



Nortel Networks Multiservice Switch

7400/15000/20000

ATM CAC and Bandwidth Fundamentals

NN10600-708

Nortel Networks Multiservice Switch 7400/15000/20000

ATM CAC and Bandwidth Fundamentals

Publication: NN10600-708

Document status: Standard

Document version: 6.1S1

Document date: August 2004

Copyright © 2004 Nortel Networks.

All Rights Reserved.

Printed in Canada

NORTEL, NORTEL NETWORKS, the globemark design, the NORTEL NETWORKS corporate logo, DPN, and PASSPORT are trademarks of Nortel Networks.

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

About this document **13**

Who should read this document and why 13

What you need to know 14

How this document is organized 14

What's new in this document 15

Text conventions 15

Related documents 16

How to get more help 17

Chapter 1

Connection admission control **19**

Types of CAC and terms of use 19

Overview of connection admission control 20

 Atmlf-CAC 20

 CAC and traffic characteristics 22

CAC on permanent connections 24

CAC on switched connections over PNNI 24

 Discussion: GCAC and ACAC in PNNI 1.0 24

 Notes on GCAC for PNNI 1.0 26

 Multiservice Switch implementation of GCAC: EGH-based GCAC 28

CAC on switched connections over UNI, IISP, and AINI 29

ECR calculation by service category 30

 Constant bit rate CAC 33

 Relationship between CBR CAC and CDVT 34

 Real-time variable bit rate CAC 34

 Non-real time variable bit rate CAC 37

- Unspecified bit rate CAC 37
- CAC for virtual path termination 38
 - Atmlf-CAC serving capacity 39
- Overview of VPT-CAC 39
- VPT-CAC serving capacity and pools 40
- Atmlf-CAC and VPT-CAC characteristics 40

Chapter 2
Bandwidth pool management **43**

- General characteristics of bandwidth pools 43
 - Bandwidth pools and VPT-CAC 47
- Bandwidth pool over- and under-subscription 48
- Bandwidth pool sharing 52
- Bandwidth pool management for point-to-multipoint SVCs 53
 - Connection bandwidth control 53

Chapter 3
Dynamic bandwidth management **55**

- Dynamic bandwidth and ATM connections 56
- Benefits of dynamic bandwidth management 57
 - Related features 58
- Dynamic bandwidth management and connection types 59
 - Non-elastic connections 59
 - Elastic connections 60
- Connection bandwidth control 61
 - Objectives of dynamic bandwidth management 64
 - Application of the CBC algorithm 64
 - CBC and point-to-multipoint connections 66
 - Rules governing CBC 66
 - How CBC handles bandwidth loss 72
 - How CBC handles bandwidth gain 75
- Examples of dynamic bandwidth management 78
- Example: non-elastic connections responding to dynamic bandwidth 78
 - Step A: losing the first link 80
 - Step B: losing the second link 81

Step C: losing the third and fourth links	81
Example: elastic connections responding to dynamic bandwidth	81
Step A: losing the first link	83
Step B: losing the second link	83
Step C: regaining a link	84
Example: mixed connections responding to dynamic bandwidth	85
Step A: losing the first link	87
Step B: losing the second link	88
Step C: regaining a link	89
CBC and bandwidth pools	89
Dynamic bandwidth and data protection	92
Coordinating dynamic bandwidth between nodes	93
VPT CAC response to dynamic bandwidth changes	93

List of figures

- Figure 1 Equivalent cell rate calculations under Atmlf-CAC 21
- Figure 2 CAC parameters 23
- Figure 3 Network node execution of GCAC and ACAC 26
- Figure 4 CAC on Multiservice Switch ATM FPs for switched connections over UNI, IISP, and AINI 30
- Figure 5 Calculation of equivalent cell rate for connection admission control 32
- Figure 6 Example of bandwidth allocation: 100 percent allocation 45
- Figure 7 Example of progression of pool allocation for a single service category 47
- Figure 8 Example of bandwidth allocation for VPTs 48
- Figure 9 Example of bandwidth allocation: over-subscription 50
- Figure 10 Example of bandwidth allocation: under-subscription 51
- Figure 11 Example of bandwidth allocation: common pool 52
- Figure 12 ATM logical trunks as elastic and non-elastic connections 61
- Figure 13 Example of dynamic bandwidth management using CBC algorithm 63
- Figure 14 How the CBC algorithm handles bandwidth loss 74
- Figure 15 How the CBC algorithm handles bandwidth gain 77
- Figure 16 How non-elastic connections in a pool respond to changes in bandwidth 79
- Figure 17 How elastic connections in a pool respond to changes in bandwidth 82
- Figure 18 How mixed connections in a pool respond to changes in bandwidth 86
- Figure 19 Example of how bandwidth loss is handled within pools 90

List of tables

Table 1	Multiservice Switch ATM supporting information	14
Table 2	Comparison of GCAC algorithms and metrics	27
Table 3	Allocation equivalent cell rate by ATM service category	30
Table 4	CBR CAC parameters for equivalent cell rate	33
Table 5	Real-time VBR CAC parameters for equivalent cell rate	35
Table 6	Transmit traffic descriptor type settings	36
Table 7	Comparison of Atmlf-CAC and VPT-CAC characteristics	40
Table 8	IMA and related Multiservice Switch features	58

About this document

This document contains information for Nortel Networks Multiservice Switch node ATM traffic management controls for connection admission control (CAC) and bandwidth management.

The following topics are discussed in this section:

- “Who should read this document and why” (page 13)
- “What you need to know” (page 14)
- “How this document is organized” (page 14)
- “What’s new in this document” (page 15)
- “Text conventions” (page 15)
- “Related documents” (page 16)
- “How to get more help” (page 17)

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 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 required for this document	Location of supporting information
For an explanation of the following FP types: <ul style="list-style-type: none"> • CQC 	see the description of ATM function processors in NN10600-700 <i>Nortel Networks Multiservice Switch 7400/15000/20000 ATM Technology Fundamentals</i>
For information about addressing, signaling, and routing	see NN10600-702 <i>Nortel Networks Multiservice Switch 7400/15000/20000 ATM Routing and Signalling Fundamentals</i>
For traffic management concepts	see NN10600-705 <i>Nortel Networks Multiservice Switch 7400/15000/20000 ATM Traffic Management Fundamentals</i>
For traffic shaping and policing	see NN10600-706 <i>Nortel Networks Multiservice Switch 7400/15000/20000 ATM Traffic Shaping and Policing Fundamentals</i>
For cell queuing and queue scheduling capabilities	see NN10600-707 <i>Nortel Networks Multiservice Switch 7400/15000/20000 ATM Queuing and Scheduling Fundamentals</i>

How this document is organized

NN10600-708 *Nortel Networks Multiservice Switch 7400/15000/20000 ATM CAC and Bandwidth Fundamentals* contains the following sections:

- “Connection admission control” (page 19)
- “Bandwidth pool management” (page 43)
- “Dynamic bandwidth management” (page 55)

These descriptions support the procedural information provided in the NN10600-710 *Nortel Networks Multiservice Switch 7400/15000/20000 ATM Configuration Management* and NN10600-715 *Nortel Networks Multiservice Switch 7400/15000/20000 ATM Fault and Performance Management*. To illustrate concepts, examples are provided as necessary.

What's new in this document

There were no new features added to this document.

Other changes made to this document include the following:

- For CR Q00881217, the section “Unspecified bit rate CAC” (page 37) was updated with a note stating that CBR calls with an equivalent cell rate (ECR) higher than a specified amount will fail.
- 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*.

Text conventions

This document uses the following text conventions:

- non-proportional spaced plain type
Non-proportional spaced plain type represents system generated text or text that appears on your screen.
- non-proportional spaced bold type
Non-proportional 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 node 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*
- NN10600-706 *Nortel Networks Multiservice Switch 7400/15000/20000 ATM Traffic Shaping and Policing Fundamentals*
- NN10600-707 *Nortel Networks Multiservice Switch 7400/15000/20000 ATM Queuing and Scheduling 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*
- NN10600-920 *Nortel Networks Multiservice Switch 7400/15000/20000 Operations: Frame Relay to ATM Interworking*
- NN10600-420 *Nortel Networks Multiservice Switch 7400/15000/20000 Operations: Trunking*
- 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 document*.

Chapter 1

Connection admission control

Connection admission control (CAC) is the traffic management technique for accepting or rejecting connections at set up. It is an access-specific control.

This chapter provides information in the following sections:

- “Types of CAC and terms of use” (page 19)
- “Overview of connection admission control” (page 20)
- “CAC on permanent connections” (page 24)
- “CAC on switched connections over PNNI” (page 24)
- “CAC on switched connections over UNI, IISP, and AINI” (page 29)
- “ECR calculation by service category” (page 30)
- “CAC for virtual path termination” (page 38)

Types of CAC and terms of use

There are two types of CAC: ATM interface CAC (AtmIf-CAC) and virtual path termination CAC (VPT-CAC). Throughout the discussion in this chapter, the term CAC is applied in the following ways:

- CAC is a generic term in this document that refers all types of CAC, usually in a high-level discussion.

- AtmIf-CAC is Nortel Networks Multiservice Switch node's CAC implementation that applies to independent connections (VCs and VPs), and to VPT VCCs that do not have configured bandwidth allocation. AtmIf-CAC includes generic CAC (GCAC) and actual CAC (ACAC). See "CAC on switched connections over PNNI" (page 24).
- VPT-CAC is Nortel Networks Multiservice Switch node CAC implementation that applies to VPTs and VPT VCCs. VPT-CAC is similar to AtmIf-CAC but applied on a per-VPT basis.

There are also terms like constant bit rate (CBR) CAC that refer to specific algorithms that apply to a specific service category.

Overview of connection admission control

CAC applies to the egress link to the network to help ensure service guarantees. With a correctly configured CAC, a node can refuse a new connection at the egress point to the network. Refusal occurs when there is a high risk that the new connection may prevent the network from maintaining service for established connections. If the node accepts the new connection, the connection setup function establishes the call across the network and updates any required state information along the path.

CAC takes advantage of statistical multiplexing of cells from different connections, taking into account all traffic descriptor parameters. CAC as a traffic management control maximizes link usage while maintaining the quality of service for all admitted connections.

AtmIf-CAC

Note: For information on VPT-CAC, see "Overview of VPT-CAC" (page 39).

The function processor uses the full link capacity as the serving capacity for all connection points admitted by AtmIf-CAC. Serving capacity is one of the traffic parameters used by AtmIf-CAC to determine if a link can accommodate a VCC or VPC connection point.

The figure “Equivalent cell rate calculations under AtmIf-CAC” (page 21) summarizes the equivalent cell rate (ECR) calculations for unspecified bit rate (UBR) with minimum cell rate (MCR), constant bit rate (CBR), real-time variable bit rate (RT-VBR), and non-real-time variable bit rate (NRT-VBR) traffic.

Figure 1
Equivalent cell rate calculations under AtmIf-CAC

ECR= 0 for a UBR connection without minimumCellRate provisioned. For example, minimumCellRate= 0.

ECR= minimum cell rate (MCR) if minimumCellRate is present. For example, minimumCellRate> 0.

$ECRCBR = \Phi \text{ServingCapacityCBR, QueueSizeCBR, CLRCBR, PCRCBR, CDVTCBR}$

$ECRRT = \Phi \text{ServingCapacityRT, QueueSizeRT, CLRRT, PCRRT, SCRRT, MBSRT}$

Note: Serving capacity is one of the traffic parameters used by the CAC algorithm to determine if the link node can accommodate a VCC or a VPC connection point.

The node uses ECR with the other controls that are configured for the ATM interface to determine admission. AtmIf-CAC admits the connection point if its ECR is less than the available bandwidth. AtmIf-CAC then deducts the ECR from the available bandwidth for the bandwidth pool

The ECR calculation is based on statistical multiplexing of all connection points that share the link bandwidth. In general, better effect of statistical multiplexing is achieved with more VBR type connection points sharing the physical link. Links with the higher cell rates provide access to more connections points. A node can admit a connection point with a smaller ECR

if it shares more bandwidth with other connections. The inverse is also true: when a number of connection points share more bandwidth, the node can admit a new connection if it has a smaller ECR.

CAC and traffic characteristics

Acceptance of the connection is based on the traffic characteristics of both the requested connection and the existing connections. The decision to accept or reject a connection also takes into account some or all of the following parameters:

- queue limit
- service category and cell loss ratio (CLR)
- serving capacity (link capacity for AtmIf-CAC and PCR for VPT-CAC)
- traffic descriptor parameters: peak cell rate (PCR) for the connection, sustained cell rate (SCR) for the connection, maximum burst size (MBS) for the connection, requested shaping rate, cell delay variation tolerance (CDVT), and minimum cell rate (MCR)
- PCR or SCR shaping rates

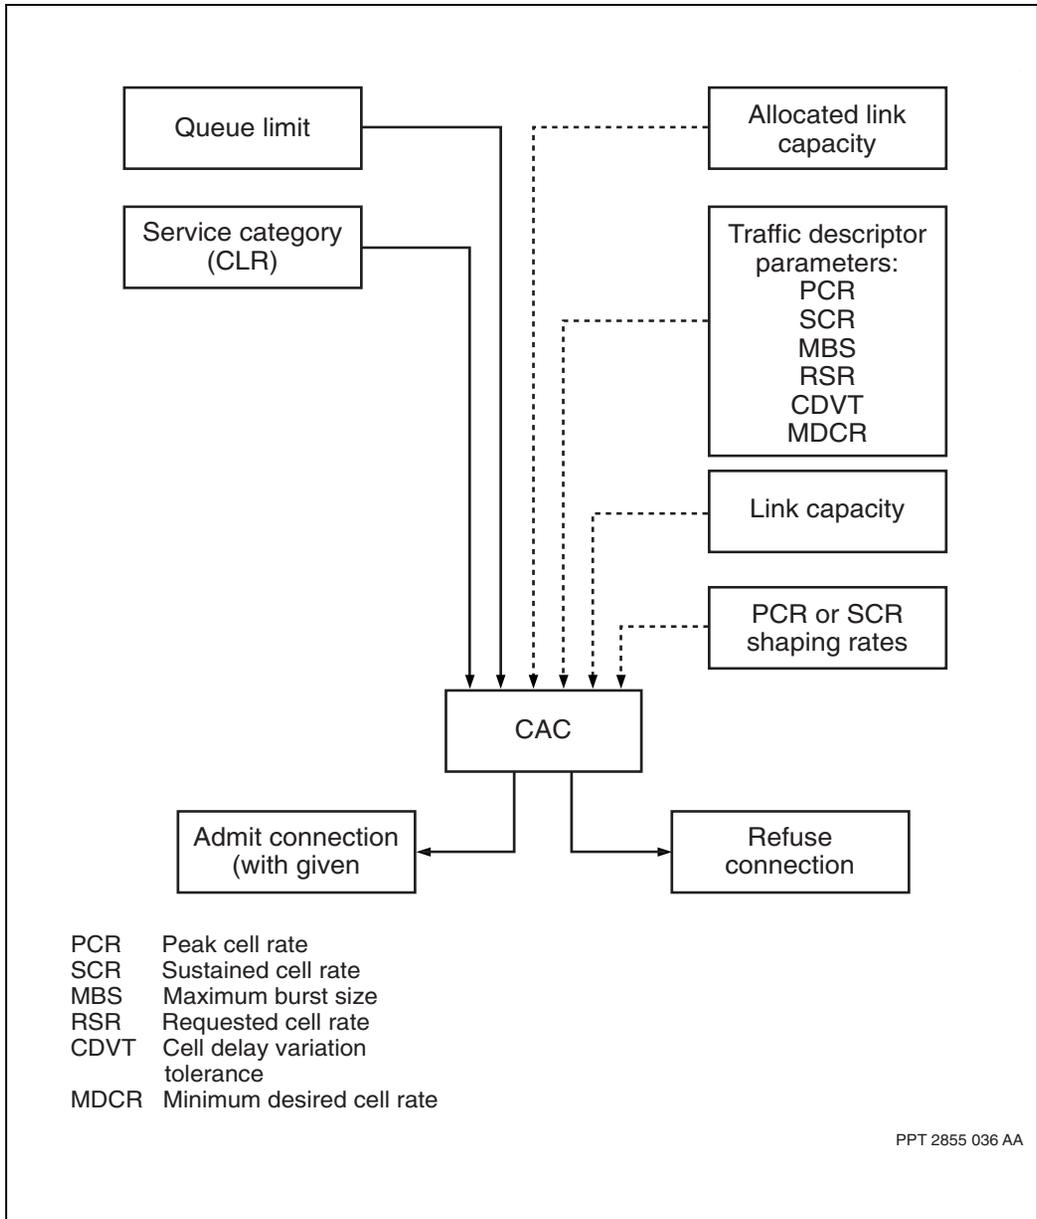
CAC considers only those links that may need capacity reservation as possible bottlenecks. The figure “CAC parameters” (page 23) summarizes these parameters. Using these parameters, the CAC algorithm determines the amount of bandwidth to reserve for the connection. Bandwidth is expressed as ECR.

Lastly, note that CAC rejects call setup requests where the call parameters fall outside of the hardware capabilities and ranges.

For more information on how these parameters influence ECR, see the following sections:

- “Constant bit rate CAC” (page 33)
- “Real-time variable bit rate CAC” (page 34)
- “Non-real time variable bit rate CAC” (page 37)
- “Unspecified bit rate CAC” (page 37)

Figure 2
CAC parameters



CAC on permanent connections

For permanent connections, Nortel Networks Multiservice Switch nodes apply AtmIf-CAC at each configured relay and end-point.

CAC on switched connections over PNNI

PNNI 1.0 includes generic connection admission control (GCAC) and actual connection admission control (ACAC).

During route determination for an incoming call setup, the PNNI node applies GCAC for route selection. GCAC calculates the expected CAC behavior of other nodes, based on the additive link metrics for that node and the QOS of the connection request. During call setup, each node along the route applies ACAC to calculate the required capacity of the connection. ACAC guarantees the QOS of the new connection and the existing connections.

The following sections provide information on:

- the differences between GCAC and ACAC in PNNI 1.0
- the GCAC options that are available in PNNI 1.0
- Multiservice Switch node implementation using the extended Gibbens-Hunt (EGH) method

Discussion: GCAC and ACAC in PNNI 1.0

When an access node receives a request for a new connection, it must select a route from source to destination. To do this, the node must have information on the required bandwidth for the connection and the available bandwidth on each link in the route. To obtain this information, ATM switches in the network advertise some information about their internal ACAC states. However, ACAC is not subject to standardization in the PNNI 1.0 protocol. Further, differences in architecture and overhead make ACAC impractical for switching systems to advertise potentially different ACAC states.

