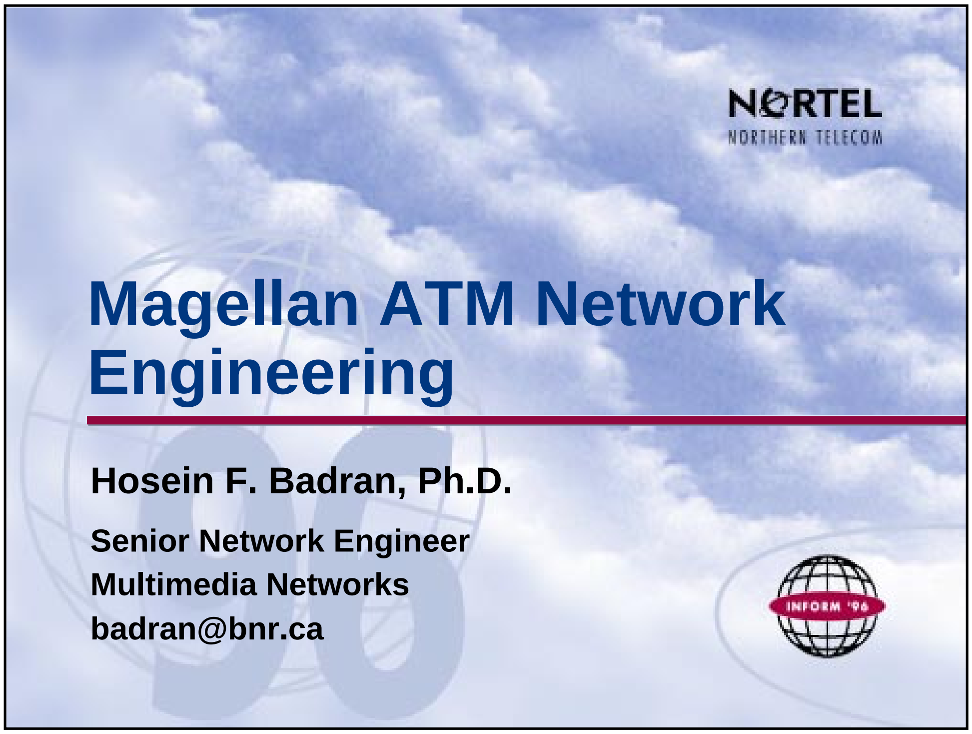

Proper engineering of ATM networks should aim at meeting the different quality of service (QOS) requirements of various service classes, such as real-time services (i.e. packetized voice, and video) and non-real time services (i.e. file transfer, and other data applications), while efficiently and cost-effectively utilizing the network resources.

ATM network engineering relies on the proper operation and effectiveness of the different mechanisms utilized to manage and control the traffic flow in the network, both at the network access, as well as within the backbone structure.

This session presents approaches for Magellan ATM network engineering. Mechanisms for achieving QOS requirements for different services will be addressed. Elements of backbone network engineering necessary to satisfy QOS requirements will be described through a network design example.

About the presenter:

Hosein Badran is a senior engineer with the Multimedia Networks Engineering group. Before joining the engineering group two years ago, he spent three years with the Magellan Traffic Management group working on different activities related to traffic management for frame relay, packetized voice, and ATM services, as well as development of the congestion management platform for the Magellan Passport product.

He obtained his Ph.D in the area of design and performance analysis of ATM networks and non-blocking switch architectures.

Dr. Badran is an active participant in the ATM Forum, and has been selected as an official ATM Forum ambassador.

Agenda

- ATM Traffic Management
- Bandwidth Allocation and Control
- ATM Network Engineering: A Case Study
- Summary

2

ATM network engineering relies on the proper operation and effectiveness of the different mechanisms utilized to manage and control the traffic flow in the network, both at the network access, as well as within the backbone structure.

This session presents approaches for Magellan ATM network engineering. Mechanisms for achieving QOS requirements for different services will be addressed. Elements of backbone network engineering to satisfy QOS requirements will be described through a network design example.

This presentation addresses the key components required to properly engineer an ATM network. First, key traffic management concepts are examined, with reference to the recent standards material. An integral step in the network design process is the allocation and enforcement of the bandwidth resources required by the different ATM connections. Key functions affecting the bandwidth allocated and controlled by the ATM network switches are presented.

A case study for the design and engineering of an ATM network carrying video, packetized voice and frame relay data services is discussed in detail.

