





Thomson Gateway



Ethernet QoS Configuration Guide

R7.4 and higher

Thomson Gateway

Ethernet QoS Configuration Guide
R7.4 and higher

Copyright

Copyright ©1999-2008 Thomson. All rights reserved.

Distribution and copying of this document, use and communication of its contents is not permitted without written authorization from Thomson. The content of this document is furnished for informational use only, may be subject to change without notice, and should not be construed as a commitment by Thomson. Thomson assumes no responsibility or liability for any errors or inaccuracies that may appear in this document.

Thomson Telecom Belgium
Prins Boudewijnlaan, 47
B-2650 Edegem
Belgium

<http://www.thomson-broadband.com>

Trademarks

The following trademarks may be used in this document:

- DECT is a trademark of ETSI.
- Bluetooth® word mark and logos are owned by the Bluetooth SIG, Inc.
- Ethernet™ is a trademark of Xerox Corporation.
- Wi-Fi® and the Wi-Fi logo are registered trademarks of the Wi-Fi Alliance. "Wi-Fi CERTIFIED", "Wi-Fi ZONE", "Wi-Fi Alliance", their respective logos and "Wi-Fi Protected Access" are trademarks of the Wi-Fi Alliance.
- UPnP™ is a certification mark of the UPnP™ Implementers Corporation.
- Microsoft®, MS-DOS®, Windows®, Windows NT® and Windows Vista® are either registered trademarks or trademarks of Microsoft Corporation in the United States and/or other countries.
- Apple® and Mac OS® are registered trademarks of Apple Computer, Incorporated, registered in the United States and other countries.
- UNIX® is a registered trademark of UNIX System Laboratories, Incorporated.
- Adobe®, the Adobe logo, Acrobat and Acrobat Reader are trademarks or registered trademarks of Adobe Systems, Incorporated, registered in the United States and/or other countries.

Other brands and product names may be trademarks or registered trademarks of their respective holders.

Document Information

Status: v1.0 (March 2008)

Reference: E-DOC-CTC-20080229-0002

Short Title: Config Guide: Ethernet QoS R7.4 and higher

About this Ethernet QoS Configuration Guide.....	1
1 Introduction.....	3
2 Ethernet QoS Overview	5
3 QoS Capabilities of the Ethernet Switch.....	7
3.1 Classification and Queuing	8
3.2 Rate Control	11
3.3 Storm Control	15
4 QoS Capabilities of the Ethernet Bridge	17
4.1 802.1p User Priority Regeneration.....	18
4.2 Ingress Classification	20
4.2.1 Medium Priority	21
4.2.2 Evaluated Priority	22
4.2.3 Internal Priority	26
4.3 Queuing	28
4.4 Priority Tagging of Egress Traffic.....	31
4.5 802.1p User Priority Preservation.....	32
5 Logical Ethernet Interfaces and QoS	33
6 Layer 2 QoS and Layer 3 QoS Interaction	35

Contents

About this Ethernet QoS Configuration Guide

Used Symbols



A **note** provides additional information about a topic.



A **caution** warns you about potential problems or specific precautions that need to be taken.

Terminology

Generally, the Thomson Gateway356 will be referred to as Wireless USB Adaptor in this Ethernet QoS Configuration Guide.

Typographical Conventions

Following typographical convention is used throughout this manual:

- **Sample text** indicates a hyperlink to a Web site.
Example: For more information, visit us at www.thomson-broadband.com.
- **Sample text** indicates an internal cross-reference.
Example: If you want to know more about guide, see "1 Introduction" on page 7".
- **Sample text** indicates an important content-related word.
Example: To enter the network, you **must** authenticate yourself.
- **Sample text** indicates a GUI element (commands on menus and buttons, dialog box elements, file names, paths and folders).
Example: On the **File** menu, click **Open** to open a file.

Documentation and software updates

Thomson continuously develops new solutions, but is also committed to improving its existing products.

For more information on Thomson's latest technological innovations, documents and software releases, visit us at <http://www.thomson-broadband.com>.

About this Ethernet QoS Configuration Guide

Overview

First, this Ethernet QoS Configuration Guide shortly lists the Ethernet QoS features of the Thomson Gateway and the relevant standards. Some background information on several Ethernet QoS-related concepts is also provided. Next, this document describes the configuration of each Ethernet QoS feature on the Thomson Gateway in detail. CLI (Command Line Interface) commands are used for the configuration.

This document is structured as follows:

Topic	Page
"1 Introduction"	3
"2 Ethernet QoS Overview"	5
"3 QoS Capabilities of the Ethernet Switch"	7
"4 QoS Capabilities of the Ethernet Bridge"	17
"5 Logical Ethernet Interfaces and QoS"	33
"6 Layer 2 QoS and Layer 3 QoS Interaction"	35

1 Introduction

Introduction

Ethernet QoS, or Layer 2 QoS, is the subject of this document. In this chapter, we shortly list the Ethernet QoS features that are supported by the Thomson Gateway, the relevant standards and related documents.

Supported Ethernet QoS features

Thomson Gateway devices support following Ethernet QoS features:

- **QoS features of the Ethernet switch**
 - ▶ Classification
 - ▶ Queuing and scheduling
 - ▶ Pause frame flow control
 - ▶ Rate control (ingress rate limiting and egress traffic shaping)
 - ▶ Storm control
- **QoS features of the Ethernet bridge**
 - ▶ 802.1p user priority regeneration
 - ▶ Ingress classification
 - ▶ Queuing
 - ▶ Priority tagging
 - ▶ 802.1p user priority preservation
- **QoS features of the logical Ethernet interfaces**
 - ▶ WAN-side priority tagging

Standards compliancy

Thomson Gateway devices are compliant with following standards, which are relevant to Ethernet, Ethernet bridging, VLAN bridging, VLAN user priorities and stacked VLANs:

- **IEEE Std 802.3 - 2000:** Part 3: Carrier sense multiple access with collision detection (CSMA/CD) access method and physical layer specifications.
- **IEEE Std 802.1D - 2004:** IEEE standard for local and metropolitan area networks - Media Access Control (MAC) bridges.
- **IEEE Std 802.1Q - 1998:** IEEE standard for local and metropolitan area networks - Virtual bridged local area networks.
- **IEEE Std 802.1p:** Traffic Class Expediting and Dynamic Multicast Filtering. This standard is merged into 802.1D-2004.
- **IEEE Std 802.1ad - 2005:** IEEE standard for local and metropolitan area networks - Virtual bridged local area networks - Amendment 4: Provider bridges.

Related documents

Other Layer 2 related features, such as Ethernet and VLANs, are described in other documents. See the "Ethernet Configuration Guide" and the "VLAN Configuration Guide" for more information.

1 | Introduction

2 Ethernet QoS Overview

Introduction

The overall QoS framework of Thomson Gateway devices consists of several *QoS elements* to support demanding applications such as Voice over IP (VoIP), IPTV, management and so on. The two QoS elements within the scope of this document are the Ethernet switch and the Ethernet bridge.

