Nortel Networks Multiservice Switch

7400/15000/20000

ATM Traffic Shaping and Policing

NN10600-706

Nortel Networks Multiservice Switch 7400/15000/20000

ATM Traffic Shaping and Policing

Publication: NN10600-706
Document status: Standard
Document version: 6.1S1
Document date: August 2004

Printed in Canada

Publication history

August 2004

6.1S1 Standard
General availability. Contains information on Nortel Networks Multiservice Switch 7400, 15000, and 20000 for the PCR6.1 release.

Contents

List of figures

List of tables

About this document

This document contains information on Nortel Networks Multiservice Switch ATM traffic management controls for shaping and policing.

The following topics are discussed in this section:

Who should read this document and why

This document is intended for persons responsible for performing the following:

- network planning and engineering
- installation and configuration
- operations
- fault management

What you need to know

Be familiar with the operating principles of the Nortel Networks Multiservice Switch system before you read this document. To fully understand the information in this guide, be familiar with ATM and the Open Systems Interconnection (OSI) model, and standards and recommendations published by the ATM Forum and the International Telecommunication Union (ITU).

Table 1
Multiservice Switch ATM supporting information

ATM supporting information for this document	Location of supporting information
For an explanation of the following FP types: • CQC, AQM, APC, GQM, PQC	see the description of ATM function processors in NN10600-700 *Nortel Networks Multiservice Switch 7400/15000/20000 ATM Technology Fundamentals*
For information about addressing, signaling, and routing	see NN10600-702 *Nortel Networks Multiservice Switch 7400/ 15000/20000 ATM Routing and Signalling Fundamentals*
For traffic management concepts	see NN10600-705 *Nortel Networks Multiservice Switch 7400/ 15000/20000 ATM Traffic Management Fundamentals*
For cell queuing and queue scheduling capabilities	see NN10600-707 *Nortel Networks Multiservice Switch 7400/ 15000/20000 ATM Queuing and Scheduling Fundamentals*
For connection admission controls and bandwidth management	see NN10600-708 *Nortel Networks Multiservice Switch 7400/ 15000/20000 ATM CAC and Bandwidth Fundamentals*

What's new in this document

There were no new features added to this document.

Other changes made to this document include the following:

- The terms Passport and PVG have been rebranded in conjunction with the new Nortel Networks' brand simplified naming format. Passport is now referred to as the Nortel Networks Multiservice Switch, and PVG is now Media Gateway 7480/15000. For more information on the product rebranding, refer to NN10600-000 *Nortel Networks Multiservice Switch 7400/15000/20000 What's New in PCR6.1*.

- added the section "Policing enforcer violation counts for GQM-based FPs" (page 91).

Text conventions

This document uses the following text conventions:

- `nonproportional spaced plain type`

  Nonproportional spaced plain type represents system generated text or text that appears on your screen.

- `nonproportional spaced bold type`

  Nonproportional spaced bold type represents words that you should type or that you should select on the screen.

- *italics*

  Statements that appear in italics in a procedure explain the results of a particular step and appear immediately following the step.

  Words that appear in italics in text are for naming.

- `[optional_parameter]`

  Words in square brackets represent optional parameters. The command can be entered with or without the words in the square brackets.

- `<general_term>`

  Words in angle brackets represent variables which are to be replaced with specific values.

- UPPERCASE, lowercase

  Nortel Networks Multiservice Switch commands are not case-sensitive and do not have to match commands and parameters exactly as shown in this document, with the exception of string options values (for example, file and directory names) and string attribute values.

- |

  This symbol separates items from which you may select one; for example, ON|OFF indicates that you may specify ON or OFF. If you do not make a choice, a default ON is assumed.

- ...

  Three dots in a command indicate that the parameter may be repeated more than once in succession.

The term absolute path name refers to the full specification of a path starting from the root directory. Absolute path names always begin with the slash (/) symbol. A relative path name takes the current directory as its starting point, and starts with any alphanumeric character (other than /).

Related documents

See the following documents for related information:

- NN10600-030 *Nortel Networks Multiservice Switch 7400/15000/20000 Overview*
- NN10600-700 *Nortel Networks Multiservice Switch 7400/15000/20000 ATM Technology Fundamentals*
- NN10600-702 *Nortel Networks Multiservice Switch 7400/15000/20000 ATM Routing and Signalling Fundamentals*
- NN10600-705 *Nortel Networks Multiservice Switch 7400/15000/20000 ATM Traffic Management Fundamentals*
- NN10600-707 *Nortel Networks Multiservice Switch 7400/15000/20000 ATM Queuing and Scheduling Fundamentals*
- NN10600-708 *Nortel Networks Multiservice Switch 7400/15000/20000 ATM CAC and Bandwidth Fundamentals*
- NN10600-715 *Nortel Networks Multiservice Switch 7400/15000/20000 ATM Fault and Performance Management*
- NN10600-720 *Nortel Networks Multiservice Switch 7400/15000/20000 AAL1 Circuit Emulation Operations*

- NN10600-730 *Nortel Networks Multiservice Switch 7400/15000/20000 Inverse Multiplexing for ATM Operations*
- NN10600-420 *Nortel Networks Multiservice Switch 7400/15000/20000 Operations: Trunking*
- NN10600-920 *Nortel Networks Multiservice Switch 7400/15000/20000 Operations: Frame Relay to ATM Interworking*
- NN10600-060 *Nortel Networks Multiservice Switch 7400/15000/20000 Component Reference*
- NN10600-500 *Nortel Networks Multiservice Switch 6400/7400/15000/20000 Alarms Reference*
- NN10600-560 *Nortel Networks Multiservice Switch 7400/15000/20000 Accounting*
- *Nortel Networks Multiservice Switch Release Notes*

For a list of related industry standards, see the NN10600-700 *Nortel Networks Multiservice Switch 7400/15000/20000 ATM Technology Fundamentals*.

How to get more help

For information on training, problem reporting, and technical support, see the “Nortel Networks support services” section in the product overview.

Chapter 1
Multiservice SwitchTraffic shaping

This chapter describes general principles of traffic shaping as implemented for Nortel Networks Multiservice Switch nodes and function processors (FPs) in an ATM network.

Information is provided in the following sections:

For information on traffic shaping that is specific to FP types, see the following chapters:

Overview of traffic shaping

Traffic shaping ensures that transmitted traffic conforms to subscribed traffic parameters by regulating the emission interval of cells transmitted on the link. Traffic shaping regulates the departure of cells for each connection. Shaping limits the burst length and provides regular spacing between cells. By pacing cell transmission, the service provider avoids cell discard at points where traffic policing applies. Shaping also allows the network manager to balance the quality of service (QoS) requirements against network costs.

Traffic shaping is required when enterprise or carrier switches are connected through another public frame relay or ATM network, and there is a need to conform to the traffic contract specified by the ATM public carrier. Shaping occurs on the transmit queues for the link. Traffic shaping is recommended for any connection to an external ATM network.

Linear traffic shaping regulates each connection with a single rate. It means the cell interval between the current cell departure and the next scheduled departure equals the reciprocal of the connection's peak cell rate (PCR).

Dual-rate inverse usage parameter control (UPC) traffic shaping regulates each connection using a PCR and a sustained cell rate (SCR) with maximum burst size (MBS). An inverse UPC shaper can serve the cells of a connection consecutively at its PCR up to its MBS provided that the service on average remains below its SCR.

The figure "Model of how traffic shaping works" (page 21) shows at a high level how traffic shaping works. In this illustration, the shaping takes unevenly spaced incoming cells from the backplane and smooths out the delay variation for transmission over the link. In effect, the shaper adds some delay as a result of spacing transmitted cells to regular intervals. However, this delay is balanced by the reduction of removal of bursty traffic which results in more efficient resource allocation across the network.

Figure 1
Model of how traffic shaping works

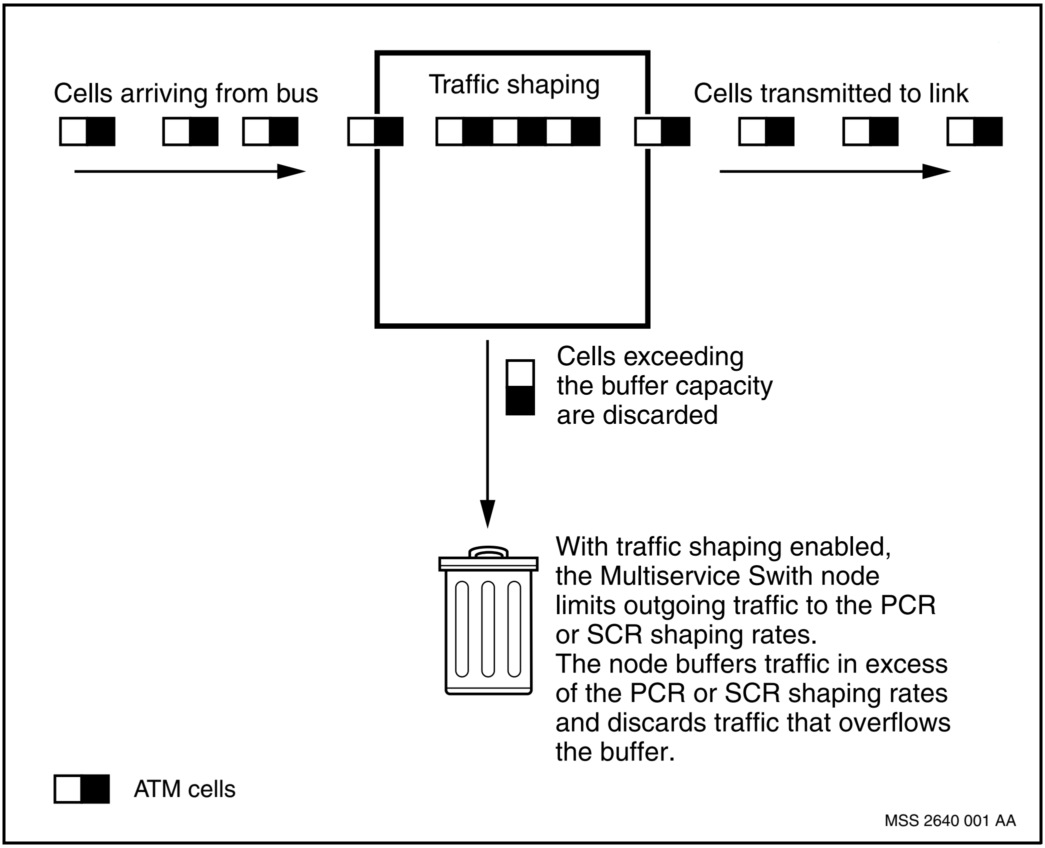

There are several applications for traffic shaping in an ATM network. The most critical application is when an enterprise switch must pace cell flow to the service provider switch. Pacing ensures that the cell arrival rate does not exceed the agreed peak cell rate (PCR) and sustained cell rate (SCR) as defined in the service contract. Exceeding PCR or SCR triggers the policing function to discard cells. If traffic on a connection exceeds the shaped PCR, traffic shaping buffers the excess cells. If the buffer fills completely (the traffic on the connection significantly exceeds the PCR for an extended period) the node discards cells. When the network transports frames over a Nortel Networks Multiservice Switch trunk over ATM and excessive congestion results in cell discard, the node can also discard entire frames. Cell and frame discard indicates either unusual fault conditions on the network or the need to re-engineer and expand services.

Traffic shaping also applies within a network as a method of smoothing out bursty traffic sources. By reducing or eliminating bursty traffic, the service provider can more efficiently allocate memory on nodes throughout the network, which results in increased network efficiency.

Shaper congestion

Shaper congestion occurs when traffic input exceeds the bandwidth that is available to the shaper. For example, the bandwidth available to a shaper could be the link rate if the shaper is set at emission priority zero (EP0). If the shaper is at an EP lower than EP1, the available bandwidth would be the bandwidth that is guaranteed for the EP out of the total bandwidth that is left over from the demands of EP0 and EP1.

On these cards, a congested shaper may show a throughput loss on shaped connections: 4pOC3SmIrAtm (NTHR21CA, NTHW15AA) and 4pOC3MmAtm (NTHR17CA, NTHW05AA). Such throughput loss tends to penalize connections with higher shaping rates more than connections with lower shaping rates. Setting the attribute *shapeRecoupPolicy* in the ATM service category to the value *maximumEfficiency* will improve the shaping throughput in most cases, but may not eliminate the throughput loss for connections with high shaping rates.

On the same cards, when there are VCs shaped in a higher priority EP and when the traffic input to a lower shaped EP is close to or exceeds the bandwidth available to this shaped lower EP, the VCs with large shaping rates in the lower EP may have a throughput loss. If the shaping rate that is configured for a connection is very high compared to the link rate, and if the connection is having a throughput loss, use of the value *maximumEfficiency* for the attribute *shapeRecoupPolicy* may not improve the throughput for this connection. (In most cases, a shaper shapes many connections whose shaping rates are low compared to the link rate.)

Summary of shaping options by function processor type

The table “Shaping options by transmit traffic descriptor type and FP type” (page 23) summarizes shaping options by function processor (FP) type. The type of shaping which is actually applied is indicated by an operational attribute under each connection.

Table 2
Shaping options by transmit traffic descriptor type and FP type

Descriptor type	ATM IP and GQM-based FPs	CQC-based FPs
1, 2	no shaping	no shaping
3, 4, 5	linear shaping	linear shaping
6, 7, 8	VBR-shaping (dual rate inverse UPC) or linear shaping (default)	linear shaping
9	linear shaping	linear shaping

Configuration parameters for traffic shaping

For ATM IP, CQC-based, and GQM-based function processors (FPs), you can enable or disable traffic shaping at two levels:

- the service category level (all FPs)
- the connection level (all FPs)

All Nortel Networks Multiservice Switch nodes support queuing strategies and traffic shaping according to the following rules:

- If enabled at the service category level, shaping can be either enabled or disabled at the connection level.
- If disabled at the service category level, shaping is disabled and cannot be enabled for any connection associated with that interface.
- If the connection is configured with traffic descriptor types 1 or 2, traffic shaping is automatically disabled.

- If a connection has traffic shaping enabled, per-VC queuing is implemented at the connection shaping rate. For a switched connection, the shaping of the connection is determined by the setting for traffic shaping at the service category level.
- If the connection has traffic shaping disabled and per-VC queuing enabled by service category, per-VC queuing is implemented by shaping the connection at the fastest shaping rate available to the ATM interface. The maximum shaping rate available to a connection depends on
  - — the number of ports being shaped on the FP
  - — the global scaling factor. See “Example: per-VC queuing enabled with shaping disabled” (page 25).
- If the connection has traffic shaping disabled and per-VC queuing disabled, common queuing is performed by enqueuing cells or frames to the appropriate emission priority common queue for the connection.
- For APC-based FPs, you need to configure the connections at emission priorities other than zero if the connection has traffic shaping disabled and the service category unshaped. However, to enable traffic shaping on a given class or connection, you need to map the entire class to EP0.
- For AQM-based ATM interfaces, unshaped connections are not allowed when shaping is enabled on an EP being used by the IMA link group. If all connections that belong to a shaped service category do not have shaping enabled, queue congestion and high cell discards may occur.
- GQM-based FPs allow dynamic connections which request a shaping rate below the minimum to pass. When the suggested shaping rate is less than the minimum supported (100 cells per second), the connection will be admitted but shaped at 100 cells per second instead of the requested rate.