Appendix A contains a glossary of terms used in this presentation.

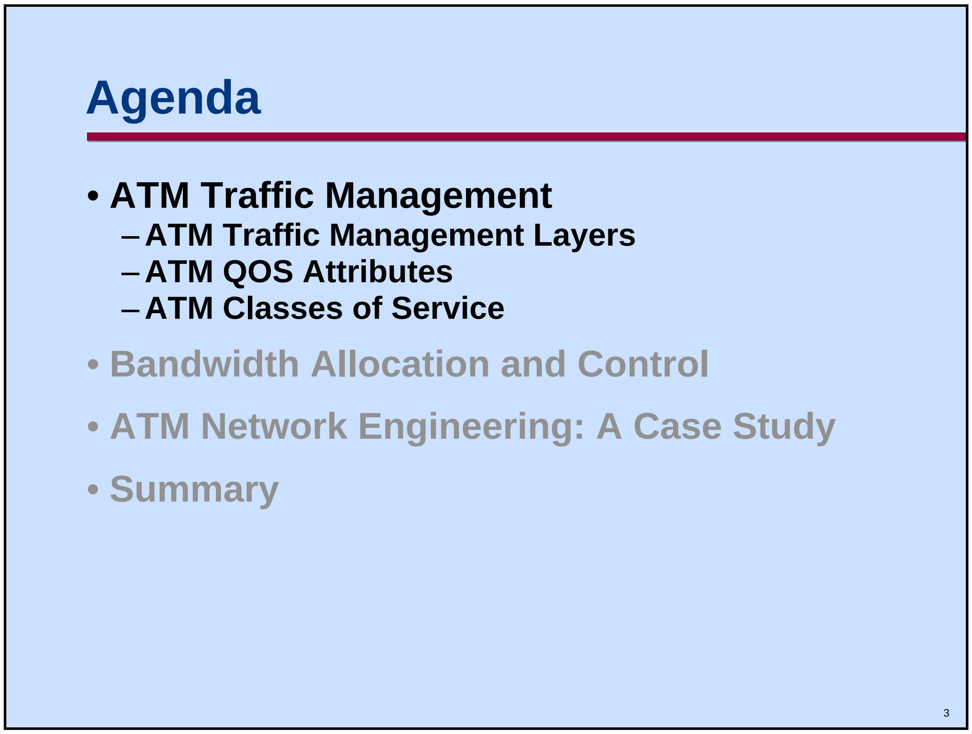

This section covers key concepts in traffic management that directly affect the design and engineering of ATM networks.

The classification of traffic management functions and schemes, divided according to the time duration in which these schemes control the traffic flow, is presented. The ATM QOS attributes, as well as the ATM service classes and QOS parameters defined by the ATM Forum, are described. Meeting these QOS requirements of the different services carried by the ATM network is an important objective of a properly designed and engineered network.

