



Technical Notes

QoS Features on the Business Ethernet Switch 50 (BES50)

Version:	NN70000-004 issue 1.00
Date:	February 3 rd , 2009
Status:	Released

Copyright © 2009 Nortel Networks. All rights reserved.

The information in this document is subject to change without notice. The statements, configurations, technical data, and recommendations in this document are believed to be accurate and reliable, but are presented without express or implied warranty. Users must take full responsibility for their applications of any products specified in this document. The information in this document is the property of Nortel Networks.

Revision History

Revision (Ver.Iss)	Issue Date	Document Status	Change Description
1.00	2009-02-03	Released	First released version.

About this Document

This application note describes the Quality Of Service (QoS) capability of BES50 family of switches. This document also outlines the procedure to properly configure different QoS features of BES50.

The document is intended for people responsible for deployment and support of the Nortel Business Ethernet Switch 50. This may include:

- Customer System Engineering
- Channel Partner Engineering
- Nortel Solutions Engineering
- Knowledge Systems / Education Services

Table of Contents

Revision History	2
About this Document.....	3
Table of Contents	4
Table of Figures.....	5
Overview.....	6
High-level Description of QoS Treatment in BES50	7
Configure BES50 Traffic Prioritization.....	14
1. Default Port Priority.....	14
2. 802.1p User Priority	14
3. Traffic prioritization based on IP DSCP	15
4. ACL Based Traffic prioritization	17
BES50 Re-Marking Capability	29
IP DSCP Re-Marking on BES50FE	29
IP DSCP Re-Marking on BES50GE	29
Difference Between BES50FE and BES50GE.....	32
1. Number of traffic class.....	32
2. Different ways to configure ACL Based traffic prioritization	32
3. Location of frame/packet re-marking.....	32
4. IP DSCP re-marking.....	32
Default QoS Behaviour of BES50FE and BES50GE	33
Default QoS Behaviour on BCM Traffic.....	33
Software Baseline	35

Table of Figures

Figure 1: BES50FE Priority to Traffic Class Mapping	8
Figure 2: BES50GE Priority to Traffic Class Mapping	8
Figure 3: Scheduler Queuing Mode	9
Figure 4: BES50FE WRR Weight Settings	10
Figure 5: BES50GE WRR Weight Settings.....	10
Figure 6: BES50FE Traffic Prioritization Flow Chart	12
Figure 7: BES50GE Traffic Prioritization Flow Chart	13
Figure 8: BES50 Default Port Priority Settings.....	14
Figure 9: BES50 IP DSCP Priority Status	16
Figure 10: BES50 DSCP to COS Mapping	17
Figure 11: BES50FE Extended ACL Rule.....	19
Figure 12: BES50FE Class Map	20
Figure 13: BES50FE Policy Setting CoS Value	22
Figure 14: BES50FE Policy Setting DSCP Value	23
Figure 15: BES50FE Apply QoS Policy to Ingress Port.....	24
Figure 16: BES50GE Extended ACL Rule	26
Figure 17: BES50GE ACL Priority Mapping.....	28
Figure 18: BES50GE IP DSCP Re-Marking Status	29
Figure 19: BES50GE Priority Index to IP DSCP Mapping	30

Overview

BES50 provides a number of features to support different Quality of Service (QoS) for different types of traffic. There are multiple levels of traffic prioritization on BES50 (both BES50FE and BES50GE). Prioritization can be based on

- (1) Default port priority
- (2) 802.1p value of the incoming tagged Ethernet frame
- (3) IP DSCP of the incoming IP packet
- (4) ACL based (using the 5-tuples in IP header: source and destination IP addresses, source and destination ports and protocol number)

BES50 has pre-defined precedence (or order) when applying these mechanisms. ACL based prioritization has the highest precedence over the others, and default port priority has the lowest precedence if nothing else has been configured.

When an incoming Ethernet frame arrives at an ingress port, BES50 first determines the traffic priority/CoS (value between 0 – 7) based on the L2/L3 information carried in the Ethernet frame and/or the IP header (IP DSCP, IP addresses, port numbers and protocol number).

Each priority/CoS (basically 802.1p bit value) maps into one of the traffic class (4 classes for BES50FE and 8 classes for BES50GE). These traffic classes are served by a scheduler that can run in either Weighted Round Robin (WRR) or Strict Priority.

BES50 also provides the capability to mark or re-mark the 802.1p value (CoS) or the DSCP value of the egress Ethernet frames and/or IP packets. The following table summarizes these two features:

	Traffic Prioritization	Egress (re) marking		
		CoS/1p (tagged Ethernet frames)	DSCP (IP packets)	
			BES50GE	BES50FE
Default port priority	Yes	Yes	Yes (optional)	No
802.1p bit (tagged Ethernet frame)	Yes	Yes	Yes (optional)	No
IP DSCP	Yes	Yes	Yes (optional)	No
ACL based	Yes	Yes	Yes (optional)	Yes (if the applied policy has an associated action to set IP DSCP value)

High-level Description of QoS Treatment in BES50

A central element for providing different QoS treatment is the concept of traffic priority and traffic classes. When an incoming Ethernet frame arrives at an ingress port, BES50 first determines the priority/CoS (value between 0 – 7) of the incoming frame based on the L2/L3 information carried in the Ethernet frame and/or the IP payload. Each priority/CoS (basically 802.1p bit value) maps into one of the traffic class (4 classes for BES50FE and 8 classes for BES50GE).

BES50 provides a default priority to traffic class mapping. The priority to traffic mapping for BES50FE and BES50GE are depicted in the following table.

Priority/CoS	Traffic Class	
	BES50FE	BES50GE
0	1	2
1	0	0
2	0	1
3	1	3
4	2	4
5	2	5
6	3	6
7	3	7

Table 1: BES50 Priority to Traffic Class Mapping

This mapping is also user configurable if the user desired to use a different setting. To review and/or configure the priority to traffic class mapping, navigate to the Applications > Priority > Traffic Classes panel

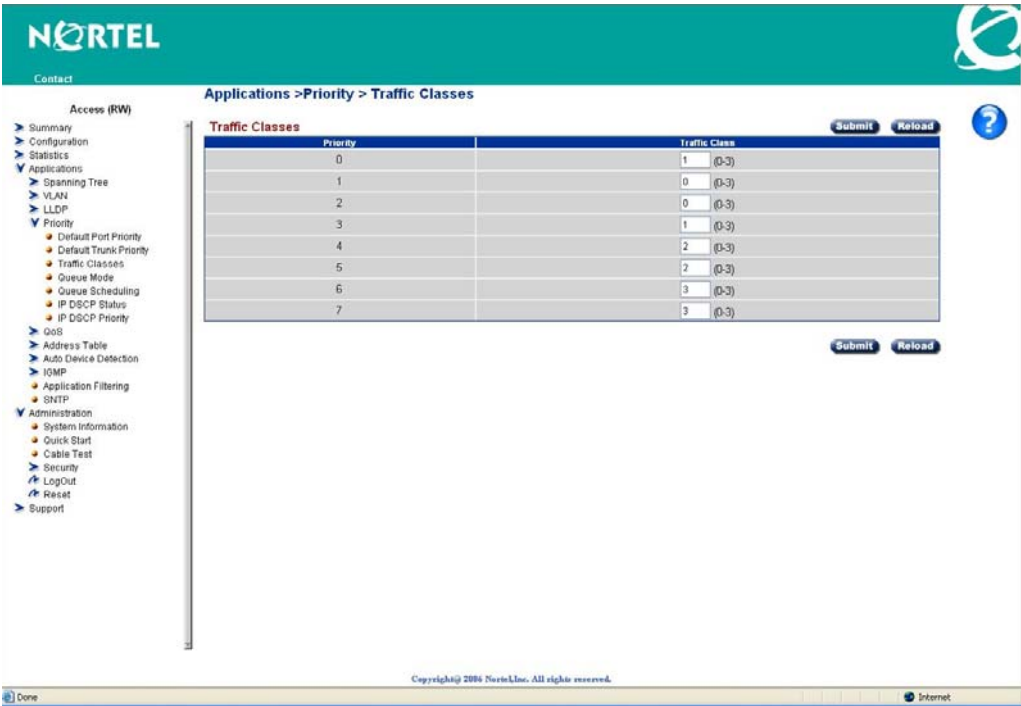


Figure 1: BES50FE Priority to Traffic Class Mapping

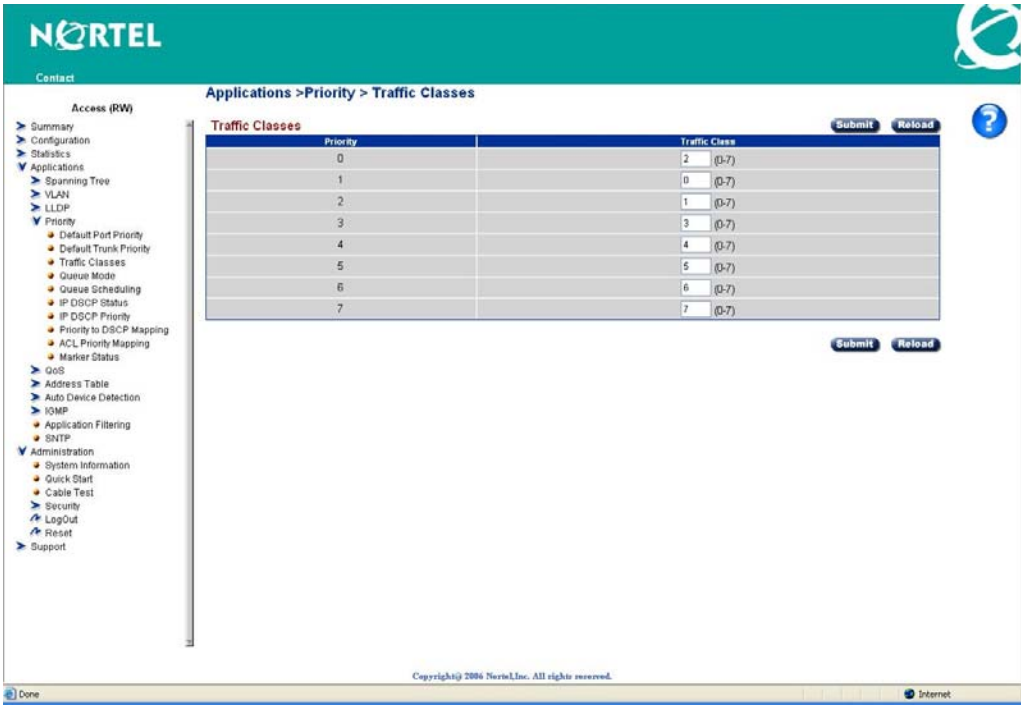


Figure 2: BES50GE Priority to Traffic Class Mapping

These traffic classes are served by a scheduler that can run in either Weighted Round Robin (WRR) or Strict Priority.