When the node determines the shaping rate based on PCR or SCR, the general rule is to round down on fractions or use the next lowest available rate. This approach ensures that the node is not sending traffic at a rate that is higher than the policer at the other end of the connection expects.

For UBR with MDCR connections, the node determines the shaping rate based only on the PCR value.

Example: per-VC queuing enabled with shaping disabled

This configuration is best illustrated by example. On a CQC 3-port E3 FP, each port has a link rate of 34.368 Mbits/sec (effectively 80 000 cells/s). If all three ports have shaping enabled, the maximum available shaping rate is 58 962 cells/s for each connection on any of these ports, assuming a global scaling factor of 1.0. See the table of defined traffic shaping rates for CQC-based FPs in “Shaping and policing on CQC-based FPs” (page 43). By configuring multiple connections per port in balance with the available shaping rate, you can be sure to use all available bandwidth on the link.

In this example, there is nothing to be gained by having only one port shaped on the CQC 3-port E3 FP. Referring to the table of defined traffic shaping rates for CQC-based FPs in “Shaping and policing on CQC-based FPs” (page 43), the next highest cell rate is 83 333 cells/s (shaping identifier 8). Since the E3 link rate is lower at 80 000 cells/s, the higher rate defined for shaping identifier 8 can never be achieved.

In this release, operational attributes do not show an available shaping rate that is higher than the link rate since a connection can never be shaped at that higher rate.

Shaping rate for bandwidth-elastic connections

Only connections that are bandwidth-elastic can decrease or increase their speed in response to changes in bandwidth available over an ATM link. When elastic connections decrease their speed, the shaping rate adjusts to the new cell rate. The idea behind traffic shaping for elastic connections is that no cells will be dropped on any traffic conforming to the new shaping rate derived for elastic connections in a state of reduced bandwidth.

> ***Note:*** To guarantee that cells are preserved, the ATM application controlling the traffic must adjust to the new cell rate.

Since PCR and SCR values do not reflect a state of reduced bandwidth, a special input parameter is used by the shaping rate algorithm to determine PCR or SCR shaping rates for elastic connections.

For elastic connections, current cell rate (CCR) is used as the input parameter to the shaping rate algorithm. CCR for a connection is determined as follows:

CCR = <ecr> - <bwReduction>

where:
<ecr> is the equivalent cell rate for the connection
<bwreduction> is the amount by which the speed of an elastic connection is reduced in cells per second (this is the value for the *bandwidthReduction* attribute for a *Vcc* component). CCR is used only when the value for <bwreduction> is greater than zero.

To determine the PCR or SCR shaping rates for an elastic connection, CCR is used as the input parameter to the shaping rate algorithm instead of PCR and SCR. The use of CCR guarantees that the PCR or SCR shaping rate is always above the value for CCR, where this relationship ensure that adequate bandwidth is available to the elastic connection. See “VBR shaping using PCR or SCR as the shaping rate” (page 75). The algorithm specifies that PCR or SCR shaping rates equals the next available shaping rate greater than TSR. In the case of elastic connections, TSR equals CCR.

Chapter 2
Multiservice SwitchTraffic policing

This chapter describes general principles of traffic policing as implemented for Nortel Networks Multiservice Switch nodes and function processors (FPs) in an ATM network. Information is provided in the following sections:

- "Overview of usage parameter controls" (page 28)
- "UPC enforcement algorithms" (page 29)
- "Generic cell rate algorithm" (page 30)
- "UPC operating characteristics" (page 38)
- "Summary of limitations on UPC" (page 38)
- "Additional considerations for UPC" (page 39)

For information on traffic policing that is specific to FP types, see the following chapters:

- "Shaping and policing on CQC-based FPs" (page 43)
- "Shaping and policing on ATM IP FPs" (page 63)
- "Shaping and policing on APC/PQC-based FPs" (page 83)
- "Shaping and policing for GQM-based FPs" (page 89)

Overview of usage parameter controls

Traffic policing is the set of operations that the node applies to monitor and control traffic at an end-system access point. Policing mechanisms evaluate both the traffic offered by end stations at the access point and the validity of the ATM connection. Nortel Networks Multiservice Switch nodes implement usage parameter controls (UPC) for all traffic policing.

For all Nortel Networks Multiservice Switch ATM FP types, UPC undertakes the following operations:

- monitors traffic arriving on a connection
- determines if the traffic conforms to the contract for the connection
- discards or tags non-conforming traffic

Through these operations, UPC protects network resources from malicious as well as unintentional misbehavior, which can affect the QOS of other already established connections, by detecting violations of negotiated parameters and taking appropriate actions. Such actions include cell discard and cell tagging.

The main purpose of UPC is to protect network resources from traffic demands in excess of those defined through the traffic contract, which can affect the service category of established connections. This purpose is achieved by detecting violations of negotiated parameters and taking appropriate actions, such as cell tagging and cell discard. Cells are tagged in congestion situations.

Tagging is done through UPC in conformance with the network signaling protocol, where the UPC changes the cell loss priority (CLP) bit in the cell header from 0 (not tagged) to 1 (tagged). Tagged cells may be discarded at any point further down the connection, where preference is given to untagged cells over tagged cells (where discard priority is enforced). Cell tagging affects discard priority of traffic.

An AAL5 partial packet discard (PPD) mechanism is also implemented in this policer. PPD is automatically activated at connection setup if the frame discard option is selected. If any cell of a frame is discarded by the policer, PPD then discards all subsequent cells of that frame up to but not including the end-of-message (EOM) cell of the frame.

UPC on Multiservice Switch nodes is based on a continuous-state leaky bucket algorithm that controls the rate and maximum burst size (MBS) that can flow into the ATM network; see "Relationship between UPC and CDVT" (page 39).

UPC functions are based on the conformance definition of a negotiated connection, which is determined by the traffic descriptor type and the ATM service category of an ATM cell stream (which are configured or derived from signaled BBC-IE). From the 5-byte ATM cell header, UPC can identify two priority levels using the CLP bit (which is defined by the network's signaling protocol). Untagged cells have CLP=0 and are viewed as more important; tagged cells have CLP=1 and are viewed as less important. Tagged cells may be discarded in preference to untagged cells. Due to UPC, CLP=0 cells may be tagged as CLP=1. UPC also includes OAM and resource management (RM) cells in its calculations.

UPC uses the traffic descriptor parameters to identify non-conforming cells. See the section on traffic contracts in NN10600-705 *Nortel Networks Multiservice Switch 7400/15000/20000 ATM Traffic Management Fundamentals* for information on traffic descriptor types and traffic parameters.

UPC enforcement algorithms

All Nortel Networks Multiservice Switch ATM FPs implement UPC enforcement through generic cell rate algorithms (GCRA). Depending on the traffic type, the node applies one or two GCRAs to identify non-conforming cells. The UPC enforcer then tags or discards these cells. The results of the algorithms are always rounded up to ensure that the policer can accept cells over a shaped connection at the expected rate (no unnecessary or unwarranted tagging or discards).

With ATM IP FPs, you have the option of counting non-conforming cells.

With the following FPs, you also have up to two GCRA enforcers in tandem for each connection so that up to two rates can be monitored and enforced upon the connection:

- APC-based
- AQM-based

- AQS-based
- GQM-based

UPC counts discarded cells for each VCC and VPC. UPC discards cells for the following reasons:

- non-conformance with UPC
- congestion on the Multiservice Switch node backplane
- AAL5 reassembly errors

UPC counts discarded cells under the same operational attribute that counts received cells that UPC discards due to backplane congestion. For data collection on the number of cells that a node tags through UPC, ATM IP FPs collect these counts. CQC-based FPs do not.

The information in this section describes how the Multiservice Switch node applies UPC enforcers depending on the traffic descriptor type. For each GCRA, see the section on traffic contracts NN10600-705 *Nortel Networks Multiservice Switch 7400/15000/20000 ATM Traffic Management Fundamentals* for a summary of the configurable parameters.

See the section on configuring traffic descriptor parameters in NN10600-705 *Nortel Networks Multiservice Switch 7400/15000/20000 ATM Traffic Management Fundamentals* for descriptions of parameter usage.

Generic cell rate algorithm

The GCRA algorithm is defined with two parameters, I and L, and is expressed as

$$\mathrm{GCRA}(I,L)$$

where:
I is the increment
L is the limit

These GCRA parameters are based on the traffic descriptor parameters configured for each connection:

- the *I* parameter is based on either PCR or SCR in traffic descriptor parameters 1 and 2
- the *L* parameter is based on MBS in parameter 3 and CDVT in parameters 2 or 4

Note: CDVT is taken from parameter 2 for traffic descriptor type 9 and for parameter 4 for all other traffic descriptor types.

Parameter 5 - RSR - does not apply to UPC. See the table on Nortel Networks Multiservice Switch node traffic descriptor types and common configurable parameter rules in the section on traffic contracts in NN10600-705 *Nortel Networks Multiservice Switch 7400/15000/20000 ATM Traffic Management Fundamentals*.

Note: For UBR with MDCR, refer to the following ATM Form document: af-tm-0150.000.

Example 1: single UPC enforcer

For example, traffic descriptor type 3 has a single UPC enforcer, which is defined as

GCRA(1/PCR, CDVT)

Because a second UPC enforcer is not required for traffic descriptor type 3, traffic descriptor parameter 2 is not used. The application of this enforcer is illustrated in the figure "Traffic descriptor type 3 UPC enforcement" (page 34).

Example 2: dual UPC enforcers

As a second example, examine the GCRAs required for traffic descriptor type 6, which are expressed as

GCRA(1/PCR, CDVT)   GCRA(1/SCR, BT+CDVT)

Because traffic descriptor parameter 2 is used for traffic descriptor type 6, a second UPC enforcer is required. Because traffic descriptor parameter 3 is used, the limit is based on burst tolerance (BT) and CDVT taken together. BT is a function of all traffic descriptor parameters, and is expressed as:

$$BT = (MBS - 1)(1/SCR - 1/PCR)$$

Note the strong position that MBS (traffic descriptor parameter 3) assumes in this relationship. The significance of the intervals expressed by 1/SCR and 1/PCR are as follows:

- As the difference between these intervals increases, burst tolerance increases.
- When the intervals are equal, burst tolerance is zero. In other words, deviation upward from the SCR is not permitted.

The application of these enforcers is illustrated in the figure “Traffic descriptor type 6 UPC enforcement” (page 36).

UPC enforcement for traffic descriptor types

The table “Summary of UPC enforcement by traffic descriptor type” (page 32) summarizes GCRAs and tag/discard actions for each traffic descriptor type.

Table 3
Summary of UPC enforcement by traffic descriptor type

<table>
<tr><th>Traffic descriptor type</th><th>GCRA 1 (see Note 1)</th><th>GCRA 2</th><th>Tag or discard</th></tr>
<tr><td>1</td><td colspan="3">indicates no traffic descriptor type for the connection</td></tr>
<tr><td>2</td><td colspan="3">traffic descriptor type not used</td></tr>
<tr><td>3</td><td>1/PCR, CDVT</td><td>n/a</td><td>discard at GCRA 1</td></tr>
<tr><td>4</td><td rowspan="2">1/PCR, CDVT</td><td rowspan="2">1/PCR0, CDVT</td><td>discard at GCRA 2</td></tr>
<tr><td>5</td><td>tag at GCRA 2</td></tr>
<tr><td>6</td><td>1/PCR, CDVT</td><td>1/SCR, BT + CDVT</td><td>discard at GCRA 2</td></tr>
<tr><td colspan="4">(Sheet 1 of 2)</td></tr>
</table>

Table 3 (continued)
Summary of UPC enforcement by traffic descriptor type

Traffic descriptor type	GCRA 1 (see Note 1)	GCRA 2	Tag or discard
7	1/PCR, CDVT	1/SCR0, BT + CDVT	discard at GCRA 2
8			tag at GCRA 2
9	1/PCR, CDVT	n/a	discard at GCRA 1
Note 1: GCRA1 always results in discard if policing action is required. ***Note 2:*** If PCR=0, UPC is disabled.			
(Sheet 2 of 2)			

The following sections describe UPC enforcement for each traffic descriptor type.

UPC enforcement for descriptor types 1 and 2

Nortel Networks Multiservice Switch nodes do not apply UPC enforcer functions for these descriptor types since traffic management is not available. Traffic descriptor type 1 is used to indicate that there is no traffic descriptor for the connection. Traffic descriptor type 2 is not used.

UPC enforcement for descriptor type 3

For connections defined with traffic descriptor type 3, a Nortel Networks Multiservice Switch node applies one enforcer, which discards CLP=0+1 cells based on the configured PCR in parameter 1. Non-violating cells are passed to the next FP for transmitting. Enforcement is illustrated in the figure “Traffic descriptor type 3 UPC enforcement” (page 34).

Figure 2
Traffic descriptor type 3 UPC enforcement

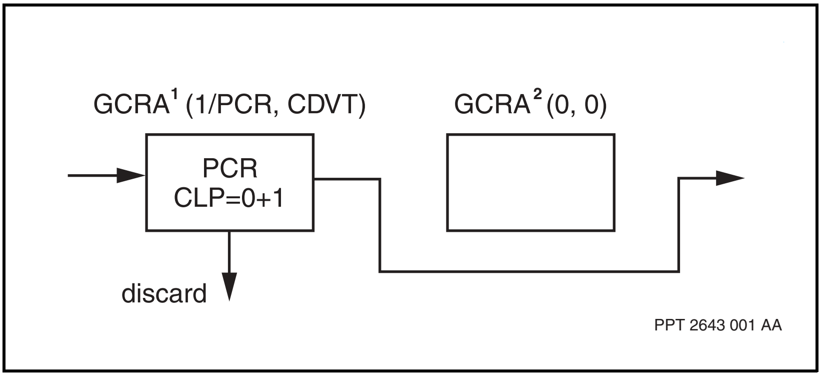

UPC enforcement for descriptor types 4 and 5

For connections defined with traffic descriptor types 4 or 5, a Nortel Networks Multiservice Switch nodes applies two enforcers. Enforcer 1 discards CLP=0+1 cells based on the configured PCR in parameter 1.

For traffic descriptor type 4, enforcer 2 discards CLP=0 cells that exceed the PCR0 rate configured in traffic descriptor parameter 2 for that connection. For traffic descriptor type 5, enforcer 2 tags CLP=0 cells (changes the CLP bit from 0 to 1) that exceed the configured PCR0 rate.

Traffic descriptor types 4 and 5, although supported, are not recommended since they are incompatible with frame relay service requirements. Further, both descriptors types specify the PCR for CLP0+1 in parameter 1 and the PCR0 in parameter 2.

