabled – on by default
  - Precedence levels 15 and 16 are permanently occupied
    - Precedence level 15 is used by SPB
    - Precedence level 16 is used by ARP
  - Certain features if enabled can take up one or more precedence levels
    - Please see Appendix for list of precedence levels used by various features
  - Up to 256 classifiers for each precedence for each ASIC
  - Up to 256 meters for each precedence for each ASIC, usable on a maximum of 8 out of the 16 available precedences
  - Up to 128 counters for each precedence for each ASIC
  - Up to 256 counters for each mask precedence for each ASIC
- Default Queue set is 2 with buffering set to large

Model	ASIC Device 1
5928GTS	Port 1–28
5928GTS-PWR+	Port 1–28
5952GTS	Port 1–52
5952GTS-PWR+	Port 1–52

## Verification

### 1 Verify QoS queue set and buffer:

5928GTS-PWR+#**show qos agent**

```
QoS Operational Mode: Enabled
QoS NVRam Commit Delay: 10 seconds
QoS Current Queue Set: 2
QoS Next Boot Queue Set: 2
QoS Current Buffering: Large
QoS Next Boot Buffering: Large
QoS UBP Support Level: Disabled
QoS Default Statistics Tracking: Aggregate
QoS DoS Attack Prevention: Disabled
    Minimum TCP Header Length: 20
    Maximum IPv4 ICMP Length: 512
    Maximum IPv6 ICMP Length: 512
Auto QoS Mode: Disabled
```

### 2 Verify QoS interface group and assignments

5928GTS-PWR+#**show qos if-group**

Role Combination	Interface Class	Capabilities	Statistics Tracking	Storage Type
allQoSPolicyIfcs	Untrusted	Input 802, Input IP Aggregate		ReadOnly
\$qosDisabledIfcs	Unrestricted	Input 802, Input IP Disabled		Other

5928GTS-PWR+#**show qos if-assign**

Unit	Port	IfIndex	Role Combination	Queue Set	Capability	DAPP Support
1	1	1	allQoSPolicyIfcs	2	Version 1,2	Yes
1	2	2	allQoSPolicyIfcs	2	Version 1,2	Yes
1	27	27	allQoSPolicyIfcs	2	Version 1,2	Yes
1	28	28	allQoSPolicyIfcs	2	Version 1,2	Yes

### 3 Verify QoS precedence levels and filters / meters used

5928GTS-PWR+#*show qos diag*

Unit/Port	Mask Precedence Usage															
	16	15	14	13	12	11	10	9	8	7	6	5	4	3	2	1
1/1	AR	SB	DH												Q	Q
1/2	AR	SB	DH												Q	Q
1/3	AR	SB	DH												Q	Q
1/27	AR	SB	DH												Q	Q
1/28	AR	SB	DH												Q	Q

AR=ARP DH=DHCP EA=EAP Q=QoS SB=SPB

Unit/Port	NonQoS NonQoS										RngChk
	Prec	Filter	Meter	Cntr	Filter	Meter	Filter	Meter	Cntr		
		Used	Used	Used	Used	Used	Total	Total	Total	Used	
1 /1 -28	16	0	0	0	29	29	256	256	128		
	15	0	0	0	28	0	256	256	128		
	14	0	0	0	112	28	256	256	128		
	13	0	0	0	3	0	256	256	128		
	12	0	0	0	0	0	256	256	128		
	11	0	0	0	0	0	256	256	128		
	10	0	0	0	0	0	256	256	128		
	9	0	0	0	0	0	256	256	128		
	8	0	0	0	0	0	256	256	128		
	7	0	0	0	0	0	256	256	128		
	6	0	0	0	0	0	256	256	128		
	5	0	0	0	0	0	256	256	128		
	4	0	0	0	0	0	256	256	128		
	3	0	0	0	0	0	256	256	128		
	2	28	0	28	0	0	256	256	128		
	1	28	0	28	0	0	256	256	128		