To solve this problem, GCAC allows switching systems to advertise ACAC information that is generic and compact but flexible enough to support any CAC. GCAC has these characteristics:

- is independent of switch-architecture
- is invoked at the access node for route selection

- calculates an estimate of required connection capacity
- uses advertised link state metrics to select links
- includes only those links with sufficient capacity

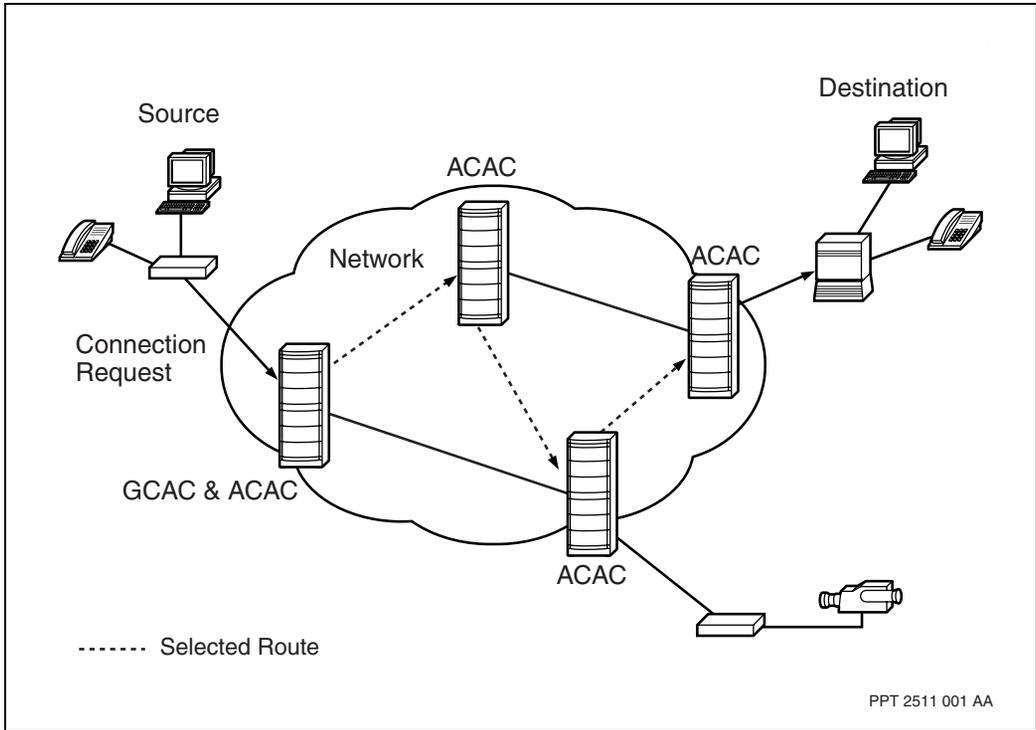
The GCAC algorithm uses the advertised parameters (available from routing database) and the characteristics of the requested connection (available from signaling) to include nodes and links in the route. If the GCAC algorithm determines that a node or link can accept the connection, it includes that node or link in the route.

The ACAC algorithm works differently from the GCAC algorithm. After the node at the network access point selects the path, each node along the selected route executes its own ACAC algorithm. This application of ACAC at the node level guarantees quality of service. ACAC uses the node's knowledge of its queuing architecture, link speed, and available serving capacity, together with the connection's characteristics and QOS requirements. See the figure "Network node execution of GCAC and ACAC" (page 26).

ACAC has these characteristics:

- is precise
- is dependent on switch architecture
- has less stringent execution time requirements than GCAC
- is fast, since it operates in real time

Figure 3
Network node execution of GCAC and ACAC



Notes on GCAC for PNNI 1.0

The PNNI 1.0 specification defines two GCAC methods: complex GCAC and simple GCAC. The table “Comparison of GCAC algorithms and metrics” (page 27) summarizes the applicable parameters for both GCAC methods.

Table 2
Comparison of GCAC algorithms and metrics

Metric	Required/ Optional	Complex GCAC	Simple GCAC
available cell rate (AvCR)	required	used	used
maximum carrying capacity (MCC)	optional	used	not used
cell rate margin (CRM)	optional	used	not used
variance factor (VF)	optional	used	not used
minimum cell rate (MCR)	optional	not used	used

For a network in which nodes advertise all four link state metrics, the service provider uses complex GCAC. Complex GCAC takes advantage of these metrics and makes a good guess about the required cell rate for a connection. However, complex GCAC requires a calculation of one equation for each link in the connection path. Further, each node must update and broadcast all four link state metrics for each link and traffic service category. The update and broadcast processes increase both node overhead and network traffic.

Nortel Networks Multiservice Switch nodes do not support complex GCAC, nor do nodes support the generation, advertisement, or maintenance of associated MCC, CRM, and VF optional parameters.

For a network in which nodes advertise available cell rate (AvCR) only, the service provider uses simple GCAC. Simple GCAC overcomes the overhead and resource demands inherent in complex GCAC. However, observation and experimentation demonstrates that the EGH method is considerably faster and more accurate than simple GCAC.

Multiservice Switch implementation of GCAC: EGH-based GCAC

Nortel Networks Multiservice Switch nodes use a GCAC algorithm based on the extended Gibbens-Hunt (EGH) method. This algorithm estimates ECR. As a result, EGH-based GCAC provides a way to predict a typical node-specific ACAC algorithm with a minimum number of link state parameters. To estimate ECR, nodes calculate the average between the minimum and maximum number of connections that the link can support. Calculations involve the peak cell rate (PCR), sustained cell rate (SCR), and the link rate.

A Multiservice Switch node calculates the minimum number of connections that it can admit without introducing nodal loss or delay as

$$\min = R/\gamma$$

where the sum of the PCRs (denoted by $\Sigma \gamma$, where $\text{PCR}=\gamma$) is less than the link rate (denoted by R).

A node calculates the maximum number of connections as

$$\max = R/\Sigma\theta\gamma$$

where, under conditions of sustained congestion on the link, the sum of the SCRs (denoted by $\Sigma \theta\gamma$, where $\text{SCR}=\theta\gamma$ and θ is the source activity defined as $\theta = \text{SCR}/\text{PCR}$) exceeds the link rate (denoted by R).

The algorithm then divides the link rate by the average of min and max to obtain an estimate of the ECR (denoted by Ω_0) as

$$\Omega_0 = \gamma((2\theta)/(1+\theta))$$

For UBR connections, the ECR value is equal to the MCR value. If MCR is not specified, then the ECR value is equal to 0.

After a node estimates the ECR value once per connection, it compares the resulting ECR value to the available cell rate (AvCR) value for each link (AvCR is in the routing database) and includes or excludes the link according to the following criteria:

if ($ECR < AvCR$) include the link; else exclude the link

For a UBR connection with a network generated MCR IE, the ECR value is always equal to 0.

CAC on switched connections over UNI, IISP, and AINI

For switched connections, Nortel Networks Multiservice Switch network applies CAC at each hop. When the originating node sets up a switched connection, the nodes in the route undertake the following steps:

- 1 At the requesting end-point, CAC admits or refuses the connection based on traffic characteristics and congestion conditions at that node.
- 2 Through signaling, the node directs call set-up to the next node in the selected route.
- 3 At each subsequent hop, the node applies the CAC algorithm to admit or refuse the connection.

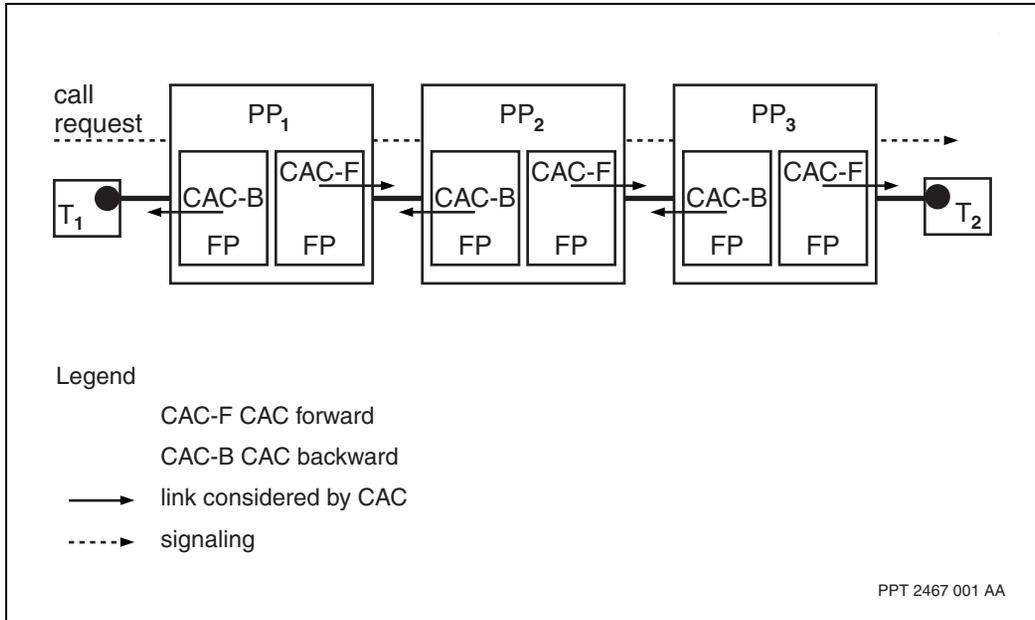
CAC determines if the node can support the requested ATM service category and if there is enough capacity left on the link to accommodate the new call without affecting those calls already in progress.

As a rule, configuration defines the same ATM service category for both directions of a connection. On the other hand, the QOS for the connection can be different for each direction. Multiservice Switch node's CAC ignores the QOS values except for the unspecified bit rate (UBR) services. For all other services (CBR, RT-VBR, and NRT-VBR), only the ATM service category determines how CAC admits the connection.

Note: For UBR services, set QOS to 0 for connections over UNI interfaces and for IISP 1.0 interfaces based on UNI signaling.

For each connection, the node invokes CAC twice, as shown in the figure "CAC on Multiservice Switch ATM FPs for switched connections over UNI, IISP, and AINI" (page 30). First, the receiving FP verifies that its transmit link can accommodate the backward direction of the requested connection. Then, the transmitting FP verifies that its transmit link can accommodate the forward direction of the requested connection. If the link can satisfy these constraints, the node signals the call setup message over the transmit interface to the next hop on the path to the destination.

Figure 4
CAC on Multiservice Switch ATM FPs for switched connections over UNI, IISP, and AINI



ECR calculation by service category

The node determines equivalent cell rate (ECR) based on the ATM service category for the incoming connection request. The traffic parameters for the traffic descriptor type also determine the value of the ECR. The table “Allocation equivalent cell rate by ATM service category” (page 30) shows these equivalents.

Table 3
Allocation equivalent cell rate by ATM service category

Service category	Allocated ECR
CBR	$PCR \leq ECR$
RT-VBR, NRT-VBR	$SCR \leq ECR \leq PCR$
UBR	$zero \leq MCR = ECR$

If a connection does not have shaping enabled, the node calculates the ECR value based on the connection PCR shaping rate. If a connection does have shaping enabled, the node calculates the ECR value based on the connection PCR shaping rate or based on PCR, depending on the type of shaping enabled:

- if linear shaping is enabled, the basis for ECR is the PCR shaping rate for the connection
- if VBR shaping is enabled, the basis for ECR is
 - the actual PCR for the connection if ECR is less than or equal to the actual PCR
 - the PCR shaping rate for the connection if ECR is equal to the actual PCR

The figure “Calculation of equivalent cell rate for connection admission control” (page 32) summarizes this relationship.

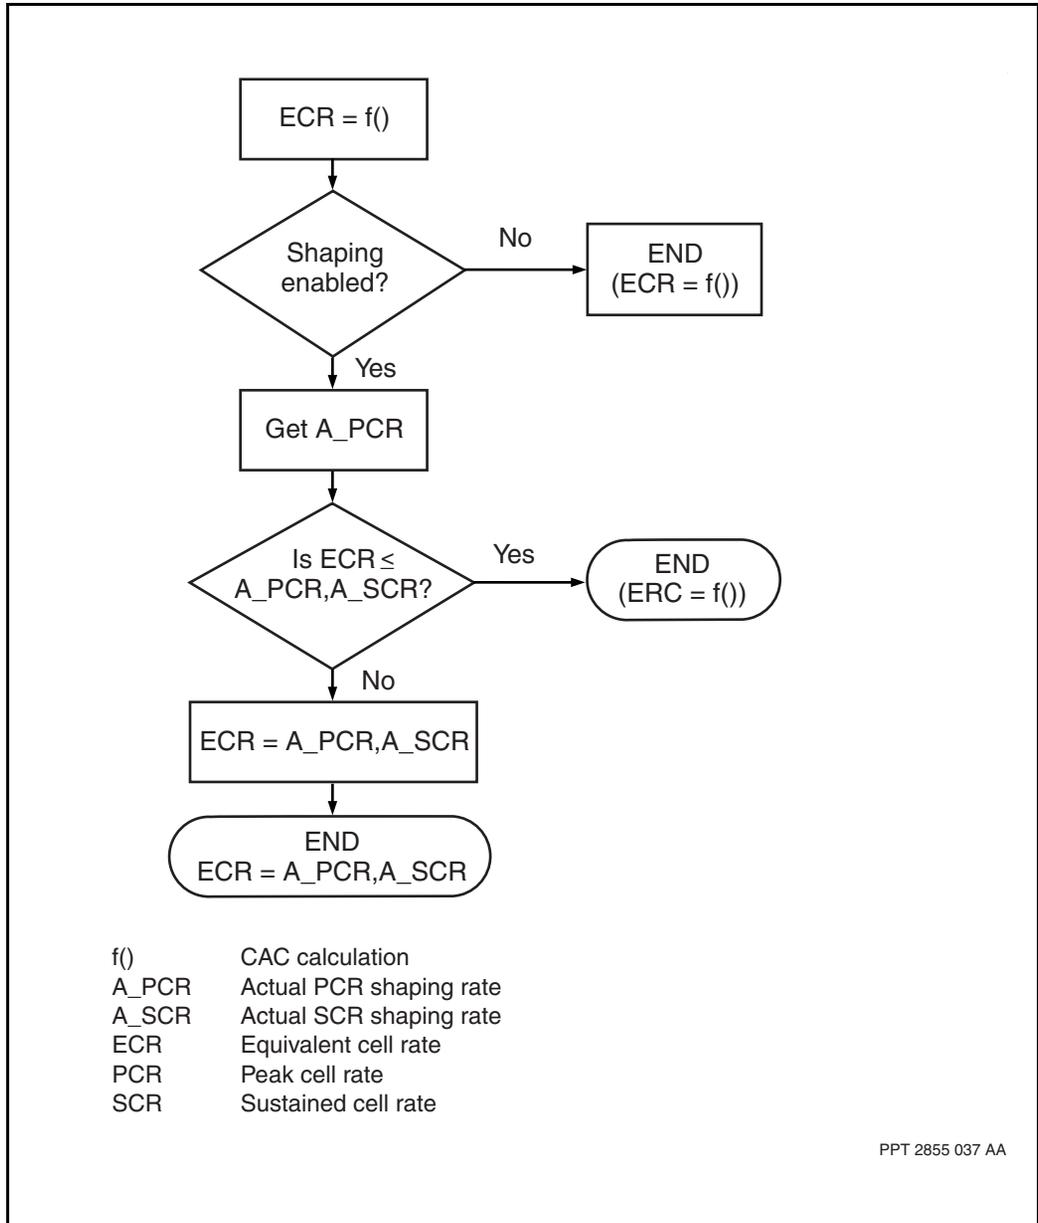
When traffic shaping and per-VC queuing is disabled, the queue sizes of the common queues are used in the ECR computation. When per-VC queuing or traffic shaping is enabled, the queue size corresponding to the shaping rate is used in the ECR computation.

The following sections provide descriptions of how Multiservice Switch nodes apply CAC to traffic in each service category:

- “Unspecified bit rate CAC” (page 37)
- “Constant bit rate CAC” (page 33)
- “Real-time variable bit rate CAC” (page 34)
- “Non-real time variable bit rate CAC” (page 37)

For information on traffic shaping, see NN10600-706 *Nortel Networks Multiservice Switch 7400/15000/20000 ATM Traffic Shaping and Policing Fundamentals*.

Figure 5
Calculation of equivalent cell rate for connection admission control



Constant bit rate CAC

Because CBR traffic maps to high emission priority, VBR traffic, which is served by a lower-level emission priority, does not affect CBR traffic. As a result, Nortel Networks Multiservice Switch nodes apply CBR CAC independently from CAC for other traffic.

The value of ECR for a CBR connection is greater than or equal to its PCR. The ECR depends on the parameters shown in the table “CBR CAC parameters for equivalent cell rate” (page 33).

Table 4
CBR CAC parameters for equivalent cell rate

Parameter	Direction change of this parameter	Effect of the change on ECR	Remarks
PCR	increasing	increases	PCR is $V_{cc} V_{cd} T_m$ $txTrafficDescParm1$
CDVT	increasing	increases	transmit CDVT is either <ul style="list-style-type: none"> • $V_{cc} V_{cd} T_m$ $txTrafficDescParm 4$ • $Cbr/0 Cdv t$
CLR	decreasing	increases	CLR is $Cbr/0 Clr$

In theory, CBR traffic is characterized by periodic cell arrival, such that ECR of a CBR connection is equal to its PCR. In practice, however, cells do not always arrive at the node equally spaced in time because of buffering or contention at the customer premises equipment (CPE) or the upstream nodes. As a result, ECR is always greater than or equal to the PCR.

CBR CAC also takes into account the CDVT of a connection and the desired maximum cell loss ratio (CLR). The CDVT value for a connection defines the expected degree of irregularity in cell arrival for CBR traffic. This irregularity may, in some cases, cause the ECR (reserved bandwidth) to be larger than the PCR. The difference is especially important when the CDVT is greater than $1/PCR$ (assuming the input link rate is much greater than PCR) or when the buffer size is not large enough to accommodate the random starting phases of different CBR connections.

Relationship between CBR CAC and CDVT

In summary, CDVT is used by CBR CAC for traffic descriptor types 3 through 8 in the transmit direction. Recall that CDVT does not apply to either RT-VBR CAC or NRT-VBR CAC.

The value defined for CDVT affects the amount of bandwidth reserved for the connection, which in turn influences how CAC evaluates the connection request. If, for example, $PCR = 1000$ cells/s and $CDVT = 0$, then $ECR = 1000$ cells/s (that is, CDVT has no effect on the equivalent cell rate and therefore the reserved bandwidth). As the CDVT is increases, the ECR value based on PCR also increases, and the net result is an increase in the reserved bandwidth for the connection. In a busy network where the available bandwidth may be low, CAC may reject a request for a connection that is defined with a high CDVT value.

The effective increase for ECR is limited to 115% of PCR. A higher ECR does not provide better service and therefore reserves unnecessary bandwidth. This unnecessary reservation of bandwidth could block admission of other connections that the link would otherwise readily support.

Real-time variable bit rate CAC

Real-time VBR traffic maps to the medium-level priority queues, and NRT-VBR traffic maps to the lower-level priority queues. The exact mapping depends on the function processor type (ATM IP or CQC). The node calculates the ECR for VBR CAC by subtracting the cumulative cell rate for CBR traffic from the serving capacity, and the ECR for NRT-VBR CAC by taking into account the remaining capacity available to NRT-VBR traffic.

The value of ECR for a RT-VBR connection lies between the values of PCR and SCR for that connection. The ECR depends on the parameters shown in the table “Real-time VBR CAC parameters for equivalent cell rate” (page 35).