Enforcement is illustrated in the figure “Traffic descriptor types 4 and 5 UPC enforcement” (page 35).

Figure 3
Traffic descriptor types 4 and 5 UPC enforcement

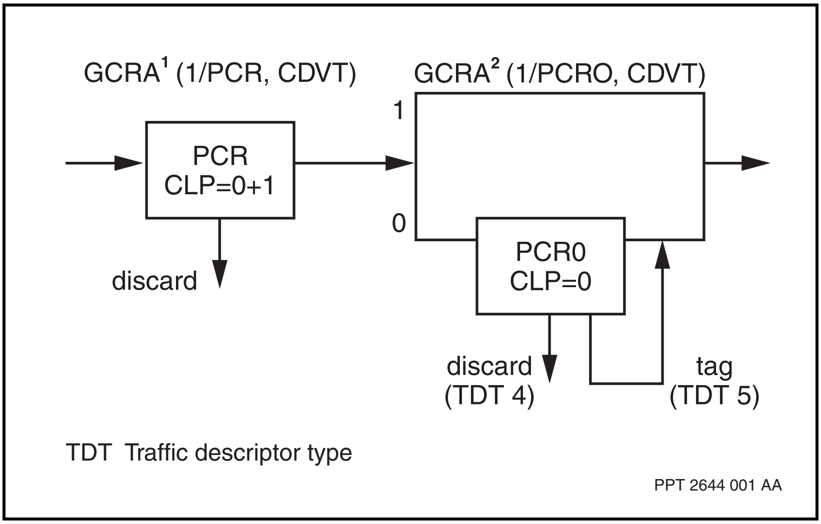

UPC enforcement for descriptor type 6

For connections defined with traffic descriptor type 6, a Nortel Networks Multiservice Switch nodes applies two enforcers. Enforcer 1 discards PCR CLP=0+1 cells based on the configured PCR in parameter 1. Enforcer 2 discards SCR CLP=0+1 cells based on the configured SCR in parameter 2. Enforcement is illustrated in the figure “Traffic descriptor type 6 UPC enforcement” (page 36).

Figure 4
Traffic descriptor type 6 UPC enforcement

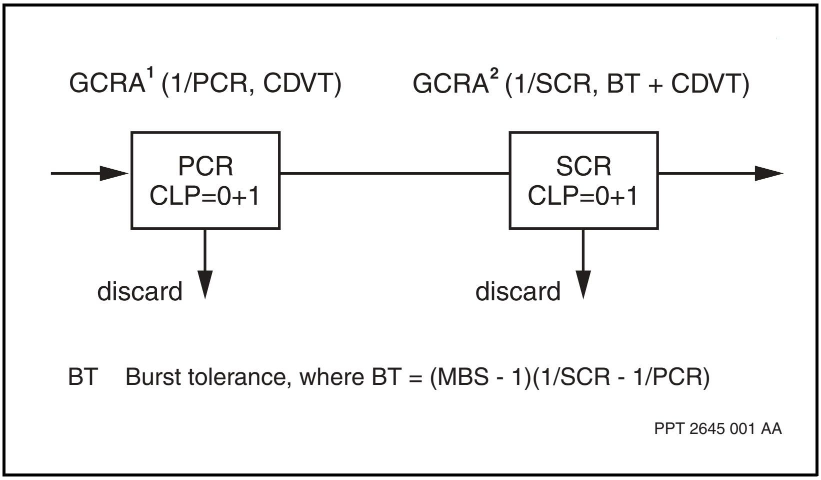

UPC enforcement for descriptor types 7 and 8

For connections defined with traffic descriptor types 7 or 8, a Nortel Networks Multiservice Switch node applies two enforcers. Enforcer 1 discards PCR CLP=0+1 cells based on the configured PCR in parameter 1. For traffic descriptor type 7, enforcer 2 discards excess cells. For traffic descriptor type 8, enforcer 2 tags CLP=0 cells (changes the CLP bit from 0 to 1).

Enforcement is illustrated in the figure “Traffic descriptor types 7 and 8 UPC enforcement” (page 37).

Figure 5
Traffic descriptor types 7 and 8 UPC enforcement

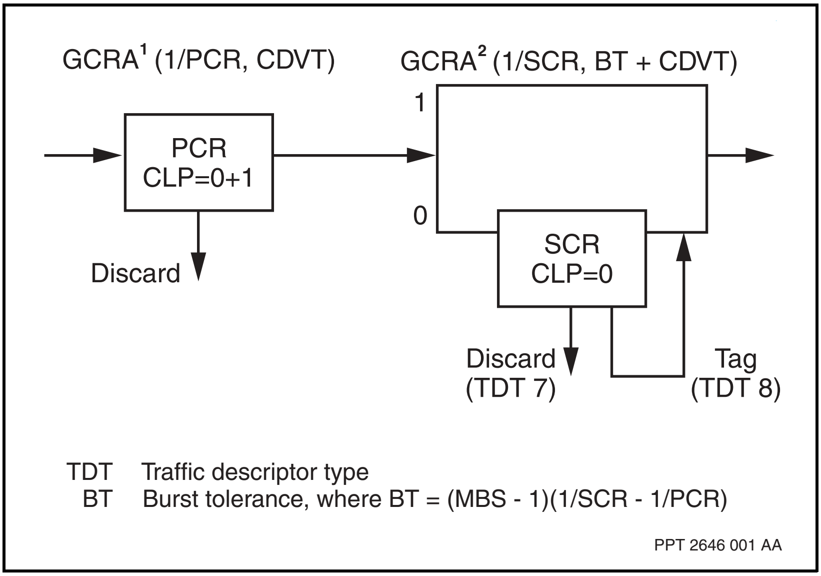

UPC enforcement for descriptor type 9

For UBR with MDCR connections defined with traffic descriptor type 9, a Nortel Networks Multiservice Switch node applies one enforcer, which discards CLP=0+1 cells based on the configured PCR in parameter 1. Non-violating cells are passed to the next FP for transmission. Enforcement is illustrated in figure “Traffic descriptor type 9 UPC enforcement” (page 38). Note that the MDCR value is not used by the enforcer.

Figure 6
Traffic descriptor type 9 UPC enforcement

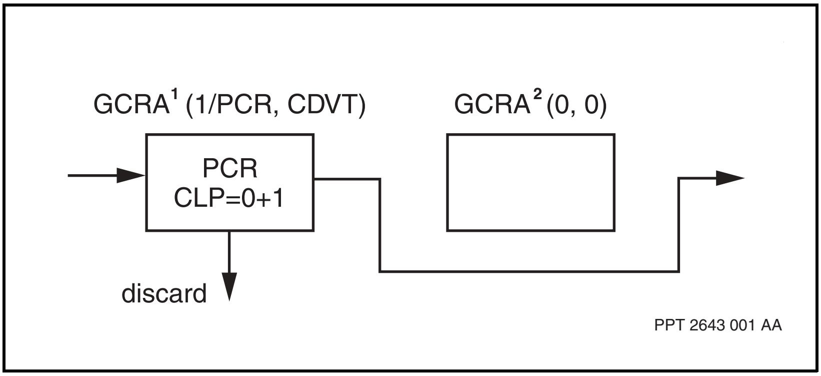

UPC operating characteristics

UPC can be enabled or disabled at the service category level for all switched and permanent connections. By default, UPC is disabled for the service category. The service category configuration can be overridden at the connection level.

Summary of limitations on UPC

UPC is not supported under the following configurations:

- The traffic descriptor type for the connection is 1 or 2.
- UPC is disabled for all service categories under the ATM interface (service category settings for UPC can be overridden at the connection level).
- UPC is disabled specifically for the connection.
- Connections are configured with traffic parameters that fall outside the hardware range. In this case, the connection is not accepted, either due to a failure through the configuration semantic check or failure on SVC call setup.

UPC may be specified for all ATM service categories, including UBR connections.

Additional considerations for UPC

This section presents additional information on UPC and its interaction with the following aspects of Nortel Networks Multiservice Switch node traffic management. Information appears in the following sections:

- “Relationship between UPC and CDVT” (page 39), which describes the requirements for balancing delay and discard
- “Relationship between UPC and linear traffic shaping” (page 40), which describes interactions between near- and far-end nodes, and how to ensure that the network does not unnecessarily tag or discard cells
- “Relationship between UPC and trunks” (page 41), which describes configuration requirements for connections over Multiservice Switch trunks over ATM
- “ATM over-policing by UPC” (page 41) which describes the behavior of the UPC policer when the configured PCR is less than 1/CDVT.

Relationship between UPC and CDVT

When UPC is enabled for a connection, care must be taken when defining the CDVT value for that connection. CDVT defines the level of tolerance for the cell delay variation on a connection. In terms of traffic management for the connection, this definition or tolerance translates into a limit at which cells are discarded as they are received from the connection. As cells are received, the inter-cell intervals are noted and Nortel Networks Multiservice Switch nodes increment an internal counter once for each interval. Multiservice Switch nodes decrease the counter as fewer cells are received over an internally defined time frame. When the limit defined by CDVT is reached, received cells are considered to be non-conforming and discarded.

For example, if traffic is received at rate that is slightly higher than the rates defined through the traffic descriptor type, cells are discarded as a result of the GCRA. However, if the CDVT is set higher, cell discards will stop.

> ***Note:*** Unexpected cell loss with a comparatively small PCR, which is in excess of the agreed upon PCR can occur when applying GCRA to CBR flow. This cell loss can be avoided by maintaining CDVT at a value greater than the emission interval (1/PCR).

UPC uses CDVT for traffic descriptor types 3 through 9.

Relationship between UPC and linear traffic shaping

Connections with the same traffic descriptor parameters on both ends (not necessarily same traffic descriptor types), and UPC enabled on the far end (the receive end), may experience cell discards or cell tagging even if traffic shaping is enabled on the near end (the transmit end). This scenario is possible is CQC-based FPs are configured on both ends. It can be avoided by using the VBR shaper that is configurable on ATM IP FPs.

Traffic shaping on Nortel Networks Multiservice Switch nodes is based on PCR CLP0+1 only, whereas UPC may be based on either PCR CLP0+1 (and possibly on PCR CLP0) or on SCR, depending on the traffic descriptor type. Traffic shaping helps only in conformance to PCR CLP0+1 at the far end.

For connections with any traffic descriptor type, the PCR/SCR shaping rate value used by the shaper cannot be higher than the PCR CLP0+1 value used by the policer to conform to UPC at the far end.

For connections with traffic descriptor type 3 or 9 at the far end, the condition outlined in the preceding paragraph is the only requirement since both traffic shaping and UPC are based only on PCR CLP0+1. For connections with traffic descriptor types 4 through 8, UPC enforces either PCR CLP0 or SCR in addition to PCR CLP0+1.

For connections with traffic descriptor types 4,6, and 7 at the far end, cell discards occur when the following conditions are true:

- the traffic offered to the shaper is in excess of its PCR CLP0 or SCR, and
- the PCR/SCR shaping rate that is used at the near end is higher than the PCR CLP0 or SCR at the far end

For connections with traffic descriptor types 5 and 8 at the far end, cell tagging occurs if

- the traffic offered to the shaper is in excess of its PCR CLP0 or SCR, and
- the PCR/SCR shaping rate that is used at the near end is higher than the PCR CLP0 or SCR at the far end

Tagging reduces the priority of the cells and makes them potential targets for discards when the network experiences congestion.

Non-conformance to UPC in each of the cases described in this section can be reduced by adjusting the RSR of the shaper at the near end to a value that is lower than the current value.

On CQC-based FPs, traffic shaping is not supported on CBR connections. As a result, if the traffic is not well-behaved and if UPC is enabled on the far end, cell discard or tagging may occur.

Relationship between UPC and trunks

For connections used as Nortel Networks Multiservice Switch trunks over ATM, if UPC is enabled on the far end (the receive end) of the connection, traffic shaping must be enabled on the near end (the transmit end) of the connection. If this condition is not satisfied, the trunk does not stage.

However, traffic shaping is not supported on connections with an *atmServiceCategory* of CBR. Therefore, when setting up logical trunks with an *atmServiceCategory* of CBR, ensure that UPC is disabled at the other end of the connection. If this condition is not satisfied, the trunk does not stage.

ATM over-policing by UPC

As documented in ITU I.371 recommendation Appendix III, when the UPC policer operates with a PCR value less than 1/CDVT, and when there is non-conforming user traffic, the policer can discard the non-conforming traffic, as well as conforming traffic. The amount of conforming traffic being discarded can be significant, depending on the PCR and CDVT values and depending on the amount of non-conforming traffic. This behavior is called over-policing because it discards not only all the non-conforming traffic, but a certain amount of conforming traffic as well.

Over-policing is intrinsic to the generic cell rate algorithm (GCRA) proposed for UPC by the ITU and by the ATM Forum. Therefore, UPC policers implemented in all Nortel Networks Multiservice Switch releases are expected to exhibit this behavior.

To avoid this type of over-policing, the signaled or provisioned CDVT (including the typical default value of CDVT, 250 microseconds) should be no less than 1/PCR, that is, the CDVT needs to be greater than or equal to 1/PCR, where the unit for PCR is in cells per second.

Note: Over-policing only happens when there is non-conforming (violating) traffic from the user, who attempts to send more traffic than is permitted by the traffic contract. If there is no violating traffic, there will be no over-policing.

Chapter 3
Shaping and policing on CQC-based FPs

This chapter described traffic shaping and policing for CQC-based function processors (FP).

Information is organized into the following sections:

Overview of traffic shaping on CQC

Nortel Networks Multiservice Switch nodes support the following traffic shaping features on CQC-based function processors:

- traffic shaping is configurable on two ports of two-port function processors, up to three ports of three-port function processors, and up to four ports of the eight-port function processors
- you enable and disable shaping for service categories
- you can disable shaping at the connection level
- a global scaling factor is configurable for each function processor on the node to provide more variety of shaping rates, which allows you to shift the available range of rates
- configurable per-VC queue congestion control thresholds

When configuring for service categories, choices are enabled, linear, and disabled. When configuring by VCCs or VPCs, choices are same as service category and disabled.

On CQC-based function processors, configuring traffic shaping to either enabled-based-on-traffic-descriptor or linear serves to enable linear traffic shaping. CQC-based function processors do not support VBR shaping.

Traffic shaping on CQC-based function processors makes no distinction between CLP=0 and CLP=1 cells once the cells have entered the shaped queues. On shaped connections, all cells are emitted at the configured rate. However, CLP=1 is discarded before CLP=0 in the event of congestion. To use traffic shaping, you must first enable per-VC queuing.

CQC shaping options

The table “Shaping options by traffic descriptor type” (page 45) summarizes shaping options by function processor type. The actual type of shaping which is being performed is indicated by an operational attribute under each connection.

Table 4
Shaping options by traffic descriptor type

Transmit traffic descriptor type	CQC-based function processors
1, 2	no shaping
3, 4, 5, 6, 7, 8, 9	linear shaping

Shaping rates for CQC-based function processors

There are 28 shaping rates at the medium scheduling priority and 28 shaping rates at the low priority (shaping is not allowed and not required at the high priority). Shaping rates are selected from a contiguous range of 33 rates, where each rate is associated with an average cell interval value. See “Defined traffic shaping rates for CQC-based function processors (available rates per shaped port)” (page 50).