- Strict Priority queuing operates in a pre-emptive fashion. Frames in the higher priority queues always take precedence over frames in lower priority queues. A low priority queue will not be serviced as long as there is a frame in one of the higher priority queues. When the highest priority queue is empty, frames from the other queues are serviced based on their precedence (i.e. from the higher priority queue to the lower priority queue). With strict priority queuing, it is possible that some queues may never be serviced, causing frames dropped and starvation. Use of Strict Priority queuing must be used with great care.
- With Weighted Round Robin queuing, each queue is assigned a weight which represents a relative proportion of time during which the queue can send packets. This technique ensures each queue gets appropriate share of bandwidth for transmitting its packets. With Weighted Round Robin, no priority is assigned to the queues. Each queue sends frames in proportion to its configured weight.

BES50 uses WRR queuing by default. If the user decided to change the scheduling algorithm, navigate to the Applications > Priority > Queue Mode panel

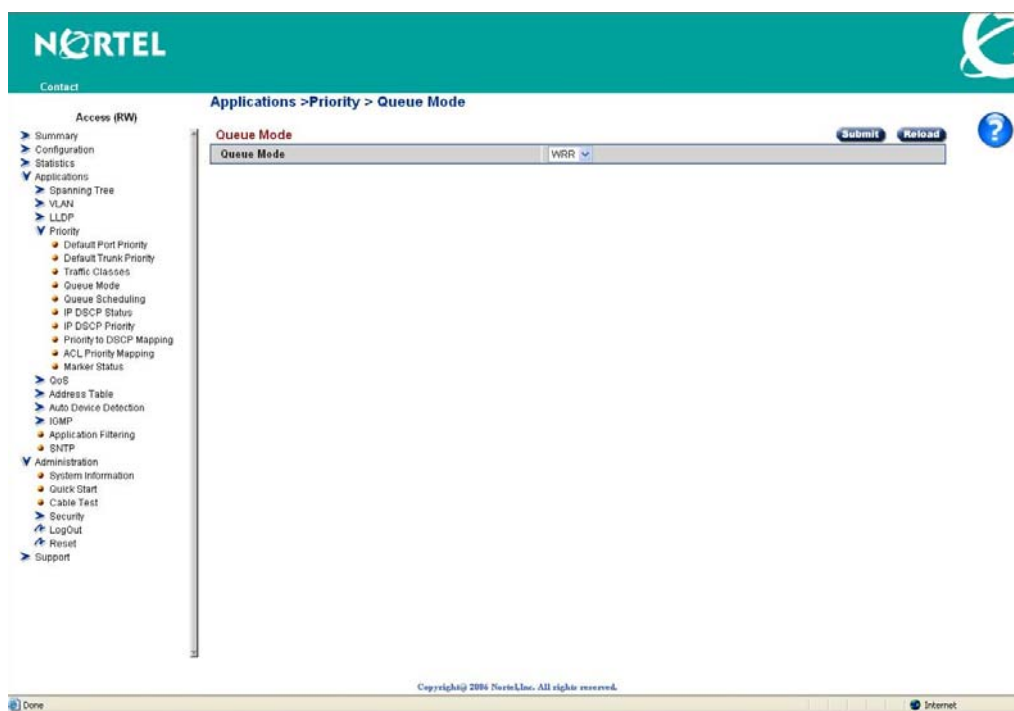


Figure 3: Scheduler Queuing Mode

When running in WRR, the weight assigned to each class are pre-defined and not user configurable. To view the assigned weight for each class, navigate to the Applications > Priority > Queue Scheduling panel

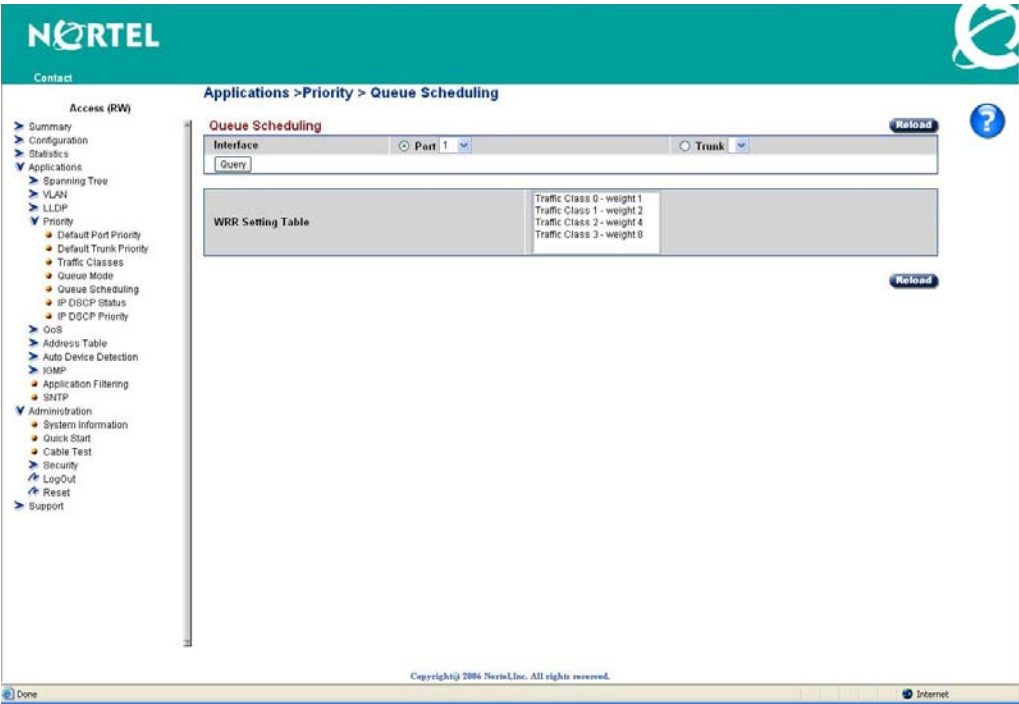


Figure 4: BES50FE WRR Weight Settings

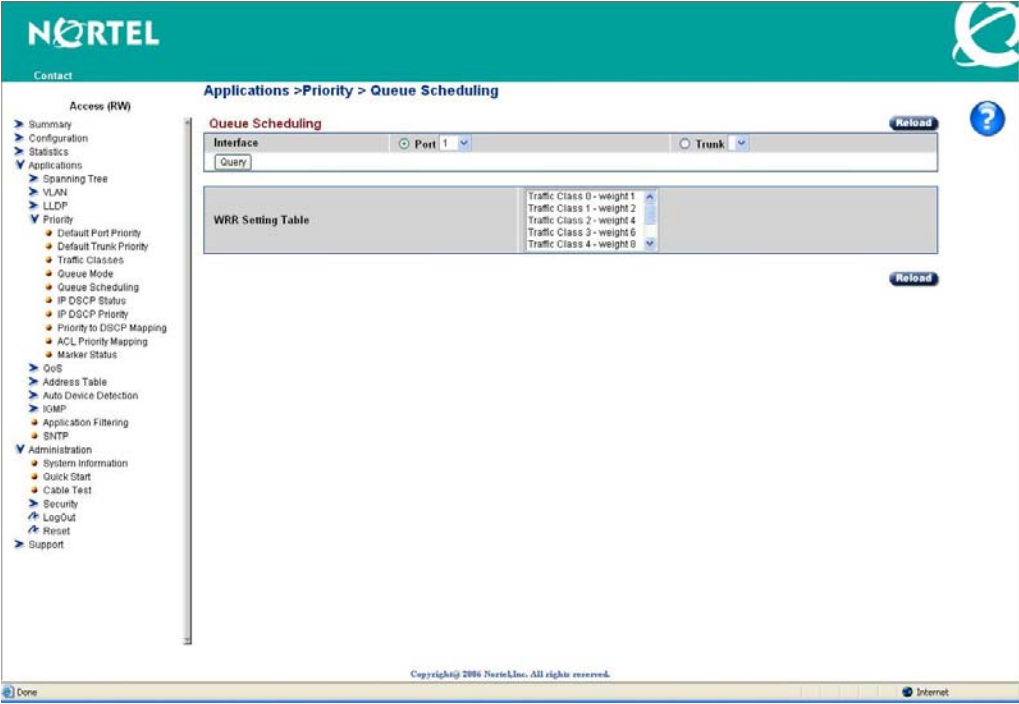


Figure 5: BES50GE WRR Weight Settings

BES50 provides multiple ways to classify different types of traffic and map the traffic to one of the 8 defined priority/CoS values. There are four different methods to classify and prioritize incoming traffic, these methods are based on:

1. Default port priority
2. 802.1p value of the incoming tagged Ethernet frame
3. IP DSCP of the incoming IP packet
4. ACL based (using the 5-tuples in IP header: source and destination IP addresses, source and destination ports and protocol number)

Traffic priority is determined by applying these mechanisms in a pre-defined precedence. ACL based prioritization has the highest precedence over the others, and default port priority has the lowest precedence.

- If ACL based prioritization is enabled and the ingress traffic matches one of the defined ACL rules, the CoS value is determined by the action associated with the configured ACL.
- If ACL based prioritization is disabled or the ingress traffic does not match any of the ACL rules, then the ingress traffic is subjected to IP DSCP based prioritization if this feature is enabled (IP DSCP priority is disabled by default). If the ingress Ethernet frame encapsulates an IP packet, BES50 uses the DSCP value contained in the IP header to determine the CoS value for the ingress traffic.
- If IP DSCP based prioritization is disabled or the ingress Ethernet frame does not encapsulate an IP packet, BES50 extracts the 802.1p value from the tagged ingress Ethernet frame to determine the appropriate CoS.
- If the ingress Ethernet frame is untagged, BES50 eventually use the user configurable default port priority to determine the CoS for the incoming traffic.

Even though both BES50FE and BES50GE follow similar framework to prioritize incoming traffic, there are minor difference in the way these mechanisms are implemented.

- BES50FE maps traffic directly into a CoS/priority value, which in turns identifies the traffic class which will be used for the frame or packet.
- BES50GE first maps traffic into a priority index, which in turns maps into a CoS/priority value to identify the traffic class which will be used for the frame or packet.

The following two figures illustrate the flow chart of how traffic prioritization is achieved in both BES50FE and BES50GE.