Table 5
Real-time VBR CAC parameters for equivalent cell rate

Parameter	Direction change of this parameter	Effect of the change on ECR	Remarks
PCR or the PCR or SCR shaping rates	increasing	increases	PCR is <i>Vcc Vcd Tm txTrafficDescParm 1</i> PCR or SCR shaping rates is <i>Vcc Vcd Tm txTrafficDescParm 5</i>
SCR	increasing	increases	SCR is <i>Vcc Vcd Tm txTrafficDescParm 2</i>
MBS	increasing	increases	MBS is <i>Vcc Vcd Tm txTrafficDescParm 3</i>
VBR capacity	decreasing	increases	VBR capacity is not a visible attribute.
queue limit	decreasing	increases	Queue size is a visible attribute. Queue limit is <i>Vcc txQueueThresholds 0</i>
CLR	decreasing	increases	CLR is either <ul style="list-style-type: none"> • <i>RtVbr/0 Clr</i> • <i>NrtVbr/0 Clr</i>

If traffic shaping is not enabled for a connection, Nortel Networks Multiservice Switch nodes calculate the ECR value based on its PCR. However, if traffic shaping is enabled, the basis for ECR on ATM IP function processors is the PCR for the connection. For information on queue limits, see the sections on queuing and scheduling in NN10600-707 *Nortel Networks Multiservice Switch 7400/15000/20000 ATM Queuing and Scheduling Fundamentals*.

When traffic shaping and per-VC queuing are disabled, the queue sizes of the common queues are used in the ECR computation. However, when per-VC queuing or traffic shaping is enabled, the queue size corresponding to the shaping rate is used in the ECR computation.

Multiservice Switch nodes also apply correction factors through which the CLR probability is determined based on the probability that the combined source rates exceed the link rate (the need for buffering) and the probability of buffer overflow. The calculations involve three steps in evaluating the requested connection:

- 1 compute the probability of cell overflow into the buffer
- 2 compute the probability of buffer overflow
- 3 compute the ECR

The connection is admitted only if the CLR probability is less than the target CLR configured for the connection. This approach yields the best link usage over all other evaluated CAC methods, with efficiency that approaches the theoretical upper bound. See *Performance Evaluation of Connection Admission Control Techniques in ATM networks* for details. (For a complete reference, see the list of references in NN10600-700 *Nortel Networks Multiservice Switch 7400/15000/20000 ATM Technology Fundamentals*.)

VBR connections are normally specified with transmit traffic descriptor type 6, 7 or 8, and a defined PCR, SCR, and MBS. These connections are admitted as defined in the preceding paragraphs. However, connections with other transmit traffic descriptor types are treated as defined in the table “Transmit traffic descriptor type settings” (page 36).

Table 6
Transmit traffic descriptor type settings

Attribute setting	Meaning
txTrafficDescriptorType 1 or 2	no bandwidth reserved (ECR=0)
txTrafficDescriptorType 3	ECR = PCR
txTrafficDescriptorType 4 or 5	ECR = PCR for CLP0

Non-real time variable bit rate CAC

Non-real time VBR CAC is based on the available serving capacity as a function of the medium priority queue. To derive a useful ECR, CAC must first calculate the equivalent of the available serving capacity to serve NRT-VBR traffic. Then, the method used for VBR CAC is applied to derive the ECR.

For NRT-VBR traffic, the VBR capacity is a function of the link rate and the combined bandwidth of admitted CBR and RT-VBR connections. NRT-VBR traffic gets bandwidth only after any CBR or RT-VBR traffic has been served.

Unspecified bit rate CAC

UBR and UBR with MDCR CAC admit all connections up to the maximum number of UBR connections defined through connection administration for the interface. The number of UBR connections may be limited by setting the *maxVccs* and *maxVpcs* attributes of the *AtmIf Ca UBR* component.

Note: UBR calls greater than 1, 639, 344 ECR will fail with cause code 37.

If no MDCR value is defined for a UBR connection, no bandwidth will be reserved for that connection and the connection is treated as if it has an equivalent cell rate (ECR) equal to zero. If an MDCR value is defined, then the MDCR value is deemed to be the ECR value and is used for reserving bandwidth for the connection. This MDCR value is removed from the available bandwidth for the UBR pool when the connection is admitted. When the UBR pool bandwidth is exhausted, no other UBR connections will be admitted.

For UBR with MDCR CAC, the FPs offer the following methods to reserve bandwidth:

- UBR service category connections may be optionally assigned to a dedicated bandwidth pool.
- the number of UBR connections may be controlled by setting the *maxVccs* attribute of the *AtmIf Ca ubr/0* component.

- for designated UBR connections, a specific MDCR may be provisioned which applies only to that connection. When that connection is activated, the provisioned MDCR is removed from the bandwidth pool for UBR connections and the current number of configurable connections is decreased by 1.

Note: VPT MDCR can be set the same as AtmIf MDCR.

For information on how an MDCR value can be defined for a UBR connection, see NN10600-710 *Nortel Networks Multiservice Switch 7400/15000/20000 ATM Configuration Management*.

CAC for virtual path termination

For VP termination, a node applies one of two levels of CAC, depending on the connection. The first level is ATM interface CAC (AtmIf-CAC) which controls the admission of VPC and VCC connection points under the ATM interface. The second level is virtual path termination CAC (VPT-CAC) which controls the admission of VCC connection points associated with a VPT.

When VPT-CAC does not apply, AtmIf-CAC treats a VPT as if it has an ECR equal to 0 and ignores VPT traffic descriptor provisioning. As a result, AtmIf-CAC does not refuse connection admission for the VPT because the VPT does not request bandwidth. AtmIf-CAC then admits each associated VPT VCC based on the VCC bandwidth requirements and the available bandwidth on the interface.

The service provider can also configure a VPT at the connection administrator level so that the node admits its VCCs under VPT-CAC. This configuration is essential for standard VPTs and virtual interfaces. Under VPT-CAC, the ATM interface admits the VPT using CAC but only if the interface has sufficient available bandwidth to accommodate the VPT ECR requirement. On VPT admission, the VPT admits its associated VCCs using VPT-CAC but only if the VPT has sufficient available bandwidth to accommodate the VCC ECR requirements.

Bandwidth pool allocation is identical to pool allocation for the ATM interface under the VPT connection administrator. For information on bandwidth pool allocation, see “Bandwidth pools and VPT-CAC” (page 47).

Under dynamic bandwidth management, when bandwidth changes dynamically over certain ATM links, the node applies connection bandwidth control (CBC). Unlike CAC, CBC defines how different ATM connections react to changes in bandwidth over an ATM link once the connections are admitted and operating. For details on CBC, see “Dynamic bandwidth management” (page 55).

AtmIf-CAC serving capacity

AtmIf-CAC handles a VPT that is configured for VPT-CAC like any other VPC connection point: it calculates the ECR based on the traffic parameters for its ATM service category.

The ECR calculation considers only the link bandwidth. With the other AtmIf-CAC parameters fixed, the computed ECR for a connection is always the same, regardless of the order of connection admission.

These characteristics of AtmIf-CAC allows higher link utilization while meeting requested QoS, and simplifies connection admission. With a larger serving capacity, the effect of statistical multiplexing of VCs on one physical link allows AtmIf-CAC to admit a VC with a smaller ECR. By using a smaller ECR, the node can admit more VCs. Moreover, with the constant serving capacity, the ECR is independent of the order of connection admission.

Overview of VPT-CAC

VPT-CAC determines admission of a VPT VCC. This determination is made when

- the PVC is configured
- the SVC requests call setup

Following successful admission of the VPT as part of the ATM interface, the node applies VPT-CAC to each VCC that requests admission as part of the VPT. As with AtmIf-CAC, the VPT-CAC first computes the ECR of the VCC, then admits the VCC if it passes the various controls established for the VPT.

The various admission controls provisioned for the VPT are similar to those provisioned by AtmIf-CAC. These controls include the following items:

- limits on the number of VCCs and the number of UBR connection points that the VPT supports

- the amount of bandwidth allotted to each of the five VPT bandwidth pools
- per-ATM service category controls for CBR, RT-VBR, and NRT-VBR (see the figure “Equivalent cell rate calculations under AtmIf-CAC” (page 21)).

VPT-CAC also includes CAC overbooking.

VPT-CAC serving capacity and pools

Serving capacity and pools for VPT-CAC has the following characteristics:

- VPT service capacity is equal to the PCR
- bandwidth pools are set up as a percentage of ECR
- VPT-CAC uses the VPT serving capacity for calculating the ECR of the VCC

Compare these characteristics with the general characteristics for AtmIf-CAC in the section “AtmIf-CAC serving capacity” (page 39)

AtmIf-CAC and VPT-CAC characteristics

The table “Comparison of AtmIf-CAC and VPT-CAC characteristics” (page 40) summarizes the characteristics for AtmIf-CAC and VPT-CAC. Most of the VPT-CAC attributes have the same default values as the AtmIf-CAC attributes.

The VPT-CAC parameters can be changed independently of the AtmIf-CAC parameters. Important aspects of the VPT-CAC are serving capacity, VPT bandwidth pools, QOS, and overbooking.

Table 7
Comparison of AtmIf-CAC and VPT-CAC characteristics

Parameter	AtmIf-CAC	VPT-CAC
Serving capacity	link rate	PCR_{VPT}
Pool capacity	% of link rate	% of ECR_{VPT}
(Sheet 1 of 3)		

Table 7 (continued)
Comparison of Atmlf-CAC and VPT-CAC characteristics

Parameter	Atmlf-CAC	VPT-CAC
maximum number of VCCs (<i>maxVccs</i>)	maximum number of connections that can be created under the given ATM interface, including VPT-VCCs	maximum number of connections that can be created under the given VPT
maximum number of UBR connections under the ATM interface (<i>maxVccs</i> under <i>Atmlf Ca Ubr</i>)	maximum number of UBR connections that can be created under the given ATM interface, including VPT-UBR VCCs	maximum number of UBR connections that can be created under the given VPT
number of used UBR connections (<i>VccUsage</i> , <i>VpcUsage</i> , and <i>VptUsage</i> under <i>Atmlf Ca Ubr</i>)	number of UBR connections under the given ATM interface	available UBR connections under the VPT
minimum cell rate (MCR)	MCR for the UBR connections under the given atm interface.	MCR for the UBR connections under the given VPT
number of permanent VCCs (<i>permanentVccs</i>)	number of permanent VCCs that are currently provisioned on this interface This number includes VCCs with <i>SrcPvc</i> subcomponents, and any VCCs that are associated with a VPT.	number of permanent VCCs that are currently associated with this VPT
number of switched VCCs (<i>switchedVccs</i>) - see Note	number of switched VCCs that are currently active on this interface, including switched VCCs associated with a VPT	number of switched Vccs that are currently associated with this VPT
number of troubled VCCs <i>troubledVccs</i> - see Note	number of troubled VCCs on the ATM interface. It does not include troubled VCCs associated with the VPT.	number of troubled VCCs that are associated with the VPT
connection pool usage (<i>connectionPoolUsage</i>)	number of enabled VCCs and VPCs under the ATM interface, including any VPTs	number of enabled VCCs associated with this VPT
(Sheet 2 of 3)		

Table 7 (continued)
Comparison of Atmlf-CAC and VPT-CAC characteristics

Parameter	Atmlf-CAC	VPT-CAC
maximum number of VPCs (<i>maxVpcs</i>)	maximum number of VPCs that can be configured on the ATM interface	not applicable
maximum number of VPTs (<i>maxVpts</i>)	maximum number of VPTs that can be configured on the ATM interface	not applicable
minimum number of multicast branches (<i>minMulticastBranches</i>)	guaranteed number of multicast branches that can be configured on the ATM interface	uses Atmlf-CAC limit
<i>maxMulticastBranches</i>	maximum number of configurable multicast branches that can be configured on the ATM interface	uses Atmlf-CAC limit
minimum value for VCI under VPI 0 (<i>minAutoSelectedVciForVpi Zero</i>)	minimum VCI value that signaling automatically allocate for a switched VCC under VPI 0	uses Atmlf-CAC limit
minimum value for VCI under VPIs other than 0 (<i>minAutoSelectedVciForNon ZeroVpi</i>)	minimum VCI value that signaling automatically allocates for a switched VCC with a non-zero VPI value	uses Atmlf-CAC limit
Note: These parameters are not specific to CAC.		
(Sheet 3 of 3)		

Chapter 2

Bandwidth pool management

You can partition the capacity on each port into bandwidth pools. Partitioning allows you to fulfill the different requirements of each service category. When you know the amount of bandwidth that a given traffic type needs, it is desirable to set aside the necessary port capacity. In this way, you protect that capacity from being used by other traffic classes.

This chapter provides information in the following sections:

- “General characteristics of bandwidth pools” (page 43)
- “Bandwidth pool over- and under-subscription” (page 48)
- “Bandwidth pool sharing” (page 52)
- “Bandwidth pool management for point-to-multipoint SVCs” (page 53)

General characteristics of bandwidth pools

Bandwidth pool management limits the maximum amount of bandwidth for connections under each service category. In connection admission control (CAC), bandwidth pool management ensures that the sum of the equivalent cell rates (ECR) of connections under a specific service category do not exceed the allocated pool capacity for that service category, adjusted by the configured overbooking factor. For example, the service provider can provide a small bandwidth pool for the CBR traffic as a way of limiting the amount of video calls in a network.

Bandwidth pools have the following general characteristics:

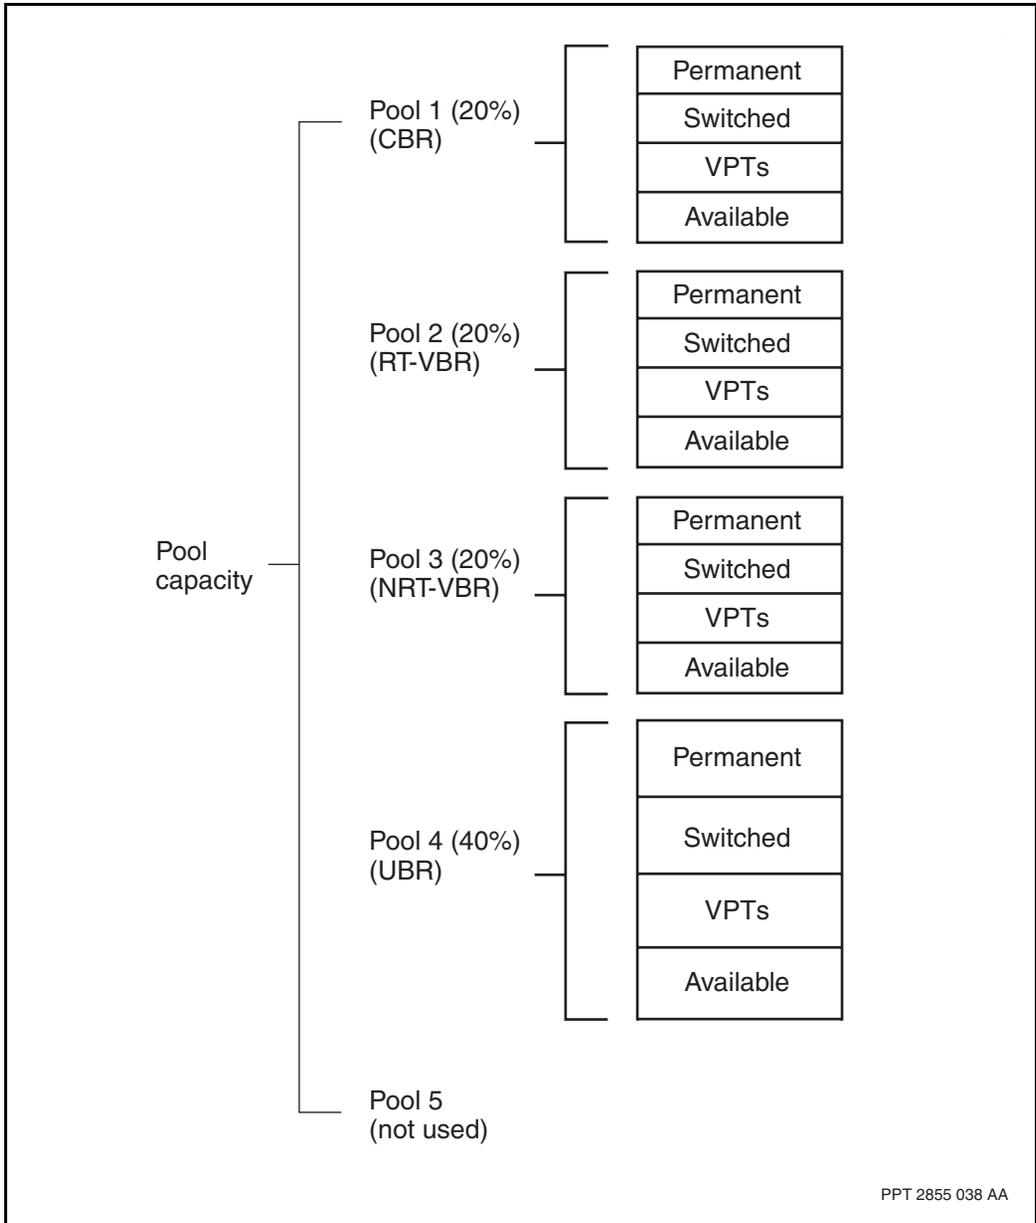
- there are five available pools, which allows mapping of each service category to a dedicated pool
- the fifth pool is available for future implementation of available bit rate (ABR)
- Nortel Networks Multiservice Switch nodes permit sharing of port capacity (and by association ATM interface capacity) between service categories
- Multiservice Switch nodes permit over-subscription for each bandwidth pool at 12,800%

Through configuration, you map each service category to a given pool and assign each pool a percentage of port capacity. A service category can be assigned to one pool only. The allocation of the pool capacity allows for the five service categories to flexibly share the link resource in a scheme that ranges from complete sharing to complete partitioning.

A sophisticated CAC mechanism is based on the calculation of the ECR for each connection request and the calculation of the ACR after the node admits a connection. The network needs bandwidth management at the node level between ATM service categories.

The figure “Example of bandwidth allocation: 100 percent allocation” (page 45) shows an example of how total pool capacity can be allocated across service categories. The figure also shows how within each pool a portion of capacity is dynamically allocated to permanent and switched connections, and to virtual path terminators (VPT).