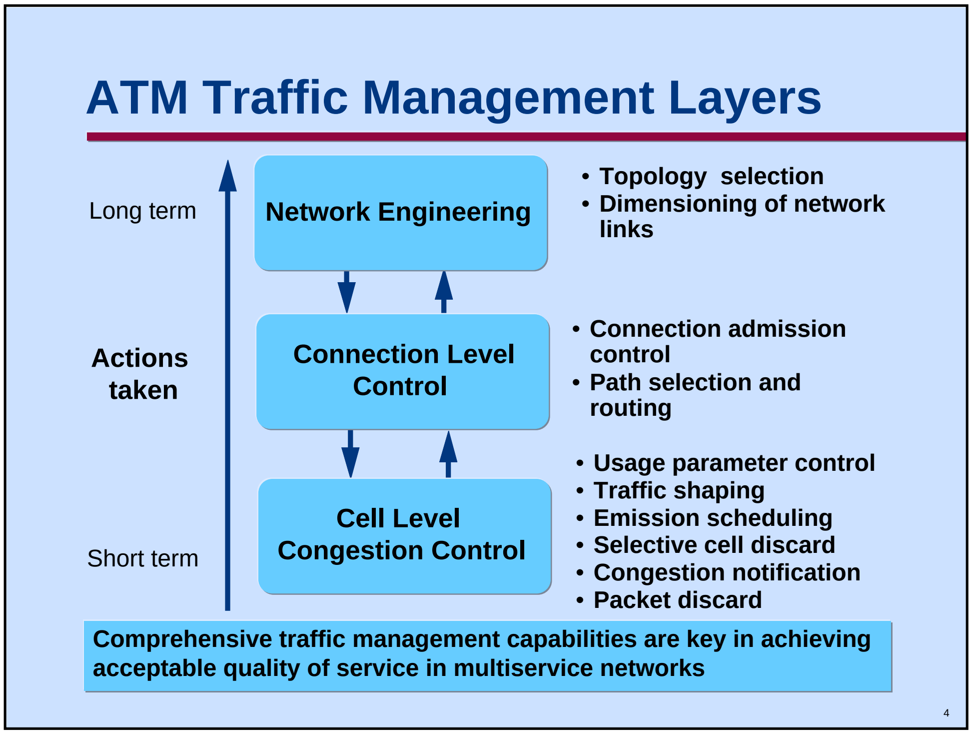

Traffic management functions and schemes in ATM networks can be classified according to the time duration in which these schemes control the network traffic flow, or equivalently, the time duration of network congestion that they are aimed at controlling.

Network engineering aims at providing customers with acceptable levels of service for expected traffic demands. Topology selection and dimensioning of internal network links to meet grade of service (GOS) and QOS requirements are the main components of proper network engineering.

For sporadic congestion, one method is to route traffic according to the load level of network links, and to reject new connections if all paths are highly loaded. This is called “connection admission control” or CAC. CAC is effective mainly for medium duration congestion, since once the connection is admitted, the congestion may persist for the duration of the connection.

For congestion lasting less than a connection duration, either end-to-end control schemes, or cell level control schemes, can be used. For example, during connection setup, the source traffic rate may be negotiated. Once a connection is accepted, the usage parameter control (UPC) device at the network access ensures that the source traffic rate stays within negotiated limits. Emission scheduling and selective cell discard are two other control mechanisms that work at controlling and alleviating cell level congestion.

An important point here is that although the mechanisms in the different layers control different network congestion scenarios, the acceptable overall level of service can be met only if each layer relies on the correct design and function of the layers below. For example, while CAC operates on the call level, it takes into consideration the cell level behavior of the network traffic. At the same time appropriate dimensioning and scheduling of the node queues affects the correct operation of the CAC algorithms. In the same sense, service engineering and network dimensioning are affected by the appropriate connection level and cell level mechanisms.

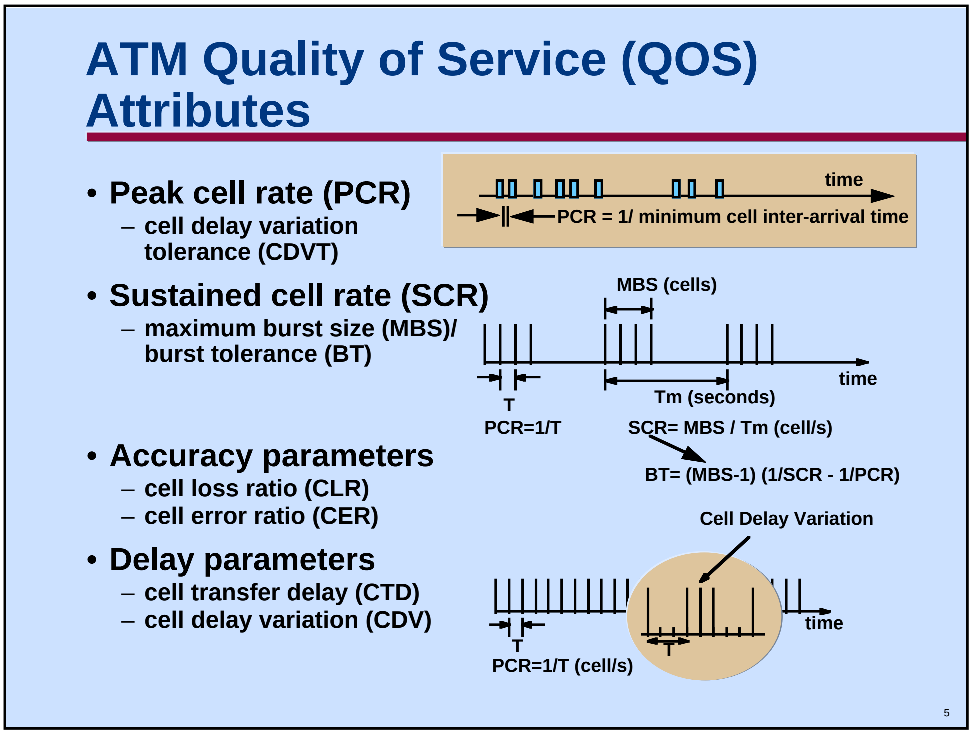

Peak cell rate (PCR): The maximum cell rate at which the user will transmit. PCR is the inverse of the minimum cell inter-arrival time.