The 28 shaping rates can be allocated across up to three ports simultaneously on most CQC-based function processors, or across up to four ports on the eight-port DS1/E1 function processors. Allocation determines the number of rates available for each port, and the range in which the rates fall, according to the following rules:

- Traffic shaping can be disabled for all ports.
- When allocated to port 0, all 28 contiguous shaping rates are available for that port.
- When allocated on ports 0 and 1, 14 contiguous shaping rates are available per port, where both ports draw from the same range of available rates.
- When allocated on ports 0, 1 and 2, seven non-contiguous shaping rates are available per port, where all ports draw from the same range of available rates.
- When allocated on ports 0, 1, 2, and 3, seven non-contiguous shaping rates are available per port, where all ports draw from the same range of available rates.
- Ports that have no shaping rates allocated to them can only support connections with common queuing.

Per-VC queuing and traffic shaping configurations

The table "Per-VC queuing and traffic shaping configurations supported on CQC-based function processors" (page 46) shows the possible configurations for enabling per-VC queuing and traffic shaping. The configurations apply whenever per-VC queuing is enabled (per-VC queuing is a requirement for traffic shaping.) For details on per-VC queuing, see NN10600-707 *Nortel Networks Multiservice Switch 7400/15000/20000 ATM Queuing and Scheduling Fundamentals*.

When per-VC queuing is enabled on the eight-port DS1/E1 function processor, only ports 0, 1, 2, and 3 can carry traffic directly. The other ports are available for use in IMA link groups. For details on using per-VC queuing and traffic shaping with the IMA feature, see NN10600-730 *Nortel Networks Multiservice Switch 7400/15000/20000 Inverse Multiplexing for ATM Operations*.

> ***Note:*** The per-VC queuing and traffic shaping configuration on ports 0, 1 and 2 is not available on the eight-port DS1/E1 function processor.

Table 5
Per-VC queuing and traffic shaping configurations supported on CQC-based function processors

Function processor types	**Per-VC queuing availability**			
	0	**0, 1**	**0, 1, 2**	**0, 1, 2, 3**
JT2 ATM	yes	yes		
3-port DS1/E1 ATM	yes	yes	yes	
8-port DS1/E1 ATM	yes	yes		yes
DS3/E3 ATM	yes	yes	yes	
OC-3 ATM	yes	yes	yes	

Note: The per-VC queuing and traffic shaping configuration on ports 0, 1 and 2 is not available on the eight-port DS1/E1 function processor.

Global scaling factors

Nortel Networks Multiservice Switch node traffic shaping uses six global scaling factors (1, 1.4, 2, 2.8, 4, and 5.6). You configure only one scaling factor for a function processor, such that all ATM interfaces on the function processor use the same scaling factor.

Application of a scaling factor does not increase the number of available rates on the interfaces, nor does it affect the granularity between rates. Instead, the scaling factor moves the available range toward the lower end of the scale. The net effect is a “sliding window” of available shaping rates, through which a fixed number of rates is always available. The scaling factor is applied to the average cell emission interval to derive the final traffic shaping characteristics.

The basic premise of this approach to allocating shaping rates is illustrated in the figure “Approach to allocating shaping rates based on the number of ports enabled” (page 49). Note the following characteristics in the allocation of shaping rates:

- When shaping is enabled on port 0 only, all 33 defined rates can be made available using the global scaling factors, with 28 rates accessible for any given scaling factor.
- When shaping is enabled on ports 0 and 1, 19 of 33 defined rates can be made available using global scaling factors, with 14 rates accessible for any given scaling factor.
- When shaping is enabled on ports 0, 1, and 2, 18 of 33 defined rates can be made available using global scaling factors, with 7 rates accessible for any given scaling factor. Observe the gap between each accessible rate within the range.
- When shaping is enabled on ports 0, 1, 2, and 3, 18 of 33 defined rates can be made available using global scaling factors, with 7 rates accessible for any given scaling factor. Observe the gap between each accessible rate within the range.
- Shaping rates 1, 2 and 3 cannot be used since the corresponding traffic characteristics are higher than those possible on the OC-3 port (149.76 Mbits/s).

The shaping rates are a function of the hardware shaper. The node cannot apply a rate that is higher than the link speed as the shaping rate. For lower- speed links, the scaling factor can be used to access different subsets of the available set of rates. In this way, shaping rates that fall within the range of the link speed can be used.

The configured scaling factor applies to all shaping stacks on the function processor: only one global scaling factor is configurable per logical processor (LP).

Figure 7
Approach to allocating shaping rates based on the number of ports enabled

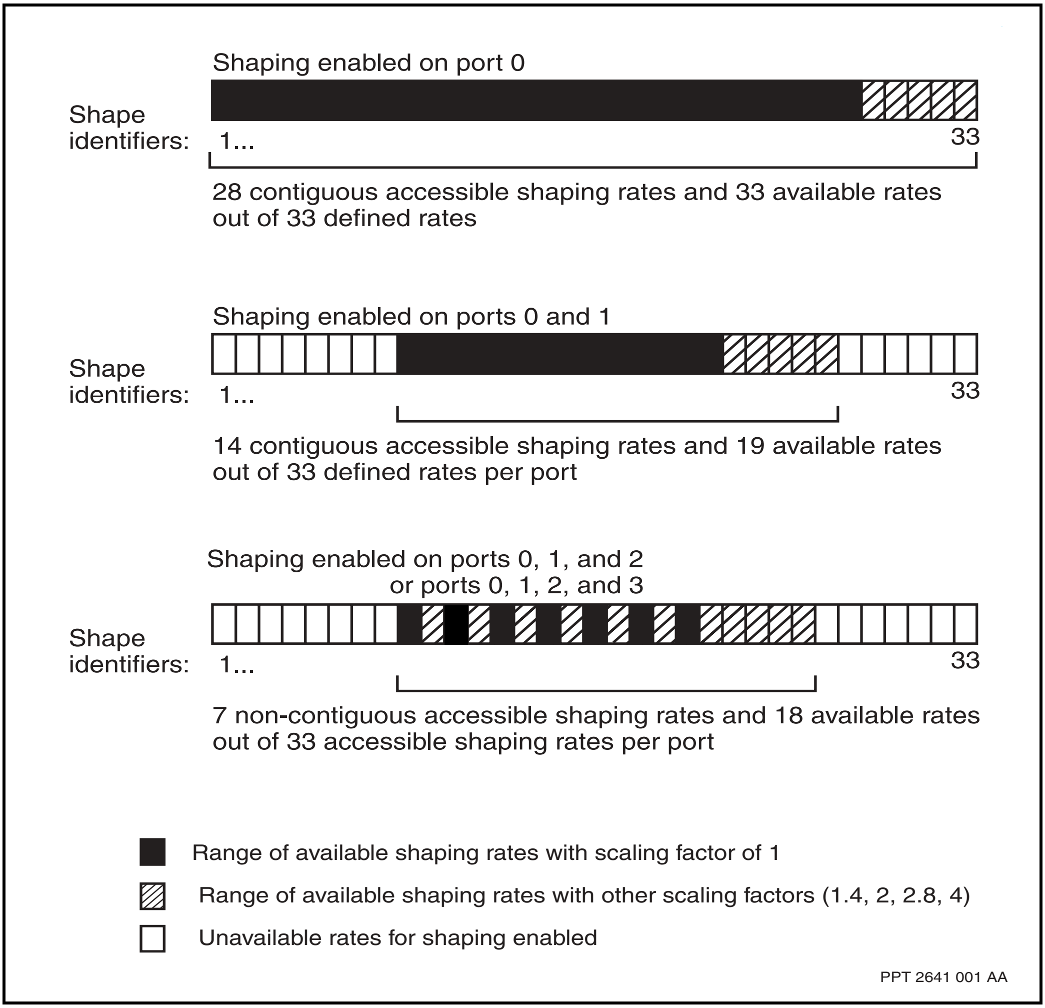

The global scaling factor applies to the average cell emission interval. How this sliding window approach applies to the actual range of defined and available shaping rates is shown in the table “Defined traffic shaping rates for CQC-based function processors (available rates per shaped port)” (page 50).

Table 6
Defined traffic shaping rates for CQC-based function processors (available rates per shaped port)

Shape identifiers				Resulting traffic characteristics			
Port 0	Ports 0, 1	Ports 0, 1, 2	Ports 0, 1, 2, 3	Avg. cell emission interval (microsec.)	Bandwidth demand	Equivalent bandwidth	Cell rate (cells/s)
1				1.06	400 Mbits/s	362 Mbits/s	943 396
2				1.50	282Mbits/s	256 Mbits/s	666 667
3				2.12	200 Mbits/s	181 Mbits/s	471 698
4				3.00	141 Mbits/s	128 Mbits/s	333 333
5				4.24	100 Mbits/s	90.6 Mbits/s	235 849
6				6.00	70.7 Mbits/s	64 Mbits/s	166 667
7				8.48	50 Mbits/s	45.3 Mbits/s	117 924
8				12.0	35.3 Mbits/s	32 Mbits/s	83 333
9	1	1	1	17.0	25 Mbits/s	22.6 Mbits/s	58 962
10	2	|	|	24.0	17.7 Mbits/s	16 Mbits/s	41 667
11	3	2	2	33.9	12.5 Mbits/s	11.3 Mbits/s	29 481
12	4	|	|	48.0	8.83 Mbits/s	8.0 Mbits/s	20 833
13	5	3	3	67.8	6.25 Mbits/s	5.66 Mbits/s	14 740
14	6	|	|	96.0	4.42 Mbits/s	4.0 Mbits/s	10 416
15	7	4	4	135	3.12 Mbits/s	2.83 Mbits/s	7370
16	8	|	|	192	2.21 Mbits/s	2.0 Mbits/s	5208
17	9	5	5	271	1.56 Mbits/s	1.42 Mbits/s	3685
18	10	|	|	384	1.10 Mbits/s	1.0 Mbits/s	2604
19	11	6	6	543	781 Kbits/s	708 Kbits/s	1842
20	12	|	|	768	552 Kbits/s	500 Kbits/s	1302
21	13	7	7	1085	391 Kbits/s	354 Kbits/s	921
(Sheet 1 of 2)							

Table 6 (continued)
Defined traffic shaping rates for CQC-based function processors (available rates per shaped port)

Shape identifiers				Resulting traffic characteristics			
Port 0	Ports 0, 1	Ports 0, 1, 2	Ports 0, 1, 2, 3	Avg. cell emission interval (microsec.)	Bandwidth demand	Equivalent bandwidth	Cell rate (cells/s)
22	14				1536	276 Kbits/s	250 Kbits/s	651
23							2171	195 Kbits/s	177 Kbits/s	460
24							3072	138 Kbits/s	125 Kbits/s	325
25							4341	97.6 Kbits/s	88.5 Kbits/s	230
26							6144	69.0 Kbits/s	62.5 Kbits/s	163
27					8683	48.8 Kbits/s	44.2 Kbits/s	115
28				12 288	34.5 Kbits/s	31.2 Kbits/s	82
29				17 366	24.4 Kbits/s	22.1 Kbits/s	58
30				24 576	17.3 Kbits/s	15.6 Kbits/s	41
31				34 732	12.2 Kbits/s	11.1 Kbits/s	29
32				49 152	8.6 Kbits/s	7.8 Kbits/s	20
33				69 464	6.1 Kbits/s	5.6 Kbits/s	14

Note 1: The shaping rates available per enabled port are identical on the three-port CQC-based function processors with all ports configured for traffic shaping, and the eight-port DS1/E1 function processors with the maximum of four ports configured for traffic shaping.

Note 2: Due to the granularity of the hardware counters, the cell rate measured on the link may be up to four percent lower than the values presented in this table and in the actual shaping rate for the VCC or VPC. See “Difference between actual and measured shaping rates” (page 57) for details on the difference between measured and stated PCR or SCR shaping rates.

(Sheet 2 of 2)

Implications of multi-port traffic shaping

Configuring more than one port to support traffic shaping has a cost: there are fewer available shaping rates per port. This cost has two implications:

- The range of shaping rates per port is reduced.

- If only port 0 is shaped, the range of available emission intervals is from 1.06 microseconds to 69 milliseconds, which translates into a range of cell rates from 943 396 cells/s down to 14 cells/s.
- If port 0 and 1 are shaped, the range of available emission intervals is reduced to values from 17 microseconds and 8.6 milliseconds, which translates into a range of cell rates from 58 962 cells/s down to 115 cells/s.
- If ports 0, 1, and 2, or ports 0, 1, 2, and 3 are shaped, the range of available emission intervals is further reduced to values from 17 microseconds and 6.1milliseconds, which translates into a range of cell rates from 58 962 cells/s down to 163 cells/s.

• The difference in shaping rates between any two shapes is increased when three or four ports have shaping ability as shown in column 3 of the table "Defined traffic shaping rates for CQC-based function processors (available rates per shaped port)" (page 50). Notice that the difference is constant for Port 0 shaped and Port 0 and 1 shaped.

When traffic shaping is enabled for a connection, all traffic received from the bus for that connection is queued on a per-VC queue and is emitted from this queue according to the shaping rate selected. Traffic shaping supports per-VC queue congestion control by providing an absolute queue length limit for each shaping stack. This length applies to each per-VC queue on that stack. For more information, see NN10600-707 *Nortel Networks Multiservice Switch 7400/15000/20000 ATM Queuing and Scheduling Fundamentals*.

Traffic shaping is supported at the low and medium link emission priority, but not at the high emission priority. The service category determines the emission priority. For more information, see the section on service category mapping priorities for CQC in NN10600-707 *Nortel Networks Multiservice Switch 7400/15000/20000 ATM Queuing and Scheduling Fundamentals*.

Shaping on CBR connections

On CQC-based function processors, a service provider can use QOS interworking over PVCs to configure shaping on CBR connections. For example, a service provider can map a CBR PVC onto a RT-VBR PVC to obtain the shaping rate that the CQC-based function processor provides. The service provider can use this configuration in situations where a carrier provides a fixed rate CBR VPC service.

Linear shaping using PCR shaping rates

All ATM function processors support linear shaping. Nortel Networks Multiservice Switch CAC nodes derive PCR shaping rates from the traffic descriptor parameters that are either configured or signaled for the VCCs and VPCs. Calculations include attributes that define the traffic descriptor type and the parameter values for that descriptor type.

The figure “Algorithm for linear shaping: all ATM function processors” (page 54) shows the algorithms that Multiservice Switch CAC nodes use to derive the PCR shaping rate for linear shaping. Essential to this algorithm is how the CQC-based function processor determines these rates. Normally, the PCR shaping rate is the next available shaping rate that is lower than or equal to the calculated target shaping rate (TSR). See the table “Defined traffic shaping rates for CQC-based function processors (available rates per shaped port)” (page 50).