Ethernet QoS functionality

The Ethernet QoS functionality mainly consists of two steps:

- At the *ingress* side, the incoming frames are *classified*. This means that an internal class, i.e. priority, is assigned to each frame.
- At the *egress* side, each frame is *queued* into the correct queue, based on the assigned internal class.

Internal class information

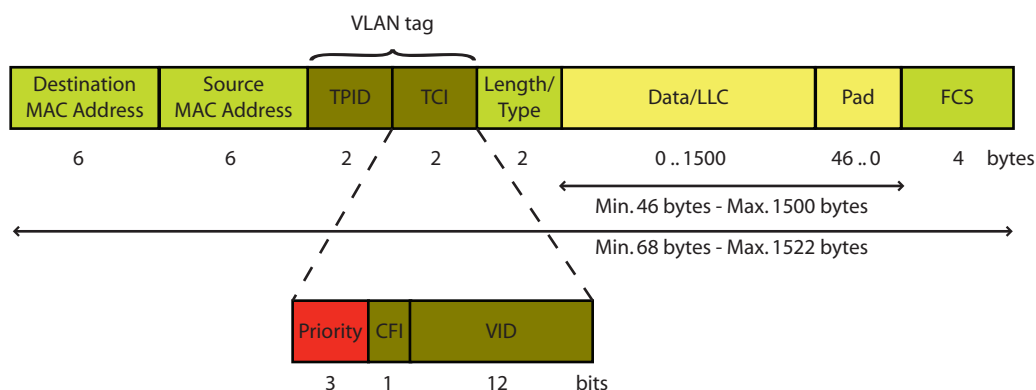
The *internal class* of a frame or packet plays an important role during both classification and queuing. There are 16 internal classes, numbered from 0 through 15.

All QoS elements that are part of the overall QoS framework work together and *share* the internal class information for each frame or packet. This means that when classification is done by the Ethernet bridge, the resulting internal class can be used or overruled by the IP router. The other way around, when classification is done by the IP router, the resulting internal class can be used to bridge the traffic to the correct ATM PVC or the correct QoS queue of a single ATM PVC.

Classification-related frame fields

In order to assign an internal class to a received frame, several information fields of the frame can be used. Following information fields will be used in this document:


- **802.1p user priority field:** this is a three-bit field in the VLAN tag of an Ethernet frame, indicated in red in following illustration:



2| Ethernet QoS Overview


- **IP precedence field:** this is a three-bit field in the header of an IP packet. It is part of the ToS (Type of Service) byte of an IPv4 header or the Traffic Class byte of an IPv6 header.

IP Precedence Field



- **DSCP (Differentiated Services Code Point) field:** this is a six-bit field in the header of an IP packet. It is part of the Differentiated Services (DS) field of the IP header. The DS field replaces the old ToS byte of the IPv4 header and the old Traffic Class byte of the IPv6 header. The values of the DSCP field provide limited backwards compatibility with the IP precedence field.

Differentiated Services Code Point Field



3 QoS Capabilities of the Ethernet Switch

Introduction

Which QoS features of the Ethernet switch are available depend on the used Thomson Gateway device. We make a distinction between residential devices and business devices. Whenever necessary, it is noted to which type of devices the feature is applicable.

Listing the QoS capabilities of the switch

To display the QoS capabilities of the Ethernet switch, execute the command `:eth switch info`.

To list the QoS capabilities of a *residential device*, for example the Thomson ST780, execute following command:

```
=>:eth switch info
Ethernet switch capabilities:
+ Number of external ports=4
+ QoS supported: number of queues=4
  + diffserv classifier supported
+ Storm control supported, possible values: 100 200 300 400
```

To list the QoS capabilities of a *business device*, for example the Thomson ST620, execute following command:

```
=>:eth switch info
Ethernet switch capabilities:
+ Number of external ports=4
+ QoS supported: number of queues=4
  + realtime queue supported
  + diffserv classifier supported
  + tos classifier supported
+ Ingress shaping supported: number of ingress shapers=2
+ Egress shaping supported: number of egress shapers=1
+ Storm control supported, possible values: 33 50 100 200
```

Overview

Following QoS features are described in this chapter:

Topic	Page
"3.1 Classification and Queuing"	8
"3.2 Rate Control"	11
"3.3 Storm Control"	15

3| QoS Capabilities of the Ethernet Switch

3.1 Classification and Queuing

Introduction

Simply explained, classification and queuing work as follows:

- At the ingress side, each frame is **classified**: the Ethernet switch assigns a priority to each received frame, based on the configured classification criterion.
- At the egress side the frame is **queued**: the Ethernet switch uses the priority of the frame to map the frame onto the correct queue. The number of queues can be configured (up to four).
- For the mapping from the classification result (priority) to the correct queue, the Ethernet switch has **internal mapping tables**. You cannot configure these mapping tables.

The CLI commands described in this section belong to the command group `:eth switch qos`.

Listing classification and queuing information

To obtain information on QoS support, the configured classification criterion and queues, execute the command `:eth switch qos list`.

By default, QoS support is disabled:

```
=>:eth switch qos list
No queues configured => QoS disabled.
```

If QoS support is enabled, the listed information consists of three parts, as shown in the following illustration:

The diagram shows the output of the `:eth switch qos list` command. It is divided into three sections with red callouts:

- Enabling QoS support command `config`**: Points to the first line of output: "QoS enabled on switch. 2 queues configured. Realtime disabled."
- Queuing command `weights`**: Points to the second table of output showing queue weights.
- Classification command `ifconfig`**: Points to the third table of output showing port classification and flow control settings.

Queue	Weight
0	33
1	67

Port	Classifier	FlowControl
1	802.1P	off
2	802.1P	off
3	802.1P	off
4	802.1P	off

Enabling QoS support

The common QoS parameters can be configured with the command `:eth switch qos config`.

Following parameter must be specified:

- **State**: this parameter is used to enable or disable QoS support on the Ethernet switch. By default, QoS support is disabled.
 - ▶ To **disable** QoS support, set this parameter to disabled.
 - ▶ To **enable** QoS support, set this parameter to enabled. In addition, define at least two queues per switch port.

3 | QoS Capabilities of the Ethernet Switch

Optionally, two other parameters can be specified:

- **NbrOfQueues:** this parameter indicates the number of queues per switch port. This is a number within the range from 0 through 4.



At least **two** queues **must** be defined in order to enable QoS support on the Ethernet switch.

- **Realtime:** the highest priority queue can be set as a real-time queue by enabling this parameter. By default, this parameter is disabled.



In case of a residential device, the parameter realtime is **not** applicable.

```
=>:eth switch qos config state=enabled nbrOfQueues=3 realtime=enabled
=>:eth switch qos list
QoS enabled on switch. 3 queues configured. Realtime enabled.
```

Queue	Weight
0	14
1	28
2	58

Port	Classifier	FlowControl
1	802.1P	off
2	802.1P	off
3	802.1P	off
4	802.1P	off

Classification

Classification can be configured for each switch port with the command `:eth switch qos ifconfig`.