Cell delay variation tolerance (CDVT): Allowable variation in the minimum cell inter-arrival time. This parameter specifies how much the user is allowed to vary around the negotiated PCR.

Sustained cell rate (SCR): This is the average rate, as measured over a long interval, in the order of the connection life time.

Burst tolerance (BT): This parameter determines the maximum burst that can be sent at the peak rate. This is the bucket size parameter for the enforcement algorithm that is used to control the traffic entering the network. The maximum burst size (MBS) is proportional to the burst tolerance BT, and is defined through the relation:

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

Cell loss ratio (CLR): Cell loss ratio is the percentage of cells that are lost in the network due to congestion and buffer overflow, and are not delivered at the destination.

Cell error rate (CER): Percentage of cells delivered at destination with payload errors. Errors are mainly due to transmission medium inadequacy.

Cell transfer delay (CTD): The delay experienced by a cell between network entry and exit points is called the cell transfer delay. It includes propagation delays, queueing delays at various intermediate switches, and service times at queueing points.

Cell delay variation (CDV): A measure of the variance of the cell transfer delay. High variation implies larger buffering for delay sensitive traffic such as voice and video.

ATM Classes of Service

- **ATM classes of service**
  - constant bit rate (CBR)
  - real-time variable bit rate (real-time VBR)
  - non-real-time variable bit rate (non-real-time VBR)
  - available bit rate (ABR)
  - unspecified bit rate (UBR)

QOS parameters to be specified for each service class:

	ATM Layer Service Categories				
Attribute	CBR	rt-VBR	nrt-VBR	UBR	ABR
Traffic Parameters:					
PCR and CDVT	specified			specified	specified
SCR, MBS, CDVT	n / a	specified		n / a	
MCR	n / a			n / a	specified
QoS Parameters:					
peak-to-peak CDV	specified	specified	unspecified	unspecified	unspecified
maximum CTD	specified	specified	unspecified	unspecified	unspecified
CLR	specified			unspecified	See Notes
Other Attributes:					
Feedback	unspecified			unspecified	specified

6

There are five classes of service defined for ATM (as per ATM Forum UNI 4.0 Specification, which is currently being finalized). The QOS parameters for these service classes are summarized in the table shown above.

1. Constant bit rate (CBR): This class is used for emulating circuit switching. The cell rate is constant with time. Applications are quite sensitive to cell delay variation, and cell loss has to be strictly controlled. Examples of applications that can use CBR are telephone traffic (i.e. nx64 Kbit/s), video conferencing, and television.

2. Variable bit rate (VBR): This class allows users to send traffic at a rate that varies with time, depending on the availability of user information to transmit. Statistical multiplexing is used to benefit from the traffic burstiness to make maximum use of network resources. Depending on whether or not the application is sensitive to cell delay variation, this class is subdivided into two categories: real-time VBR, and non-real-time VBR. While the cell transfer delay parameter is specified for both VBR categories, cell delay variation is specified only for real-time VBR. Examples for real-time VBR are voice with speech activity detection (SAD) and interactive compressed video. For non-real-time VBR multimedia e-mail is an example.

3. Available bit rate (ABR): This class of ATM services is currently being specified by the standardization bodies, such as the ITU-T and the ATM Forum. It is aimed at data traffic such as file transfer and e-mail. Although the standard does not require the cell transfer delay and cell loss ratio to be guaranteed or minimized, it is desirable for switches to minimize the delay and loss as much as possible.

Depending upon the congestion state of the network, the source is required to control its rate. The users are allowed to declare a minimum cell rate, which is guaranteed to the connection by the network.

4. Unspecified bit rate (UBR): This class is designed for those data applications that want to use any left-over capacity and are not sensitive to cell loss or delay. Such connections are not rejected by the network on the basis of bandwidth shortage (no connection admission control) and not policed for their usage. During congestion, the cells are lost but the sources are not expected to reduce their rate. Instead, these applications may have their own higher-level cell loss recovery and retransmission mechanisms. Examples of applications that can use this service are e-mail, file transfer, news feed, etc. Of course, if a more reliable service is desired, these applications can use the ABR service instead.