0 /32

## 7.4 Default Settings: VSP 7000

### Default Settings

- The VSP 7000 supports 10 precedences used by QoS and non-QoS
  - Precedence levels 1 to 7 are not used by default
  - Precedence level 8 is used by DHCP Relay
    - Can be cleared up if IP DHCP Relay is disabled – on by default
  - Precedence levels 9 and 10 are permanently occupied
    - Precedence level 9 is used by SPB
    - Precedence level 10 is used by ARP
  - Certain features if enabled can take up one or more precedence levels
    - Please see Appendix for list of precedence levels used by various features
  - Up to 128 filter components for precedences 1–4 and 256 filter components for precedences 5–10
  - Up to 128 meters per precedence per hardware device group, usable on a maximum of 8 out of the 10 available precedences
  - Up to 64 counters for precedences 1–4 and 128 counters for precedences 5–10.
  - Up to 32 TCP/UDP port range checkers
- Default Queue set is 2 with buffering set to large

### Verification

#### 1 Verify QoS queue set and buffer:

```
7024XLS#show qos agent
QoS Operational Mode: Enabled
QoS NVRam Commit Delay: 10 seconds
QoS Current Queue Set: 2
QoS Next Boot Queue Set: 2
QoS Current Buffering: Large
QoS Next Boot Buffering: Large
QoS Default Statistics Tracking: Aggregate
```

#### 2 Verify QoS interface group and assignments

```
7024XLS#show qos if-group
```

Role	Interface	Capabilities	Statistics	Storage
Combination	Class		Tracking	Type
allQoSPolicyIfcs	Trusted	Input 802, Input IP	Aggregate	ReadOnly
\$qosDisabledIfcs	Unrestricted	Input 802, Input IP	Disabled	Other

```
7024XLS#show qos if-assign
```

```
Unit Port IfIndex Role Combination Queue Set Capability
-----
1 1 1 allQoSPolicyIfcs 2 Version 1,2
1 2 2 allQoSPolicyIfcs 2 Version 1,2
|
|
1 31 31 allQoSPolicyIfcs 2 Version 1,2
1 32 32 allQoSPolicyIfcs 2 Version 1,2
```

### 3 Verify QoS precedence levels and filters / meters used

```
7024XLS#show qos diag
```

```
Unit/Port          Mask Precedence Usage
                   10  9  8  7  6  5  4  3  2  1
-----
1/1                AR  SB  DH
1/2                AR  SB  DH
|
|
1/31               AR  SB  DH
1/32               AR  SB  DH
AR=ARP DH=DHCP SB=SPB
```

```

NonQoS NonQoS
Unit/Port  Prec Filter Meter Cntr Filter Meter Filter Meter Cntr RngChk
          Used  Used  Used  Used  Used  Total  Total  Total  Used
-----
1 /1 -32  10   0    0    0   32   32   256   256   128
          9    0    0    0   32    0   256   256   128
          8    0    0    0  128   32   256   256   128
          7    0    0    0    0    0   256   256   128
          6    0    0    0    0    0   256   256   128
          5    0    0    0    0    0   256   256   128
          4    0    0    0    0    0   128   256    64
          3    0    0    0    0    0   128   256    64
          2    0    0    0    0    0   128   256    64
          1    0    0    0    0    0   128   256    64
```

0 / 32



## 7.5 QoS Precedence Levels – Masks used by various applications

Masks	Rules	Application	Code
1	1	Broadcast ARP / ARP Inspection	AR
2	2	QoS (default untrusted policy)	Q
2	10	IGMP	na
1	2	Port Mirroring	PM
1	2	DHCP Replay / DHCP Snooping / NSNA DHCP	DH
1	32	BaySecure (port)	na
1	10	BaySecure (MAC-DA)	BS
1	32	EAP MHMA Allowed Cients	EA
1	1	IPFIX	IF
1	32	NSNA MAC	NS
5	8	NSNA (R/Y/G filters)	Q
1	1	ADAC	AD
1	1	RIP	RI
1	1	UDP Bcast	UB
1	3	VRRP	DR
1	3	OSPF	DR
1	10	IP Source Guard	IS
1	3	PIM	DR
1	1	PIM	PI
3	1	Content-based Forwarding	CF

## 8. Cisco to Avaya Comparison

This section compares various Cisco IOS ACL commands into their Avaya equivalents.