Two parameters must be specified:

- **Port:** this parameter specifies the switch port to be configured, indicated by a number from 1 through 4.
- **Mode:** this parameter indicates the used **classification criterion**.
 - ▶ **None:** all traffic will be classified as low priority.



At the egress switch ports, the classification criterion can be set to none as all classification is done at the ingress side.

- ▶ **High:** all traffic will be classified as high priority.
- ▶ **802.1p:** classification is based on the value of the 802.1p user priority field in the VLAN tag of a VLAN-tagged or priority-tagged frame. This is the default classification criterion.
- ▶ **Diffserv:** classification is based on the value of the DSCP field in the IP header of a packet.
- ▶ **ToS:** classification is based on the value of the IP precedence field in the IP header of a packet.



In case of a residential device, the ToS mode is **not** supported.

Optionally, following parameter can be specified:

- **Flowcontrol:** this parameter is used to enable or disable pause frame flow control. By default, flow control is disabled.

A flow control mechanism can be used when the data rate sent by a device exceeds the input data rate of the receiver. In this case, the receiver sends a pause frame which stops the transmission of the sender for a specified period of time. A pause frame is a MAC control frame.

3| QoS Capabilities of the Ethernet Switch

For example, configure the classification as follows:

```
=>:eth switch qos ifconfig port=2 mode=high flowcontrol=enabled
=>:eth switch qos list
QoS enabled on switch. 3 queues configured. Realtime enabled.

Queue      Weight
0          14
1          28
2          58

Port  Classifier FlowControl
1     802.1P  off
2     High   on
3     802.1P  off
4     802.1P  off
```



QoS support **must** be enabled on the Ethernet switch before the classification parameters can be configured.

Queuing

At the **egress side**, frames are queued. The used scheduling algorithm is **WFQ (Weighted Fair Queuing)**. This means that the queues are being served based upon their weight. A queue with a weight of W% will achieve an average data rate of W% of the link data rate. The higher the weight, the higher the priority.

The weight of each queue can be configured with the command **:eth switch qos weights**.

For each queue, the weight must be specified:

- **QueueX**: this parameter indicates the weight of queue X in WFQ, expressed in percent. The weight can range from 0% through 100%. By default, the weights are as follows:
 - ▶ **2 queues**: 33% and 67%.
 - ▶ **3 queues**: 14%, 28% and 58%.
 - ▶ **4 queues**: 6%, 13%, 27% and 54%.

```
=>:eth switch qos weights queue0=24 queue1=31 queue2=45
=>:eth switch qos list
QoS enabled on switch. 3 queues configured. Realtime enabled.

Queue  Weight
0      24
1      31
2      45

Port  Classifier FlowControl
1     802.1P  off
2     High   on
3     802.1P  off
4     802.1P  off
```



The sum of the weights of the queues **must** be equal to **100**.



QoS support **must** be enabled on the Ethernet switch and the correct number of queues must be defined before the weights can be configured.

3.2 Rate Control

Introduction

Rate control is used to meter and limit the rate of a data stream. Rate control can be performed through rate limiting or shaping:

■ *Ingress rate limiting*

At the ingress side, rate limiting is used to restrict the bit rate of specific traffic types. Excess traffic is dropped.

Two ingress traffic shapers, also referred to as buckets, can be defined. Specific traffic types can be assigned to one of the two traffic shapers. Per switch port, a speed and burst size can be specified.

■ *Egress shaping*

At the egress side, traffic shaping is used: traffic is buffered and then sent at a predefined rate. In this case, excess traffic will not be dropped as long as the buffer is not congested.

Only one shaper or bucket is used. Thus, all egress traffic must pass the egress rate control checking. Per switch port, a speed and burst size can be specified.

The CLI commands described in this section belong to the command group `:eth switch shaper`.



Rate limiting and shaping are *not* supported by residential devices. The description in this section applies to business devices.

How does a token bucket algorithm work?

Rate control is typically achieved using one or more *token bucket algorithms*. The operation of such a token bucket algorithm can be explained as follows:

- The bucket or traffic shaper is characterized as follows:
 - ▶ Tokens are added to the bucket at a given *speed*.
 - ▶ The bucket can contain only a limited number of tokens, defined by the *size* of the bucket.
- Every time a frame arrives at the bucket, a token is taken from the bucket before the frame is accepted and forwarded.
- If a frame arrives at the bucket and no tokens are available (the bucket is empty), the frame is dropped.

As a result, the speed tokens are added to the bucket limits the *rate* of the data frames and the size of the bucket defines the maximum *burst size* of the data stream. Both the speed limit and burst size are configurable.

3| QoS Capabilities of the Ethernet Switch

Listing the rate control configuration

To look up the rate control configuration, execute following command:

```
=>:eth switch shaper iflist
Port      Direction Shaper  State   Speed   BurstSize  Flags
1         egress    0      off     0       0          nnnnnn
1         ingress   0      off     0       0          -----
1         ingress   1      off     0       0          -----
2         egress    0      off     0       0          nnnnnn
2         ingress   0      off     0       0          -----
2         ingress   1      off     0       0          -----
3         egress    0      off     0       0          nnnnnn
3         ingress   0      off     0       0          -----
3         ingress   1      off     0       0          -----
4         egress    0      off     0       0          nnnnnn
4         ingress   0      off     0       0          -----
4         ingress   1      off     0       0          -----
State: configurable per port and shaper type (ingress or egress).
Speed & BurstSize: configurable per port and shaper. Only limited set of values.
Input is rounded.
Flags: configurable per ingress shaper: u(nicast), m(ulticast), b(roadcast),
c(ontrol), a(ddress unknown), d(iscard), n(ot applicable).
The same frametype can not be set on different ingress shapers.
```

Configuring the ingress traffic shapers

Two different *ingress* traffic shapers can be defined with the command `:eth switch shaper config`. These two traffic shapers exist for each switch port and can not be configured per port.

One parameter must be specified:

- **Shaper**: this parameter specifies the traffic shaper instance to be configured. Only two traffic shaper instances exist, indicated with a number (0 or 1).

Optionally, different parameters can be specified to configure the selected traffic shaper instance.

- The **traffic types** to be shaped must be specified. For each traffic type, the corresponding parameter can be enabled or disabled. By default, all traffic types are disabled.
 - ▶ **Unicast**: unicast traffic
 - ▶ **Multicast**: multicast traffic
 - ▶ **Broadcast**: broadcast traffic
 - ▶ **Control**: MAC control traffic
 - ▶ **Unknown**: traffic that resulted in a destination look up failure (DLF). This means that the destination address is unknown.



You can not enable the same traffic type for both traffic shaper instances.