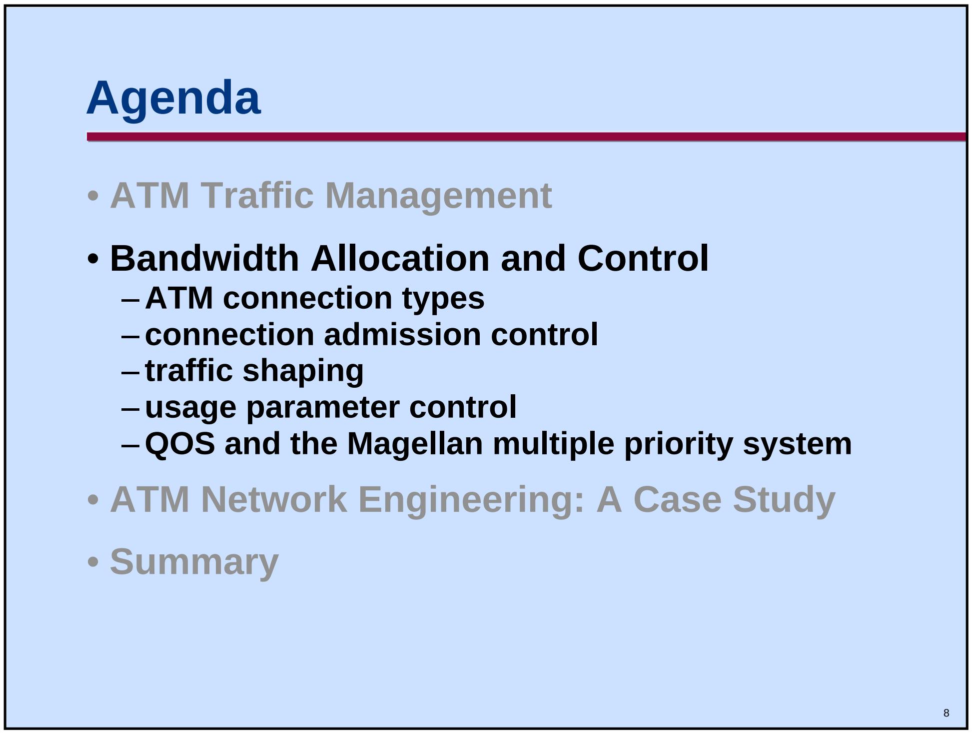

The specification of the bandwidth requirements of incoming connections as well as the enforcement and control of the declared traffic characteristics are key elements in the design of an ATM network.

In this section the two types of ATM connections available to carry the end-user traffic across an ATM network are reviewed. Three main aspects for the connections bandwidth allocation and control are considered, namely call admission control (CAC), traffic shaping, and usage parameter control (UPC). The relationship between the available ATM QOS classes and the Magellan Multiple Priority System, in terms of the multiple emission and discard priorities, is highlighted.

ATM Connection Types

- **Virtual channel connections (VCCs)**
  - end-to-end
  - used to transfer
    - user information
    - signaling information
    - management information

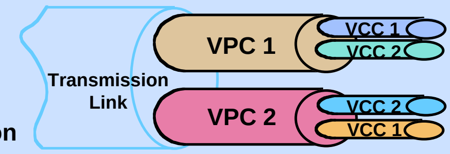

- **Virtual path connections (VPCs)**
  - end-to-end
  - or network segment wide

VPCs introduce the concept of virtual trunks, where multiple VCCs are grouped together into a VPC

9

Transport function and capability in an ATM network are structured hierarchically. Two levels of connections are included in ATM. A virtual channel connection is the basic unit, which carries a single stream of cells in order, from user to user.

Moreover, a collection of virtual channel connections can be bundled together into a virtual path connection. A virtual path connection can be created from end-to-end across an ATM network. In this case, the ATM network does not route cells belonging to a particular virtual channel, but cells of all the virtual channels belonging to a particular virtual path are routed the same way through the ATM network.

But an ATM network is also able to make use of the concept of a virtual path for purposes of bundling virtual circuits together between switches. Two ATM switches may have many different virtual channel connections between them belonging to different users. These can be bundled by the two ATM switches into a virtual path connection. This can serve the purpose of a virtual trunk between the two switches. This virtual trunk can then be handled as a single entity by perhaps multiple intermediate virtual path cross connects between the two virtual circuit switches.