Config Task Description	Cisco IOS	Avaya ACLI
Deny src-ip 192.168.1.0/24 Allow everyone else	access-list <1> deny <192.168.1.0> <0.0.0.255> access-list <1> permit any	qos traffic-profile classifier name <one> src-ip <192.168.1.0/24> drop-action <enable> block <one>
Apply ACLs to interfaces	interface range <gi1/0/1 - 20> ip access-group <1> in exit	qos traffic-profile set port <1-20> name <one> show qos traffic-profile set
Permit SMTP connection from any host to mail server	access-list <100> permit tcp any host <10.1.1.1> eq smtp	qos traffic-profile classifier name <one> dst-ip <10.1.1.1/32> protocol <6> dst-port-min <25> dst- port-max <25> block <one>
Permit TCP connections to port range between 1024 and 2024 to subnet 10.1.1.0/24	access-list <100> permit tcp any <10.1.1.0> <0.0.0.255> range <1024> <2024>	qos traffic-profile classifier name <one> src-ip 10.1.1.1/32> protocol <6> src-port-min <1024> src-port-max <2048> block <one>
Permit UDP connections to port range between 1024 and 2024 to subnet 10.1.1.0/24	access-list <100> permit udp any <10.1.1.0> <0.0.0.255> range <1024> <2024>	qos traffic-profile classifier name <one> src-ip 10.1.1.1/32> protocol <17> src-port-min <1024> src-port-max <2048> block <one>
Bootps	10 permit udp any any eq bootps	qos traffic-profile classifier name <one> addr-type <ipv4> protocol <17> dst-port-min <67> dst-port- max <67> block one
Bootpc	20 permit udp any any eq bootpc	qos traffic-profile classifier name <one> addr-type <ipv4> protocol <17> dst-port-min <68> dst-port- max <68> block one
IP Range	30 deny ip 10.57.240.0 0.0.1.255 host 128.1.0.1	qos traffic-profile classifier name <one> src-ip <10.57.240.0/24> dst-ip <128.1.0.2/32> drop- action enable block <one>  qos traffic-profile classifier name <one> src-ip <10.57.241.0/24> dst-ip <128.1.0.2/32> drop- action enable block <one>
IP Range	110 permit ip host 10.57.240.1 10.0.0.0 0.255.255.255	qos traffic-profile classifier name <one> src-ip <10.57.240.1/32> dst-ip <10.0.0.0/8> block <one>

## 9. Reference Documentation

Document Title	Publication Number	Description
Shortest Path Bridging (802.1aq) for ERS 8800 and VSP 9000 Technical Configuration Guide	NN48500-617	
Basic SPBM Configuration	NN48500-632	
Migrating to a Virtual Services Fabric using Shortest Path Bridging Technical Configuration Guide	NN48500-622	
Avaya Virtual Services Platform 7000 Series Configuration - SPBM	NN46202-510	
Configuring Quality of Service on Avaya Ethernet Routing Switch 5900 Series	NN47211-504	

© 2015 Avaya Inc. All Rights Reserved.

Avaya and the Avaya Logo are trademarks of Avaya Inc. and are registered in the United States and other countries. All trademarks identified by ®, TM or SM are registered marks, trademarks, and service marks, respectively, of Avaya Inc. All other trademarks are the property of their respective owners. Avaya may also have trademark rights in other terms used herein. References to Avaya include the Nortel Enterprise business, which was acquired as of December 18, 2009.