31 QoS Capabilities of the Ethernet Switch

- The way **excess traffic** is handled must be specified:
 - ▶ **Discard**: if this parameter is enabled, excess frames are discarded. If the parameter is disabled, flow control is used (if available) and a pause frame is sent out. By default, this parameter is disabled.

```
=>:eth switch shaper config shaper=0 unicast=enabled control=enabled unknown=enabled
discard=enabled
=>:eth switch shaper config shaper=1 multicast=enabled broadcast=enabled discard=enabled
=>:eth switch shaper iflist
```

Port	Direction	Shaper	State	Speed	BurstSize	Flags
1	egress	0	off	0	0	nnnnnn
1	ingress	0	off	0	0	u--cad
1	ingress	1	off	0	0	-mb--d
2	egress	0	off	0	0	nnnnnn
2	ingress	0	off	0	0	u--cad
2	ingress	1	off	0	0	-mb--d
3	egress	0	off	0	0	nnnnnn
3	ingress	0	off	0	0	u--cad
3	ingress	1	off	0	0	-mb--d
4	egress	0	off	0	0	nnnnnn
4	ingress	0	off	0	0	u--cad
4	ingress	1	off	0	0	-mb--d

State: configurable per port and shaper type (ingress or egress).
Speed & BurstSize: configurable per port and shaper. Only limited set of values.
Input is rounded.

Flags: configurable per ingress shaper: u(nicast), m(ulticast), b(roadcast), c(ontrol), a(ddress unknown), d(iscard), n(ot applicable).
The same frametype can not be set on different ingress shapers.

The egress traffic shaper

Only one **egress** traffic shaper exists for each switch port. This egress traffic shaper can not be configured:

- **All traffic types** must pass through the traffic shaper.
- **Excess traffic** is not discarded as long as the queue that contains this traffic is not congested.

Configuring per port rate control

Rate control can be configured on a **per port** basis. To configure rate control on a specific switch port, execute the command **:eth switch shaper ifconfig**.



Before you configure rate control on a switch port, disable the corresponding bridge port first.

Per port and per traffic direction the traffic shaper can be selected. Therefore, three parameters must be specified:

- **Port**: this parameter specifies the switch port to be configured, indicated by a number from 1 through 4.
- **Ingress**: this parameter is used to select the traffic direction.
 - ▶ If the parameter is **enabled**, the ingress traffic direction is selected.
 - ▶ If the parameter is **disabled**, the egress traffic direction is selected.
- **Shaper**: this parameter selects the (ingress) traffic shaper instance. Only two ingress traffic shaper instances exist, indicated with a number (0 or 1). If the egress traffic direction is selected, the value of this parameter is not relevant and set to 0.

Once the combination of port, traffic direction and traffic shaper is selected, several optional parameters can be set:

- **State**: this parameter specifies whether rate control is used or not. By default, rate control is disabled. For the ingress traffic direction, rate control is always enabled or disabled for **both** traffic shaper instances of the selected port.

3| QoS Capabilities of the Ethernet Switch

- **Speed:** this parameter specifies the speed limit, expressed in bits per second (bps). The value can only be changed in discrete steps:
 - ▶ Steps of 64 000 bps within the range from 64 000 bps through 1 792 000 bps.
 - ▶ The next possible value after 1 792 000 bps is 2 000 000 bps.
 - ▶ Steps of 1 000 000 bps within the range from 2 000 000 bps through 100 000 000 bps.

As only discrete steps are possible, rounding is used.



Each value is mapped to the *next* (higher) possible discrete value. For example, 64 001 bps is mapped to 128 000 bps.

- **Burstsize:** this parameter specifies the maximum burst size, expressed in kilobytes (KB). This value can only be changed in discrete steps within the range from 6 KB through 126 KB. The possible values are 6 KB, 10 KB, 18 KB, 34 KB, 66 KB and 126 KB. Only these values can be specified during configuration (no rounding is used).

```
=>:eth bridge ifconfig intf=ethport2 portstate=disabled
=>:eth switch shaper ifconfig port=2 ingress=disabled shaper=0 state=enabled speed=90000000
burstsize=18
=>:eth switch shaper ifconfig port=2 ingress=enabled shaper=0 state=enabled speed=30000000
burstsize=10
=>:eth switch shaper iflist
Port      Direction  Shaper    State    Speed    BurstSize  Flags
1         egress     0         off      0         0          nnnnnn
1         ingress    0         off      0         0          u--cad
1         ingress    1         off      0         0          -mb--d
2         egress     0         on       90000000  18         nnnnnn
2         ingress    0         on       30000000  10         u--cad
2         ingress    1         on       100000000 126        -mb--d
3         egress     0         off      0         0          nnnnnn
3         ingress    0         off      0         0          u--cad
3         ingress    1         off      0         0          -mb--d
4         egress     0         off      0         0          nnnnnn
4         ingress    0         off      0         0          u--cad
4         ingress    1         off      0         0          -mb--d
State: configurable per port and shaper type (ingress or egress).
Speed & BurstSize: configurable per port and shaper. Only limited set of values.
Input is rounded.
Flags: configurable per ingress shaper: u(nicast), m(ulticast), b(roadcast),
c(ontrol), a(ddress unknown), d(iscard), n(ot applicable).
The same frametype can not be set on different ingress shapers.
=>:eth bridge ifconfig intf=ethport2 portstate=forwarding
```

3.3 Storm Control

Introduction

A **traffic storm** occurs when frames flood the LAN, creating excessive traffic and degrading the network performance.

The **storm control** feature prevents the switch ports from being disrupted by a broadcast, multicast, or unicast traffic storm. Storm control is thus rather a security mechanism than a QoS feature.

The CLI commands described in this section belong to the command group **:eth switch storm**.

Listing the storm control configuration

To look up the storm control configuration, execute following command:

```
=>:eth switch storm iflist
Port      State      Rate(0.1%) BurstSize  Flags
1         off        0          0         ---
2         off        0          0         ---
3         off        0          0         ---
4         off        0          0         ---
Rate & BurstSize: only limited set of values. Input is rounded.
Flags: b(roadcast), m(ulticast), a(ddress unknown).
```

Configuring per port storm control

Ingress storm control can be activated per port with the command **:eth switch storm ifconfig**.

The port must be specified:

- **Port:** this parameter specifies the switch port to be configured, indicated by a number from 1 through 4.

Optionally, following parameters can be specified:

- **State:** this parameter specifies whether storm control is enabled or not. By default, storm control is disabled.
- **Rate:** this parameter specifies the maximum allowed rate of the selected traffic types, expressed in tenths of one percent (0.1%). Only specific values can be used during configuration (no rounding is used).
 - ▶ In case of a business device, the only possible values are 33, 50, 100 and 200.
 - ▶ In case of a residential device, the only possible values are 100, 200, 300 and 400.
- **Burstsize:** this parameter specifies the maximum burst size, expressed in kilobytes (KB). This value can only be changed in discrete steps of 2 KB within the range from 2 KB through 8 KB. As only discrete steps are possible, rounding is used.



Each value is mapped to the **next** (higher) possible discrete value. For example, 3 KB is mapped to 4 KB.

3| QoS Capabilities of the Ethernet Switch

Storm control can be used to control bursts of different *traffic types*. For each traffic type, the corresponding parameter can be enabled or disabled. By default, all traffic types are disabled.