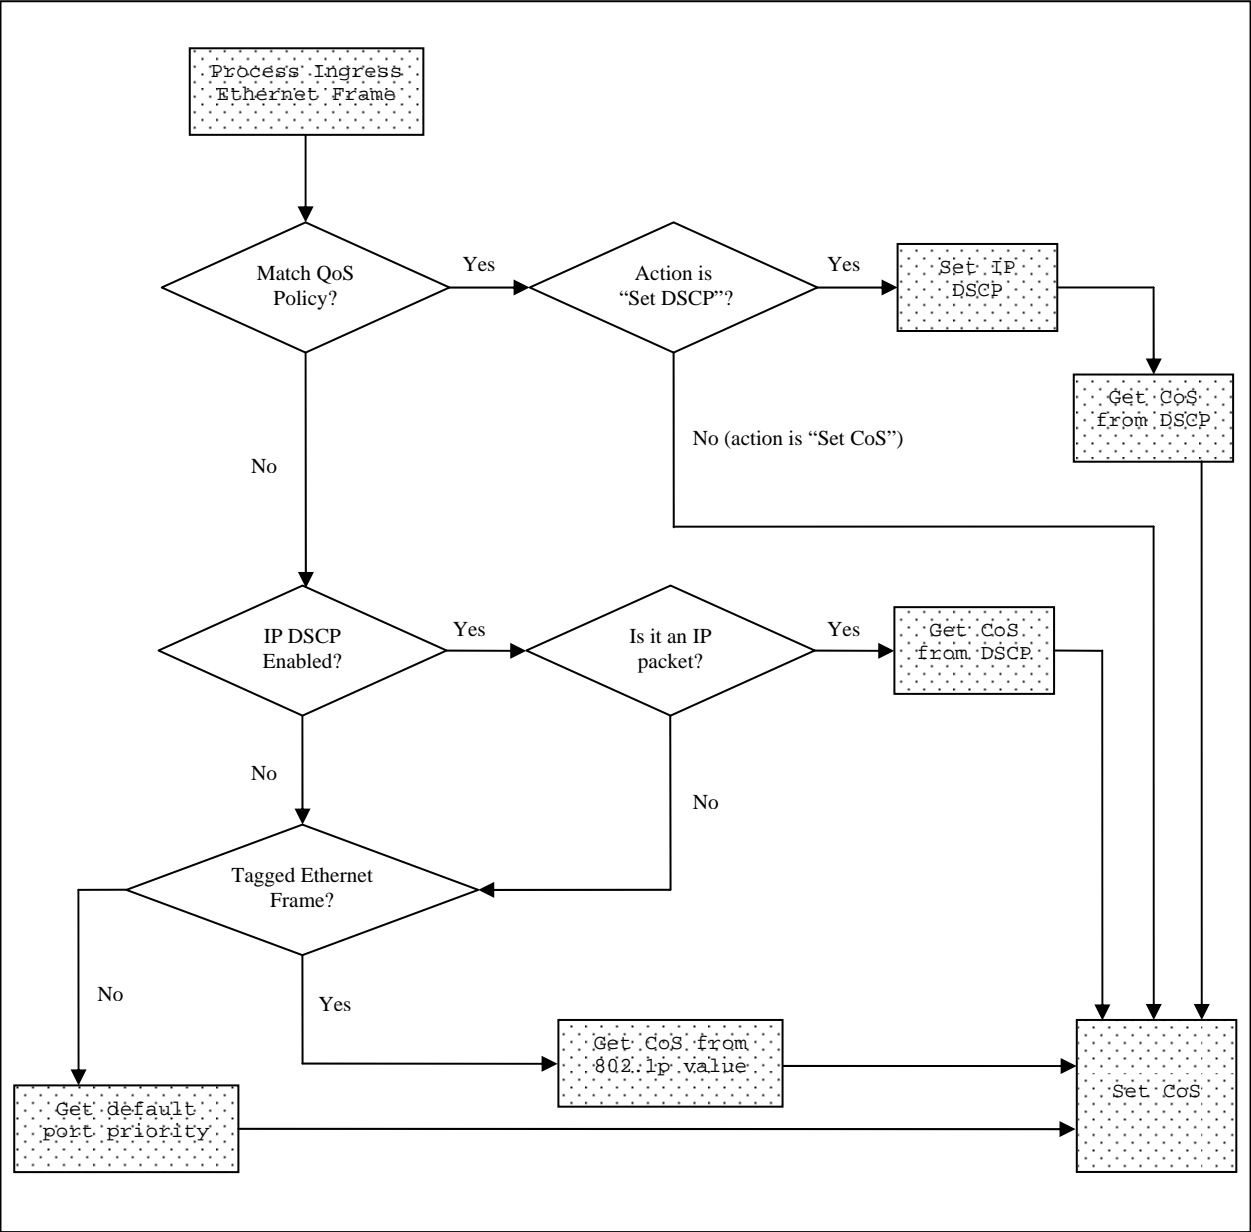


Figure 6: BES50FE Traffic Prioritization Flow Chart

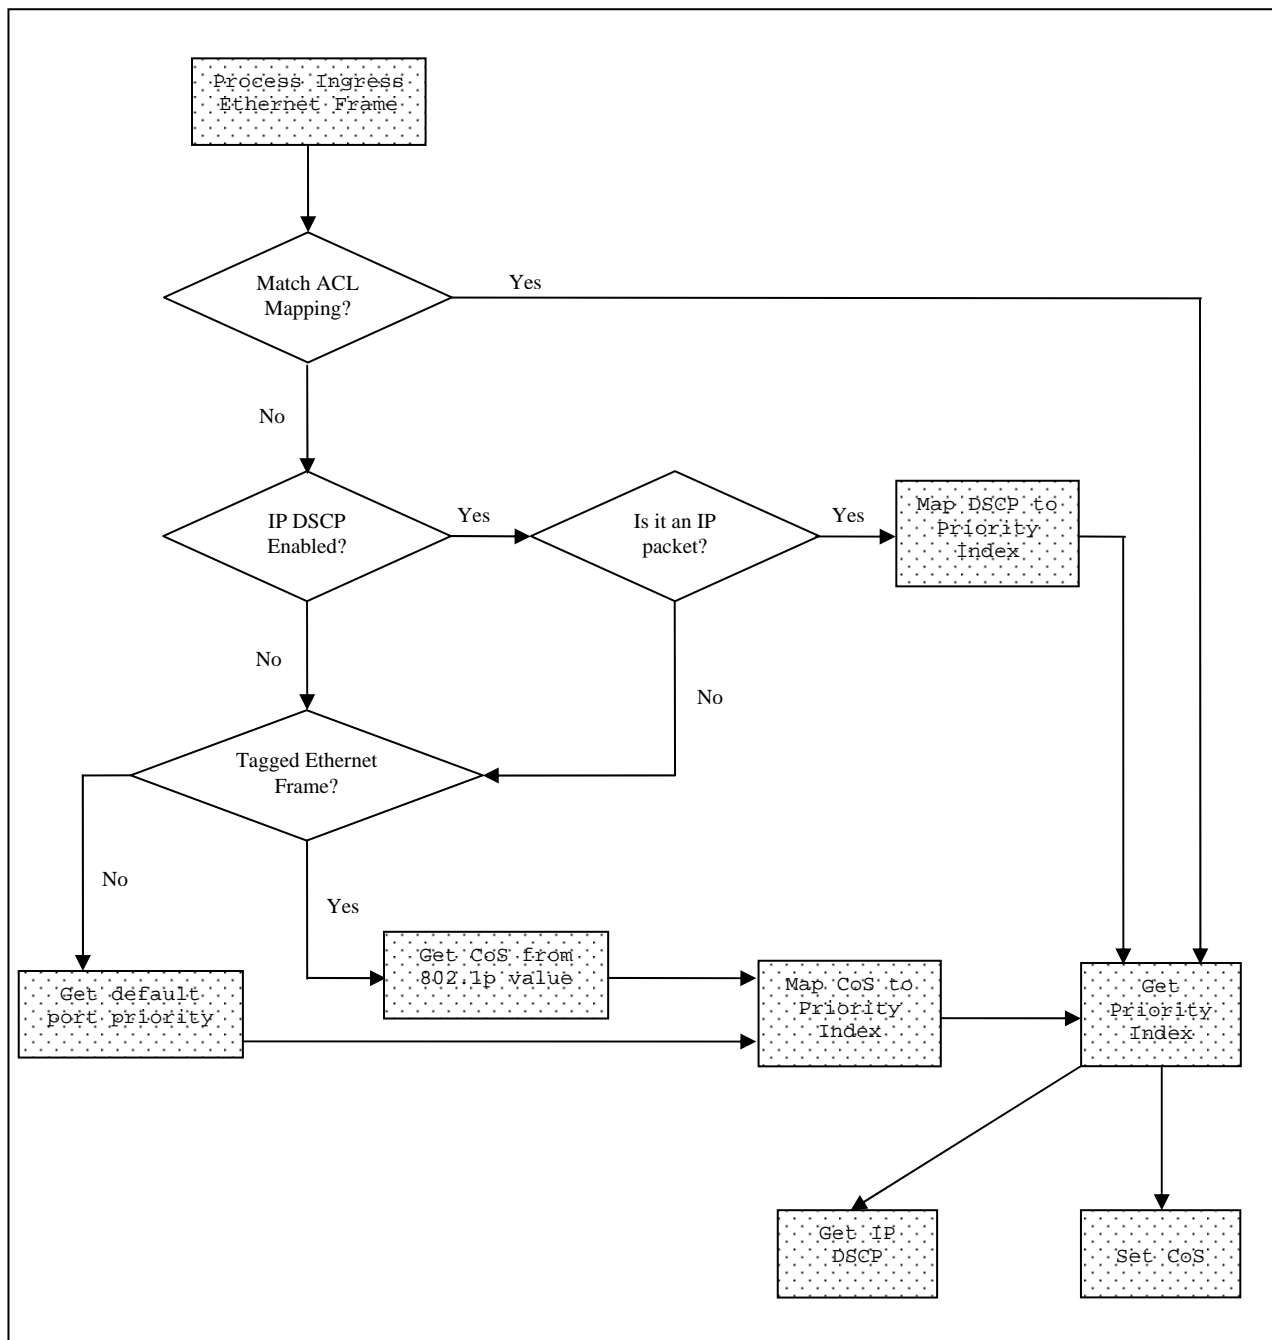


Figure 7: BES50GE Traffic Prioritization Flow Chart

The resulting priority/CoS (basically 802.1p bit value) maps into one of the traffic class (4 classes for BES50FE and 8 classes for BES50GE). These traffic classes are served by a scheduler that can run in either Weighted Round Robin (WRR) or Strict Priority.

To review and/or configure the traffic priority to traffic class mapping, navigate to the Applications > Priority > Traffic Classes panel (see “Figure 1: BES50FE Traffic Classes” and “Figure 2: BES50GE Traffic Classes”).

Configure BES50 Traffic Prioritization

1. Default Port Priority

A default port priority can be specified for each port on the BES50. All untagged packets entering the BES50 are tagged with the specified default port priority, and then sorted into the appropriate priority queue at the output port.