The actual values used are fixed values that cannot be adjusted either up or down. The shaping rate used is applied relative to the PCR that is configured for the connection through the traffic descriptor parameters. A small increase in the configured PCR value does not necessarily result in a change of shaping rate.

Figure 8
Algorithm for linear shaping: all ATM function processors

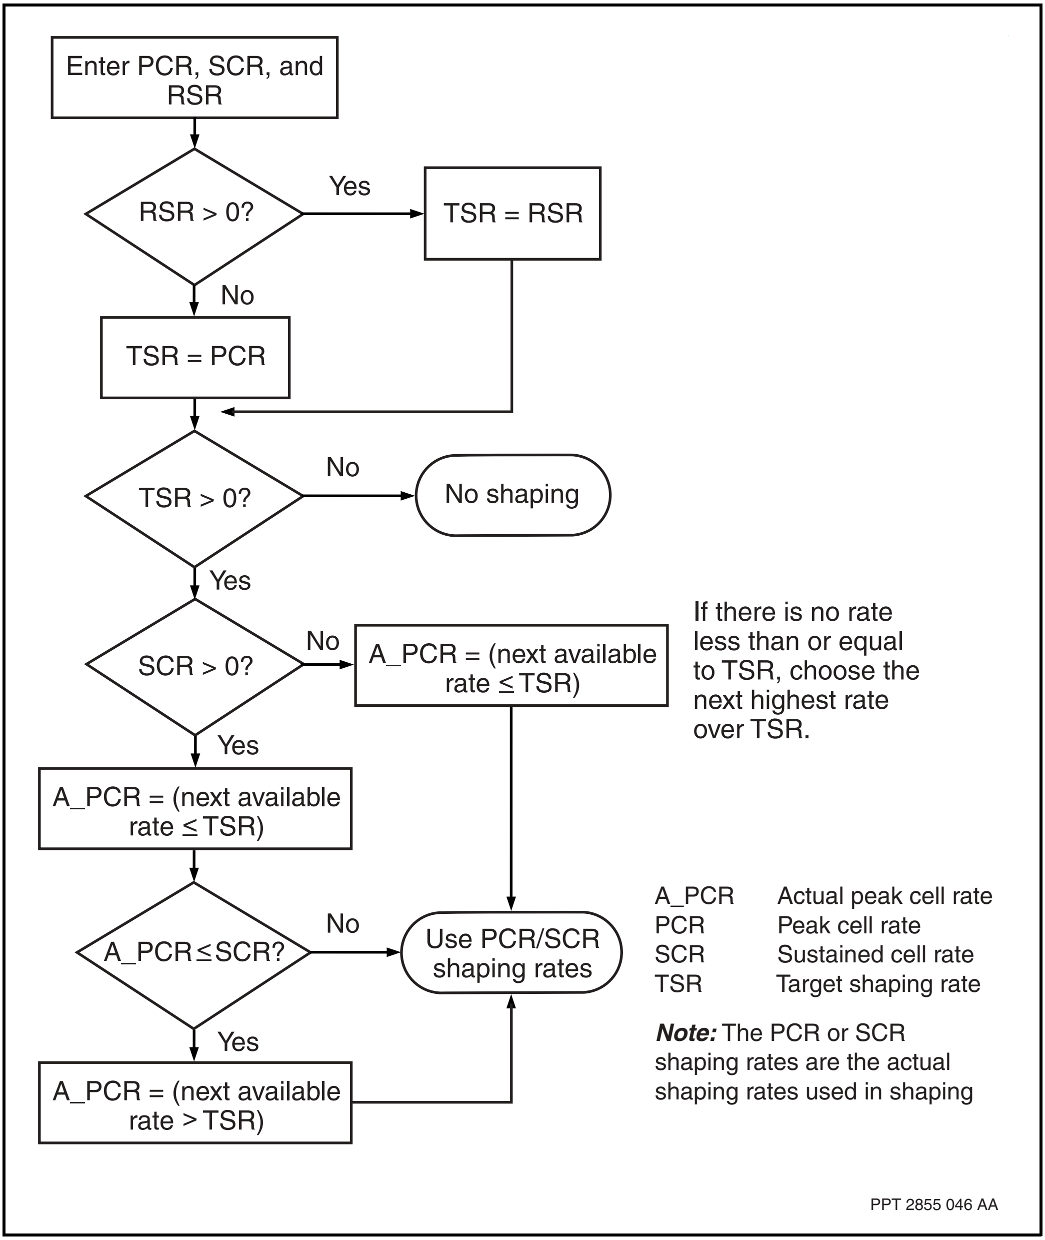

Configuration for traffic shaping

In addition to configuration at the service category and connection levels, CQC-based function processors also permit configuration at the ATM resource control level.

Examples of shaping on CQC-based function processors

Given the function processor type, the configured scaling factor for the LP, and the number of ports shaped on the LP, it is possible that the calculated TSR is not available. In this case, the lowest available shaping rate is used. For example, if

scaling factor = 1.0
traffic shaping is activated on port 0
calculated TSR = 50

then the next available rate of 41 cells/s is not available. The lowest available shaping rate is 82 cells/s, which is the rate used as the PCR or SCR shaping rates.

The table "Examples of shaping rate selection for scaling factors 1.4 and 2" (page 55) summarizes the results of the shaping rate selection process. Compare the results with the defined traffic shaping rates provided in the table "Defined traffic shaping rates for CQC-based function processors (available rates per shaped port)" (page 50).

Table 7
Examples of shaping rate selection for scaling factors 1.4 and 2

Target Shaping Rate	PCR or SCR shaping rates scaling factor 1.4			PCR or SCR shaping rates scaling factor 2		
	28 shapes	14 shapes	7 shapes	28 shapes	14 shapes	7 shapes
300 cell/s	230 cells/s	460 cells/s	651 cells/s	230 cells/s	325 cells/s	460 cells/s
10 000 cell/s	7370 cells/s	7370 cells/s	5208 cells/s	7370 cells/s	7370 cells/s	7370 cells/s
(Sheet 1 of 2)						

Table 7 (continued)
Examples of shaping rate selection for scaling factors 1.4 and 2

Target Shaping Rate	PCR or SCR shaping rates scaling factor 1.4			PCR or SCR shaping rates scaling factor 2		
	28 shapes	14 shapes	7 shapes	28 shapes	14 shapes	7 shapes
14 000 cell/s	10 416 cells/s	10 416 cells/s	10 416 cells/s	10 416 cells/s	10 416 cells/s	7370 cells/s
21 000 cell/s	20 833 cells/s	20 833 cells/s	20 833 cells/s	20 833 cells/s	20 833 cells/s	14 740 cells/s
(Sheet 2 of 2)						

The table “Examples of shaping rate algorithms and actual shaping rates” (page 56) summarizes the shaping rate algorithm where PCR, SCR, and RSR are available as traffic descriptor parameters. Again, compare the results with the defined traffic shaping rates provided in the table “Defined traffic shaping rates for CQC-based function processors (available rates per shaped port)” (page 50).

Table 8
Examples of shaping rate algorithms and actual shaping rates

Peak cell rate (PCR)	Sustained cell rate (SCR)	Requested shaping rate (RSR)	PCR or SCR shaping rates are based on	Resulting PCR or SCR shaping rates
10 000 cells/s	7000 cells/s	10 000 cells/s	RSR	7370 cells/s
10 000 cells/s	7000 cells/s	not specified	SCR	7370 cells/s
10 000 cells/s	7500 cells/s	not specified	PCR	10 416 cells/s
10 000 cells/s	not specified	not specified	PCR	7370 cells/s

Note 1: The PCR, SCR, and RSR parameters and the PCR or SCR shaping rates, given a global scaling factor of 2.

Note 2: RSR applies to PVCs only.

When both PCR and SCR are greater than zero, traffic is shaped at the highest available rate that lies between these two values. If there is no defined shaping rate between these values, the next higher defined shaping rate above PCR is used.

Difference between actual and measured shaping rates

The table "Difference between actual and measured shaping rates" (page 57) provides details on the difference between the actual shaping rate and the measured shaping rate. This difference decreases as the scaling factor increases.

Table 9
Difference between actual and measured shaping rates

Shaping scaling factor	Difference between measured and stated actual shaping rate
1.0	measured is 2 percent to 4 percent lower than the PCR or SCR shaping rates
1.4	measured is 2 percent to 4 percent lower than the PCR or SCR shaping rates
2.0	measured is 1 percent to 2 percent lower than the PCR or SCR shaping rates
2.8	measured is 1 percent to 2 percent lower than the PCR or SCR shaping rates
4.0	measured is the same as the PCR or SCR shaping rates
5.6	measured is the same as the PCR or SCR shaping rates

As an example of the difference between actual and measured shaping rates, the table "Measured shaping rates for Multiservice Switch OC-3 function processors" (page 58) summarizes the shaping rates for the OC-3 function processors for scaling factors 1 to 5.6.

Table 10
Measured shaping rates for Multiservice Switch OC-3 function processors

Expected shaping rate	Scaling factor = 1	Scaling factor = 1.4	Scaling factor = 2.8	Scaling factor = 4	Scaling factor = 5.6
15					-14
20				20	20
29			-28	-28	-28
41			-40	-40	-40
58		-57	-57	-57	-57
82	-78	-81	-81	-81	-81
115	-113	115	115	115	115
163	-156	-162	-162	-162	-162
230	-226	-229	-229	-229	+231
325	-312	325	325	325	325
460	-452	+462	+462	+462	+462
651	-626	651	651	651	651
921	-904	921	921	921	+925
1302	-1252	-1252	1302	1302	1302
1842	-1808	1842	1842	1842	+1851
2604	-2504	-2504	2604	2604	2604
3658	-3616	+3685	+3684	+3684	+3702
5208	-5008	-5008	5208	5208	5208
7370	-7233	7370	7370	7370	+7405
10 416	-10 016	-10 016	10 416	10 416	10 416
14 740	-14 467	14 740	14 740	14 740	+14 810
20 833	-20 032	-20 032	20 833	20 833	20 833
29 481	-28 935	29 481	29 481	29 481	+29 621
(Sheet 1 of 2)					

Table 10 (continued)
Measured shaping rates for Multiservice Switch OC-3 function processors

Expected shaping rate	Scaling factor = 1	Scaling factor = 1.4	Scaling factor = 2.8	Scaling factor = 4	Scaling factor = 5.6
41 667	-40 064	-40 064	41 667	41 667	41 667
58 962	-57 870	+58 963	+58 963	+58 963	+59 242
83 333	-80 129	-80 129	+83 334	+83 334	+83 334
117 924	-115 742	+117 926	+117 926	+117 926	+118 484
166 667	-160 257	-160 258	+166 668	+166 668	+166 668
235 849	-231 484	+235 852	+235 852	+235 852	
333 333	-320 517	-320 517	+333 337		
Note: A "-" sign indicates shaping rate lower than expected and a + sign indicates a shaping rate higher than expected. No marking indicates that the shaping rate is as expected.					
(Sheet 2 of 2)					

UPC characteristics for CQC-based function processors

For traffic descriptor types 6, 7, or 8, the following restrictions apply

- For SCR $\leq$380 cells/s
  (X = 1/SCR + CDVT + BT) < 1.342 s
  where X is the leaky bucket in the GCRA.
  If the check is not met, the connection request is rejected.
- For SCR > 380 cells/s
  (L = CDVT + BT) < 1.339 s
  If the check is not met, the connection request is rejected.

When you enable UPC on a connection, the node can use one or two enforcers. I1 and L1 relate to the first enforcer, and I2 and L2 relate to the second enforcer.

The following statements are true for I1, I2, L1, and L2

- For any traffic descriptor type (except 1 and 2), I1 = 1/PCR, L1 = CDVT.

- For traffic descriptor type 3, the Nortel Networks Multiservice Switch system does not use I2 and L2.
- For traffic descriptor types 4 and 5, $I2 = 1/PCR_0$ and $L2 = CDVT$.
- For traffic descriptor types 6, 7, and 8, I2 = 1/SCR and L2 = CDVT + BT, where BT = (MBS - 1) * (1/SCR - 1/PCR).

If either PCR or SCR is greater than 48 828 cells/s, then the following conditions must be true:

1 $I + L \leq 167\ 769.60$ microsecond

2 $L \leq 167\ 444.48$ microsecond

3 PCR and SCR $\geq$ 24 cells/s

Otherwise, the following conditions must be true:

1 $I + L \leq 1\ 342\ 156.80$ microsecond

2 $L \leq 1\ 339\ 555.8$ microsecond

3 PCR and SCR $\geq$ 3 cells/s

You need two sets of conditions because the UPC mechanism can use four different units of time. You select the time units through UPC parameter configuration. For the largest time unit (20.48 microsecond), PCR or SCR must be greater than 48 828 cells/sec.

Cell counting

CQC-based function processors do not collect counts of UPC cell discards or tagging. CQC-based function processors can count receive cell discards, but not as recognized UPC violations.

CDVT setting for UPC at the far end (shaping on CQC)

On CQC-based function processors, when traffic shaping is used over a connection and UPC is turned on for that connection at the far end, it is necessary to incorporate into the CDVT the credit of the two user cells. This credit is issued by the Nortel Networks Multiservice Switch shaper. The following equation determines the CDVT for accommodating clumped cells and for restricting the CDVT to a reasonable value for setting the UPC.

CDVT ≥ 2 * (inverse of PCR/SCR shaping rate)

If you prefer a lower CDVT value, set the CDV attenuation under the ATM resource control to a value of on. This setting causes the shaper to issue one user cell. However, this setting also reduces the transmit bandwidth utilization of the three-port DS3 ATM and three-port OC-3 ATM function processors. By default, the CDV attenuation is off. Setting it to on changes these relationships to the following equation and provides much lower CDVT values:

CDVT ≥ 1* (inverse of PCR/SCR shaping rate)

These equations specify CDVT due to traffic shaping. The equations do not include the cell rate de-coupling effect. The maximum value of CDVT factor is one cell time. One cell time is equal to the time required to transmit one cell at the link rate.

See the *Nortel Networks Multiservice Switch Release Notes* for more information.

Relationship between UPC and linear traffic shaping

On CQC-based function processors, traffic shaping is not supported on CBR connections. As a result, if the traffic is not well-behaved and if UPC is enabled on the far end, cell discard or tagging may occur.

Chapter 4
Shaping and policing on ATM IP FPs

This chapter describes traffic shaping characteristics that are specific to ATM IP function processors (FPs). Information is organized into the following sections:

Traffic shaping on ATM IP function processors

The traffic shaper on an ATM IP function processor (FP) uses a shaped fair queuing (SFQ) mechanism. This control has the following general characteristics:

- permits shaping of queues in up to two emission priorities
- can apply to any service category
- supports linear shaping
- supports VBR shaping (also called inverse UPC shaping)
- supports increased range and granularity (within 1.0% of requirements for the traffic contract) for both linear and VBR shaping and does not require global scaling factors
- provides shaping on all ports with full shaping rate range and granularity (no restrictions related to per-VC queuing and available ports)
- provides cell delay variation (CDV) that conforms to ITU I.371

If shaping is required on connections for three of more service categories, you must map multiple service categories to a single emission priority. For example, if CBR, RT-VBR, and NRT-VBR connections require shaping, map RT-VBR and NRT-VBR to the same emission priority.