- **Broadcast**: broadcast traffic
- **Multicast**: multicast traffic
- **Unknown**: traffic that resulted in a destination look up failure (DLF). This means that the destination address is unknown.

```
=>:eth switch storm ifconfig port=1 state=enabled rate=33 burstsize=5 broadcast=enabled
=>:eth switch storm iflist
Port      State      Rate(0.1%) BurstSize  Flags
1        on        33       6        b--
2         off        0          0          ---
3         off        0          0          ---
4         off        0          0          ---
Rate & BurstSize: only limited set of values. Input is rounded.
Flags: b(roadcast), m(ulticast), a(ddress unknown).
```



Storm control can only be enabled if both the rate and burst size are different from zero and at least one traffic type is selected.

4 QoS Capabilities of the Ethernet Bridge

Overview

Following QoS features are described in this chapter:

Topic	Page
"4.1 802.1p User Priority Regeneration"	18
"4.2 Ingress Classification"	20
"4.3 Queuing"	28
"4.4 Priority Tagging of Egress Traffic"	31
"4.5 802.1p User Priority Preservation"	32

4| QoS Capabilities of the Ethernet Bridge

4.1 802.1p User Priority Regeneration

Introduction

The VLAN tag of VLAN-tagged frames contains a three-bit user priority field. This field indicates the **802.1p user priority**, which is also referred to as the Priority Code Point (PCP) or the Class of Service (CoS).

802.1p user priority regeneration maps the 802.1p user priority value of a received frame to the **regenerated user priority**, also referred to as the VLAN user priority. The mapping is based on the user priority regeneration table.

The regenerated user priority corresponds to a fixed **internal class**.

User priority regeneration table

The **user priority regeneration table** maps the 802.1p user priority of a received frame to a regenerated user priority. The table contains eight entries, one entry for each possible value of the 802.1p user priority.

To look up the regeneration table of a bridge port, execute following command:

```
=>:eth bridge iflist intf=ethport1
ethport1 : dest : ethif1
          Connection State: connected   Retry: 10
          Priority Tagging: NA (destination switch interface)
          Port: ethport1      PortNr: 1      PortState: forwarding Interface: up
          Multicast filter: disabled      Dynamic VLAN : disabled
          WAN : disabled
          IGMP snooping : enabled
          Transparent Prio: disabled
          BPDU Filtering : disabled
          Extra Tagging : none
          VLAN: Default VLAN: default Ingressfiltering: disabled Acceptvlanonly: disabled
          VLAN: Priority: disabled IP Prec: disabled Priority: 0
          Regeneration table: 0 1 2 3 4 5 6 7
          RX bytes: 49421 frames: 14
          TX bytes: 10512 frames: 22 dropframes: 0
```

The regeneration table is indicated by a **string**. The default table is represented by the string 0 1 2 3 4 5 6 7. This must be interpreted as follows:

- 802.1p user priority 0 is mapped to the regenerated user priority 0.
- 802.1p user priority 1 is mapped to the regenerated user priority 1.
- 802.1p user priority 2 is mapped to the regenerated user priority 2.
- 802.1p user priority 3 is mapped to the regenerated user priority 3.
- 802.1p user priority 4 is mapped to the regenerated user priority 4.
- 802.1p user priority 5 is mapped to the regenerated user priority 5.
- 802.1p user priority 6 is mapped to the regenerated user priority 6.
- 802.1p user priority 7 is mapped to the regenerated user priority 7.

Configuring the regeneration table

The user priority regeneration table can be configured independently for each bridge port and for each value of the 802.1p user priority. For the configuration of an entry, the full range of values can be used for the regenerated user priority.

If you want to configure the regeneration as follows:

- 802.1p user priority 0 is mapped to the regenerated user priority 7.

4| QoS Capabilities of the Ethernet Bridge

- 802.1p user priority 1 is mapped to the regenerated user priority 6.
- 802.1p user priority 2 is mapped to the regenerated user priority 5.
- 802.1p user priority 3 is mapped to the regenerated user priority 4.
- 802.1p user priority 4 is mapped to the regenerated user priority 3.
- 802.1p user priority 5 is mapped to the regenerated user priority 2.
- 802.1p user priority 6 is mapped to the regenerated user priority 1.
- 802.1p user priority 7 is mapped to the regenerated user priority 0.

Then the regeneration table must be configured with the string 7 6 5 4 3 2 1 0.

To configure the regeneration table, execute following command:

```
=>:eth bridge ifconfig intf=ethport1 regenprio=76543210
```



Type the string *without* interspacing.



In case of a business device, the *same* user priority regeneration table is applied to *all bridge ports* that are *connected to a switch port*.

From regenerated user priority to internal class

Each regenerated user priority is mapped to one of the sixteen internal classes. This mapping is *fixed*. Following table shows this mapping:

Regenerated user priority	Internal class
7	15
6	14
	13
	12
	11
5	10
	9
4	8
	7
3	6
	5
0 (Default)	4 (Default)
	3
2	2
	1
1	0

4.2 Ingress Classification

Introduction

During ingress classification, an internal class is assigned to a received frame, based on a frame's medium priority and evaluated priority:

- **The medium priority:** this priority is based on the medium (physical interface) the received frame is entering on.
- **The evaluated priority:** this priority is based on the information fields in the received frame. The used information can be:
 - ▶ 802.1p user priority field
 - ▶ port default priority
 - ▶ IP precedence field
 - ▶ DSCP field

The resulting internal class assigned to the frame is also referred to as the final priority or **internal priority**.

4.2.1 Medium Priority

Medium priority

The medium priority is an internal class. Hence, the value can range from 0 through 15.

Based on *the medium (physical interface) the frame is entering on*, a medium priority is assigned to each incoming frame:

- For *Ethernet*, the medium priority is 4.
- For *ATM*, the medium priority depends on the ATM QoS profile set on the interface (ATM PVC). An ATM QoS profile consists of a CTD (Connection Traffic Descriptor) for the upstream and downstream direction. For each CTD, an ATM QoS category was specified. Following table shows the mapping between the ATM QoS category and the medium priority:

ATM QoS category	Internal class
CBR	15
VBR-rt	14
VBR-nrt (low CDVT)	13
GFR (low CDVT)	12
VBR-nrt (high CDVT)	11
GFR (high CDVT)	10
	9
	8
UBR BCS 7	7
UBR BCS 6, ABR	6
UBR BCS 5, UBR-mdcr	5
UBR BCS 4, UBR	4 (default)
UBR BCS 3	3
UBR BCS 2	2
UBR BCS 1	1
UBR BCS 0	0

4| QoS Capabilities of the Ethernet Bridge

4.2.2 Evaluated Priority

Evaluated priority

The evaluated priority is an internal class. Hence, the value can range from 0 through 15.

Based on *the information fields in a received frame*, an evaluated priority is assigned to the frame. The information fields that can be used are:

- The **802.1p user priority field** in the VLAN tag of a VLAN-tagged frame.
- The **IP precedence field** in the IP header of a packet.
- The **DSCP field** in the IP header of a packet.

During configuration, one information field must be selected. The value of the selected field is translated to the evaluated priority.



For more information on these information fields, see “Classification-related frame fields” on page 5.