The default priority applies for an untagged frame received on a port set to accept all frame types (i.e., can receive both untagged and tagged frames). The default port priority does not apply to IEEE 802.1Q VLAN tagged frames. If the incoming frame is an IEEE 802.1Q VLAN tagged frame, the IEEE 802.1p User Priority bits contained in the Ethernet frame will be used for traffic prioritization. If the egress port is an untagged member of the associated VLAN, the egress frames are stripped of all VLAN tags prior to transmission.

To review and/or modify the default port priority settings, navigate to the Applications > Priority > Default Port Priority panel.

Port	Default Priority (0-7)	Number of Egress Traffic Classes	Trunk
1	0	4	
2	0	4	
3	0	4	
4	0	4	
5	0	4	
6	0	4	
7	0	4	
8	0	4	
9	0	4	
10	0	4	
11	0	4	
12	0	4	

Figure 8: BES50 Default Port Priority Settings

2. 802.1p User Priority

BES50 can use the IEEE 802.1p user priority embedded in the incoming tagged Ethernet frames to determine the appropriate user traffic priority. Up to 8 traffic priorities are defined in IEEE 802.1D. BES50 provides a default user priority to traffic class mapping in accordance with Annex G of the IEEE 802.1D (keep in mind there are 4 traffic classes for BES50FE and 8 traffic classes for BES50GE).

This mapping is also user configurable if the user desired to use a different mapping. To review and/or configure the traffic priority to traffic class mapping, navigate to the Applications > Priority > Traffic Classes panel (see Figure 1: BES50FE Priority to Traffic Class and Figure 2: BES50GE).

3. Traffic prioritization based on IP DSCP

This section focuses on prioritizing traffic based on the IP DSCP value in the incoming IP packets. In brief, the user first needs to enable IP DSCP prioritization (it is disabled by default). When an incoming IP packet arrives at an ingress port, BES50 first maps the DSCP value in the incoming IP packet to a traffic priority/CoS (value between 0 – 7). Each traffic priority/CoS is mapped to one of the traffic classes, each traffic class will be serviced by the scheduler.

Keep in mind that the IP DSCP field of an IP packet is 6 bits while the 802.1p field of an Ethernet frame is only 3 bits, there are a lot more IP DSCP value combination than the 802.1p value. BES50 provides a default DSCP to CoS mapping as depicted in the following table. The DSCP to CoS mapping is also user configurable if the user desired to use a different setting.

IP DSCP Value	CoS Value
0	0
8	1
10, 12, 14, 16	2
18, 20, 22, 24	3
26, 28, 30, 32, 34, 36	4
38, 40, 42	5
48	6
46, 56	7

Table 2: BES50 IP DSCP to CoS Mapping

E.g.

- DSCP value of 40 or 0x28 maps into CoS 5, and then maps into traffic class 2 in BES50FE or traffic class 5 in BES50GE
- DSCP value of 46 or 0x2e maps into CoS 7, and then maps into traffic class 3 in BES50FE or traffic class 7 in BES50GE.

NOTE:

When IP DSCP prioritization is enabled, in the event that the incoming Ethernet frame is also tagged with an 802.1p value, BES50 will use the DSCP value in the IP packet for traffic prioritization instead of the 802.1p value in the Ethernet frame.

Step (1) Enable IP DSCP Prioritization

BES50 has the capability to prioritize traffic based on DSCP of the incoming IP packet. By default this feature is disabled, the user has to enable it if desired.

Each DSCP is mapped into a CoS/priority, which in turns in mapped into one of the traffic class.

Navigate to the Applications > Priority > IP DSCP Status panel

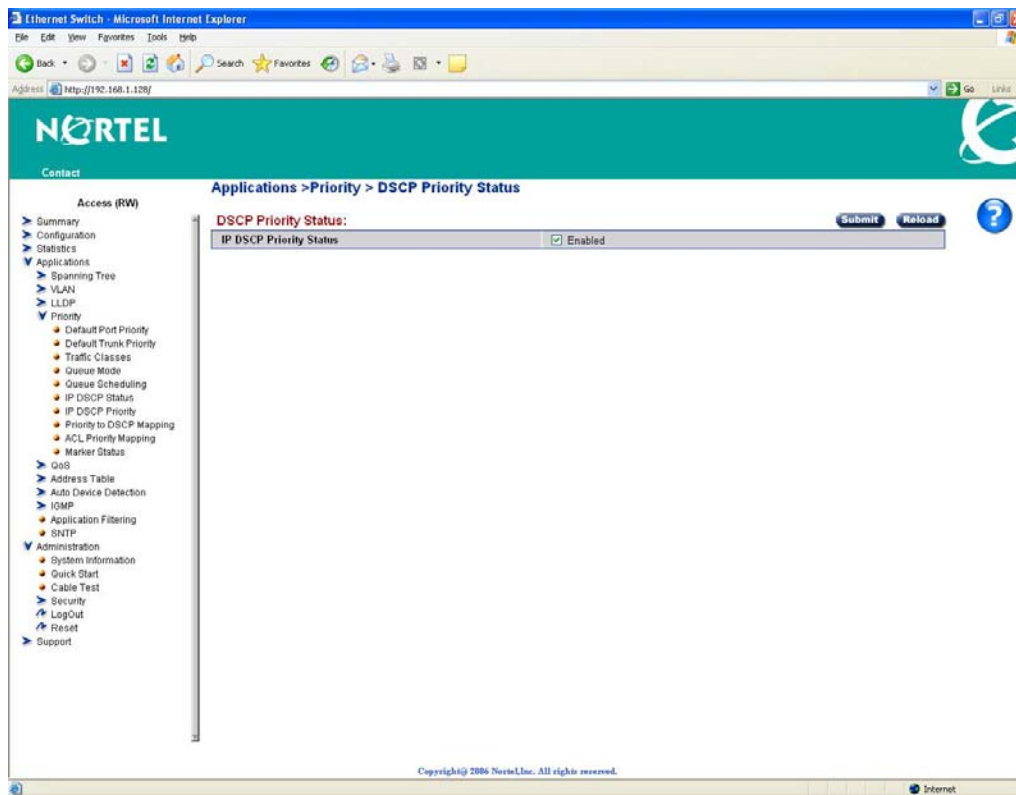


Figure 9: BES50 IP DSCP Priority Status

Step (2) Configure/Check DSCP to CoS mapping

BES50 has a predefined DSCP to CoS mapping. However the users can still modify the mapping if they desired. Navigate to the Applications > Priority > IP DSCP Priority panel

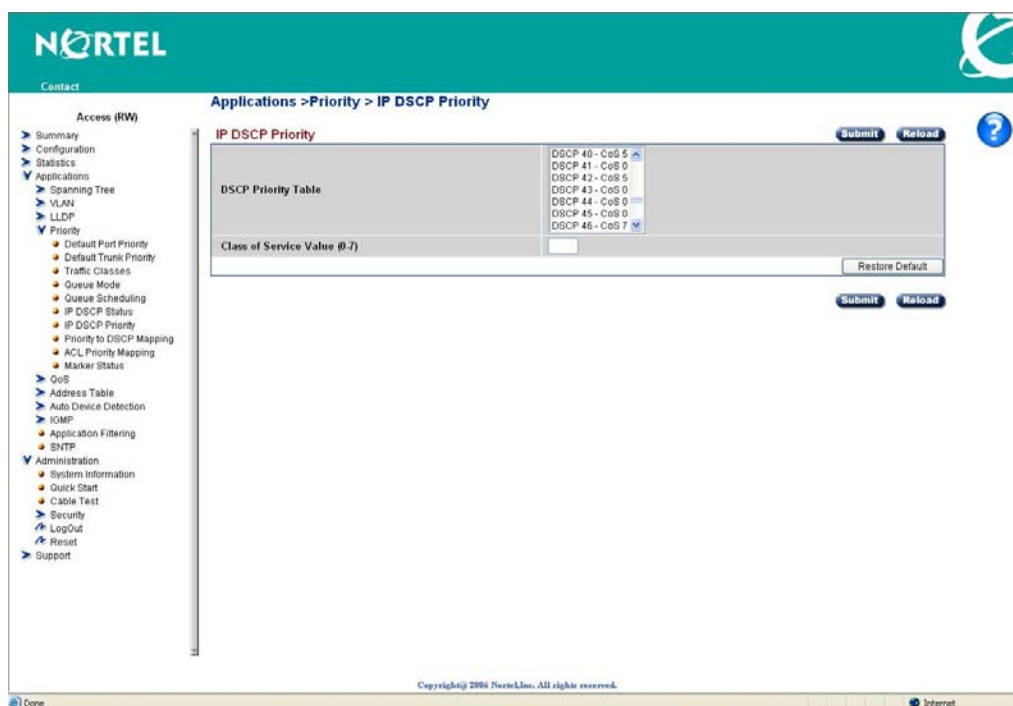


Figure 10: BES50 DSCP to COS Mapping

4. ACL Based Traffic prioritization

Using the ACL classification rules, BES50 can also prioritize traffic based on the following IP header information of the incoming IP packets:

- Source IP address
- Destination IP address
- Source port number
- Destination port number
- Protocol ID

Even though both BES50FE and BES50GE support traffic prioritization based on these five tuples in the IP header, the configuration process on these two units are different from each other.

BES50FE Configuration

Step (1) Define the traffic classifier in an extended Access Control List (ACL)

This step describes the classification criteria to identify a subset of incoming traffic for the desired QoS treatment. The criteria are based on the five tuples in the IP header.

NOTE:

The action associated with the ACL rule MUST be set to “Permit”.

NORTEL

Contact

Access (RW)

Summary

Configuration

Statistics

Applications

Spanning Tree

VLAN

LLDP

Priority

QoS

Address Table

Auto Device Detection

IGMP

Application Filtering

SNTP

Administration

System Information

Quick Start

Cable Test

Security

Users Accounts

Authentication Settings

HTTP Settings

Port Security

802.1X

ACL

ACL Configuration

Port Binding

LogOut

Reset

Support

Administration > Security > ACL Configuration

ACL Configuration

SubmitReload

Type	Name	Remove	Edit
Add ACL			
Name	ACL001		
Type	Extended		

SubmitReload

Copyright © 2006 Nortel Inc. All rights reserved.

DoneInternet

NORTEL

Contact

Access (RW)

Summary

Configuration

Statistics

Applications

Spanning Tree

VLAN

LLDP

Priority

QoS

Address Table

Auto Device Detection

IGMP

Application Filtering

SNTP

Administration

System Information

Quick Start

Cable Test

Security

Users Accounts

Authentication Settings

HTTP Settings

Port Security

802.1X

ACL

ACL Configuration

Port Binding

LogOut

Reset

Support

Administration > Security > ACL Configuration - Extended ACL

Extended ACL

SubmitReloadBack

Action	Source IP Address/ Mask	Destination IP Address/ Mask	Prot- ocol	Source Port/ Bit Mask	Destination Port/ Bit Mask	Remove
Add ACE						
Action	Permit			Source Port (0-65535)		
Source Address Type	Any			Source Port Bit Mask (0-65535)		
Source IP Address	0.0.0.0			Destination Port (0-65535)		
Source Subnet Mask	0.0.0.0			Destination Port Bit Mask (0-65535)		
Destination Address Type	Any					
Destination IP Address	0.0.0.0					
Destination Subnet Mask	0.0.0.0					
Protocol	<input type="radio"/> TCP (6) <input type="radio"/> UDP (17) <input checked="" type="radio"/> Others 1					

SubmitReloadBack

Copyright © 2006 Nortel Inc. All rights reserved.