Figure 6
Example of bandwidth allocation: 100 percent allocation



The pool allocation values are a percentage of the total pool capacity for the interface. In the example in the figure “Example of bandwidth allocation: 100 percent allocation” (page 45), the allocation percentages are arbitrary for this example only. The defaults values for bandwidth pool allocation are:

- pool 1: 100%
- pool 2: 0%
- pool 3: 0%
- pool 4: 0%
- pool 5: 0%

This approach allows you to set limits on the capacity that the node allocates to a given service category. For example, if the traffic mix consists of

- 40% constant bit rate (CBR)
- 30% real-time variable bit rate (RT-VBR)
- 30% non-real-time variable bit rate NRT-VBR

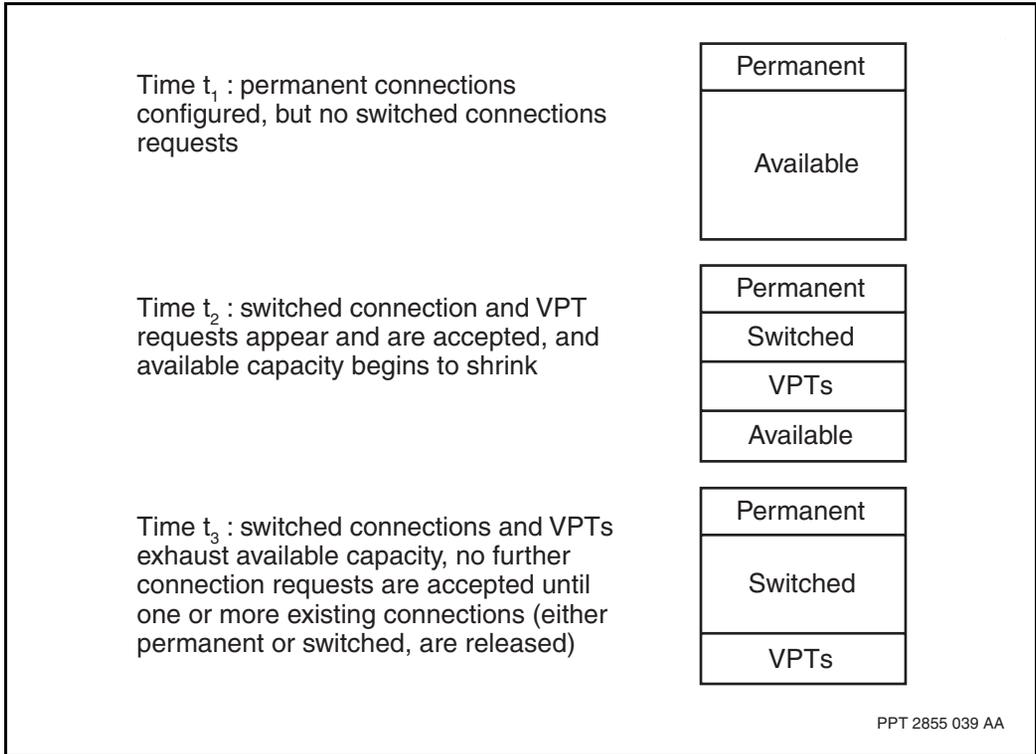
you can configure pool 1, 2, and 3 to have 40%, 30%, and 30% of the port capacity respectively, and do not allocate pools 4 and 5. In addition, you configure CBR, RT-VBR, and NRT-VBR to reserve bandwidth through CAC from Pool1, Pool2, and Pool3 respectively.

The capacity that you allocate to each service category accommodates permanent and switched connections. When the node receives a request for a new connection, CAC computes the equivalent cell rate (ECR) in real time and decides on the admittance of the connection. The node dynamically updates the available cell rate (ACR) for each pool to reflect the effects of the new connection.

For the duration of a connection, the node allocates pool capacity to that connection. The node can assign any remaining available pool capacity to new connections until all available capacity is exhausted. The figure “Example of progression of pool allocation for a single service category” (page 47) shows how the pool capacity allocated to switched connections and

VPTs fluctuates and the allocation for permanent connections remains stable. This example assumes that the service provider does not configure any new permanent connections.

Figure 7
Example of progression of pool allocation for a single service category

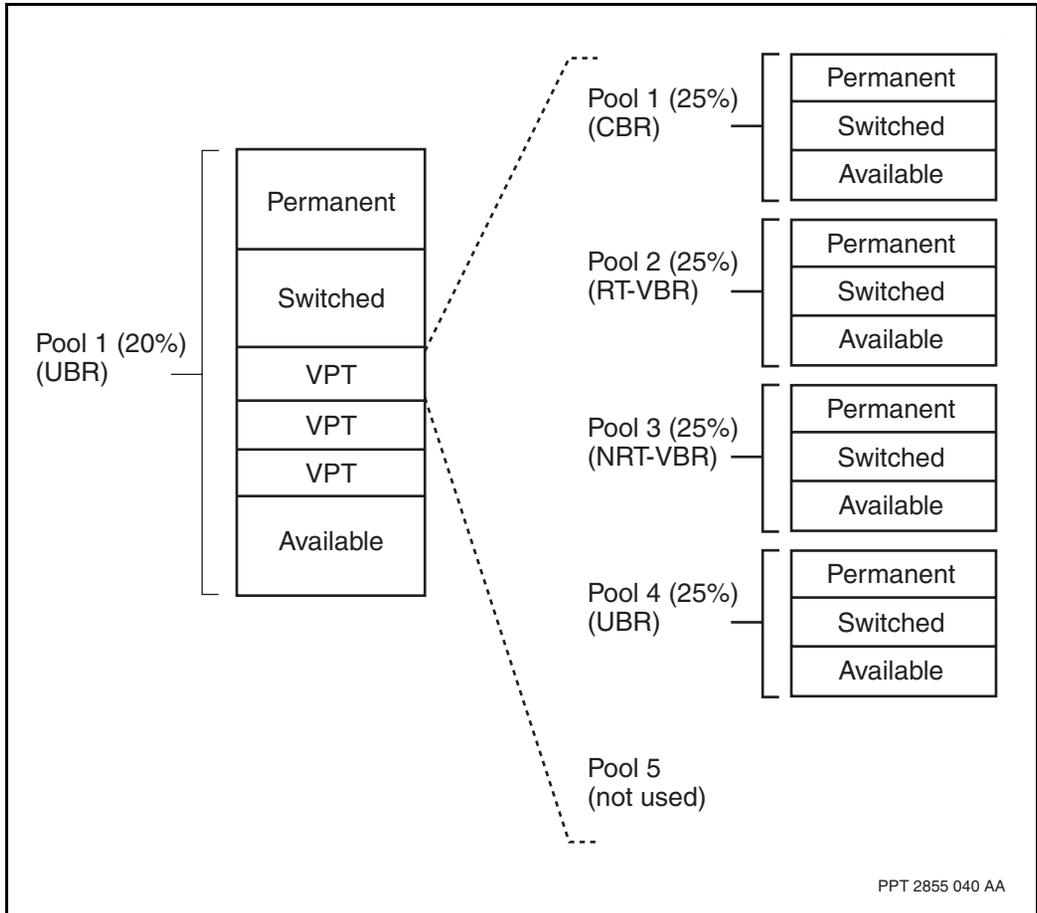


Bandwidth pools and VPT-CAC

Through configuration at the connection administrator level, you configure VPT bandwidth allocation using the same set of parameters that you use for the ATM interface. As the node receives call setup requests for VPTs, it allocates bandwidth on a per VPT basis.

The figure “Example of bandwidth allocation for VPTs” (page 48) shows how the node achieves bandwidth allocation for VPTs. Note how a single VPT can support VCCs from multiple service categories.

Figure 8
Example of bandwidth allocation for VPTs



Bandwidth pool over- and under-subscription

The percentages assigned to the five bandwidth pools do not have to add up to 100 per cent. This flexibility allows for over- and under-subscription of the port as shown in the figures “Example of bandwidth allocation: over-subscription” (page 50) and “Example of bandwidth allocation: under-subscription” (page 51). CAC uses this flexibility to increase or decrease port usage.

Nodes can use under-subscription (percentages totalling less than 100 per cent) to implicitly reserve bandwidth for UBR connections when such connections are not assigned a separate bandwidth pool. Nodes can use over-subscription (percentages totalling more than 100 per cent of port capacity) to maximize node and facility usage where traffic patterns and end-user requirements permit loss of lower priority traffic to make way for high-priority traffic.

Figure 9
Example of bandwidth allocation: over-subscription

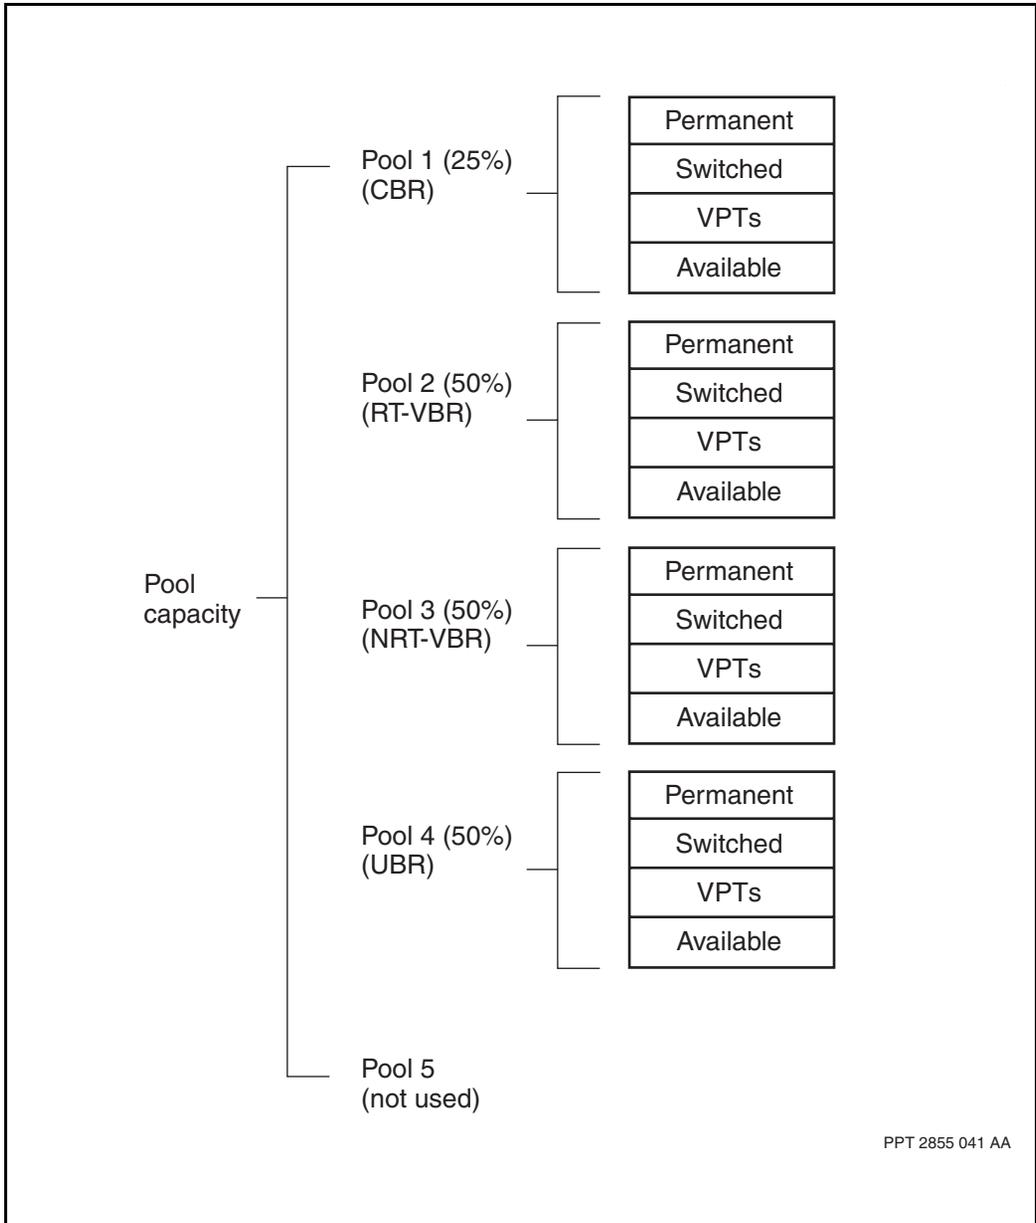
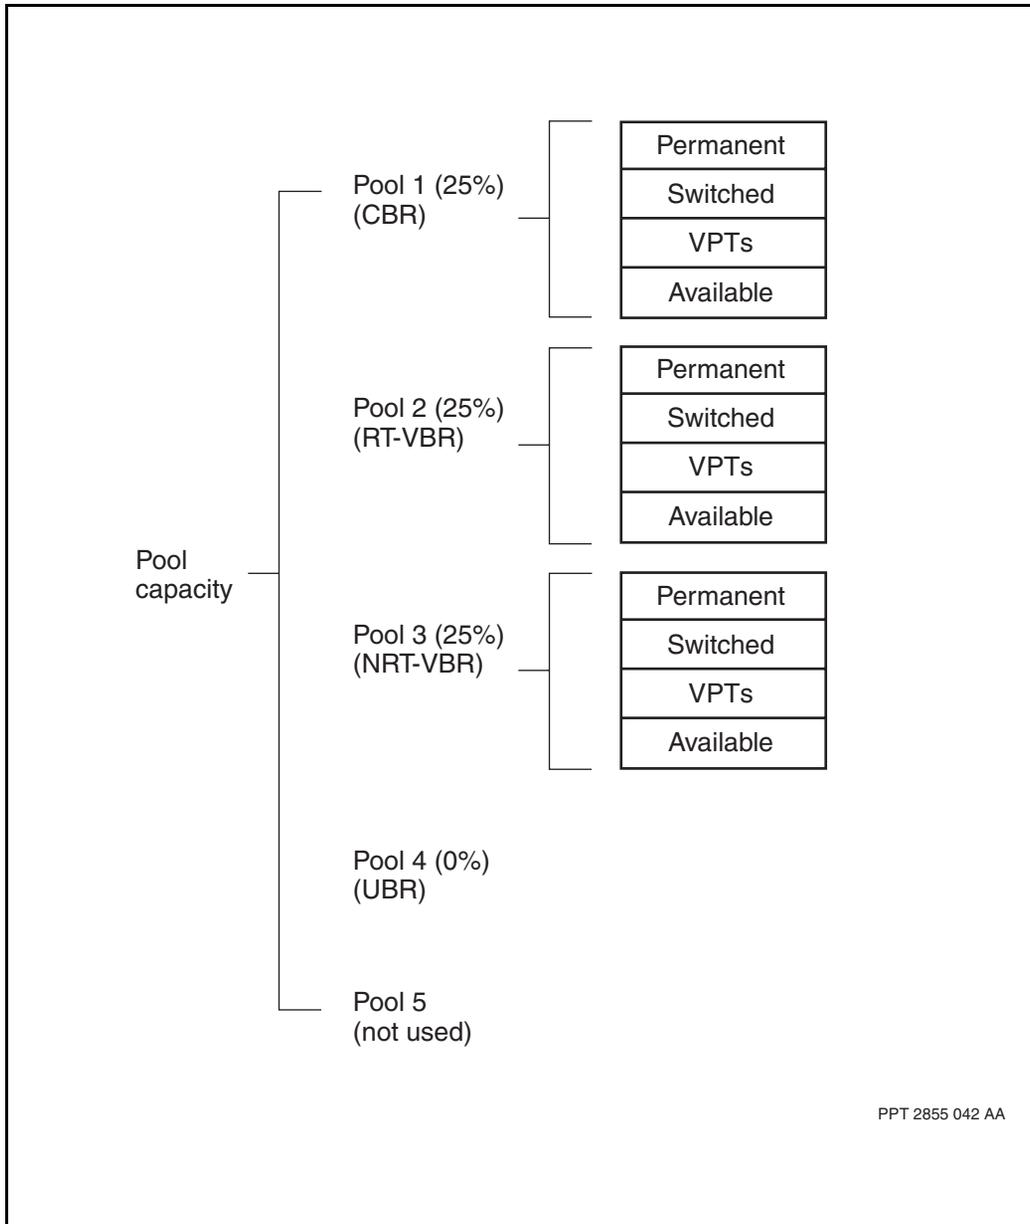


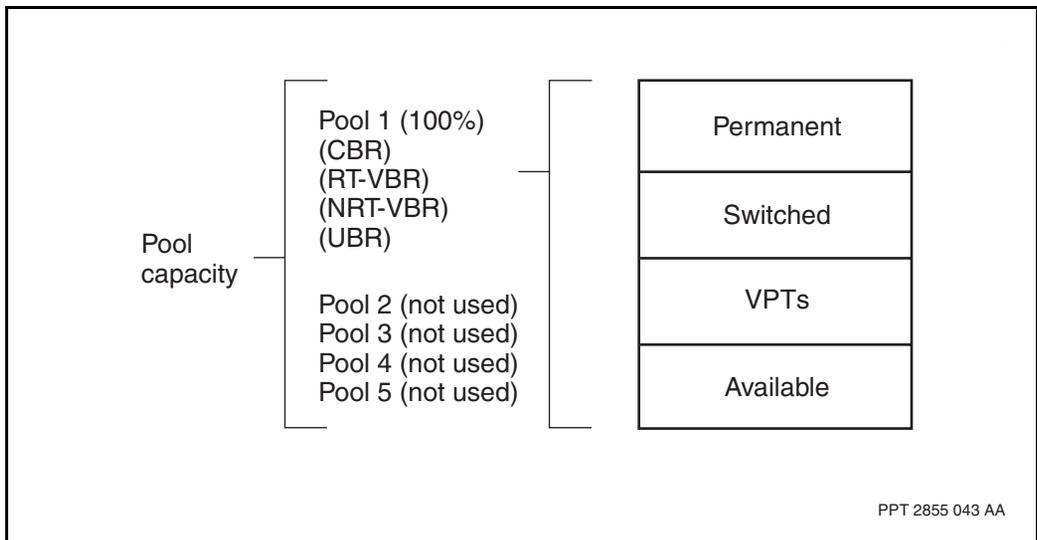
Figure 10
Example of bandwidth allocation: under-subscription



Bandwidth pool sharing

This bandwidth management strategy is flexible enough to permit full sharing of port capacity without pre-set partitions between ATM service categories. Configure this level of port sharing by assigning 100 percent (or more) of the capacity to one common bandwidth pool (pool1) from which all connections reserve their required ECR regardless of the ATM service category. See the figure “Example of bandwidth allocation: common pool” (page 52). In this configuration, the entire port capacity is available to any ATM service category on a first-come first-served basis. The node maintains one available cell rate counter to account for bandwidth allocation and de-allocation.

Figure 11
Example of bandwidth allocation: common pool



Even if no ports are active on a function processor or if links are out of service, the *poolAvailableBandwidth* attribute initializes to one port capacity. This initialization occurs to permit greater efficiency and allows the registration of permanent connections before the ports or links recover. This initialization also provides CAC and connection bandwidth control (CBC) with information about the pool requested bandwidth. This information speeds up the connection admission process. Registered permanent connections can then be enabled as soon as the port or link returns to service.

Bandwidth pool management for point-to-multipoint SVCs

Nortel Networks Multiservice Switch nodes point-to-multipoint SVCs have no specific traffic management requirements. The QOS and forward bandwidth of the point-to-multipoint leaf nodes are identical to the root node. The leaf nodes have zero return bandwidth.

For CAC, bandwidth reservation applies to transmission onto the link (transmit direction only). The transmit bandwidth reservation of the node at the far end of the link determines the receive direction (return bandwidth). Since the leaf nodes have zero return bandwidth, nodes do not require CAC on the receive port.