Configuring the classification criterion

Only one information field can be used to determine the evaluated priority. The information fields can not be combined. The **classification criterion** indicates which information field must be used.

The parameter **ipprec** of the command **:eth bridge ifconfig** indicates the classification criterion. The value of this parameter can be set to:

- **Disabled**: in this case, the 802.1p user priority is used as classification criterion. The 802.1p user priority is translated to the internal class that corresponds to the regenerated user priority.
- **Precedence**: in this case, the IP precedence value is used as classification criterion. If no IP packet is found in the frame, the classification criterion “disabled” is used.
- **DSCP**: in this case, the DSCP value is used as classification criterion. If no IP packet is found in the frame, the classification criterion “disabled” is used.

The parameter can be configured per bridge port. By default, the parameter is set to disabled. To change the selected classification criterion, modify the parameter **ipprec** as follows:

```
=>:eth bridge ifconfig intf=ethport1 ipprec=precedence
=>:eth bridge iflist intf=ethport1
ethport1 : dest : ethif1
          Connection State: connected  Retry: 10
          Priority Tagging: NA (destination switch interface)
          Port: ethport1      PortNr: 1      PortState: forwarding  Interface: up
          Multicast filter: disabled      Dynamic VLAN : disabled
          WAN : disabled
          IGMP snooping : enabled
          Transparent Prio: disabled
          BPDU Filtering : disabled
          Extra Tagging : none
          VLAN: Default VLAN: default  Ingressfiltering: disabled  Acceptvlanonly: disabled
          VLAN: Priority: overwrite  IP Prec: precedence  Priority: 2
          Regeneration table: 0 1 2 3 4 5 6 7
          RX bytes: 49421      frames: 14
          TX bytes: 10512     frames: 22      dropframes: 0
```



When the 802.1p user priority is selected as classification criterion, frames received by the Thomson Gateway will have a VLAN header. It is required to configure explicitly that the bridge is **not VLAN aware**. If the bridge is VLAN aware, you must define all the VLANs that have to pass through the Thomson Gateway. Otherwise, unknown VLANs will be dropped.

4| QoS Capabilities of the Ethernet Bridge

Configuring the default user priority

If the 802.1p user priority field is selected as classification criterion and the received frame is untagged, a default user priority is used. This default user priority can be configured per bridge port.

By default, the default user priority is set to 0. To change the default user priority of a bridge port, execute following command:

```
=>:eth bridge ifconfig intf=ethport1 priority=2
=>:eth bridge iflist intf=ethport1
ethport1 : dest : ethif1
          Connection State: connected   Retry: 10
          Priority Tagging: NA (destination switch interface)
          Port: ethport1      PortNr: 1      PortState: forwarding Interface: up
          Multicast filter: disabled      Dynamic VLAN : disabled
          WAN : disabled
          IGMP snooping : enabled
          Transparent Prio: disabled
          BPDU Filtering : disabled
          Extra Tagging : none
          VLAN: Default VLAN: default Ingressfiltering: disabled Acceptvlanonly: disabled
          VLAN: Priority: overwrite IP Prec: precedence Priority: 2
          Regeneration table: 0 1 2 3 4 5 6 7
          RX bytes: 49421      frames: 14
          TX bytes: 10512     frames: 22      dropframes: 0
```



The user priority regeneration table is only used to remap the 802.1p user priority of a VLAN-tagged frame. The table is **not** used to remap the default user priority assigned to untagged frames.

IP precedence mapping table

The IP precedence mapping table defines the translation from the IP precedence value to an internal class. The following table shows this mapping:

IP precedence (label)	Internal class
6 (Internetwork Control), 7 (Network Control)	15
5 (CRITIC/ECP)	14
4 (Flash Override)	13
	12
3 (Flash)	11
	10
2 (Immediate)	9
	8
1 (Priority)	7
	6
	5
0 (Routine)	4 (default)
	3
	2

4| QoS Capabilities of the Ethernet Bridge

IP precedence (label)	Internal class
	1
	0

Configuring the IP precedence mapping table

The mapping from the IP precedence value to an internal class is *not fixed*. The mapping table can be configured for the whole bridge (not per bridge port) using the parameter **precedencemap** of the command **:eth bridge config**:

```
=>:eth bridge config
Ageing : 300
Filter : no_WAN_broadcast
VLAN   : disabled
IPQoS precedence map for TOS:
  IP priority  QoS internal class
  0            4
  1            7
  2            9
  3            11
  4            13
  5            14
  6            15
  7            15

=>:eth bridge config precedencemap=1,3,5,7,9,11,13,15
=>:eth bridge config
Ageing : 300
Filter : no_WAN_broadcast
VLAN   : disabled
IPQoS precedence map for TOS:
  IP priority  QoS internal class
  0            1
  1            3
  2            5
  3            7
  4            9
  5            11
  6            13
  7            15
```



The precedence map is a string that consists of exactly eight internal classes, separated by **commas**.

DSCP mapping table

Some DSCP values have been standardized by the IETF. The translation from the DSCP value to an internal class is defined by the DSCP mapping table, which is depicted below (EF = Expedited Forwarding, CS = Class Selector, AF = Assured Forwarding):

DSCP decimal value (name)	Internal class
48 (CS6), 56 (CS7)	15
40 (CS5), 46 (EF)	14
32 (CS4), 34 (AF41)	13

4| QoS Capabilities of the Ethernet Bridge

DSCP decimal value (name)	Internal class
36 (AF42), 38 (AF43)	12
24 (CS3), 26 (AF31)	11
28 (AF32), 30 (AF33)	10
16 (CS2), 18 (AF21)	9
20 (AF22), 22 (AF23)	8
8 (CS1), 10 (AF11)	7
12 (AF12), 14 (AF13)	6
	5
0 (CS0, Best-effort)	4 (default)
	3
	2
	1
	0



You can *not* configure the DSCP mapping table.

4| QoS Capabilities of the Ethernet Bridge

4.2.3 Internal Priority

Internal priority

The internal priority is an internal class. Hence, the value can range from 0 through 15.

This priority is the final priority assigned to a received frame, based on a frame's medium priority and evaluated priority.

Configuring the internal priority mechanism

The mechanism to determine the internal priority is configured by the parameter `prioconfig`, which has one of the following values:

- **Disabled:** in this case, the evaluated priority is not used, the medium priority is assigned to the internal priority.
- **Overwrite:** in this case, the medium priority is not used, the evaluated priority is assigned to the internal priority.
- **Increase:** in this case, the medium priority is compared with the evaluated priority:
 - ▶ If the evaluated priority is better (higher value), the internal priority is set to the evaluated priority.
 - ▶ In all other cases, the internal priority is simply the medium priority.