ATM IP FPs support two types of shaping:

- linear shaping, which shapes connections using a shaper bound on PCR
- VBR shaping, which shapes connections using inverse UPC shaping based on PCR and SCR

For PCR, SFQ shapes at rates from 10 cell/s up to the link rate with accuracy of 1.0% or better of requested cell rates. For SCR, SFQ shapes at the following rates (dependent on FP port type), with accuracy within 0.6%:

- 44 cells/s for OC3 FPs (actual shaping rate is ±264 cells/s)
- 13 cells/s for DS3 FPs (actual shaping rate is ±078 cells/s)
- 10 cells/s for E3 FPs (actual shaping rate is ±06 cells/s)
- 176 cells/s for OC12 FPs (actual shaping rate is ±1056 cells/s)

Shaping is configured for each service category as:

- disabled
- enabled (based on the traffic descriptor, linear shaping, or VBR shaping)
- enabled (linear shaping)

Shaping at each configured connection can be disabled or use the same configuration as the service category. If a service category is shaped but a particular connection has shaping disabled, the connection is actually shaped at the maximum shaping rate available on the interface.

ATM IP shaping process

The class scheduler is based on a cell departure table, in which each connection reserves a scheduling opportunity. The reservation is based on the configured PCR and SCR values for the connection. The algorithm also allows the node to retrieve bandwidth for any connection that experiences delay. Connections with shaping disabled at the connection level receive that same fairness as shaped connections.

Under congestion, when there is temporary over booking of the bandwidth allocated to all shaped connections, the scheduler undertakes a cell-by-cell polling of all shaped connections for the emission priority. This approach ensures fairness between all shaped connections under this emission priority. For connections that have not achieved the configured emission timings, this approach also recoups lost emission opportunities up to the CDVT value for the connection. For more information on recoup of cell service opportunities, see “Recovery of shaped cell opportunities” (page 71).

When the rate of cells arriving on a shaped connection exceeds the shaping rate, the connection queue builds up. Discards occur as the queue threshold is reached.

If the volume of CBR traffic for a link is high relative to traffic under other service categories, shaped connections may not be able to transmit at their configured shaping rate. Even if a shaped connection is well-behaved, the high volume of CBR traffic starves the emission priority and cells collect in the connection queue.

If there are several shaped connections that are congested by traffic of higher priority, the remaining bandwidth for the shaped connections is allocated uniformly in a round-robin manner to each connection. If the remaining bandwidth is less than the sum of the shaping rates for all connections, then some connections cannot achieve their shaping rate and some discards occur. Higher speed connections are affected most be this behavior.

Consider the following example on an OC3 link with a cell rate 353 207 cell/s):

- connection C is CBR at 253 207 cell/s
- connection N1 is nrt-VBR shaped at 10 000 cell/s
- connection N2 is nrt-VBR shaped at 20 000 cell/s
- connection N3 is nrt-VBR shaped at 50 000 cell/s
- connection N4 is nrt-VBR shaped at 50 000 cell/s

There is 100 000 cell/s left over for the four shaped connections but together they require 130 000 cell/s (the sum of their shaping rates). After the queues build up to the threshold, the node must discard 30 000 cell/s. In this congestion situation, the shaped connections are served in round-robin manner up to their requested shaping rate. As a result each connection achieves the following throughput:

- connection C achieves 253 207 cell/s (no discards)
- connection N1 achieves 10 000 cell/s (no discards)
- connection N2 achieves 20 000 cell/s (no discards)
- connection N3 achieves 35 000 cell/s (15 000 cell/s discards)
- connection N4 achieves 35 000 cell/s (15 000 cell/s discards)

In this example, the corrective action is to reduce the amount of traffic flowing through this interface so that the sum of bandwidth for all connections is less than the link rate.

ATM IP shaping options

You can configure the following shaping options:

- a single-rate linear shaper for CBR, VBR, and UBR, which conforms to the PCR enforcer (plus inclusion of the CDVT value). For VBR, this mode is useful in the following cases:
  - if the carrier tags but does not discard excess traffic
  - if the connection carries non-critical traffic (for example, the higher level protocol can recover from occasional traffic loss) and operators wish to take advantage of excessive bandwidth capability
- a VBR shaper for VBR, also called inverse UPC shaper, which conforms to the PCR and SCR dual-rate enforcer (plus inclusion of MBS and CDVT) value and applies automatically according to traffic descriptor type as follows:
  - VBR.1, which conforms to PCR0+1 and the SCR0+1 dual-rate generic cell rate algorithm (GCRA) enforcer
  - VBR.2 and VBR.3, which conforms to the PCR0+1 and SCR0 dual-rate GCRA enforcer, especially if the carrier discards non-conforming traffic
- flexibility between a strict CDV and extra efficiency of the shaper

The figure “General principles of VBR shaping” (page 69) shows how VBR shaping works in a simple example. This example illustrates a connection with a transmit rate that varies over time. After the initial burst that reaches a maximum at the PCR, the connection transmit rate fall below the contracted SCR. This low rate is recorded as a credit; that is, the node recognizes that the connection is not using is allocated bandwidth to the fullest. The connection can apply this credit to future spikes in its transmission rate, up to the PCR and subject to burst tolerance.

For more information on strict CDV (also called minimum CDV) and extra efficiency (also called maximum efficiency) shaping options, see “Recovery of shaped cell opportunities” (page 71). Depending on the user traffic descriptor type and user preference, individual connections can be shaped using either linear shaping or VBR shaping (also called inverse UPC shaping). Strict CDV and extra efficiency options can be selected on a per service category basis, independent of linear shaping or VBR shaping.

Strict CDV shaping schedules the next cell departure time exactly one shaped cell interval behind the current cell departure time. The length of a shaped cell interval is equal to the reciprocal of the desired shaping rate. If the departure schedule of one connection coincides with the time slot already scheduled for another connection, the scheduled departure for one of the connections must be delayed to the next slot. Delay of departure schedules, due to increased number of connections and heavy traffic, may impact shaper efficiency.

VBR shaping will record delayed departures as credits counted towards future departures. These departures can be scheduled at PCR to burst out at a length up to the number of cells indicated by the MBS. The effect of one delayed departure can be absorbed by subsequent departures, see "General principles of VBR shaping" (page 69). VBR shaped connections will maintain better traffic throughput than linear shaped connections.

Figure 9
General principles of VBR shaping

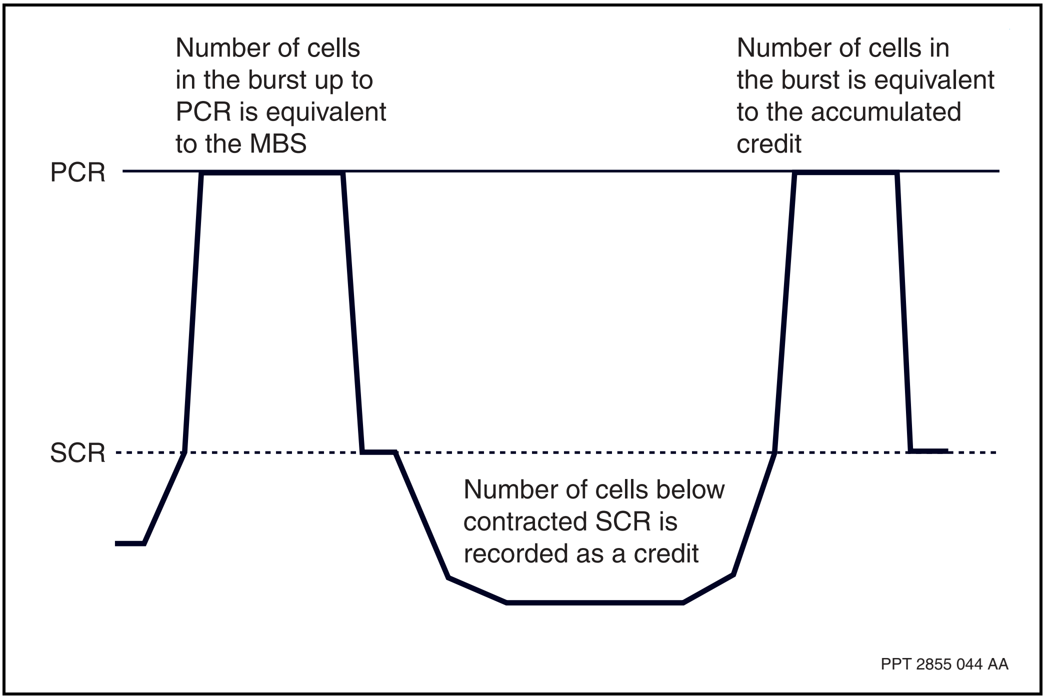

Shaping and traffic descriptor types

Shaping also depends on the traffic descriptor type for the VCC or VPC. Traffic descriptor types 1 and 2 cannot be shaped, regardless of the shaping configuration for the ATM interface. Traffic descriptor types 3 and 4 can only be shaped using a linear shaper. Traffic descriptor types 5, 6, 7, and 8 can be shaped either using linear shaping or VBR shaping.

If you configure traffic shaping as enabled based on traffic descriptor, traffic descriptor types 3, 4, and 5 provide linear shaping, types 6, 7 and 8 provide VBR shaping, and types 1 and 2 disable shaping. For traffic descriptor type 6, the VBR shaper shapes for $PCR_{0+1}$and SCR_{0+1}. For traffic descriptors 7 and 8, the VBR shaper shapes for PCR_{0+1} and SCR_0. The traffic descriptor type configured for a shaper at one node must be the same as the descriptor type being policed at the next node.

Different from UPC policing, where a policer can be configured to tag non-conforming cells (such as policing for VBR Traffic Descriptor Type 8), VBR shaping does not tag cells. Excess user cells on a connection shaped by a VBR shaper will cause the per-connection buffer to overflow and the excess cells will be discarded.

The table “Shaping options by traffic descriptor type and FP type” (page 70) summarizes shaping options by FP type.

Table 11
Shaping options by traffic descriptor type and FP type

Transmit traffic descriptor type	ATM IP with shaping enabled	ATM IP with inverse UPC shaping
1, 2	no shaping	no shaping
3, 4, 5	linear shaping	linear shaping
6, 7, 8,	VBR-shaping or linear shaping (default)	VBR shaping
9	linear shaping	linear shaping

The actual type of shaping which is being performed is indicated by an operational attribute under each connection.

To compare options for shaping and policing by traffic descriptor type, see the summary of UPC enforcement in “Multiservice SwitchTraffic policing” (page 27).

Shaping on CBR connections

On ATM IP FPs, CBR traffic may be shaped as with any other service category. While shaping on “constant” connections may seem unnecessary, this capability has the following applications:

- to return CBR traffic to an orderly flow after some CDV has been introduced
- to adjust the flow of bursty CBR traffic (for example, if voice is served on CBR, compression and speech activity detection algorithms may introduce burstiness)

- to ensure that CBR traffic can pass through a stringent policer (for example, in a carrier situation)
- to enhance conformance to the limitations of an end-system playback device which may have a restricted buffer

Recovery of shaped cell opportunities

Nortel Networks Multiservice Switch node software includes an option for each shaped service category that provides a choice between a strict CDV or extra shaping efficiency. For CBR and RT-VBR, the default is strict CDV, and for NRT-VBR and UBR the default is extra shaping efficiency.

The figure “Recovering shaped cell opportunities” (page 72) summarizes how these choices affect actual cell rates.

If the shaper tries to send a cell, c1, at a certain time, t, but that time is already used by a cell from another connection, the shaper must schedule cell c1 later at a time $t+i$ where i is the imposed delay. Normally, the shaper must schedule the next cell, c2, at a time $t+d+i$, where d is the inter-cell interval for the connection.

The shaper on the ATM IP FPs has the option to recover the lost time, i, so that cell c2 is sent at time $t+d$. This method to recover the shaped cell opportunities is called recouping. This option improves the efficiency of the shaper where there is congestion for cell opportunities. Improved efficiency comes at the cost of increased CDV. Use this option where a more lenient policer is in use which allows more cells to pass. You can use the strict CDV option where traffic policing discards cells, or for connections that require strict conformance to a specific CDV value.

Back-to-back cells may occur only in the following situations:

- when the shaping rate is greater than one half the link rate
- when using VBR shaping, and there is a burst of cells to be sent
- when enabling the extra shaping efficiency option

Improvement of traffic throughput using the extra efficiency option is more evident to connections using linear shaping than those using VBR shaping. See “ATM IP shaping options” (page 67) for information on delayed cell departure.

Figure 10
Recovering shaped cell opportunities

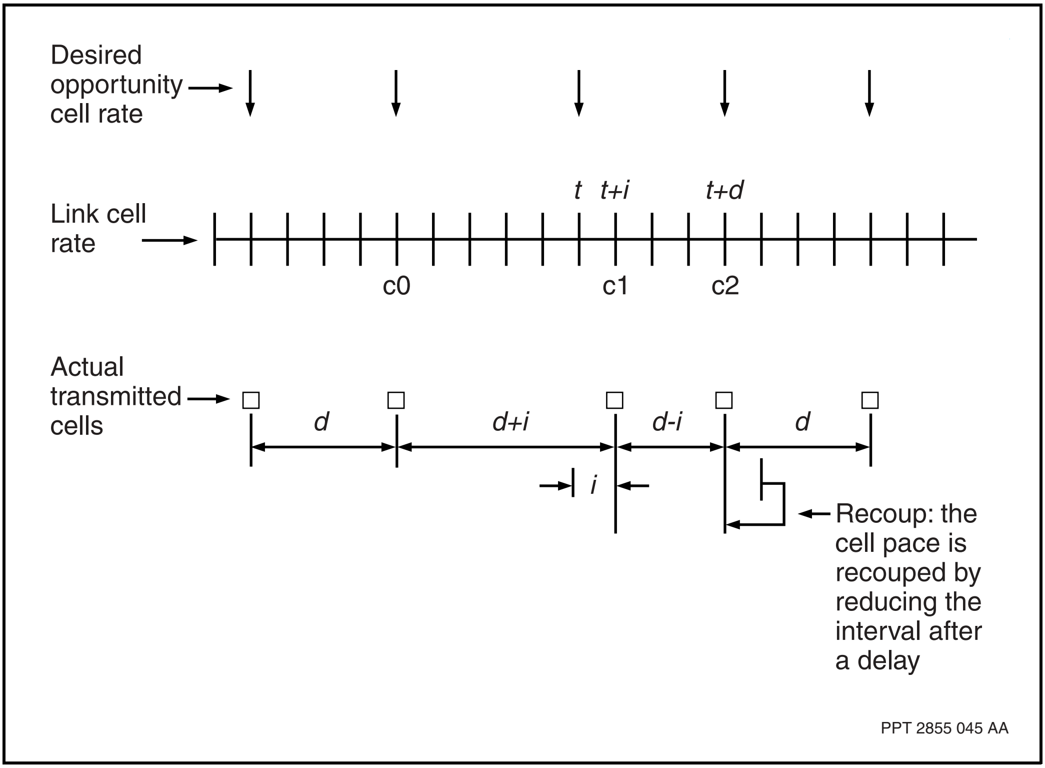

Linear shaping using PCR as the shaping rate

Nortel Networks Multiservice Switch CAC derives PCR shaping rates from the traffic descriptor parameters that are either configured or signaled for the VCCs and VPCs. Calculations include attributes that define the traffic descriptor type and the parameter values for that descriptor type.