Multiservice Switch nodes reject add-party requests to leaf nodes that cannot satisfy the QOS requirement of the root node. You must carefully plan and configure point-to-multipoint SVCs so that the interface can fulfill QOS requirements for the add-party request, with sufficient bandwidth to handle all traffic.

Connection bandwidth control

Nortel Networks Multiservice Switch nodes ATM point-to-multipoint SVC feature extends the dynamic bandwidth capability of the inverse multiplexing over ATM (IMA) feature to point-to-multipoint SVCs. The feature does not modify the connection bandwidth control (CBC) algorithm. Instead, you configure the SVC point-to-multipoint holding priority under existing CBR, RT-VBR, -time variable bit rate), and NRT-VBR configurations. You use these attributes to establish the relative holding priorities between point-to-point SVCs and point-to-multipoint SVCs.

The range of values and behavior of the SVC point-to-multipoint holding priority is equivalent to the existing SVC holding priority under the CBR, RT-VBR, and NRT-VBR. When a port or link loses bandwidth, and there is insufficient bandwidth for existing connections, the CBC algorithm reduces bandwidth of existing connections on the basis of the following connection characteristics:

- elasticity
- holding priority
- connection type

- connection number

Ports or links can lose bandwidth because of resource failure, removal, recovery, or addition. By decreasing the SVC point-to-multipoint holding priority relative to the SVC holding priority, a node releases point-to-point connections ahead of point-to-multipoint connections for a given bandwidth pool.

Chapter 3

Dynamic bandwidth management

Dynamic bandwidth management permits Nortel Networks Multiservice Switch nodes to respond to the changing bandwidth requirements of ATM connections. This result is achieved through connection bandwidth control functions, which determine how different types of ATM connections react to changes in bandwidth over an ATM link.

This chapter provides information in the following sections:

- “Dynamic bandwidth and ATM connections” (page 56)
- “Benefits of dynamic bandwidth management” (page 57)
- “Dynamic bandwidth management and connection types” (page 59)
- “Connection bandwidth control” (page 61)
- “Examples of dynamic bandwidth management” (page 78)
- “Example: non-elastic connections responding to dynamic bandwidth” (page 78)
- “Example: elastic connections responding to dynamic bandwidth” (page 81)
- “Example: mixed connections responding to dynamic bandwidth” (page 85)
- “CBC and bandwidth pools” (page 89)
- “Dynamic bandwidth and data protection” (page 92)
- “Coordinating dynamic bandwidth between nodes” (page 93)

- “VPT CAC response to dynamic bandwidth changes” (page 93)

Dynamic bandwidth and ATM connections

The ATM connections that are part of an ATM interface are subject to conditions of changing bandwidth over certain types of ATM links. When bandwidth changes, there may not be enough bandwidth for all currently admitted connections. If no action is taken, congestion will occur at the ATM interface, and the target cell loss ratio (CLR) for each connection may no longer be met (performance may not meet expectations). Nortel Networks Multiservice Switch approach is to identify lower priority connections that may be released (terminated), or connections that may undergo bandwidth reduction. This permits a high degree of user control over the ATM connections subject to dynamic bandwidth. The result is that all currently active connections fit within available bandwidth, and the user can predict how the connections will react to bandwidth changes.

A mechanism is required to manage how ATM connections respond to changes in bandwidth. Connection admission control (CAC) manages which ATM connections are admitted over an ATM link. See “Connection admission control” (page 19). By contrast, connection bandwidth control (CBC) defines how different ATM connections react to changes in bandwidth over an ATM link once the connections are admitted and operating.

Currently, bandwidth over Multiservice Switch node ATM links can change only when using an inverse multiplexing for ATM (IMA) link group. Dynamic bandwidth can occur for a number of reasons:

- One or more physical links that are part of an IMA link group may go down.
- Physical links may be removed from an IMA link group as network traffic is routed elsewhere.
- Physical links may also resume operation in an IMA link group after failing.
- Physical links may be added to an IMA link group as traffic requirements grow.

In the case of bandwidth loss, the ATM interface must re-evaluate which connections can continue at the present data rate, which ones have to be released, and which ones can stay active at a reduced rate. In the case of bandwidth gain, the ATM interface must re-evaluate which connections can be brought back to full bandwidth, and which ones can be readmitted.

The connection bandwidth control (CBC) algorithm provides the mechanism by which the ATM connections react to ongoing changes in bandwidth.

For the purposes of dynamic bandwidth, all ATM connections fall into one of two mutually exclusive categories: elastic or non-elastic connections. An elastic connection can respond to changes in bandwidth by decreasing or increasing its data rate, whereas a non-elastic connection cannot.

Benefits of dynamic bandwidth management

Dynamic bandwidth management technology on Nortel Networks Multiservice Switch nodes, offers a number of benefits, including the following:

- connection maintenance

Many connections transmitting ATM user data can continue to operate, even when one or more physical links that are part of a link group go down.

- predictability

You can predict how ATM connections will react to changes in bandwidth and, specifically, which ATM connections would be affected by a loss of bandwidth over a link. Also, QOS guarantees are maintained for different types of ATM traffic based on the orderly and predictable reallocation of available bandwidth.

- user control

When you configure your ATM connections, you can control which data is protected in the event of a reduction in available bandwidth.

For details on protecting certain types of data over links subject to dynamic bandwidth, see “Dynamic bandwidth and data protection” (page 92).

Related features

The table “IMA and related Multiservice Switch features” (page 58) shows the features that closely interwork with dynamic bandwidth management on Nortel Networks Multiservice Switch ATM nodes. These features are described elsewhere in the documentation suite.

Table 8
IMA and related Multiservice Switch features

Feature	Document in which feature is described	Description of feature
Inverse multiplexing for ATM	NN10600-730 <i>Nortel Networks Multiservice Switch 7400/15000/20000 Inverse Multiplexing for ATM Operations</i>	Describes IMA feature, benefits, functionality, configuration, monitoring, troubleshooting.
Core services connection bandwidth control (CBC)	NN10600-700 <i>Nortel Networks Multiservice Switch 7400/15000/20000 ATM Technology Fundamentals</i>	Describes how the CBC defines the response of different ATM connections to a change in bandwidth over an IMA link group.
Dynamic trunk speed changes	NN10600-420 <i>Nortel Networks Multiservice Switch 7400/15000/20000 Operations: Trunking</i>	Describes ATM direct logical trunk connections; these can be elastic connection types that are able to maintain operation when a bandwidth change occurs on an IMA link group.
Point-to-multipoint connections	NN10600-700 <i>Nortel Networks Multiservice Switch 7400/15000/20000 ATM Technology Fundamentals</i>	Describes CBC requirements for point-to-multipoint connections

The features shown in the table “IMA and related Multiservice Switch features” (page 58) interact with each other in the following ways:

- 1 The IMA feature supports fluctuations in aggregate bandwidth over an ATM link group.
- 2 Connection bandwidth control defines how all types of ATM connections, including direct ATM logical trunks, respond to changes in dynamic bandwidth in different scenarios and combinations.
- 3 Dynamic trunk speed changes enable a direct ATM logical trunk to adapt to bandwidth changes over an IMA link and remain operational.

- 4 Priority of point-to-point connections relative to point-to-multipoint connections.

Dynamic bandwidth management and connection types

For the purpose of dynamic bandwidth, there are two types of ATM connections:

- non-elastic

These are connections that are either on or off (that is, up and running or released). They are released whenever there is insufficient bandwidth in a pool to support the requirements of the connection. They can also be reactivated when bandwidth is regained.

These connections have an associated holding priority that can be configured explicitly for each PVC, and configured as a general default for SVCs (based on ATM service category).

- elastic

These are connections that can lose and regain a portion of their bandwidth, in proportion to what is available to a pool, when a change occurs to the bandwidth available to support the connections at their current data rate.

Each ATM connection falls into one of these two categories. Currently, the only type of connection that may be bandwidth elastic is an ATM logical trunk over a permanent single hop PVC. The holding priority associated with non-elastic connections is important in specifying the importance of a connection relative to other non-elastic connections. For all elastic connections, the *holdingPriority* operational attribute has the value of “not applicable”.

Non-elastic connections

All ATM connection types other than direct ATM logical trunks fall into the non-elastic category. The non-elastic ATM connections currently supported on Nortel Networks Multiservice Switch nodes are

- ATM bearer service (ABS)
- ATM logical trunks (AAL5)

- AAL1 (circuit emulation)
- ATM multi-protocol encapsulation (MPE)
- frame relay over ATM (FR-ATM)
- integrated local management interface (ILMI) and interim inter-switch signaling protocol (IISP) applications
- all switched connections
- loop component (ATM VCC/VPC)

Holding priority values indicate the relative importance of different ATM connections. When a reduction in bandwidth occurs, non-elastic connections are released in the following order:

- 1 holding priority 4
- 2 holding priority 3
- 3 holding priority 2
- 4 holding priority 1
- 5 holding priority 0

Elastic connections

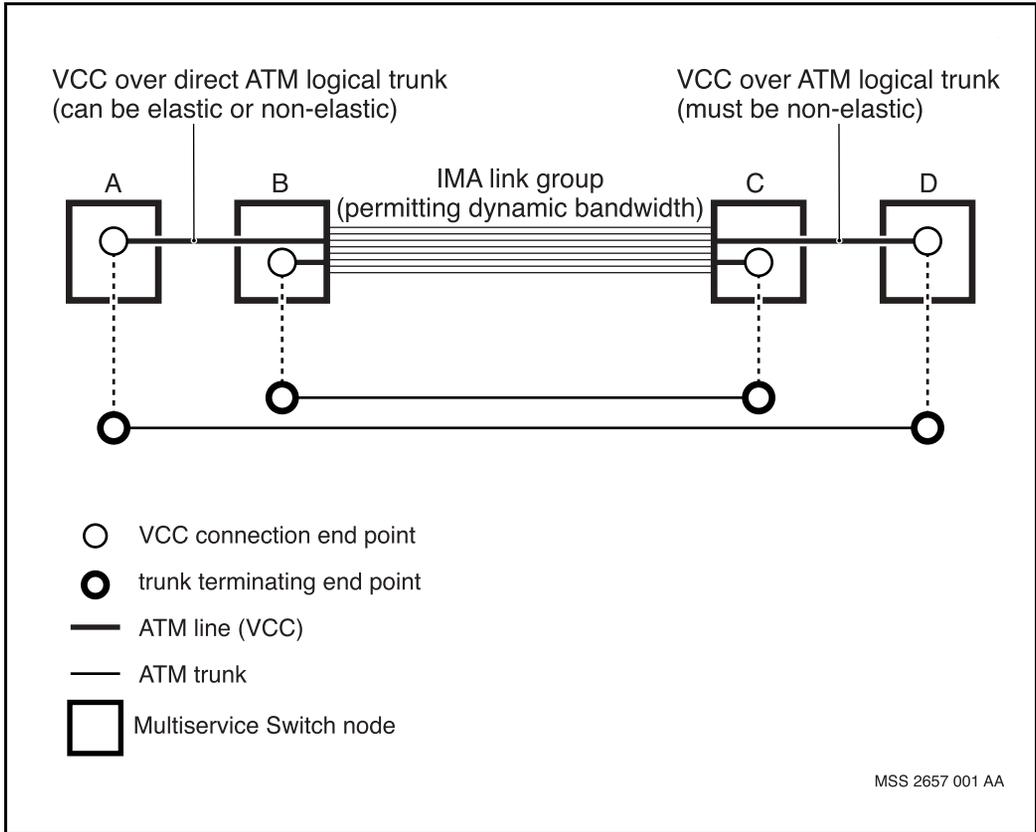
Currently, only direct ATM logical trunks (using AAL5) composed of single-hop permanent virtual circuits (PVC) can be configured as bandwidth elastic. The figure “ATM logical trunks as elastic and non-elastic connections” (page 61) shows an example of different types of ATM trunks using bandwidth over an IMA link group (Nortel Networks Multiservice Switch ATM node link subject to variations in bandwidth). In the figure, an IMA link composed of eight DS1/E1 physical links is configured between node B and node C. A bandwidth elastic direct logical trunk VCC is configured between nodes. By contrast, the ATM logical trunk VCC between node A and node D can only be a non-elastic connection type.

Note 1: Direct ATM logical trunks (on single-hop PVCs) can also be configured as non-elastic connections. However, ATM logical trunks over multi-hop VCCs must always be non-elastic connections.

Note 2: Bandwidth-elastic connection types can only use bandwidth over an IMA link group.

For more details on this type of VCC, see NN10600-420 *Nortel Networks Multiservice Switch 7400/15000/20000 Operations: Trunking*.

Figure 12
ATM logical trunks as elastic and non-elastic connections



Connection bandwidth control

Connection bandwidth control (CBC) defines how ATM connections react to dynamic bandwidth available over an ATM link. The inverse multiplexing feature can create scenarios whereby ATM connections have to respond to changing bandwidth. When configuring the IMA feature, an ATM interface (*AtmIf* component) is linked to an IMA link group so that the connections under that ATM interface are transmitted over the physical links of the IMA

group. For details on configuring the links between an ATM interface and an IMA link group, see NN10600-730 *Nortel Networks Multiservice Switch 7400/15000/20000 Inverse Multiplexing for ATM Operations*.

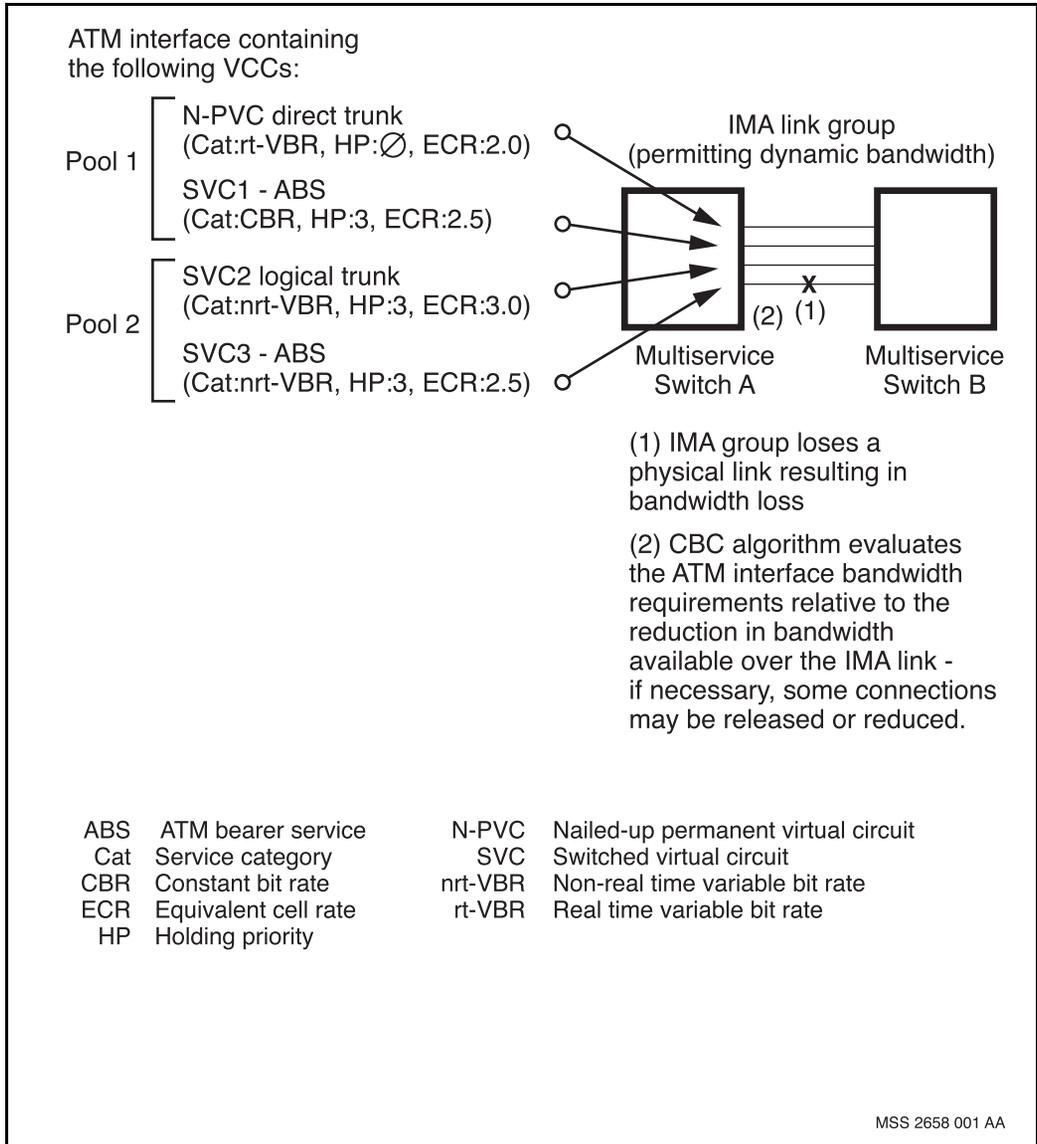
A collection of DS1/E1 links using the IMA process incorporate a high degree of robustness and reliability such that physical links can be removed or added without tearing down the link group or the ATM interface being served by the link. The failure of one physical link has no effect on the remaining physical links forming a link group, except that overall throughput is reduced over the link.

When an ATM interface receives an indication from the IMA group that the group bandwidth has changed, Nortel Networks Multiservice Switch nodes invoke the CBC algorithm to preserve the existing traffic contracts. Some non-elastic connections may be released (cleared) according to a pre-defined holding priority. Elastic connections can compensate for changes in the IMA group bandwidth and remain operational at a reduced data rate.

The figure “Example of dynamic bandwidth management using CBC algorithm” (page 63) shows an example of four connections that are part of an ATM interface using an IMA link group (composed of eight physical links) between nodes A and B. Dynamic bandwidth management is called upon in the following sequence:

- 1 IMA link group loses one of its physical links, resulting in a loss of bandwidth over the group.
- 2 CBC algorithm evaluates the ATM interface bandwidth requirements against the bandwidth available over the IMA link and makes any necessary adjustments by dropping non-elastic connections and, if necessary, reducing the bandwidth available to any elastic connections.

Figure 13
Example of dynamic bandwidth management using CBC algorithm



Objectives of dynamic bandwidth management

The primary objective of dynamic bandwidth management using the CBC algorithm is to ensure that traffic contracts are met at all times, even when IMA link group bandwidth fluctuates resulting in a congested state. If the traffic contract requirements are not adjusted, the congestion caused by the bandwidth change may have undesirable effects. By making the outcome of the bandwidth change predictable, users can choose the priority by which connections must be maintained in the event of a partial link failure.

When bandwidth is lost, the objective of the algorithm is to reduce bandwidth requirements of existing connections until a stable state is reached (that is, available bandwidth is greater than, or equal to, requested bandwidth).

A reduction in bandwidth required by the ATM interface is achieved by two integrated mechanisms:

- releasing non-elastic connections (such as ABS)
- decreasing the bandwidth allocated to elastic connections (direct logical trunks)

When bandwidth is recovered, the CBC algorithm is triggered to distribute the additional bandwidth to connections that require it most.