DoneInternet

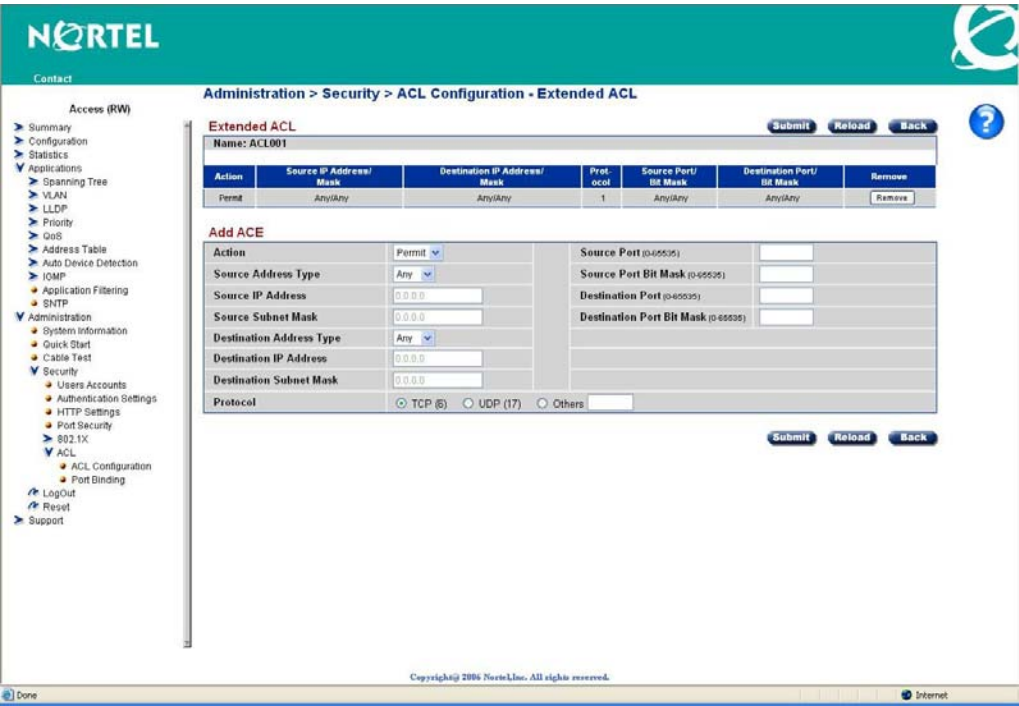
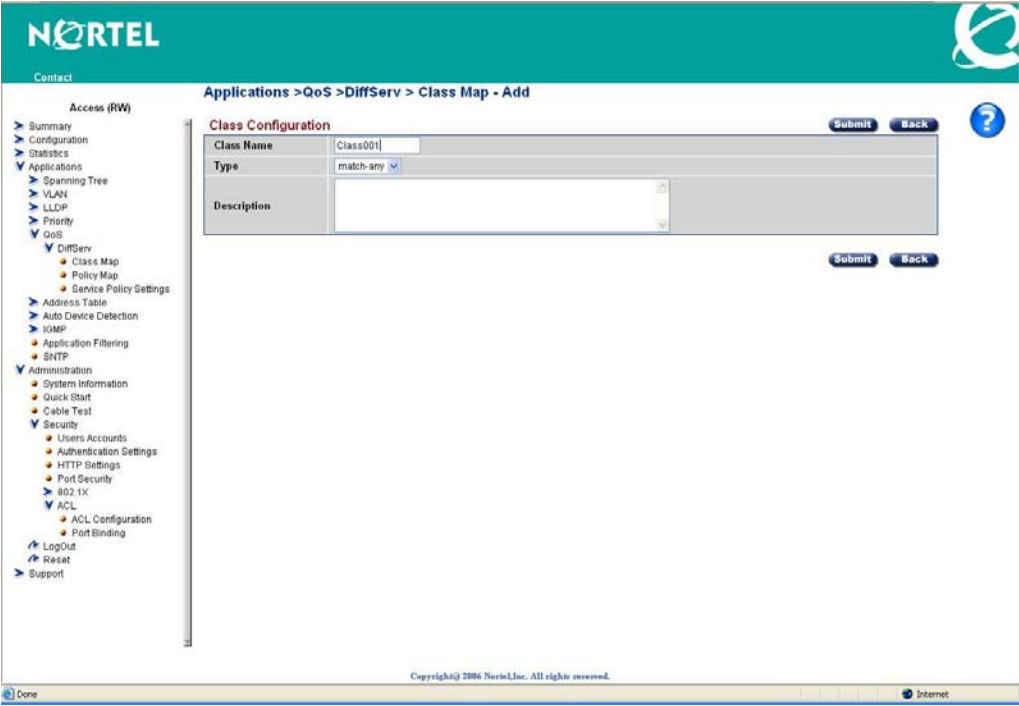


Figure 11: BES50FE Extended ACL Rule

Step (2) Define a class map that matches the classification of the ACL

This step identifies the matching traffic as a traffic stream. All packets within the same traffic stream are subject to the same QoS policy.



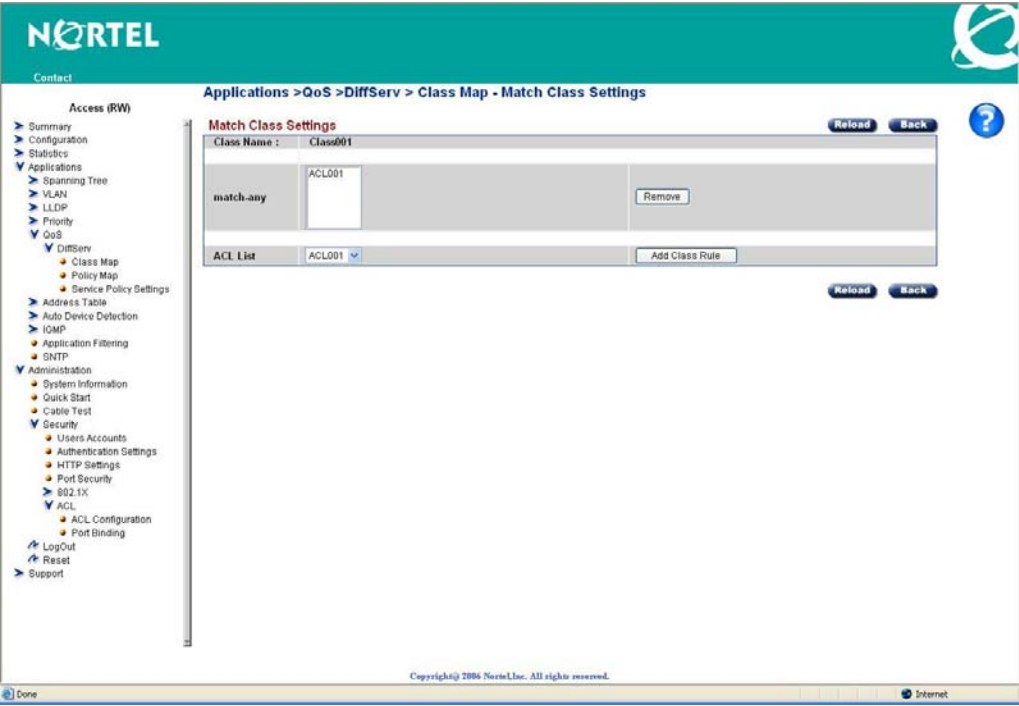
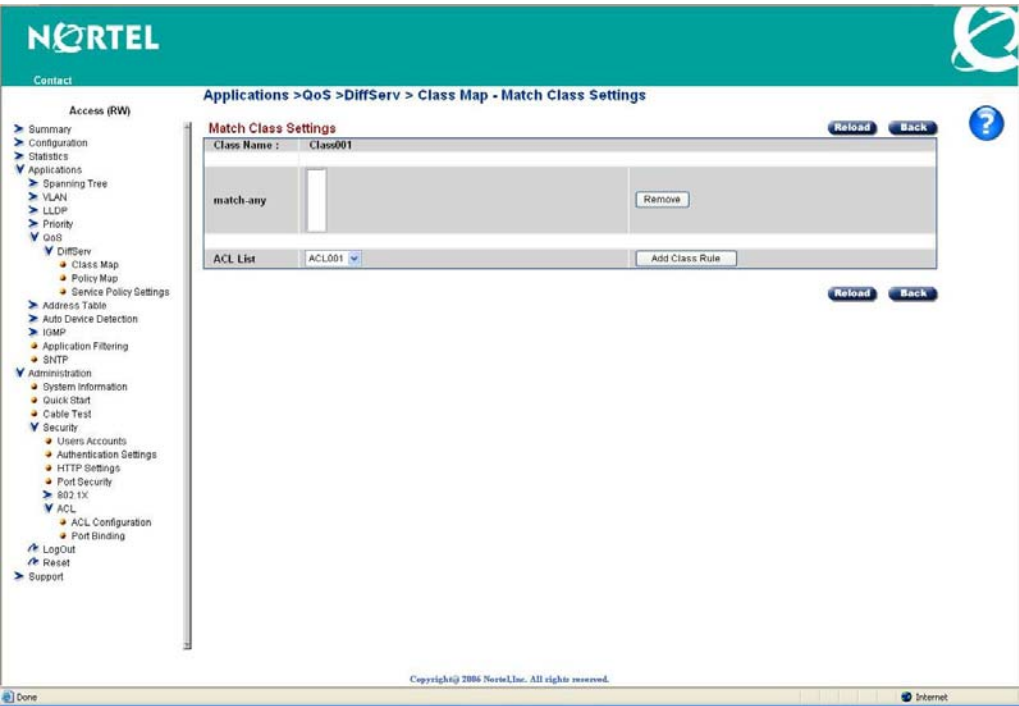