The parameter `prioconfig` of the command `:eth bridge ifconfig` can be configured for each bridge port. By default, the parameter is disabled. To modify the configuration, execute following command:

```
=>:eth bridge ifconfig intf=ethport1 prioconfig=overwrite
=>:eth bridge iflist intf=ethport1
ethport1 : dest : ethif1
          Connection State: connected   Retry: 10
          Priority Tagging: NA (destination switch interface)
          Port: ethport1      PortNr: 1      PortState: forwarding Interface: up
          Multicast filter: disabled      Dynamic VLAN : disabled
          WAN : disabled
          IGMP snooping : enabled
          Transparent Prio: disabled
          BPDU Filtering : disabled
          Extra Tagging : none
          VLAN: Default VLAN: default Ingressfiltering: disabled Acceptvlanonly: disabled
          VLAN: Priority: overwrite IP Prec: precedence Priority: 2
          Regeneration table: 0 1 2 3 4 5 6 7
          RX bytes: 49421 frames: 14
          TX bytes: 10512 frames: 22 dropframes: 0
```

Summary: priority mapping table

Several mapping tables were described in the previous subsections. The following table summarizes these mapping tables:

Regenerated user priority	IP precedence (ToS)	DSCP (DiffServ)	ATM QoS category	Internal class
7	6, 7	CS6, CS7	CBR	15
6	5	CS5, EF	VBR-rt	14
	4	CS4, AF41	VBR-nrt (low CDVT)	13

4| QoS Capabilities of the Ethernet Bridge

Regenerated user priority	IP precedence (ToS)	DSCP (DiffServ)	ATM QoS category	Internal class
		AF42, AF43	GFR (low CDVT)	12
	3	CS3, AF31	VBR-nrt (high CDVT)	11
5		AF32, AF33	GFR (high CDVT)	10
	2	CS2, AF21		9
4		AF22, AF23		8
	1	CS1, AF11	UBR BCS 7	7
3		AF12, AF13	UBR BCS 6, ABR	6
			UBR BCS 5, UBR-mdcr	5
0	0	BE, CS0	UBR BCS 4, UBR	4 (default)
			UBR BCS 3	3
2			UBR BCS 2	2
			UBR BCS 1	1
1			UBR BCS 0	0

4| QoS Capabilities of the Ethernet Bridge

4.3 Queuing

Introduction

The use of Ethernet QoS on the Thomson Gateway only improves the performance of egress traffic if the egress interfaces are QoS enabled.

Based on the assigned internal class, a frame can be queued into the *correct queue* at the egress side. In this subsection, we briefly describe two types of egress interfaces:

- *Wireless interface*
- *ATM interface (single PVC)*

The internal class can also be used to map the frames to the *correct ATM PVC*. To this end, following concept is implemented:

- *ATM interface bundle (multiple PVCs)*

QoS on the wireless interface

In case of a wireless interface, WMM (Wi-Fi Multi Media) supports QoS:

- First, the internal class assigned to a frame is mapped to the corresponding 802.1p user priority.
- Next, the 802.1p user priority is mapped to a WMM AC (Access Category), which corresponds to a queue.

To enable WMM QoS support on the wireless interface, execute following command:

```
=>:wireless qos config mode=wmm
```

The class-to-queue mapping is shown in following table. The four access categories correspond to four queues.

Internal class	VLAN user priority	Access category (AC)
15	7	AC_VO (Voice)
14	6	AC_VO (Voice)
13		
12		
11		
10	5	AC_VI (Video)
9		
8	4	AC_VI (Video)
7		
6	3	AC_BE (Best-effort)
5		
4	0	AC_BE (Best-effort)
3		
2	2	AC_BK (Background)

4| QoS Capabilities of the Ethernet Bridge

Internal class	VLAN user priority	Access category (AC)
1		
0	1	AC_BK (Background)



You cannot configure the class-to-queue mapping, which is *fixed*.

QoS on an ATM interface

IP QoS can be enabled per ATM interface, thus per PVC.

- If IP QoS is *disabled*, no IP QoS queues will be instantiated. As a result, all data (either classified or unclassified) is sent to a fixed destination queue (the AAL5 segmentation queue) for that PVC.
- If IP QoS is *enabled*, all six IP QoS queues are instantiated on the PVC. A packet queuing mechanism, which is active on ATM interface level is enabled. This queuing mechanism handles incoming packets according to their internal class. If no internal class is assigned to the data, the data is forwarded to the lowest priority best-effort queue.

To enable IP QoS on an ATM interface with destination **phonebookentry**, execute following command:

```
=>:ipqos config dest=phonebookentry state=enabled
```

The class-to-queue mapping is shown in following table. By default, the first queue (EF) has strict priority over the four queues (AF) with WFQ scheduling. These four queues have strict priority over the best-effort queue.

Internal class	QoS queue (traffic type)
15	5 (Real-time EF data)
14	
13	4 (AF Class 4)
12	
11	3 (AF Class 3)
10	
9	2 (AF Class 2)
8	
7	1 (AF Class 1)
6	

4| QoS Capabilities of the Ethernet Bridge

Internal class	QoS queue (traffic type)
5	0 (Best-effort)
4 (default)	
3	
2	
1	
0	



You cannot configure the class-to-queue mapping, which is *fixed*.

ATM interface bundles

ATM interface bundles can be used to forward frames over multiple PVCs based on the internal class assigned to each frame.

If the different PVCs would be connected to the bridge via different bridge ports, the bridge learning process would learn the same remote MAC address from different bridge ports. To prevent this, the PVCs are bundled into an ATM interface bundle, which is connected to the bridge via a single bridge port. Up to 6 PVCs can be added to a single ATM interface bundle.

To forward frames over multiple PVCs based on their internal class, the used ATM interface bundle policy is the *priority mapping policy*. Within the priority mapping policy, every PVC that is a member of the ATM interface bundle can have a *priority range* configured (as low and high selector values). All frames with a priority within the first matching member priority range will be forwarded to this PVC.



The ATM interface bundles feature is only available on selected products.

4.4 Priority Tagging of Egress Traffic

What is priority tagging?

Priority tagging means that a VLAN tag is added to a frame with the 802.1p user priority field set to a given value and the VLAN ID equal to 0. The value that is assigned to the 802.1p user priority field is *based on the internal class* of the frame.

This feature is intended to be used by:

- Bridge ports located at the *WAN side*.
- Bridge ports that are *untagged* member of a VLAN on the bridge.

Configuring priority tagging

The priority tagging feature can be configured *per bridge port*.

By default, priority tagging is disabled. To enable the feature, execute following command:

```
=>:eth bridge ifconfig intf=eth_data priotag=enabled
```



The priority tagging feature is *not* available (NA) for bridge ports that are connected to a switch port.

Listing information

To check the state of the priority tagging feature, execute following command:

```
=>:eth bridge iflist intf=eth_data
eth_data : dest : atm_0.38
          Connection State: connected  Retry: 10
          Priority Tagging: Enabled
          Port: wan0          PortNr: 6          PortState: forwarding Interface: up
          Multicast filter: disabled          Dynamic VLAN : disabled
          WAN : enabled
          IGMP snooping : enabled
          Transparent Prio: disabled
          BPDU Filtering : disabled
          Extra Tagging : none
          VLAN: Default VLAN: default Ingressfiltering: disabled Acceptvlanonly: disabled
          VLAN: Priority: disabled IP Prec: disabled Priority: 0
          Regeneration table: 0 1 2 3 4 5 6 7
          RX bytes: 49421          frames: 14
          TX bytes: 10512          frames: 22          dropframes: 0
```

4| QoS Capabilities of the Ethernet Bridge

4.5 802.1p User Priority Preservation

What is 802.1p user priority preservation?