Application of the CBC algorithm

The connection bandwidth control (CBC) algorithm is triggered whenever the bandwidth changes over an IMA link group. A change in bandwidth may be caused by:

- failure of a physical link in an IMA group
- removal of a physical link from an IMA group
- recovery of a physical link in an IMA group
- addition of a physical link to an IMA group

The scenarios in this section describe different events that may occur on an IMA link group causing the CBC algorithm to be invoked.

- All physical links were up, but some have gone down.

If available bandwidth is not sufficient to meet requested bandwidth, the CBC algorithm is invoked to determine the response of the ATM interface. Otherwise, there is no change to the status of ATM connections.

If available bandwidth drops below the required bandwidth, some connections may be released (cleared) and some connections may have their bandwidth reduced. The CBC algorithm is activated to determine which connections are affected.

Unspecified bit rate (UBR) connections may be affected even if the available bandwidth is greater than the required bandwidth after the loss of a physical link. This is because UBR connections use any excess bandwidth available on a link. If the excess bandwidth is reduced, UBR connections may be affected.

- Some physical links were down, but all have now recovered.

If available bandwidth is more than the requested bandwidth, all elastic connections are adjusted to their full requested bandwidth, and then any non-elastic connections can be readmitted. If available bandwidth is not enough, the elastic connections are upgraded first, until the available bandwidth is used.

- Some physical links were down, and now more have come up or gone down.

The CBC algorithm is activated to determine the new set of reduced and released connections. Non-elastic connections are not admitted or readmitted as long as there are any elastic connections in reduced bandwidth condition.

- Some connections were reduced and now some connections have cleared.

Any bandwidth available due to the release of connections is distributed among the remaining connections.

- A new connection is added.

Non-elastic connections are only admitted at full requested bandwidth. Elastic trunk connections, however, are admitted even at reduced bandwidth. If the available bandwidth in the pool is less than the requested bandwidth, the connection is admitted with whatever bandwidth is available. If there are already some elastic connections admitted, all elastic connections are reduced by the same amount to admit the new connection.

CBC and point-to-multipoint connections

CBC also applies to point-to-multipoint SVCs. The CBC algorithm functions in the same way for both point-to-point and point-to-multipoint SVCs.

However, you can set a holding priority for point-to-multipoint connections that is different from the holding priority for point-to-point connections. You define values for the constant bit rate (CBR), real-time variable bit rate (RT-VBR), and non-real-time variable bit-rate (NRT-VBR) service categories. The range of values and behavior of the holding priority for point-to-multipoint connections is equivalent to the behavior for the holding priority for point-to-point connections.

These attributes permit you to set the priorities for point-to-point and point-to-multipoint connections relative to each other. Resolution of bandwidth allocation is the same for implementations with one or both types of SVCs.

Rules governing CBC

The rules shown in this section are the building blocks of the CBC algorithm. These rules are provided as a reference source.

Rule 1: applying change in the IMA link group capacity

A change (decrease or increase) in the IMA link group capacity is proportionally applied to each bandwidth pool based on its configured percentage of the ATM interface bandwidth. The CBC algorithm is then applied within each pool.

Bandwidth pools are designed to reserve bandwidth for connections within a given ATM service category as a percentage of the link bandwidth. The use of bandwidth pools guarantees service categories such as CBR, RT-VBR, and NRT-VBR even access to ATM links based on ECR. For more information, see “Bandwidth pool management” (page 43).

For an example of how bandwidth pools affect the outcome of the CBC algorithm, see “CBC and bandwidth pools” (page 89).

Rule 2: specifying the importance of non-elastic connections

The importance of non-elastic connections is specified through a configured holding priority.

The order in which these connections are released within a bandwidth pool is based on their holding priority (HP). Connections with a HP 4 are the first to be released and those with a HP 0 are the first to be readmitted.

A distinct holding priority can be set for each configured connection. The holding priority for switched connections is configured as a default value for each ATM service category.

Rule 3: releasing non-elastic connections

Even non-elastic connections with a holding priority of 0 may be released.

Normally, holding priority 0 connections are the last connections to be released. However, due to the timing of bandwidth changes and connection activation, even a holding priority 0 connection can be released. This can also happen when there is a mix of elastic and non-elastic connections.

To provide the best chance that a holding priority 0 connection is not released, ensure that the sum of the bandwidth requirements for all holding priority 0 connections and all bandwidth-elastic connections fits within the bandwidth of a single link (E1Atm or DS1Atm depending on the function processor type).

Rule 4: classifying connections

All connections are classified as either bandwidth-elastic or bandwidth non-elastic.

The application associated with a connection signals whether or not the connection is bandwidth elastic based on the value of an operational attribute under the VCC.

A direct logical trunk (a single-hop PVC with end-points terminating on both sides of an IMA interface) can accommodate changes in bandwidth over an ATM link. The trunk application signals that it is prepared to use an elastic connection. The VCC then informs the trunk application of any bandwidth decrease or increase over a link group after the trunk signals that it is bandwidth-elastic.

Rule 5: releasing non-elastic switched before non-elastic configured connections

Non-elastic switched connections are released before non-elastic configured connections.

When a switched connection fails, it can be re-established on a different path if a path with sufficient bandwidth exists. By contrast, when a configured connection fails, it must be manually rerouted. Because switched connections incorporate this added dimension, they are released first within a given holding priority.

Rule 6: releasing non-elastic connections in order of VPI.VCI

After non-elastic connections are sorted based on their holding priority and whether they are dynamic or configured, connections are released based on descending VPI.VCI value. Connections are readmitted based on increasing VPI.VCI value.

VPI.VCI numbers are used because they provide an orderly, predictable method of changing the status of connections. These values also are the same on both sides of an ATM interface. Therefore, the same set of connections are chosen for release or bandwidth increase/decrease on both ends of the IMA interface. In addition, if the configured connections are assigned based on VPI.VCI values, then a LIFO (FIFO) order of release or bandwidth change is achieved.

Rule 7: UBR connections

Unspecified bit rate (UBR) connections are not reduced or released.

The UBR ATM service category is by definition a best effort service and is not allocated any link bandwidth. In this case, the equivalent cell rate (ECR) for the connection is 0. Therefore, these connections are not reserved any bandwidth by the connection admission control (CAC) algorithm and are admitted regardless of bandwidth availability on the link. The multi-priority system (MPS) in Nortel Networks Multiservice Switch ATM nodes ensures that UBR traffic is the first to be discarded under congestion, minimizing its impact on other traffic.

When MDCR is greater than zero for UBR with MDCR, this service category will not use best effort and will allocate bandwidth with an MDCR value for the CAC algorithm.

Rule 8: distributing bandwidth loss in a pool

Rule 8: bandwidth loss is distributed between non-elastic and elastic connections in a pool.

The amount of the remaining bandwidth allocated to non-elastic connections is proportional to the amount of bandwidth used by non-elastic connections before link failure. The same is the case for elastic connections. The total bandwidth usage is defined as the sum of the ECR for each non-elastic connections, plus the sum of the current cell rate (CCR) for each elastic connection.

Note: When the connection is admitted, CCR is equal to ECR. Also, for non-elastic connections, CCR is always equal to ECR.

Since non-elastic connections operate at either 100 percent or are released completely, the total amount of bandwidth dropped over a link as a result of releasing non-elastic connections often exceeds the amount required to be released when link capacity is reduced.

Rule 9: releasing non-elastic connections before elastic connections

Non-elastic connections are released first, before bandwidth-elastic connections are reduced.

If a physical link fails and bandwidth must be recovered by releasing some non-elastic connections, those connections are released first, providing immediate relief of congestion. When all connections have been released in proportion to the amount of bandwidth that must be given up by non-elastic connections, and additional bandwidth must be relinquished, the CBC algorithm investigates reducing elastic connections.

More bandwidth is made available when relatively large non-elastic connections (those consuming a large amount of bandwidth) are released. In such cases, a higher proportion of bandwidth loss is absorbed by the release of non-elastic connections. This tends to decrease the proportion of bandwidth loss incurred by elastic connections.

Rule 10: reducing bandwidth for all elastic connections

When elastic connections are forced to lose bandwidth, all elastic connections are reduced by the same percentage within a bandwidth pool.

When there is a bandwidth deficit within a pool, the *poolAvailableBandwidth* attribute appears as a negative number. This indicates that an elastic connection must be brought up to full bandwidth by the amount shown before any other non-elastic connections can be admitted. Elastic connections are admitted even at reduced bandwidth. The proportional distribution of bandwidth among elastic connections is designed to be fair and predictable, and promotes the attainment of a stable state.

Rule 11: increasing bandwidth for all elastic connections

Whenever any bandwidth becomes available, all elastic connections in a bandwidth pool are increased by the same percentage.

Any gain or loss of bandwidth is assigned uniformly to all bandwidth-elastic connections running in decreased mode within the same bandwidth pool.

When the bandwidth changes are due to the failure or reinstatement of a physical link, the bandwidth is first distributed among the bandwidth pools, and is then uniformly distributed to all elastic connections within each pool.

When the bandwidth changes are due to the release of a connection (making available additional bandwidth), the connection returns its allocated bandwidth only to the bandwidth pool to which it belongs. This event triggers the CBC algorithm to distribute the gained bandwidth equally to all elastic connections within the pool.

For UBR with MDCR, an increase in MDCR will cause the reallocation of bandwidth and affect the CBC algorithm.

Rule 12: admitting or re-admitting connections

When the newly available bandwidth has been distributed to all elastic connections, and they have returned to their full requested bandwidth (that is, CCR equals ECR), any extra bandwidth is then available to admit or re-admit connections.

First, any new elastic connections are admitted, followed by any new or released non-elastic connections. Non-elastic connections are admitted according to their holding priority, whether the connection is configured or dynamic, and increasing VPI and VCI values. Connection admission is also done according to a first-fit algorithm. This allows a small holding priority 4 connection that fits the bandwidth available to be admitted ahead of a larger holding priority 0 connection that does not fit.

Therefore, the priority for newly available bandwidth is as follows:

- Any reduced elastic connections are increased.
- Any new elastic connections are admitted.
- Any non-elastic connections are admitted or readmitted.

For UBR with MDCR, when the bandwidth is not exhausted, UBR with MDCR can be admitted. When the MDCR value changes, it causes readmitting of the UBR with MDCR connection by using the new MDCR value with the CBC algorithm. When the UBR with MDCR connection is released, it causes some connections in the *poolWaitAdmitConnections* attribute to decrease and the *UBRpoolAdmittedConnections* attribute to increase until the bandwidth is exhausted again.

Rule 13: recording bandwidth changes of less than 1%

If the bandwidth change for an elastic connection is less than one percent of the connection's total bandwidth, the change appears in the operational attributes for the VCC but is not passed on to the application running the VCC.

To limit unnecessary overhead involved in notifying the application of bandwidth changes, any change of less than one percent may be held at the VCC. If the cumulative amount of several changes exceeds one percent, then the application is informed. Therefore, the bandwidth visible to an application such as Trunks may differ by up to one percent from the current bandwidth of the VCC.

Rule 14: reporting bandwidth to a trunk

The bandwidth reported to a trunk is always less than or equal to the PCR. The ECR for a connection is the amount of bandwidth reserved to guarantee that the application can send traffic at PCR.

For CBR traffic, ECR is typically higher than PCR since ECR accounts for CDVT. However, the node can guarantee only PCR (and not the higher ECR) for a trunk application. If the trunk is elastic, with a bandwidth reduction that is less than the difference between the ECR and PCR, the trunk shows a measured speed that is equal to the VCC PCR. The trunk measured speed reflects either the value for PCR or the value for ECR minus bandwidth reduction, whichever is less.

How CBC handles bandwidth loss

When a reduction in bandwidth available to an ATM interface occurs due to the failure or removal of one or more physical links in an IMA group, the CBC algorithm reduces the bandwidth requirements of the ATM interface based on the following criteria:

- bandwidth pool, depending on ATM service category
- VCC type (bandwidth-elastic or non-elastic)
- holding priority
- connection type (dynamic or configured)
- VPI.VCI values

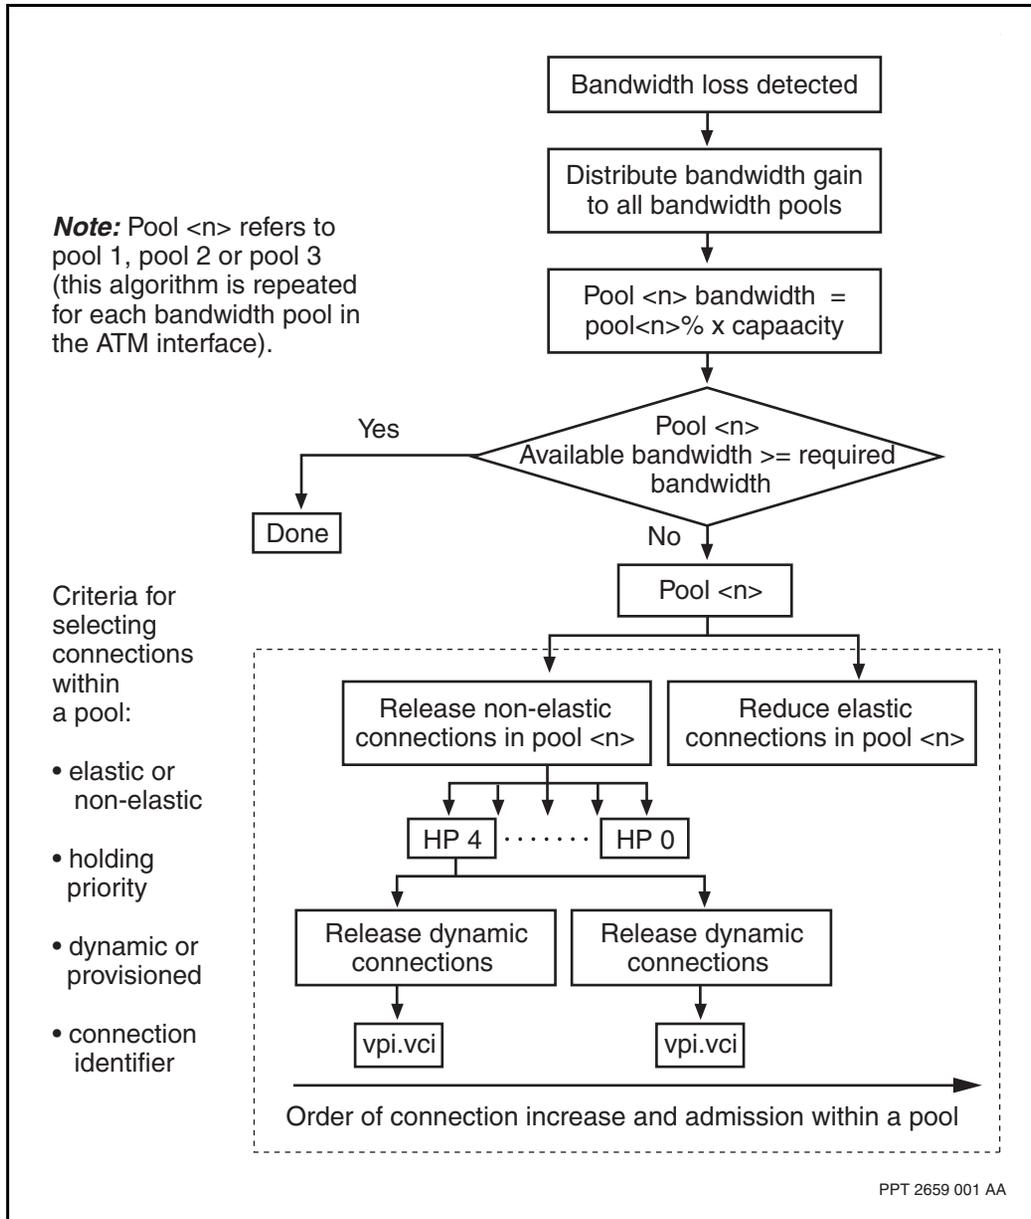
The figure “How the CBC algorithm handles bandwidth loss” (page 74) shows how the CBC algorithm handles bandwidth loss. In general, the algorithm uses the following process:

- 1 Any bandwidth loss is distributed proportionally to each pool. For each pool that has a bandwidth deficiency (available bandwidth is less than required bandwidth), the algorithm proceeds to reduce the requirement of the ATM interface such that the connections in the pool can be accommodated.
- 2 Within each pool, non-elastic and elastic connection types are assigned the bandwidth loss in proportion to the amount of bandwidth each type of connection currently has in the pool.
- 3 Non-elastic connections are released first to reduce bandwidth requirement, according to their assigned portion of the loss. They are released by increasing holding priority.
- 4 Within each holding priority, switched connections (SVCs) are released first, and then configured connections (PVCs).
- 5 Within each connection type, SVCs or PVCs are released in descending order of VPI and VCI values.
- 6 Steps 4 and 5 are repeated for each holding priority (in ascending order) until a stable state is attained; that is, required bandwidth is less than or equal to available bandwidth.
- 7 Finally, if the removal of the non-elastic connections’ portion of the loss does not satisfy the requirement for bandwidth reduction, the capacity allocated to elastic connections within a pool is reduced. To achieve a stable state (available bandwidth is greater than or equal to required bandwidth), all currently active elastic connections are reduced by the same percentage within a bandwidth pool.

All released connections are terminated; a switched connection (SVC) is torn down and must be re-established. If the connection attempts to be readmitted over the same IMA link group, bandwidth will not be available. In the case of configured connections (PVC), the connection is broken at the ATM interface in a manner similar to what happens when the interface goes down.

Connections are admitted when new bandwidth becomes available.