Figure 12: BES50FE Class Map

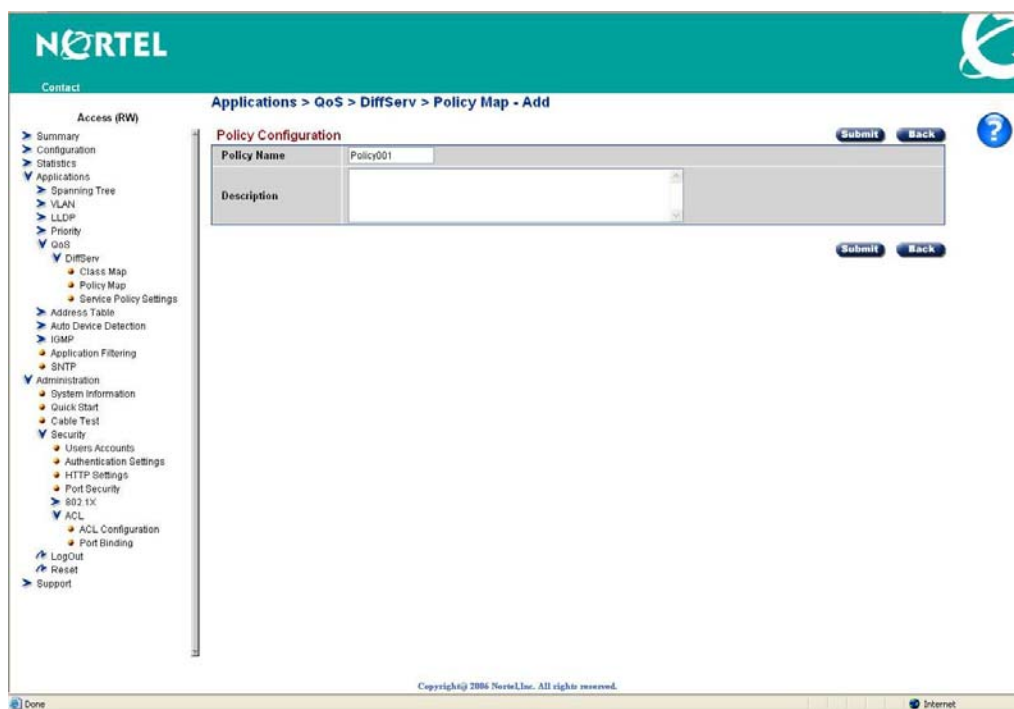
Step (3) Define a QoS policy for the selected traffic stream

A QoS policy determines the course of action applying to incoming IP packets that belong to the same traffic stream. The action can be mark/re-mark IP DSCP or mark/re-mark the CoS priority of the incoming Ethernet frames. BES50FE will provide QoS treatment based on the new IP DSCP or CoS value:

- CoS: The new CoS value maps into one of the traffic class based on the CoS to traffic class mapping (see Figure 1: BES50FE Priority to Traffic Class).

NOTE:

Using a QoS policy to set the new CoS value will not modify the DSCP of the affected IP packet. Only the 802.1p value of the egress Ethernet frame (that encapsulated the affected IP packet) will be modified.



Policy Rule Settings

Policy Name: Policy001

Class Name	Action	Meter	Rate (kbps)	Burst(byte)	Exceed Action
Class001	Set	CoS (0-7)	7		

Remove Class

Class Name: Class001

Action: Set CoS (0-7) 7

☐ Meter

Rate (1-100000) kbps

Burst (64-1522) byte

Exceed: Set IP DSCP (0-63)

Add Class

Reload Back

Policy Rule Settings

Policy Name: Policy001

Class Name	Action	Meter	Rate (kbps)	Burst(byte)	Exceed Action
Class001	Set	CoS (0-7)	7		

Remove Class

Class Name: Class001

Action: Set CoS (0-7) 7

☐ Meter

Rate (1-100000) kbps

Burst (64-1522) byte

Exceed: Set IP DSCP (0-63)

Add Class

Reload Back

Figure 13: BES50FE Policy Setting CoS Value

- IP DSCP: The new DSCP value first maps into a CoS priority value based on a defined DSCP-to-CoS mapping (see Figure 10: BES50 DSCP to COS Mapping). The new CoS value is then mapped into one of the traffic class based on the CoS to traffic class mapping (see Figure 1: BES50FE Priority to Traffic Class).

NOTE:

Using a QoS policy to set the new IP DSCP will modify the DSCP value of the target IP packet as well as the 802.1p value of the egress Ethernet frame (that encapsulated the target IP packet).

Policy Rule Settings

Policy Name: Policy001

Class Name	Action	Meter Rate (kbps)	Burst(byte)	Exceed Action
Class001	Set IP DSCP (0-63) 40	40 kbps	40 byte	Set IP DSCP (0-63) 40

Buttons: Reload, Back, Add Class, Remove Class

Policy Rule Settings

Policy Name: Policy001

Class Name	Action	Meter Rate (kbps)	Burst(byte)	Exceed Action
Class001	Set CoS (0-7) 46	40 kbps	40 byte	Set IP DSCP (0-63) 40

Buttons: Reload, Back, Add Class, Remove Class

Figure 14: BES50FE Policy Setting DSCP Value

Step (4) Apply the QoS policy to an ingress port

This step applies the QoS policy to an ingress port. Traffic classification and marking/re-marking are carried out at the ingress port where the incoming IP packets enter the BES50FE switch.

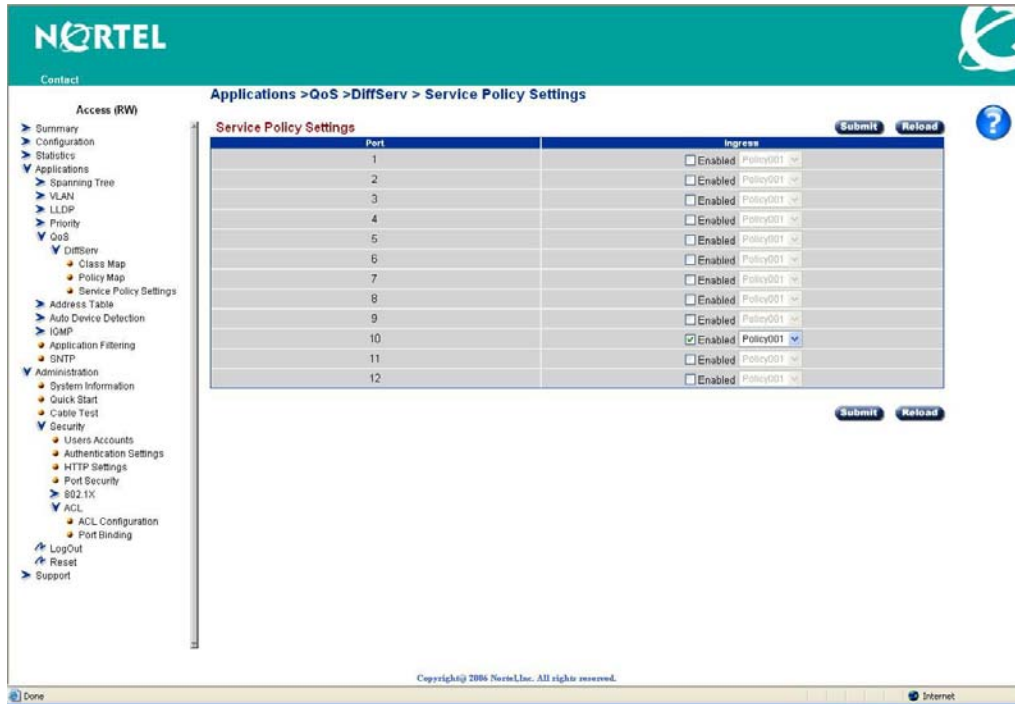


Figure 15: BES50FE Apply QoS Policy to Ingress Port

BES50GE Configuration

Step (1) Define the traffic classifier in an extended Access Control List (ACL)

This step describes the classification criteria to identify a subset of incoming traffic for the desired QoS treatment. The criteria are based on the five tuples in the IP header.

NOTE:

The action associated with the ACL rule MUST be set to “Permit”.

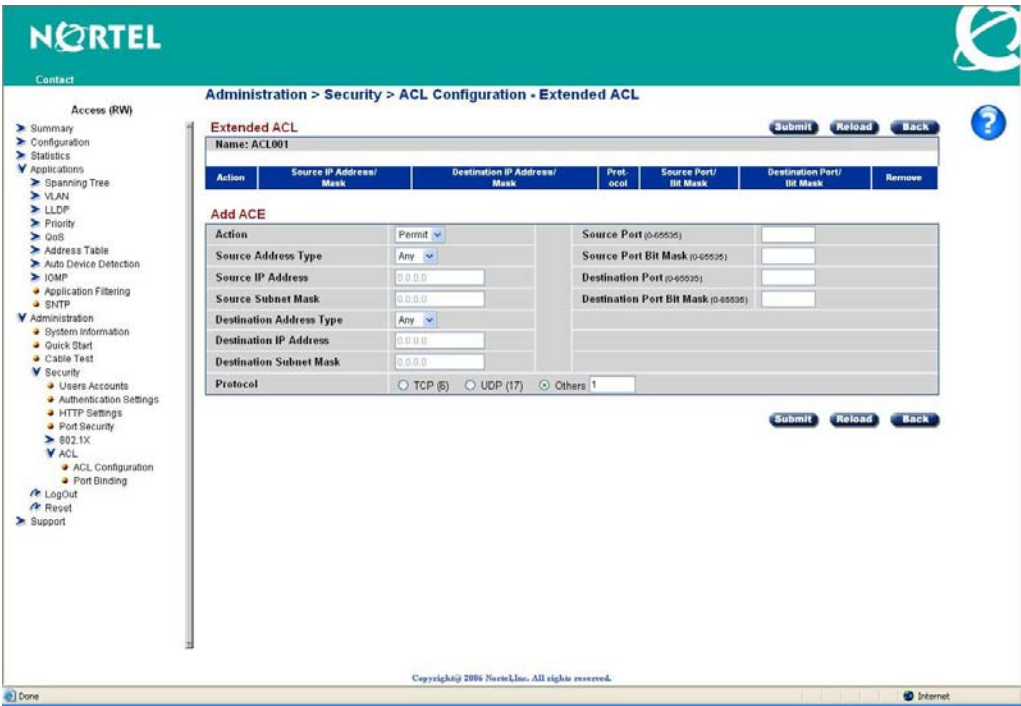
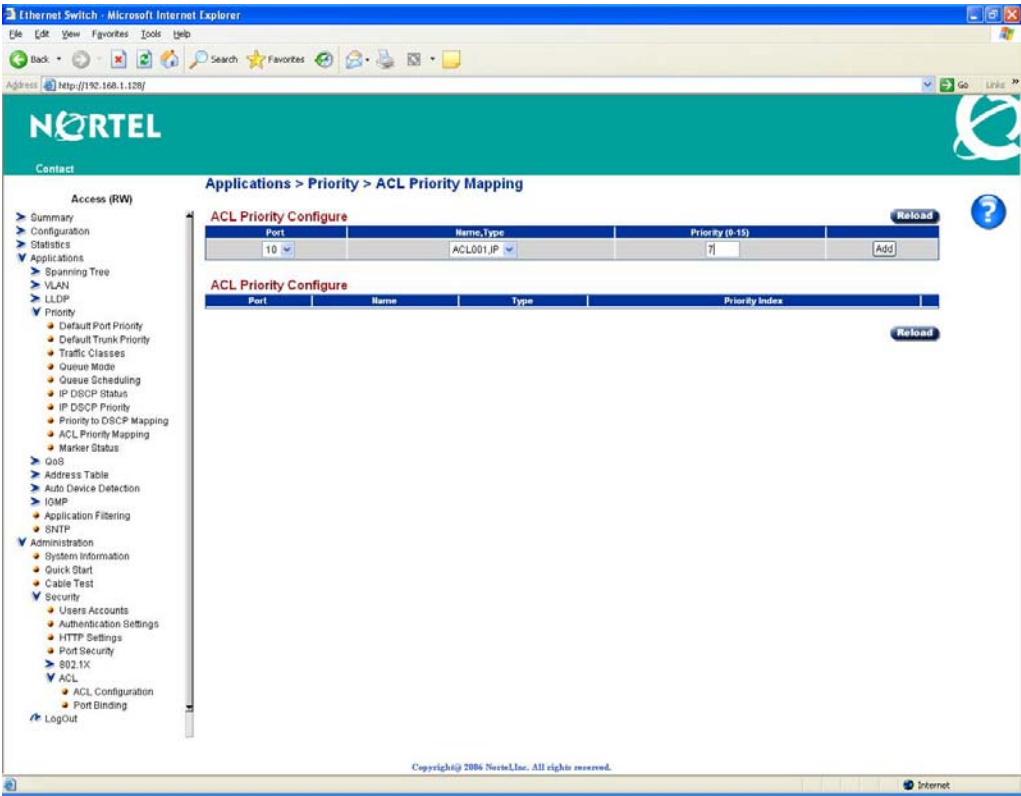