The actual values used are fixed and cannot be adjusted either up or down. The shaping rate used is applied relative to the PCR that is configured for the connection through the traffic descriptor parameters. A small increase in the configured PCR value does not necessarily result in a change of shaping rate. See the figure “Algorithm for linear shaping: all ATM FPs” (page 74) for the algorithm used to derive the PCR or SCR shaping rates.

Figure 11
Algorithm for linear shaping: all ATM FPs

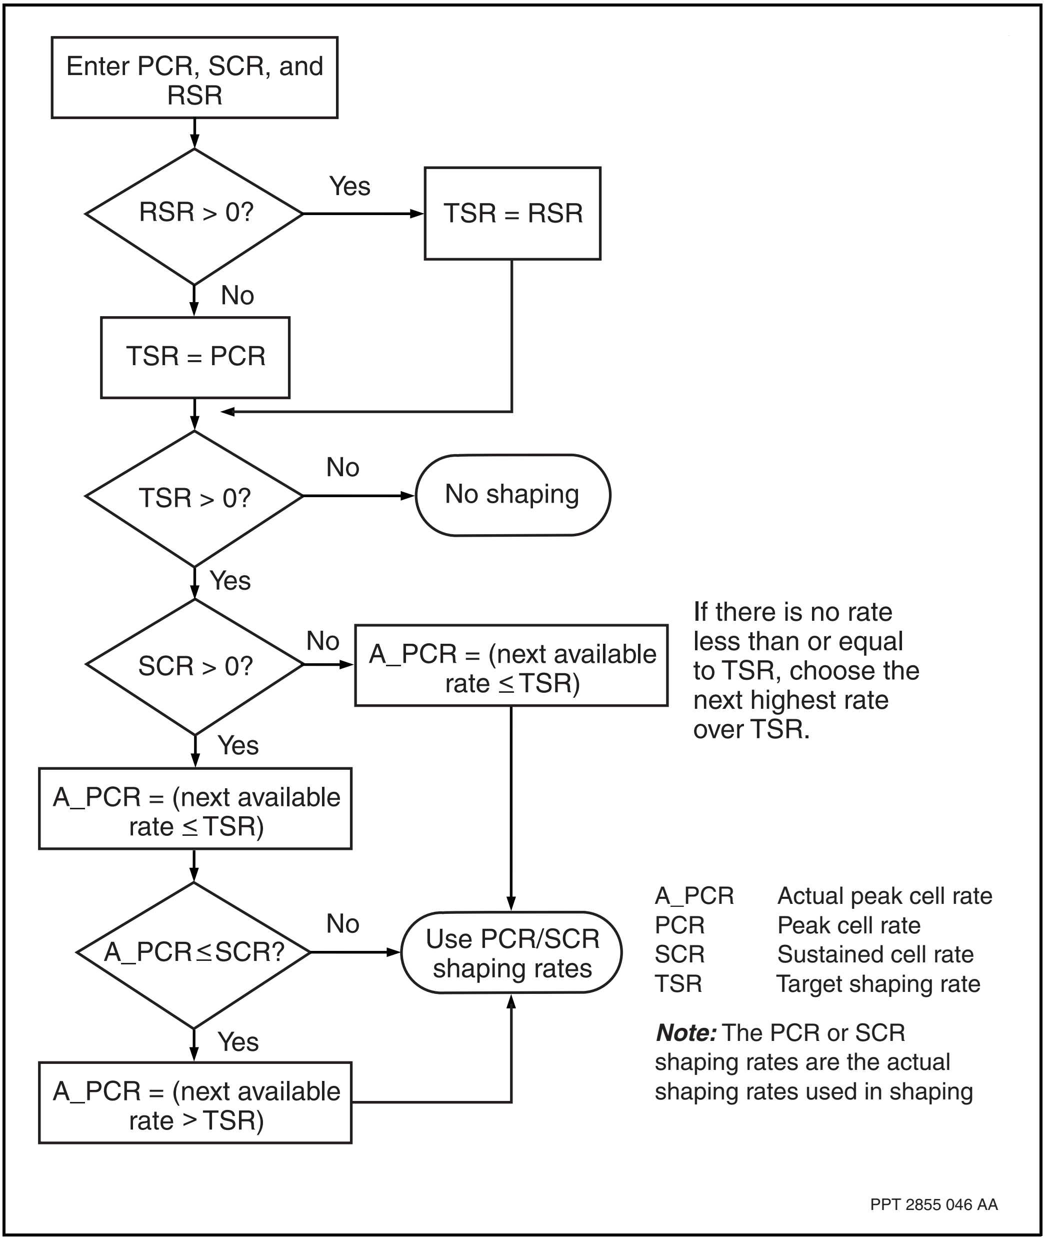

VBR shaping using PCR or SCR as the shaping rate

For ATM IP FPs, the PCR or SCR shaping rates are values that are:

$$P\_PCR * 0.99 \leq A\_PCR \leq P\_PCR$$
$$P\_SCR * 0.99 \leq A\_SCR \leq P\_SCR$$

The PCR or SCR shaping rates for VBR shaper can be expressed as follows:

1 If requested shaping rate (RSR) $\neq 0$, then PCR = RSR.

2 Set the VBR shaper values for PCR, SCR, and MBS to PCR minus up to 1%, SCR minus up to 1%, and the existing value for MBS. These values may differ from the requested PCR and SCR but never by more than 1%.

3 In the case of granularity where the exact rate is not available, use the next lower rate (that is, truncate the fraction).

This process ensures that the VBR shaper does not shape at a rate that is higher than the policer expects at the other end of the link.

The figure “Algorithm for VBR shaping: ATM IP FPs” (page 76) shows the algorithms that Nortel Networks Multiservice Switch CAC nodes use to derive the PCR and SCR shaping rates for VBR shaping.

For linear shaping, the rule is as follows:

$$(PCR \text{ or } RSR) * 0.99 \leq R\_PCR \text{ shaping rate} \leq (PCR \text{ or } RSR) * 1$$

For VBR shaping, the rule is as follows:

$$(C\_PCR \text{ or } RSR) * 0.99 \leq A\_PCR \leq (C\_PCR \text{ or } RSR) * 1$$
$$C\_SCR * 0.99 \leq A\_SCR \leq C\_SCR * 1$$

PCR shaping rates display as A_PCR
SCR shaping rates display as A_SCR

$$A\_MBS \leq C\_MBS$$

where:
A_MBS is the actual MBS
C_MBS is the configured MBS

C_PCR is the configured PCR
C_SCR is the configured SCR
R_PCR is the requested PCR shaping rate

Figure 12
Algorithm for VBR shaping: ATM IP FPs

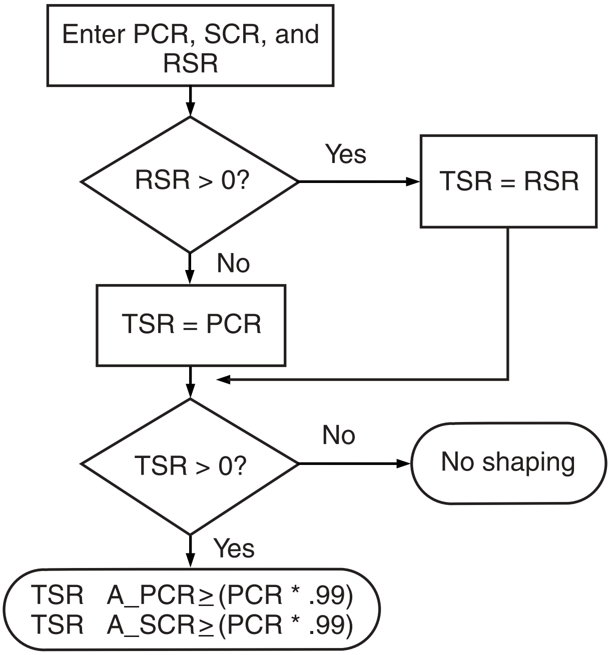

A_PCR	Actual peak cell rate
PCR	Peak cell rate
SCR	Sustained cell rate
TSR	Target shaping rate

Note: The PCR or SCR shaping rates are the actual shaping rates used in shaping

PPT 2855 047 AA

UPC and GCRA operating characteristics for ATM IP

ATM IP function processors (FPs) offer the following enhancements to UPC over CQC-based FPs:

- For GCRA 1, max(I1+L1) $\leq 2^8$ * I1 (see the table “Sample UPC operating parameters affecting maximum burst size” (page 78))
- For GCRA 2, max(I2+L2) $\leq 2^{16}$ * I2 (see the table “Sample UPC operating parameters affecting maximum burst size” (page 78))
- From the GCRA algorithm, the range for X = I+L is

  1.19 nanoseconds $\leq$ X $\leq$ 10.24 s

  for the range of UPC rates from 1 cell/s up to link rate. Compare this definition to the ranges for CQC-based FPs described in the traffic policing section in “Multiservice SwitchTraffic policing” (page 27).

- increased granularity of traffic contract (within +1.0% of traffic contract), with a worst case granularity error of 6.3% for MBS
- separate statistics for GCRA1 and GCRA2
- monitoring of GCRA violations

The separate statistics for GCRA1 and GCRA2 allows operators to determine if UPC discards are due to PCR/CDV, or SCR/BT violations. If the count of received cell discards or tagging is higher than the GCRA1 plus GCRA2 counts, the difference is typically the count of cells discarded due to PPD.

Table 12
Sample UPC operating parameters affecting maximum burst size

link per FP configuration	PCR (cell/s)	SCR (cell/s)	CDVT (cell/s)	Maximum MBS
One or more DS1 over DS3	106 132	3622	25 microsec	67872 cells + 6.3% accuracy error
One or more DS1 over OC3	353 207	3622	250 microsec	66 235 cells + 6.3% accuracy error
One or more DS3 over OC3	353207	106 132	250 microsec	93643 cells + 6.3% accuracy error
Note: For the given PCR, SCR, CDVT values, the maximum configurable MBS is the value shown in the far right column. Due to the accuracy error, the node does not discard or tag traffic even if the traffic exceeds this MBS by up to 6.3%.				

Semantic checks are as follows:

- for enforcer 1: max(I1+L1) < 2^8 * I1 (I1= 1/PCR, L1 = CDVT)
- for enforcer 2: max(I2+L2) < 2^16 * I2 (I2 = 1/SCR, L2 = BT+CDVT)
- PCR < 1/610ns
- SCR < 1/610ns
- MBS > 0

Cell counting

UPC counts discarded cells under the same operational attribute that counts received cells that UPC discards due to bus congestion. For data collection on the number of cells that a node tags through UPC, ATM IP FPs collect these counts.

ATM IP FPs also count non-conforming cells.

Interaction: ATM IP and CQC shaping and policing

Shaping and policing are complementary mechanisms. Subscribers use shaping to avoid exceeding rate contracts, while carriers use policing to prevent customers from exceeding rate contracts. At the edge of an ATM network, the ATM IP FP is the recommended FP for these reasons:

- overall throughput rate is sufficient to shape any traffic service category, including CBR
- shaping granularity is very fine, giving accurate rate choices

The figure “Traffic shaping options for ATM IP and CQC-based FPs” (page 80) shows a simple interaction scenario.

Figure 13
Traffic shaping options for ATM IP and CQC-based FPs

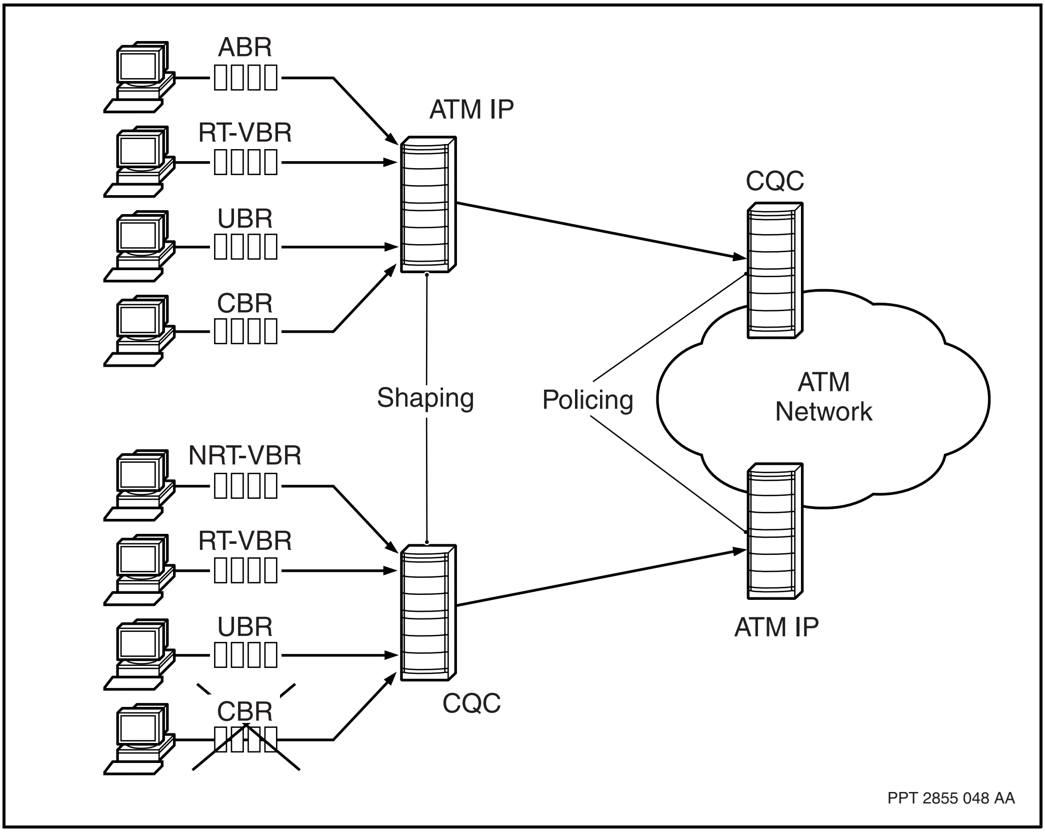

Relationship between UPC and linear traffic shaping

Connections may experience cell discards or cell tagging under these conditions:

- both ends have the same traffic descriptor parameters (not necessarily the same traffic descriptor types)
- UPC is enabled on the far end (the receive end)

Discards and tagging may occur even if traffic shaping is enabled on the near end (the transmit end). This scenario is possible if CQC-based FPs are configured on both ends. It can be avoided by using ATM IP FPs configured for VBR-ABR shaping.