Figure 14
How the CBC algorithm handles bandwidth loss



How CBC handles bandwidth gain

Additional bandwidth may be made available to an ATM interface due to:

- the recovery of a physical link in an IMA group
- the addition of a physical link to an IMA group
- the release of a connection

The CBC algorithm distributes additional bandwidth based on the following criteria:

- how the bandwidth is made available - link recovery or releasing a connection
- bandwidth pool - depending on ATM Service Category
- VCC type - bandwidth elastic or non-elastic
- holding priority
- VPI.VCI values

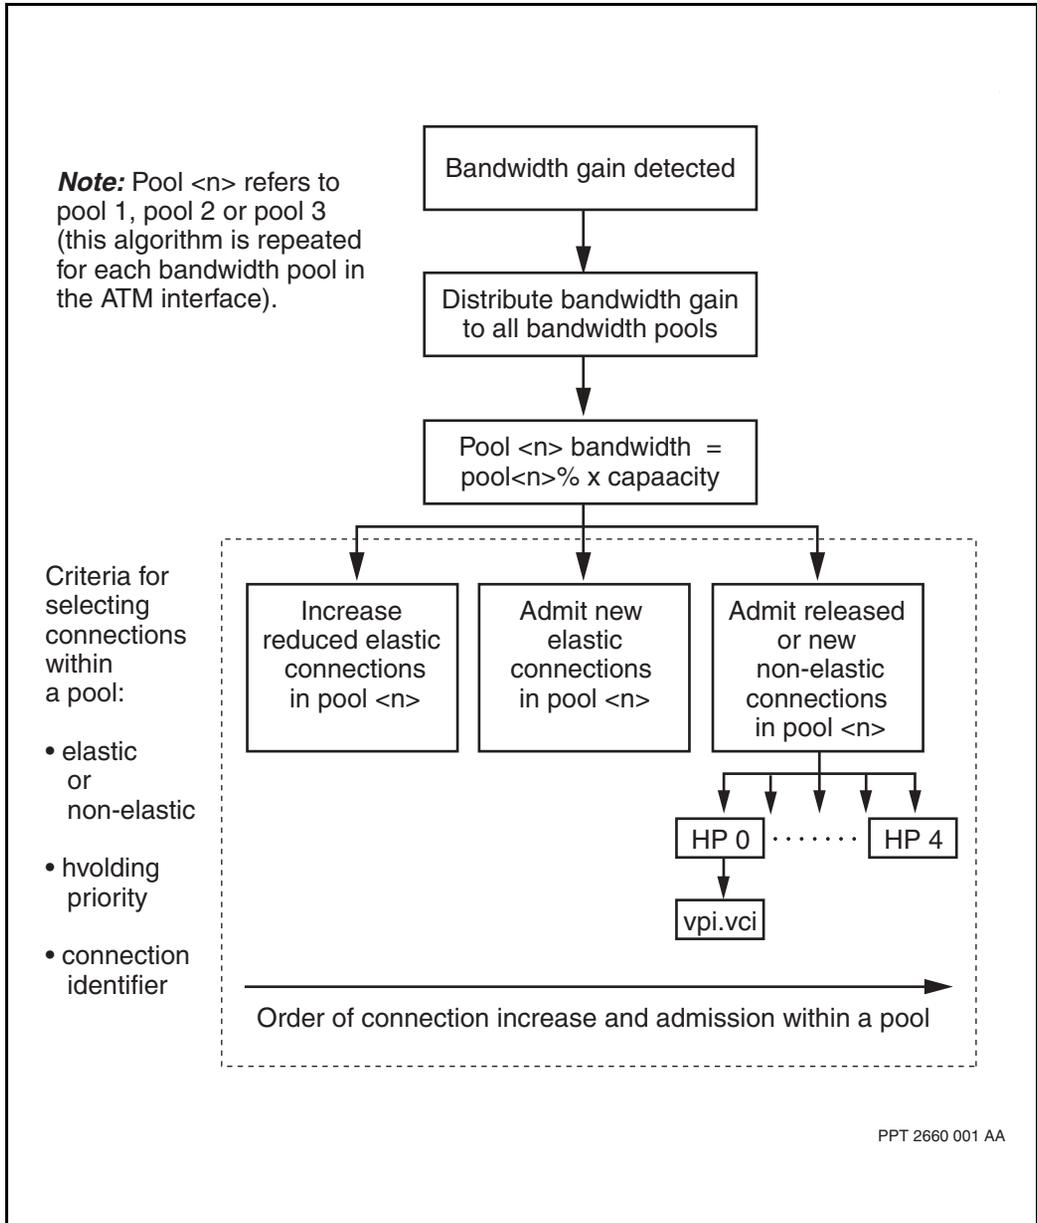
Note: For UBR with MDCR, if you change the MDCR value to a value that is less than the value that is already in place, the bandwidth will be obtained from the CBC.

The figure “How the CBC algorithm handles bandwidth gain” (page 77) shows how the CBC algorithm handles bandwidth gain. In general, the algorithm uses the following process:

- 1 Any bandwidth gain due to the recovery of a physical link is distributed proportionally to each pool. However, any bandwidth gain due to the release of a connection is distributed only within the pool to which the released connection belongs.
- 2 Within each pool, all elastic connections in a reduced state receive added bandwidth proportional to their combined equivalent cell rate (ECR). Elastic connections are brought up to their full bandwidth allotment as soon as additional bandwidth is available. Any newly added elastic connections are admitted at reduced bandwidth as soon as they are enabled.

- 3 When all elastic connections in a pool are operating at full bandwidth, any released or new non-elastic connections can be admitted to the pool. Non-elastic connections are admitted by descending order of holding priority (those with a holding priority of 0 are admitted first, and those with a holding priority of 4 are admitted last).
- 4 Within each connection type, SVCs or PVCs are admitted in ascending order of VPI and VCI values.
- 5 Step 4 is repeated for each holding priority (in descending order) until any available bandwidth is distributed.

Figure 15
How the CBC algorithm handles bandwidth gain



Examples of dynamic bandwidth management

This section provides examples of how ATM connections in a bandwidth pool respond to dynamic bandwidth. It includes the following types of examples:

- How non-elastic connections respond to dynamic bandwidth
- How elastic connections respond to dynamic bandwidth
- How mixed connections respond to dynamic bandwidth

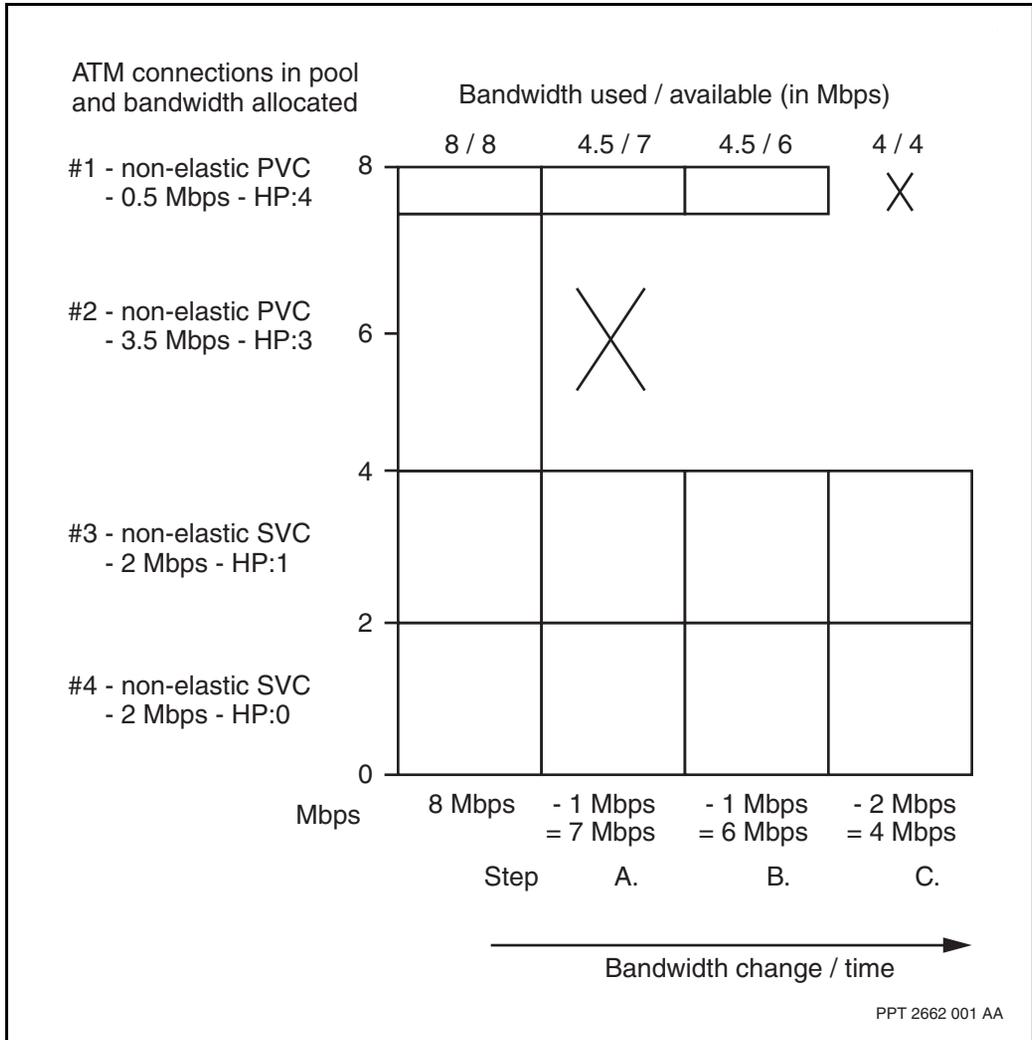
The examples shown in this section make the following assumptions:

- For ease of reference, all bandwidth is expressed in Mbits/s.
- The IMA link group in these examples begins with eight E1 physical links = 15.24 Mbits/s.
- The pool used in these examples is allotted 50 percent of IMA group bandwidth = 7.62 Mbits/s. The starting value of bandwidth allotted to the pool is shown as 8 Mbits/s to simplify the examples.
- One physical link lost or added to the IMA group = 1.905 Mbits/s (total), or 0.9525 Mbits/s lost or added to the pool. This value is shown as 1 Mbits/s to simplify the examples.

Example: non-elastic connections responding to dynamic bandwidth

The figure “How non-elastic connections in a pool respond to changes in bandwidth” (page 79) shows four non-elastic connection types in a bandwidth pool starting with 8 Mbits/s.

Figure 16
How non-elastic connections in a pool respond to changes in bandwidth



The connections in this example have the following characteristics:

- 1 - non-elastic PVC - 0.5 Mbits/s - HP:4
- 2 - non-elastic PVC - 3.5 Mbits/s - HP:3

- 3 - non-elastic SVC - 2 Mbits/s - HP:1
- 4 - non-elastic SVC - 2 Mbits/s - HP:0

The following sequence describes the effects of bandwidth changes shown in the example in the figure “How non-elastic connections in a pool respond to changes in bandwidth” (page 79).

Step A: losing the first link

The IMA link group loses one physical link, invoking the CBC bandwidth loss algorithm (see the figure “How the CBC algorithm handles bandwidth loss” (page 74)). There are now seven active links.

After proportionately distributing the bandwidth loss to the pool, approximately 1 Mbits/s must be removed, reducing the bandwidth available to the pool from 8 to 7 Mbits/s. Since there are no elastic VCCs in this example, some connections must be released in order to reduce the bandwidth requirement.

According to the algorithm, VCC 1 must be released because it has a lower holding priority than the other connections in the pool. Since VCC 1 uses only 0.5 Mbits/s, VCC 2, with a holding priority of 3, must also be released.

This leaves only 4 Mbits/s used out of a possible 7 Mbits/s bandwidth still available in the pool. At this point, the CBC bandwidth gain algorithm is invoked to investigate whether the leftover bandwidth can be used. As a result, VCC 1 is readmitted to the pool. Invoking the bandwidth gain algorithm ensures that no connections are released unnecessarily. VCC 1 was momentarily released, but is readmitted immediately; if VCC 1 were an SVC, it could not be readmitted.

Therefore, after Step A, VCC 2 is removed leaving a total of 4.5 Mbits/s bandwidth used out of 7 Mbits/s available to the pool.

Note: In this example, it makes no difference that VCCs 1 and 2 are PVCs (unlike VCCs 3 and 4, which are SVCs). This is because each connection has a different holding priority.

Step B: losing the second link

The IMA group loses another physical link, invoking the CBC bandwidth loss algorithm. There are now six active links.

An additional 1 Mbits/s must be removed from the pool, reducing the bandwidth available to the pool from 7 to 6 Mbits/s.

According to the algorithm, since the bandwidth available to the pool (6 Mbits/s) is greater than or equal to the current bandwidth requirement (4.5 Mbits/s), the status quo is maintained.

Step C: losing the third and fourth links

The IMA group loses two more physical links, invoking the CBC bandwidth loss algorithm. There are now four active links.

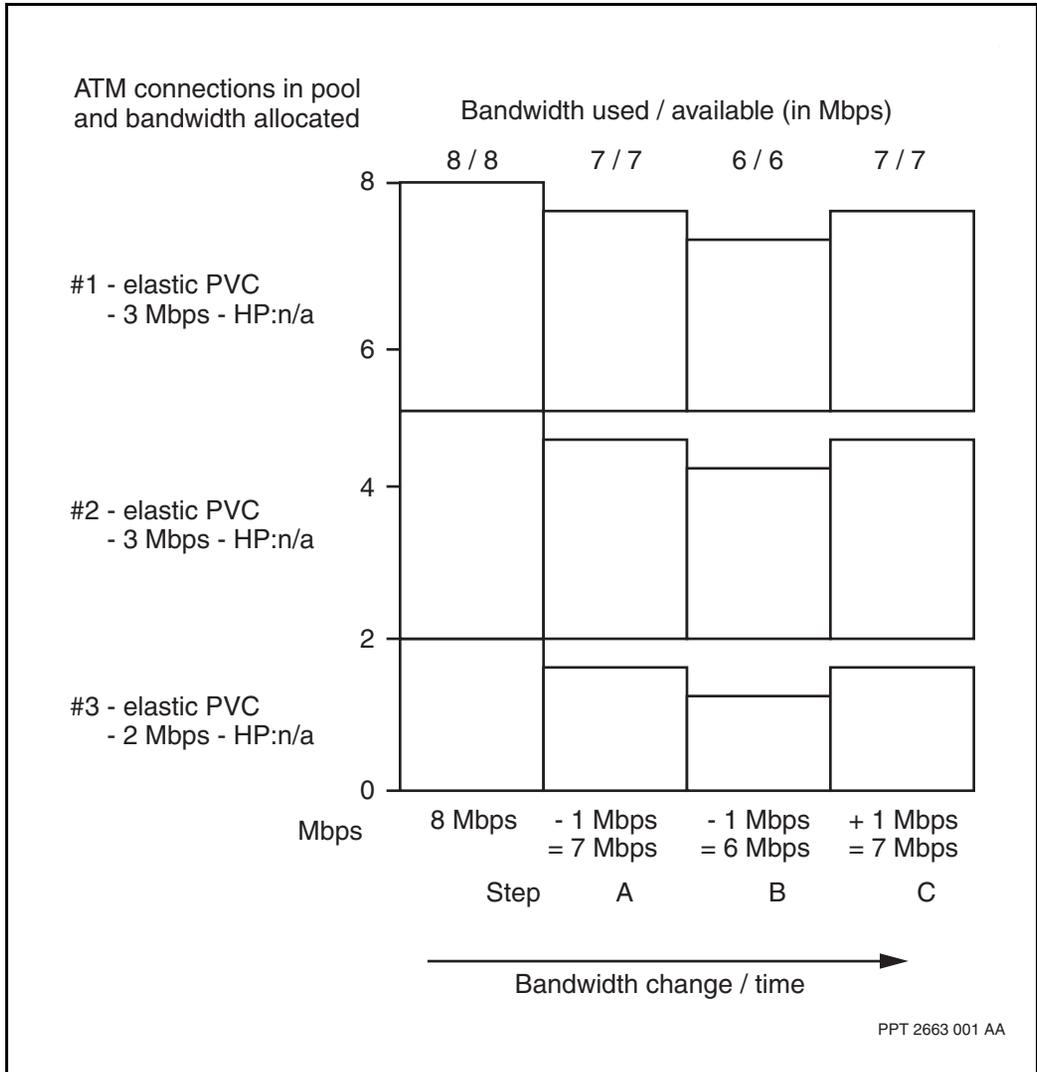
An additional 2 Mbits/s must be removed from the pool, reducing the bandwidth available to the pool from 6 to 4 Mbits/s.

According to the algorithm, VCC 1 must be released because it has a lower holding priority than the connections remaining in the pool. This leaves all 4 Mbits/s available to the pool used.

Example: elastic connections responding to dynamic bandwidth

This section describes how elastic connections behave alone in a bandwidth pool. The figure “How elastic connections in a pool respond to changes in bandwidth” (page 82) shows three elastic connection types in a bandwidth pool starting with 8 Mbits/s.

Figure 17
How elastic connections in a pool respond to changes in bandwidth



The connections in this example have the following characteristics:

- 1 - elastic PVC - 3 Mbits/s - HP:n/a
- 2 - elastic PVC - 3 Mbits/s - HP:n/a

- 3 - elastic PVC - 2 Mbits/s - HP:n/a

Currently, all elastic connections consist of permanent single-hop ATM logical trunks. The following sequence describes the effects of bandwidth changes shown in the example in the figure “How elastic connections in a pool respond to changes in bandwidth” (page 82).

Step A: losing the first link

The IMA link group loses one physical link, invoking the CBC bandwidth loss algorithm (see the figure “How the CBC algorithm handles bandwidth loss” (page 74)). There are now seven active links.

After proportionately distributing the bandwidth loss to the pool, approximately 1 Mbits/s must be removed, reducing the bandwidth available to the pool from 8 to 7 Mbits/s. Since there are no non-elastic VCCs in this example, the entire bandwidth loss must be absorbed by the elastic VCCs.

According to the algorithm, when elastic connections are forced to lose bandwidth, each elastic VCC must share its portion of the bandwidth loss. See “Rule 10: reducing bandwidth for all elastic connections” (page 70). Therefore, each VCC absorbs a portion of the loss.

- VCC 1: $3 - 0.375 = 2.625$ Mbits/s (87.5 percent of original bandwidth)
- VCC 2: $3 - 0.375 = 2.625$ Mbits/s (87.5 percent of original bandwidth)
- VCC 3: $2 - 0.25 = 1.75$ Mbits/s (87.5 percent of original bandwidth)

After Step A, all 7 Mbits/s bandwidth available to the pool is being used. In this example, unlike the example shown in the figure “How non-elastic connections in a pool respond to changes in bandwidth” (page 79), each connection continues to operate.

Step B: losing the second link

The IMA group loses another physical link, invoking the CBC bandwidth loss algorithm. There are now six active links.

An additional 1 Mbits/s must be removed from the pool, reducing the bandwidth available to the pool from 7 to 6 Mbits/s.

Each connection must share the same bandwidth loss experienced in “Step A: losing the first link” (page 83).

- VCC 1: $2.625 - 0.375 = 2.25$ Mbits/s (75 percent of original bandwidth)
- VCC 2: $2.625 - 0.375 = 2.25$ Mbits/s (75 percent of original bandwidth)
- VCC 3: $1.75 - 0.25 = 1.5$ Mbits/s (75 percent of original bandwidth)

After step B, all 6 Mbits/s bandwidth available to the pool is being used.

Step C: regaining a link

The IMA group regains one of the physical links it had lost, invoking the CBC bandwidth gain algorithm. There are now seven active links.

Bandwidth equal to 1 Mbits/s must be added to the pool, changing the bandwidth available to the pool from 6 to 7 Mbits/s.

According to the algorithm, when elastic connections gain bandwidth, each elastic VCC must benefit from its portion of the increase; see “Rule 11: increasing bandwidth for all elastic connections” (page 70). Therefore, each VCC receives a portion of the increase.

- VCC 1: $2.25 + 0.375 = 2.625$ Mbits/s (87.5 percent of original bandwidth)
- VCC 2: $2.25 + 0.375 = 2.625$ Mbits/s (87.5 percent of original bandwidth)
- VCC 3: $1.5 + 0.25 = 1.75$ Mbits/s (87.5 percent of original bandwidth)

After Step C, all 7 Mbits/s bandwidth available to the pool is being used, and the connections are operating with the same bandwidth as they were after Step A.