Figure 16: BES50GE Extended ACL Rule

Step (2) Bind ACL classification rule to an ingress port

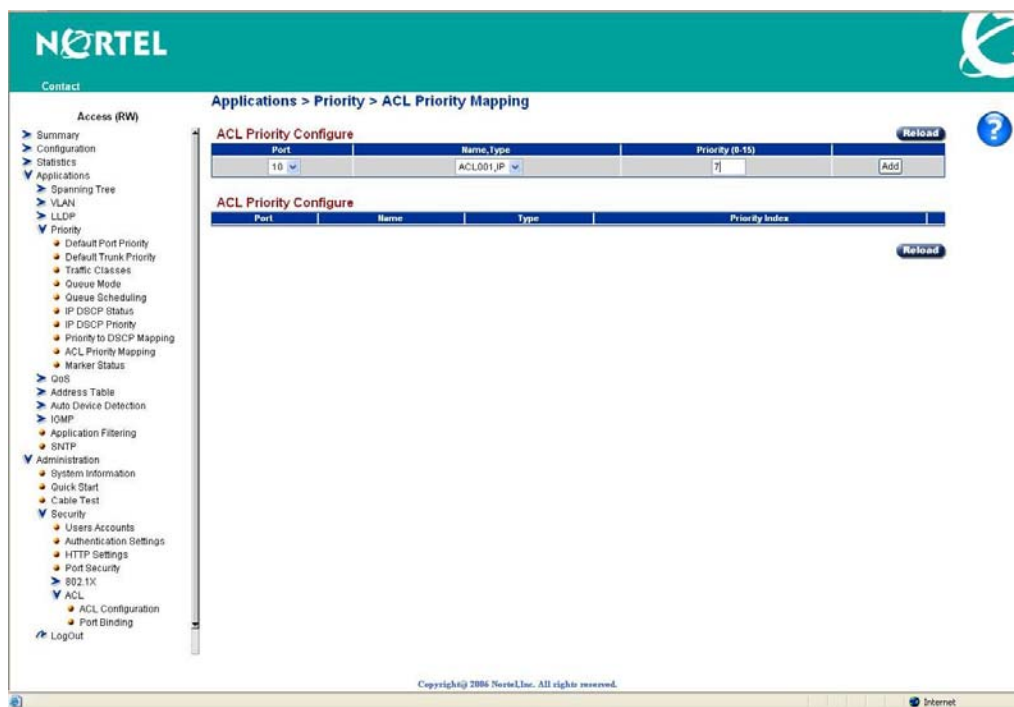
This step binds the ACL classification rule to an ingress port. Traffic classification is carried out at the ingress port where the incoming IP packets enter the BES50GE switch. All packets that match the classification rule are considered as belonging to the same traffic stream.¹ A priority index is assigned to the traffic stream which is then mapped into one of the 8 CoS value.

Since IP DSCP is a 6 bits value, there are more than 8 different combinations of valid DSCP value. A one-to-one priority index to CoS mapping limits the egress DSCP marking/remarking to only 8 DSCP values. BES50 defines 16 priority index and a non-configurable Priority Index to CoS mapping, which are used in both QoS treatment and IP DSCP marking/re-marking.

¹ Unlike BES50FE, Ethernet frame and IP packet marking/re-marking is performed on the egress side of BES50GE when the Ethernet frame is transmitted.

Priority Index	CoS Value
0	0
1	1
2	2
3	3
4	4
5	5
6	6
7	7
8	0
9	1
10	2
11	3
12	4
13	5
14	6
15	7

Table 3: BES50GE Priority Index to CoS Mapping



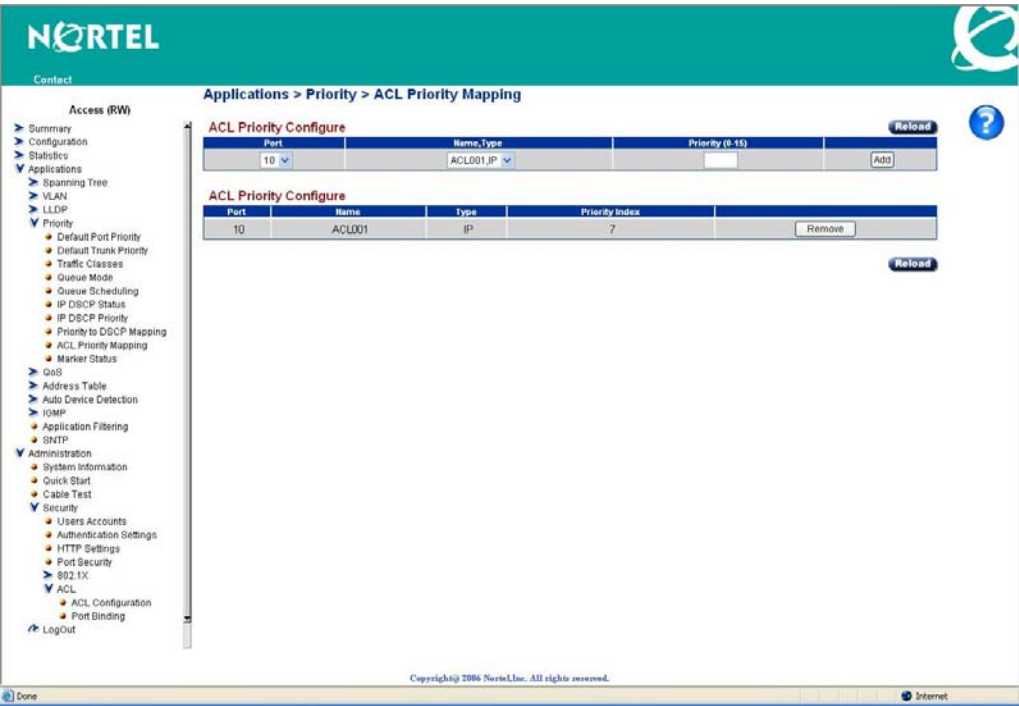


Figure 17: BES50GE ACL Priority Mapping

BES50 Re-Marking Capability

In addition to traffic prioritization, BES50 also provides the capability to re-mark the IP DSCP value of an outgoing IP packet and the 802.1p value (CoS) of an outgoing Ethernet frame.

IP DSCP Re-Marking on BES50FE

Figure 6: BES50FE Traffic Prioritization Flow Chart depicts the flow chart of how CoS/priority is determined by BES50FE. Re-marking in BES50FE is a simple extension to this process. BES50FE always re-marks the 802.1p value of the tagged egress Ethernet frame with the determined CoS/priority value. On the other hand, BES50FE will only re-mark IP DSCP of the outgoing IP packet if the matched QoS Policy has an associated “Set IP DSCP” action (see Figure 14: BES50FE Policy Setting DSCP Value).

IP DSCP Re-Marking on BES50GE

Figure 7: BES50GE Traffic Prioritization Flow Chart depicts the flow chart of how the priority index and CoS/priority are determined by BES50GE. Just like BES50FE, BES50GE always re-marks the 802.1p value of the tagged egress Ethernet frame with the determined CoS/priority value. Re-marking IP DSCP in BES50GE is slightly different from BES50FE.

By default, IP DSCP re-marking is disabled on all the ports. The user must manually enable IP DSCP re-marking on the egress port (i.e. the port in which IP packet is emitted out from BES50GE) if they desired to do so. To enable/disable IP DSCP re-marking on an egress port, navigate to the Applications > Priority > Marker Status panel.

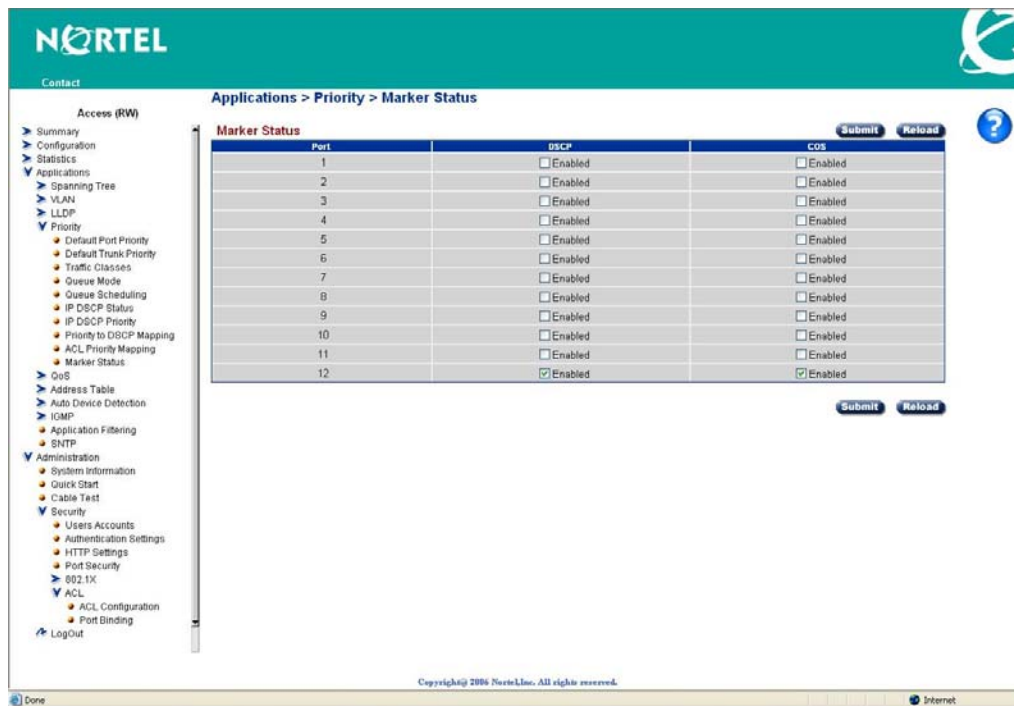


Figure 18: BES50GE IP DSCP Re-Marking Status

BES50GE provides a default priority index to IP DSCP mapping, and this mapping is also user configurable if they desired to use a different setting. To review and/or configure the traffic index to IP DSCP mapping, navigate to the Applications > Priority > Priority to DSCP panel.