Chapter 5
Shaping and policing on APC/PQC-based FPs

This chapter describes traffic shaping and policing requirements that are specific to APC/PQC-based function processors (FP).

Information is organized into the following sections:

Overview of shaping for APC/PQC-based FPs

Linear rate shaping on APC/PQC-based function processors is available through the premium emission priority, EP0. In linear rate shaping, a single shaping rate is used for each connection mapped out to EP0. The four regular emission priorities (EP2, EP3, EP4, and EP7) are unshaped.

Shaping options on EP0

Connections that use EP0 at each sub-port (ATM interface) can use 31 programmable shaping rates. The shaping rate ranges from 2 Kbits/s to 623.5 Mbits/s. There are two options available on a per-ATM service category basis:

- disabled
- enabled (linear rate shaping)

You can configure shaping at each connection to be disabled or to use the same default as the service category. If the ATM service category is shaped (CBR) but a particular connection has shaping disabled, the connection is actually shaped at the link rate because per-connection override prevails.

Shaping rate selection on EP0

As virtual circuits are configured, the shaping rates are automatically used as elements of the 31 shaping rates and are displayed under the *opShapingRate* attribute. When a connection is completed and if no other connection is using the rate, the *shapingUsage* counter for the rate becomes 0.

The default shaping rate is equal to the ATM interface rate.

You can display the number of virtual circuits using each of the 31 shaping rates through the *shapeUsage* attribute.

Shaping on CBR connections

On APC/PQC-based function processors, CBR traffic may be shaped as with any other service category. While shaping on “constant” connections may seem unnecessary, this capability has the following applications:

- to return CBR traffic to an orderly flow after some CDV has been introduced
- to adjust the flow of bursty CBR traffic (for example, if voice is served on CBR, compression and speech activity detection algorithms may introduce burstiness)
- to ensure that CBR traffic can pass through a stringent policer (for example, in a carrier situation)
- to enhance conformance to the limitations of an end-system playback device which may have a restricted buffer

Linear shaping using PCR as the shaping rate

Nortel Networks Multiservice Switch CAC nodes derive PCR shaping rates from the traffic descriptor parameters that are either configured or signaled for the VCCs and VPCs. Calculations include attributes that define the traffic descriptor type and the parameter values for that descriptor type.

The actual values used are fixed and cannot be adjusted either up or down. The shaping rate used is applied relative to the PCR that is configured for the connection through the traffic descriptor parameters. A small increase in the configured PCR value does not necessarily result in a change of shaping rate. See the figure “Algorithm for linear shaping” (page 86) for the algorithm used to derive the PCR or SCR shaping rates.

Figure 14
Algorithm for linear shaping

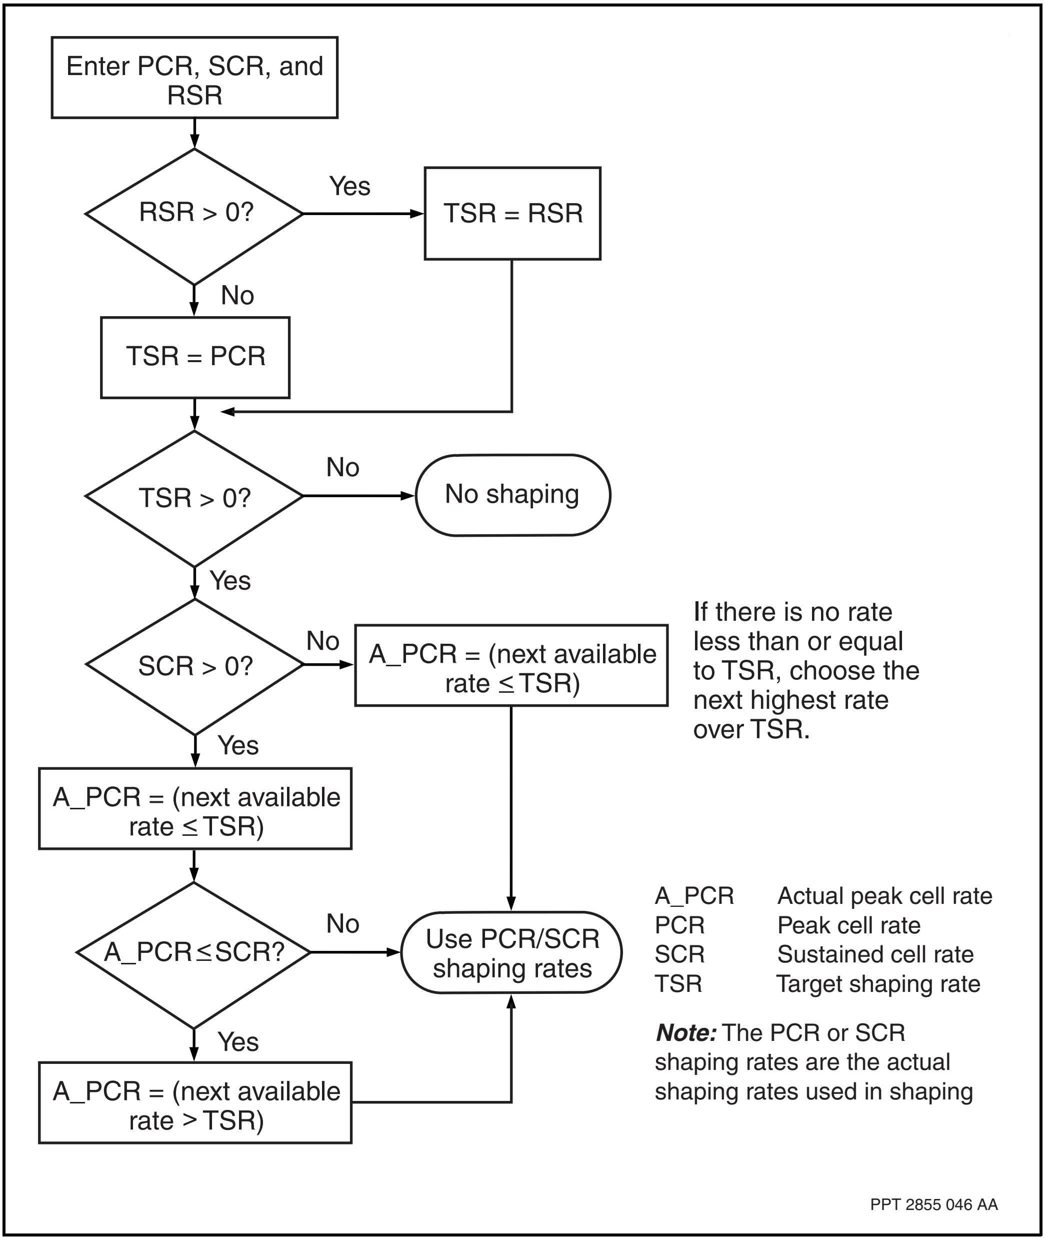

Shaping virtual circuits of multiple ATM service categories

You can map virtual circuits of two or more ATM service categories to EP0 and enable linear shaping but you must ensure that the shaping rates for the virtual circuits of different ATM service categories are properly configured. In these cases, follow these guidelines:

- do not over-book the shaping rates
- only ATM service categories with comparable QOS requirements should be mapped to EP0 to be linearly shaped together.

Traffic policing for APC/PQC-based function processors

This section describes the characteristics of traffic policing that are specific to APC/PQC-based function processors. For a general overview of traffic policing, see “Multiservice SwitchTraffic policing” (page 27).

Usage parameter control (UPC) monitors traffic for traffic contract violations on a per-VC basis and performs policing by discarding or tagging violating cells. The algorithm used to perform UPC policing is the generic cell rate algorithm (GCRA).

GCRA1 and GCRA2 work in tandem for each connection so up to two rates can be monitored and enforced upon the connection.Note that GCRA2 is disabled for all non-VBR connections. For a VBR connection, GCRA1 is used to police the peak cell rate (PCR) with a tolerance equal to the cell delay variation tolerance (CDVT). The VBR connection’s GCRA2 polices according to the sustained cell rate (SCR) with a tolerance of (CDVT + BT). BT is the burst tolerance defined as

(Maximum Burst Size - 1) * (1/SCR - 1/PCR)

The maximum burst size is the largest number of consecutive cells allowed to arrive back-to-back at the PCR rate.

Refer to the table “UPC conformance definitions for APC/PQC-based FPs” (page 88).

Table 13
UPC conformance definitions for APC/PQC-based FPs

Conformance definition	PCR policing (GCRA1)		SCR policing (GCRA2)	
	CLP type	Policing action	CLP type	Policing action
CBR	0 + 1	discarding	N/A	N/A
VBR.1	0 + 1	discarding	0 + 1	discarding
VBR.2	0 + 1	discarding	0	discarding
VBR.3	0 + 1	discarding	0	tagging
UBR.1	0 + 1	discarding	N/A	N/A
UBR.2	0 + 1	tagging/over-writing	N/A	N/A

Chapter 6
Shaping and policing for GQM-based FPs

This chapter describes traffic shaping and policing requirements that are specific to GQM-based function processors (FP).

Information is organized into the following sections:

Overview of shaping and policing for GQM-based FPs

The GQM-based FPs provide the same shaping and policing as other ATM FPs except for the additional capabilities identified in this chapter.

Types of shaping for GQM-based FPs

The types of traffic shaping for the GQM-based FPs is the same as AQM-based FPs (as described in “Shaping and policing on ATM IP FPs” (page 63)) except that up to four instead of two emission priorities (EPs) can be shaped.

Emission priorities for GQM-based FPs

The GQM-based FPs have eight emission priorities (EP) with two having absolute priority. Shaping is supported on any four of the eight emission priorities which allows any or all ATM service categories to be shaped while maintaining a separate EP for each ATM service category. The eight EPs include:

- absolute priority for EP0 and EP1 with minimum delay and cell delay variation (CDV)
- EPs 2 to 7 with minimum bandwidth guarantees (MBGs) to avoid starvation (for details, see the section for configuring MBGs for GQM-based FPs in NN10600-707 *Nortel Networks Multiservice Switch 7400/15000/20000 ATM Queuing and Scheduling Fundamentals*)
- work-conserving allocation of unused bandwidth for high utilization

As with existing Nortel Networks Multiservice Switch node ATM FPs, unshaped connections are allowed on a shaped EP. These connections are shaped at the link rate which implies that the per-VC queues are required and the individual connection weights are not used to schedule the traffic. The GQM-based FPs support the existing shaping options of minimizing cell delay variation (CDV) or maximizing bandwidth use per ATM service category through the attribute *shapeRecoupPolicy*.

Shaping rate for GQM-based FPs

The shaping rate for GQM-based FPs differs slightly from other ATM FPs, but the accuracy is the same.

Policing for GQM-based FPs

GQM-based FPs use UPC, PCR, and SCR the same as APC-based and PQC-based FPs, except the conformance to UBR.2 is only to discard cells of the connection on egress. Refer to “Traffic policing for APC/PQC-based function processors” (page 87).

Policing rate for GQM-based FPs

The maximum allowed delay variation is 17.18 seconds. The maximum is enforced by a semantic check for provisioned connections. Dynamic connections which request a cell delay variation (CDV) above the maximum are refused. The policing accuracy is 0.57% of the requested rate.

Policing enforcer violation counts for GQM-based FPs

GQM-based FPs use two enforcers per connection to police up to two rates. The non-conforming cells observed by the first and second enforcer, are counted as *rxUpcViolationOnEnforcer1* and *rxUpcViolationOnEnforcer2*, respectively. However, in certain situations, the violations counted by the enforcers on GQM-based FPs can be different from those counted by the enforcers on other FPs.

On FPs that are not GQM-based, the two enforcers perform conformance checks in a sequential manner. That is, each cell arrival is subject to a conformance check by Enforcer 1 followed by a conformance check by Enforcer 2. As a result, violations seen by Enforcer 1 will be filtered out and and will not be seen by Enforcer 2 and violations seen by Enforcer 2 will not impact the operation of Enforcer 1 in any way. This is because a cell arrival must have passed the check by Enforcer 1 before it can be checked by Enforcer 2.

In contrast, on GQM-based FPs, the two enforcers perform conformance checks simultaneously, and can influence one another in terms of violation counts. Under certain circumstances, this influence can result in readings of *rxUpcViolationOnEnforcer1* and *rxUpcViolationOnEnforcer2* on the GQM-based FPs that are different from the readings on other FPs.

GQM-based FPs have the following unique enforcer operations:

- When both enforcers are configured to discard nonconforming cells, a cell arrival seen by either one of the enforcers will advance the violation count of the enforcer by one. At the same time, the enforcer observing the violation will inform the other enforcer to disregard the cell arrival. This applies to Traffic Descriptor 4, 6, 7. Depending on enforcer parameters and traffic arrival profiles, certain violations on a connection that would have been counted as *rxUpcViolationOnEnforcer1* on other FPs can be counted as *rxUpcViolationOnEnforcer2* on GQM-based FPs.

When Enforcer 2 on a GQM-based FP detects a violation, Enforcer 2 will discard this cell and inform Enforcer 1 to disregard the cell arrival. This is different from the situation on other FPs where Enforcer 2 only checks arrivals filtered through by Enforcer 1 and in no way influences the operation of Enforcer 1.

- When Enforcer 1 is configured to discard and Enforcer 2 to tag nonconforming cells, Enforcer 1 informs Enforcer 2 about its violations but not the other way around. In this configuration, Enforcer 1 will inform Enforcer 2 to disregard a cell arrival if Enforcer 1 detects the cell as nonconforming. However, Enforcer 2 will not inform Enforcer 1 to disregard a cell arrival if Enforcer 2 sees a violation. This applies to Traffic Descriptor types 5 and 8. Depending on enforcer parameters and traffic arrival profiles, certain violations on a connection can be redundantly counted both as *rxUpcViolationOnEnforcer1* (for example, cells that are discarded) and as *rxUpcViolationOnEnforcer2* (for example, cells that are tagged) on GQM-based FPs. Redundant violation count happens on GQM-based FPs when a cell is seen by Enforcer 2 as nonconforming (and counted as an Enforcer 2 violation) but at the same time is also seen by Enforcer 1 as nonconforming. In this case, Enforcer 2 will count the cell as a tagging violation but cannot prevent Enforcer 1 from counting it again as a violation and thus discarding it. This is different from the situation on other FPs, where Enforcer 2 only checks arrivals filtered through by Enforcer 1 and thus only the cells that conform to Enforcer 1 will be checked by Enforcer 2.

For more information about the differences between the way enforcers monitor cell arrivals on GQM-based FPs and other FPs, see Multiservice Switch Engineering Guidelines: ATM Traffic Management.

Nortel Networks Multiservice Switch 7400/15000/20000

ATM Traffic Shaping and Policing

Release 6.1

Publication: NN10600-706
Document status: Standard
Document version: 6.1S1
Document date: August 2004
Printed in Canada

NORTEL NETWORKS