Example: mixed connections responding to dynamic bandwidth

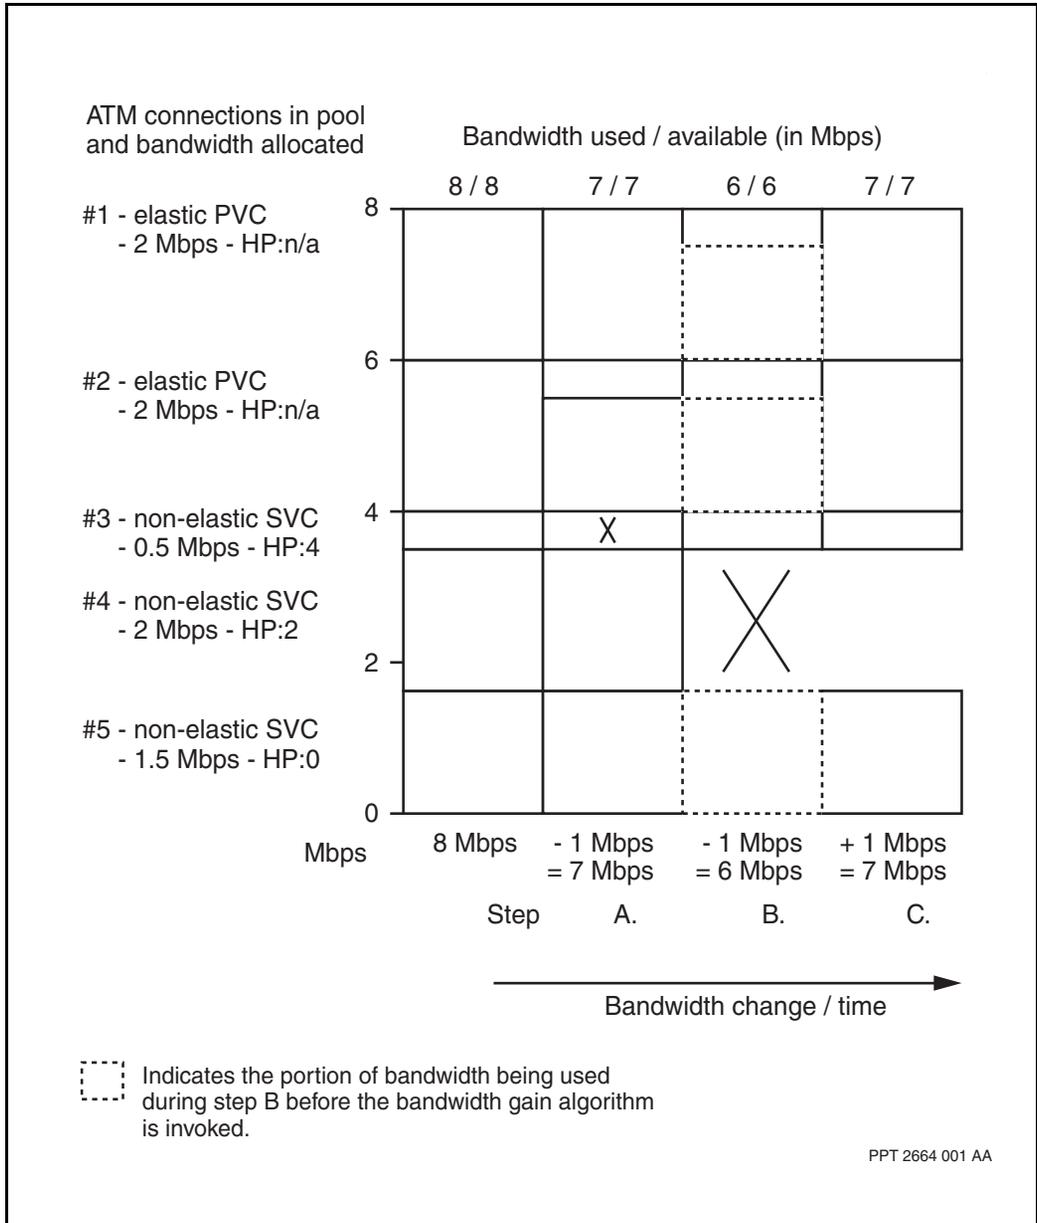
This section describes how elastic and non-elastic connections behave together in the same bandwidth pool.

In a competitive situation between elastic and non-elastic connections, elastic connections are

- the last to lose bandwidth
- the first to regain free bandwidth

The figure “How mixed connections in a pool respond to changes in bandwidth” (page 86) shows two elastic and three non-elastic connection types in a bandwidth pool starting with 8 Mbits/s.

Figure 18
How mixed connections in a pool respond to changes in bandwidth



The connections in this example have the following characteristics:

- 1 - elastic PVC - 2 Mbits/s - HP:n/a
- 2 - elastic PVC - 2 Mbits/s - HP:n/a
- 3 - non-elastic SVC - 0.5 Mbits/s - HP:4
- 4 - non-elastic SVC - 2 Mbits/s - HP:2
- 5 - non-elastic SVC - 1.5 Mbits/s - HP:0

The following sequence describes the effects of bandwidth changes shown in the example in the figure “How mixed connections in a pool respond to changes in bandwidth” (page 86).

Step A: losing the first link

The IMA link group loses one physical link, invoking the CBC bandwidth loss algorithm; see the figure “How the CBC algorithm handles bandwidth loss” (page 74). There are now seven active links.

After proportionately distributing the bandwidth loss to the pool, approximately 1 Mbits/s must be removed, reducing the bandwidth available to the pool from 8 to 7 Mbits/s. The loss must then be distributed proportionally between elastic and non-elastic connections. See “Rule 8: distributing bandwidth loss in a pool” (page 69). Since, at the start, elastic and non-elastic connections each use half of the bandwidth in the pool, each category of connections is assigned a loss of 0.5 Mbits/s.

According to the algorithm, the non-elastic connections must absorb their share of the loss before the elastic connections. See “Rule 9: releasing non-elastic connections before elastic connections” (page 69). VCC 3 must be released because it has the lowest holding priority. Since VCC 3 uses 0.5 Mbits/s, it is just enough to absorb the loss assigned to the non-elastic category. The elastic connections in the pool must also lose a total of 0.5 Mbits/s. As a result, each elastic VCC absorbs the loss of 0.25 Mbits/s.

Therefore, after Step A, all 7 Mbits/s bandwidth remaining in the pool is being used. Three connections in the pool have been affected by the loss of one physical link: one non-elastic connection has been released, and two elastic connections have been reduced.

Step B: losing the second link

The IMA group loses another physical link, invoking the CBC bandwidth loss algorithm. There are now six active links.

An additional 1 Mbits/s must be removed from the pool, reducing the bandwidth available to the pool from 7 to 6 Mbits/s. As in “Step A: losing the first link” (page 87), the loss must be distributed proportionally between elastic and non-elastic connections. Since again, elastic and non-elastic connections each use half of the bandwidth in the pool, each category of connections is assigned a loss of 0.5 Mbits/s.

According to the algorithm, VCC 4 must be released because it has a lower holding priority than the other non-elastic connection remaining. When VCC 4 is released, an additional 1.5 Mbits/s is released beyond the requirement for non-elastic connections. At this point, the pool is using a total of 5 out of 6 Mbits/s available to the pool. Therefore, there is no need to lose any bandwidth in the elastic category (the 0.5 Mbits/s owed by the elastic connections is disregarded). The 5 Mbits/s currently in use is highlighted using dashed lines in the figure “How mixed connections in a pool respond to changes in bandwidth” (page 86).

With 1 Mbits/s available, the CBC bandwidth gain algorithm is activated to distribute the free bandwidth among the connections in the pool; see the figure “How the CBC algorithm handles bandwidth gain” (page 77). According to the algorithm, any reduced elastic connections must regain their full speed when bandwidth is available; see “Rule 11: increasing bandwidth for all elastic connections” (page 70). Therefore, VCC 1 and VCC 2 are each brought back to 2 Mbits/s, consuming 0.5 Mbits/s of the excess bandwidth of 1 Mbits/s. When all the elastic connections are back at full speed, and there is bandwidth remaining, the algorithm can investigate reinstating non-elastic connections; see “Rule 12: admitting or re-admitting connections” (page 71). Since only VCC 3 fits the amount of 0.5 Mbits/s bandwidth available, it is readmitted based on the first fit rule.

At this point, all of the 6 Mbits/s bandwidth available is being used; only VCC 4 is not active.

Step C: regaining a link

The IMA group regains one of the physical links it had lost, invoking the CBC bandwidth gain algorithm. There are now six active links.

Bandwidth equal to 1 Mbits/s must be added to the pool, changing the bandwidth available to the pool from 6 to 7 Mbits/s.

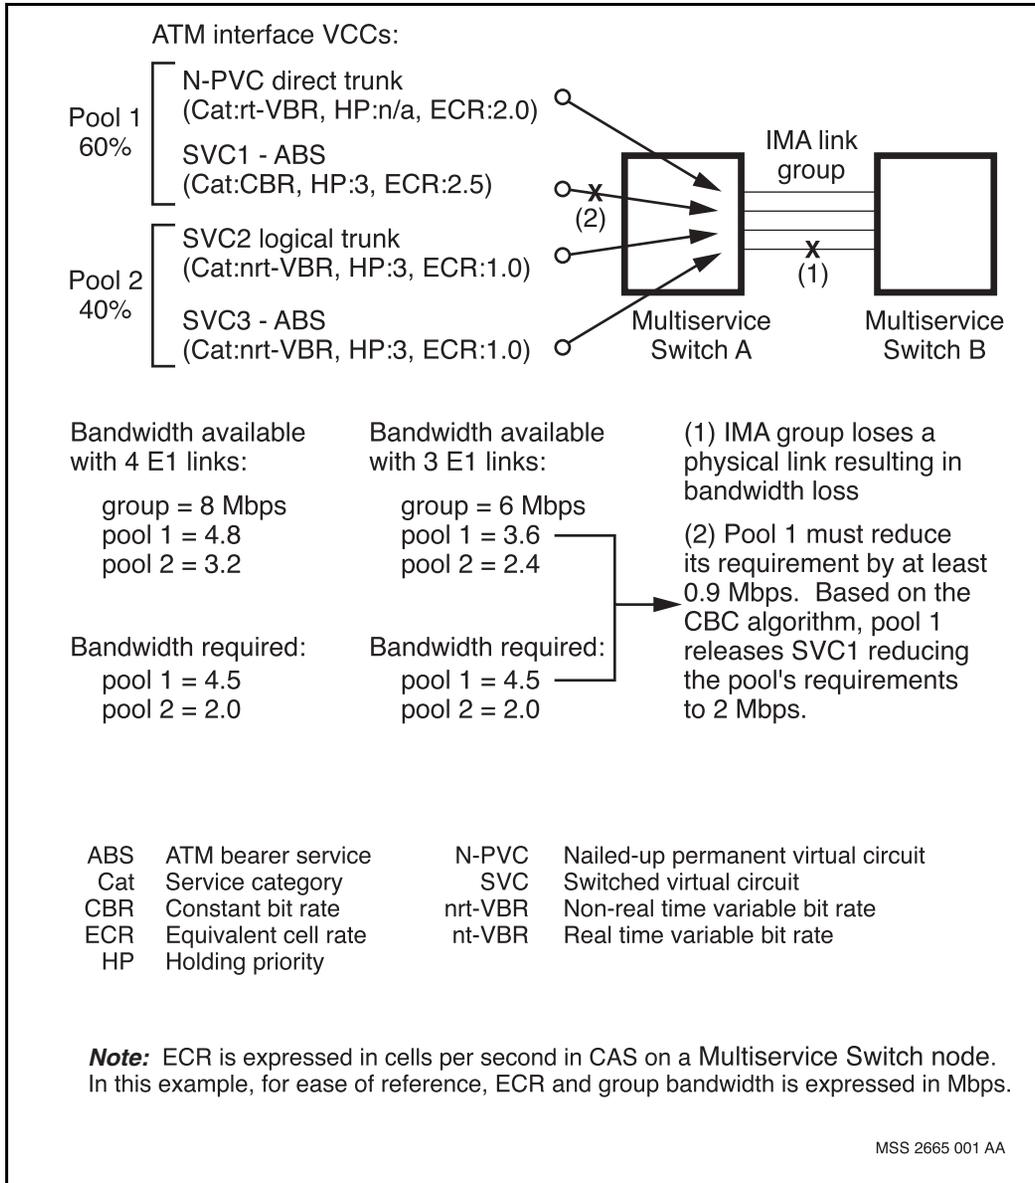
Since all elastic connections are at full speed, and the only non-elastic connection that is currently not part of the pool is VCC 4, a 2 Mbits/s connection, the extra bandwidth cannot be used.

Therefore, after Step C, 6 Mbits/s is being used out of the total of 7 Mbits/s available to the pool.

CBC and bandwidth pools

The figure “Example of how bandwidth loss is handled within pools” (page 90) shows how bandwidth pooling can affect how a bandwidth loss is handled.

Figure 19
Example of how bandwidth loss is handled within pools



In the figure “Example of how bandwidth loss is handled within pools” (page 90), one of four physical links in an IMA group fails, reducing the bandwidth over the group from 8 Mbits/s to 6 Mbits/s.

The numbers provided in the figure are rounded-off to simplify the example. They do not take into account the overhead on the physical link, or the overhead associated with the IMA process.

As a result, each pool has less bandwidth available to it in proportion to its allotment of the entire interface. The bandwidth available to pool 1 (60 percent of the interface) is reduced from 4.8 Mbits/s to 3.6 Mbits/s; the bandwidth available to pool 2 (40 percent of the interface) is reduced from 3.2 Mbits/s to 2.4 Mbits/s. Pool 3 is not used in this example.

After the adjustment, pool 2 is still able to handle its requirement of 2.0 Mbits/s.

Note: A pool’s bandwidth requirement is based on the sum of the equivalent cell rates (ECRs) for all admitted connections in the pool.

Pool 1, however, can no longer support its requirement of 4.5 Mbits/s. Based on the bandwidth loss part of the CBC algorithm, pool 1 releases SVC1 and maintains the PVC elastic connection at full bandwidth. Although SVC1 has the same holding priority as SVC2 and SVC3 and is assigned a higher profile service category (CBR), it is released because it is part of pool 1, the pool that can no longer meet its bandwidth requirements.

Without the bandwidth pools, it would make sense to release SVC2 or SVC3 in pool2. These connections are assigned a lower profile service category (NRT-VBR) and demand less bandwidth beyond the reduction requirement. Specifically, SVC2 and SVC3 each require 1.0 Mbits/s; releasing this amount of bandwidth would be closer to the reduction requirement of 0.5 Mbits/s.

Without categorizing connections into pools, the actual bandwidth reduction requirement in the figure “Example of how bandwidth loss is handled within pools” (page 90) example is the difference between 6.5 Mbits/s (total requirement) and 6.0 Mbits/s (total available).

Using bandwidth pools guarantees a percentage of bandwidth for each of the pools. This ensures that there is always some serving capacity for connections belonging to the service category associated with a pool.

However, the use of bandwidth pools may provide unwanted protection in certain instances. For example, if protecting NRT-VBR connections and releasing a CBR connection is not a desired outcome, the alternative is to place all ATM service categories in the same bandwidth pool and use holding priority to differentiate between different priorities of ATM service category. Assigning CBR connections a higher holding priority than NRT-VBR connections will result in all NRT-VBR connections being cleared before any CBR connections. In this scenario it is also possible to make specific configured NRT-VBR connections a higher holding priority.

Dynamic bandwidth and data protection

There are several items to keep in mind when configuring connections over a link group supporting dynamic bandwidth:

- Dynamic bandwidth technology, such as IMA, provides a mechanism whereby many connections transmitting ATM user data can continue to operate even when one or more physical links (that are part of a link group) go down.
- When one or more physical links go down causing a reduction in bandwidth such that insufficient bandwidth is available over a link group, data loss will occur. When you configure your ATM connections, you can control which data is protected in the event of a reduction in available bandwidth.
- Your traffic contract provides the principal mechanism for protecting your data. See the section on traffic contracts in *NN10600-705 Nortel Networks Multiservice Switch 7400/15000/20000 ATM Traffic Management Fundamentals*.

The use of bandwidth elastic and non-elastic connections provides additional flexibility to those configuring Nortel Networks Multiservice Switch ATM networks. You can protect critical connections transmitting OAM data, network management or signaling data, by configuring such connections as either bandwidth-elastic or bandwidth-non-elastic (with a holding priority of 0).

The measure of safety built into either type of connection depends on how your ATM interface is organized. For example, if your signaling data is configured on a non-elastic connection with a holding priority of 0, and this is the only non-elastic connection in a bandwidth pool, the data being transmitted over this connection would be more vulnerable than it would be if it were in a pool with several (lower priority) non-elastic connections, as in the example of VCC 4 in the figure “How non-elastic connections in a pool respond to changes in bandwidth” (page 79). The presence of other non-elastic connections will protect the connection(s) with a holding priority of 0 in the event of a reduction in available bandwidth.

Unlike non-elastic connections, all elastic connections are subject to partial bandwidth loss. Although elastic connections are the last to lose and first to gain bandwidth, they are also vulnerable to losing parts of their bandwidth, a circumstance which may adversely affect certain types of data running on a connection.

Coordinating dynamic bandwidth between nodes

Because the CBC algorithm is executed independently on both sides of an IMA interface, it may release different connections at each end of an IMA group depending on how the connections are configured and how bandwidth pools are organized at each end of the link.

There is no explicit coordination between two adjacent nodes linked by an IMA group. If each connection has a symmetrical bandwidth requirement on the transmit and receive directions, then each connection is equally affected on both sides of the IMA link when the CBC algorithm is executed. However, if the bandwidth requirements between connections at both ends are asymmetrical, then one side of the IMA interface will release, reinstate, reduce, or increase the bandwidth of more (or less) connections than the other side of the link when a change occurs.

VPT CAC response to dynamic bandwidth changes

Virtual path termination CAC (VPT-CAC) supports the same dynamic bandwidth control behavior as the ATM interface CAC (AtmIf-CAC).

The node views the following as non-elastic connection points:

- Any VPT with VPT-CAC that is admitted using AtmIf-CAC

- any VCC admitted using VPT-CAC

If an ATM interface cannot provide a VPT with the required bandwidth, the node releases the VPT and its associated VCC connection points. The node re-establishes the VPT and its associated VCCs when bandwidth is available. As with other ATM connection points, you can provision the holding priority of the VPT so that it is more or less susceptible to ATM interface bandwidth fluctuations relative to the other connection points using the interface.

During dynamic bandwidth changes, the node ignores any VPT that is not configured with a VPT-CAC. The ATM interface functions as if the VPT was not present. Instead, the node releases the VPT VCCs (and later re-admitted) on an individual basis rather than as a group according to characteristics for elasticity and holding priority. These VPT VCCs can later be re-admitted.

Nortel Networks Multiservice Switch 7400/15000/20000 ATM CAC and Bandwidth Fundamentals

Release 6.1

Copyright © 2004 Nortel Networks.
All Rights Reserved.

NORTEL, NORTEL NETWORKS, the globemark design, the NORTEL NETWORKS corporate logo, DPN, and PASSPORT are trademarks of Nortel Networks.

Publication: NN10600-708
Document status: Standard
Document version: 6.1S1
Document date: August 2004
Printed in Canada