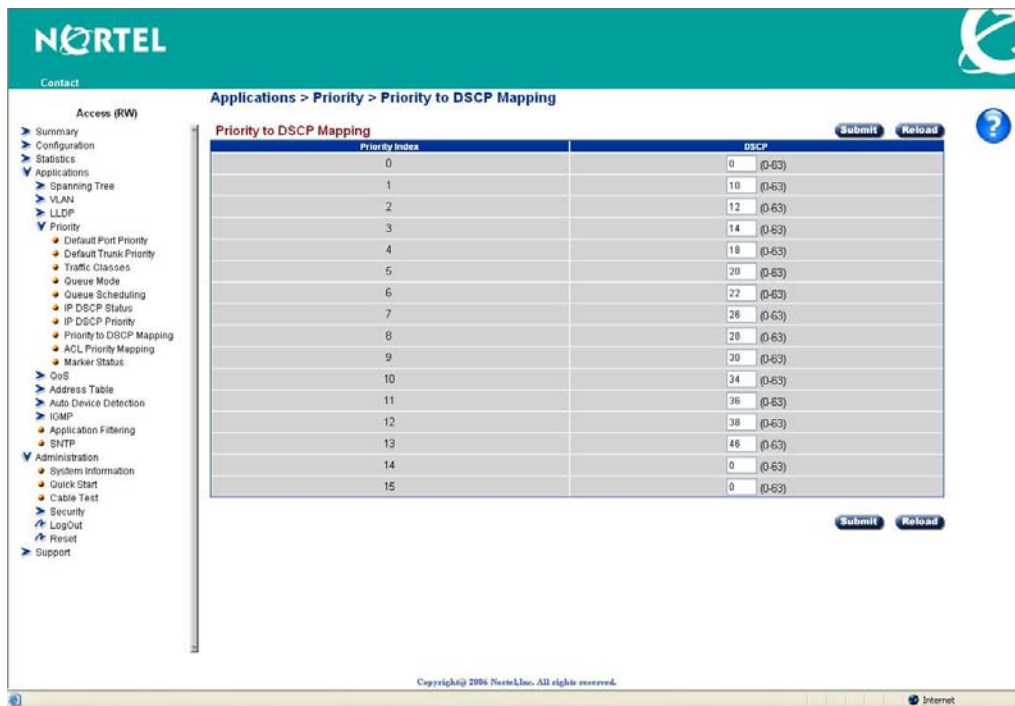


Figure 19: BES50GE Priority Index to IP DSCP Mapping

When IP DSCP re-marking is enabled, BES50GE always re-mark the DSCP value of all IP packets emitted from that egress port based on the priority index determined by BES50GE. The re-marking process is summarized as follows:

- If the egress port is not configured to re-mark IP DSCP or the outgoing Ethernet frame does not encapsulate an IP packet, then BES50GE will skip the IP DSCP re-marking process.
- If the egress port is configured to re-mark IP DSCP and the outgoing Ethernet frame encapsulates an IP packet, BES50GE will use the resulting priority index associated with the Ethernet frame to determine the corresponding DSCP value. The outgoing IP packet will be re-mark with this new DSCP value.

The following table combines the default Priority Index to CoS mapping and the default Priority Index to IP DSCP to illustrate the respective values used in CoS re-marking and IP DSCP re-marking mapped from the same priority index.

Priority Index	CoS Value	IP DSCP
0	0	0
1	1	10
2	2	12
3	3	14
4	4	18
5	5	20
6	6	22
7	7	26
8	0	28
9	1	30
10	2	34
11	3	36
12	4	38
13	5	46
14	6	0
15	7	0

Table 4: BES50GE Priority Index to CoS and IP DSCP Mapping

A word of caution when IP DSCP based traffic prioritization is enabled. Since the DSCP value is 6 bits and there are only 8 CoS/priority value, enabling IP DSCP remarking may alter IP DSCP to an undesirable value. The following table combines the default DSCP to CoS mapping with Table 4: BES50GE Priority Index to CoS and IP DSCP Mapping.

Ingress IP DSCP Value	CoS Value	Priority Index	Egress IP DSCP Value
0	0	0	0
8	1	1	10
10, 12, 14, 16	2	2	12
18, 20, 22, 24	3	3	14
26, 28, 30, 32, 34, 36	4	4	18
38, 40, 42	5	5	20
48	6	6	22
46, 56	7	7	25

E.g.

If IP DSCP traffic prioritization is enabled and the incoming IP packet with DSCP value of 36 does not match any one of the ACL rules, BES50GE will map the traffic to the traffic class 4 with a corresponding priority index value 4. When IP DSCP re-marking is enabled at the egress port, BES50 will map the priority index value 4 to IP DSCP value 18 and re-mark the outgoing IP packet with this value.

Difference Between BES50FE and BES50GE

1. Number of traffic class

BES50FE has four traffic classes, while BES50GE has eight traffic classes.

2. Different ways to configure ACL Based traffic prioritization

BES50FE uses the “QoS Policy” to facilitate ACL based traffic prioritization (see Figure 15: BES50FE Apply QoS Policy to Ingress Port), while BES50GE uses the “ACL Priority Mapping” to map traffic matching the ACL to a specified priority index (see Figure 17: BES50GE ACL Priority Mapping).

3. Location of frame/packet re-marking

BES50FE classifies and re-mark CoS/DSCP at the ingress port (i.e. the port in which the Ethernet frame enters BES50FE) of the switch. In contrast, BES50GE classifies traffic at the ingress port but re-mark CoS/DSCP at the egress port (i.e. the port in which the Ethernet frame is emitted from BES50GE) of the switch.

4. IP DSCP re-marking

Both BES50FE and BES50GE support re-marking of the DSCP value of an IP packet. However, there are subtle differences in terms of how the IP packet is re-marked:

- BES50FE re-marks packets on a per traffic stream level (i.e. all traffic that matches the same QoS policy configured on a specific ingress port).
- BES50GE re-marks packets on a per port basis, specifically IP packets that are emitted from an egress port with IP DSCP re-marking enabled.
- BES50FE only re-marks DSCP value of the IP packets that match a “QoS Policy” with an associated “Set IP DSCP” action.
- If re-marking is enabled on an egress port, BES50GE always re-mark the DSCP value of all IP packets emitted from that egress port based on the priority index determined by BES50GE.

Default QoS Behaviour of BES50FE and BES50GE

BES50 provides a set of default (out of the box) settings that are captured in the following table:

Traffic Prioritization Method	Default Status	Default CoS/Priority Value
Default port priority	Enabled	0
802.1p bit (tagged Ethernet frame)	Enabled	Taken from 802.1p value of the incoming tagged Ethernet frames
IP DSCP	Disabled	N/A
ACL based	Disabled	N/A

Table 5: Default QoS Settings of BES50

Without modifying the default settings, the followings describe the default QoS behaviour of both BES50FE and BES50GE:

- If the incoming Ethernet frame is a tagged frame, BES50 will take the 802.1p value specified in the VLAN tag as the CoS/priority value. This CoS value is used to put the incoming frame into one of the traffic classes (see Table 1: BES50 Priority to Traffic Class Mapping).
- If the incoming Ethernet frame is an untagged frame, BES50 will use the default port priority value '0' as the CoS/priority and maps the frame into one of the traffic classes (see Table 1: BES50 Priority to Traffic Class Mapping).
 - For BES50FE, CoS value of '0' maps into traffic class 1. For BES50GE, CoS value of '0' maps into traffic class 2.
 - This means that only tagged Ethernet frames with 802.1p value of either 1 or 2 will have lower priority than the untagged Ethernet frames.
- The IP DSCP value specified in the IP header has no effect on the default QoS behaviour of BES50 switches. Incoming traffic is prioritized based on either the 802.1p value of the tagged Ethernet frames or the default priority value of '0' for all the untagged Ethernet frames.
- Since both BES50GE and BES50FE do not alter the CoS value, there is no 802.1p re-marking of any Ethernet frames switched through BES50.
- By default, IP DSCP re-marking is disabled on both BES50FE and BES50GE.

Default QoS Behaviour on BCM Traffic

Since both BCM50 and BCM450 do not have VLAN capability, all the Ethernet frames transmitted from the BCM must be untagged Ethernet frames. When a BCM (either BCM50 or BCM450) is connected to a BES (either BES50FE or BES50GE) with default QoS settings, all the traffic transmitted from BCM will map to the same traffic class. Effectively there is no service differentiation among the traffic transmitted from the BCM.

In the presence of Ethernet traffic from other L2 devices, only tagged Ethernet frames with 802.1p value of either 1 or 2 will have lower priority than the traffic transmitted from BCM.

Default QoS treatment on the Ethernet frame sent to BCM is solely depending on whether the frame is tagged when it arrives at the ingress port of the BES50. The BES50 behaviour is captured in the section titled "Default QoS Behaviour of BES50FE and BES50GE".

When BES50 is used in a customer premises network with BCM50/BCM450 and QoS treatment is required, it is recommended to use the IP DSCP traffic prioritization method (see “3. Traffic prioritization based on IP DSCP” on page 15) on BES50 since BCM does not tag outgoing Ethernet frames. It is also recommended not to enable IP DSCP re-marking on BES50 so all QoS capable devices within the customer premise network will apply the same DiffServ QoS behaviour to the IP packets passing through the network. Typically DSCP is re-marked if necessary when IP packets traverse across DiffServ boundary. BES50 is a layer 2 device, re-marking DSCP is better off performed by a layer 3 router (or gateway) locating at the edge of two DiffServ boundary.

Software Baseline

This application note draws upon the capabilities available from the following BES50 software releases:

- BES50FE v1.0.3.0 (Release 1.1)
- BES50GE v1.0.5.0 (Release 1.1)