At the egress side, each bridge port can be configured with an 802.1p user priority preservation flag:

- If user priority preservation is **enabled**, the bridge port tags each frame with the **original 802.1p user priority** as present in that frame when it was received at the ingress side. Therefore, the priority preservation feature is also referred to as the priority transparency feature.
- If user priority preservation is **disabled**, the priority in each tagged frame is based on the **internal class** of the frame.

This feature is intended to be used by:

- Bridge ports located at the **WAN side**.
- Bridge ports that are **tagged** member of a VLAN on the bridge.

Configuring 802.1p user priority preservation

The user priority preservation feature can be configured **per bridge port**.

By default, user priority preservation is disabled. To enable the feature, execute following command:

```
=>:eth bridge ifconfig intf=eth_data priotransparent=enabled
```

Listing information

To check the state of the user priority preservation feature, execute following command:

```
=>:eth bridge iflist intf=eth_data
eth_data : dest : atm_0.38
          Connection State: connected   Retry: 10
          Priority Tagging: Disabled
          Port: wan0                    PortNr: 6           PortState: forwarding Interface: up
          Multicast filter: disabled     Dynamic VLAN      : disabled
          WAN                          : enabled
          IGMP snooping : enabled
          Transparent Prio: enabled
          BPDU Filtering : disabled
          Extra Tagging  : none
          VLAN: Default VLAN: default   Ingressfiltering: disabled Acceptvlanonly: disabled
          VLAN: Priority: disabled IP Prec: disabled Priority: 0
          Regeneration table: 0 1 2 3 4 5 6 7
          RX bytes: 49421 frames: 14
          TX bytes: 10512 frames: 22 dropframes: 0
```

Applications

802.1p user priority preservation enables classification without changing the original 802.1p user priority:

- On the one hand, an incoming frame is assigned to an internal class, for example based on the value of the DSCP field. This internal class is used to queue the frame in the correct queue.
- On the other hand, the original 802.1p user priority of the (tagged) frame is not overwritten. If priority tagging would be based on the internal class assigned to the frame during classification, the original 802.1p user priority may be overwritten.

5 Logical Ethernet Interfaces and QoS

Introduction

Logical Ethernet interfaces can be configured to support following Ethernet QoS feature:

- **Priority tagging** of traffic at the WAN-side

Priority tagging parameter

Priority tagging by a logical Ethernet interface can be enabled using the **priotag** parameter of the **:eth ifconfig** command.

- If the parameter is **disabled**, priority tagging is not supported. By default, priority tagging of egress traffic is disabled.
- If the parameter is **enabled**, priority tagging is supported.

```
=>:eth ifconfig intf=eth_data priotag=enabled
=>:eth iflist intf=eth_data
eth_data      : Dest: atm_data
               Connection State: not-connected  Retry: 10
               WAN: Enabled
               Priority Tagging: Enabled
               PortNr: (unassigned)
               VLAN: default
```



The priority tagging feature of logical Ethernet interfaces is very similar to the priority tagging feature of bridge ports. For more information on the priority tagging feature of bridge ports, see “4.4 Priority Tagging of Egress Traffic” on page 31.

Configuring WAN-side priority tagging

To create a logical Ethernet interface that transmits priority-tagged frames, execute following steps:

- 1 Create a logical Ethernet interface:

```
=>:eth ifadd intf=eth_data
```

- 2 Configure the logical Ethernet interface to support WAN-side priority tagging:

```
=>:eth ifconfig intf=eth_data dest=atm_data wan=enabled priotag=enabled
```

- 3 Activate the logical Ethernet interface:

```
=>:eth ifattach intf=eth_data
```


6 Layer 2 QoS and Layer 3 QoS Interaction

Introduction

At the *ingress side*, incoming frames are classified at Layer 2, as described in this document. This means that an internal class is assigned to each received Ethernet frame, thus also to the IP packet in the payload of the frame. Following questions arise with regard to this assigned internal class:

- What is the internal class of a packet that also passes through the IP router of the Thomson Gateway?
- What is the internal class of a packet that is originated by the Thomson Gateway?

At the *egress side*, WAN-side priority tagging can be used at Layer 2. Following question arises:

- Is Layer 3 QoS influenced by Layer 2 QoS and vice versa?

This chapter shortly answers these questions.

Internal class after Layer 3 classification

When a packet passes through the IP router of the Thomson Gateway, the packet is classified again. This is done in three steps:

- 1 The Layer 3 classification assigns a *label* to each packet. This label is only of internal significance and can be used in packet forwarding (routing label) and QoS definition (QoS label).
- 2 The label is mapped to the *packet class* associated with it, configured by the `defclass` parameter of the `:label modify` command:
 - ▶ **0...15**: in this case, the parameter value indicates the packet class.
 - ▶ **Dscp**: if this value is used, the packet class is set to the internal class corresponding with the value of the DSCP field of the packet.
 - ▶ **Default**: if the parameter value is set to default, the Thomson Gateway's default value 4 is used.
- 3 The packet class is mapped to the *Layer 3 internal class*. By configuring the `classification` parameter of the `:label modify` command, it is determined how the Layer 2 internal class is taken into account:
 - ▶ **Ignore**: in this case, the assigned packet class is ignored. The internal class assigned by Layer 2 is not overwritten.
 - ▶ **Overwrite**: in this case, the internal class assigned by Layer 2 is ignored. The internal class is set to the packet class.
 - ▶ **Increase**: in this case, the label classification only sets the internal class if the packet class value is higher than the internal class value assigned by Layer 2.

Internal class of Thomson Gateway originated packets

Packets that are originated by the Thomson Gateway itself also receive a label. The corresponding packet class of the label is the internal class assigned to the packet.

- The system service manager can be used to explicitly assign a QoS label to an application. If QoS label assignment is used for an application, no label classification is done.
- Application helpers can be used to establish signalling (parent)/ data (child) relationships. Label inheritance means that when parent connections get labels associated, these labels are also automatically assigned to the child connections. It is also possible to define separate child QoS labels.

6| Layer 2 QoS and Layer 3 QoS Interaction

IP marking and priority tagging

While Layer 2 supports priority tagging, IP marking is used at Layer 3:

- **IP marking:** at Layer 3, IP marking can be enabled per label using the `tosmarking` parameter of the `:label modify` command. IP marking can be used to set the IP ToS byte, IP precedence field or DSCP field of egress packets, regardless the internal class of the packet.
- **Priority tagging:** at Layer 2, priority tagging can be used. The 802.1p user priority is based on the internal class of the egress frames. Depending on the configuration, this internal class can be influenced by the Layer 3 classification.

THOMSON Telecom Belgium

Prins Boudewijnlaan 47
2650 Edegem

www.thomson-broadband.com

© Thomson 2008. All rights reserved.
E-DOC-CTC-20080229-0002 v1.0.