Advantages of VPCs include bandwidth consolidation between multiple VCCs, less network update frequency, faster routing capability, traffic shaping and enforcement on a per VPC basis, and virtual trunking between multiple sites requiring dedicated capacity.

Connection Admission Control (CAC)

Function

- **To establish, at call set-up phase, whether a connection can be accepted or rejected, while maintaining QOS objectives**

Benefit

- **Guarantee QOS requirements to new connections, as well as to already established connections**

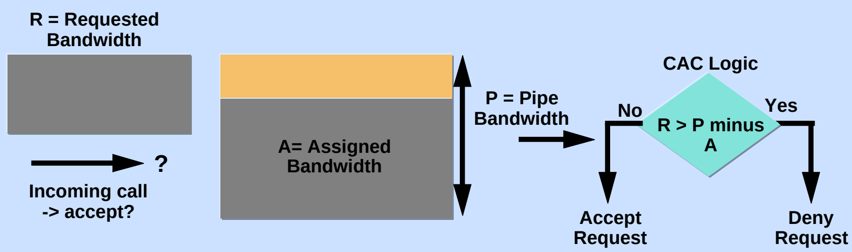

10

On receiving a request for a new connection, consisting of a request for a certain service category and related traffic parameters, the network must decide whether or not to admit the call. This aspect of the connection establishment procedure is called connection admission control. In this phase, the network checks whether it has sufficient resources that it can allocate to meet the traffic contract requested by the connection without violating existing contracts. If the network has enough resources, the allocation is “accepted”. Otherwise the connection is “rejected”.

A CAC algorithm for CBR traffic stream takes into consideration the peak cell rate required by the connection, and the cell delay variation tolerance to allow variation around the declared peak cell rate. The QOS requirements of the connection in terms of cell loss, cell delay, and cell delay variation are also taken into account.

VBR connections are admitted to meet the QOS requirements on the basis of their required bandwidth. The CAC algorithm for VBR connections takes into account the connection traffic parameters (PCR, SCR, MBS) and the QOS requirements of the VBR connection (in terms of cell loss, and cell delay).

Example of CAC Algorithms

Traffic Category	Bandwidth allocated by CAC
CBR	PCR
VBR	Equivalent Bandwidth (EB); EB = function (PCR, SCR, MBS, buffer size, cell loss target)
ABR	MCR
UBR	zero (admit all connections)

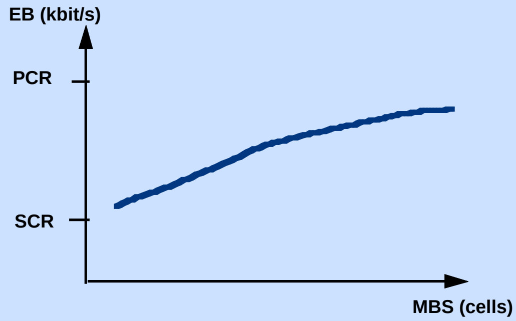

11

The chart above presents the different CAC algorithms supported by Magellan ATM switches, taking the Magellan Vector as an example. In general the CAC logic is different depending on the requested class of service:

CAC for CBR traffic:

The PCR of the new connection is compared to the bandwidth available on the outgoing ink. The connection is accepted if PCR is less than the bandwidth available . Otherwise the connection is rejected.

CAC for VBR traffic:

The CAC for VBR is more complex. The complexity lies in determining what bandwidth to allocate to the new connection. Allocating its PCR is too conservative and does not take advantage of possible statistical multiplexing gain. On the other hand, allocating SCR is too aggressive, and will result in large cell losses if multiple connections happen to send a burst of cells. To address this problem, an equivalent bandwidth (EB) for the new connection is calculated using the fluid-flow analysis and a Gaussian approximation for the queue length distribution. The required bandwidth for the connection request is estimated by the minimum value obtained by the two methods. The effective bandwidth (EB) of a single ON-OFF (VBR) source is calculated as follows:

EB=PCR*(Y-X+SQRT((Y-X)^2+4*X*(SCR/PCR)*Y))/(2*Y)
where
Y=LN(1/CLR)*(MBS/SCR)*(PCR-SCR)
CLR is hard-coded to 10^(-8)
X=Dedicated Buffer Space in cells

Effect of MBS:

The sustainable cell rate is calculated using SCR= MBS/Tm where Tm is the mean burst inter-arrival time in seconds and MBS is the maximum number of cells that can be sent at the peak cell rate. It determines an upper bound on the length of a burst transmitted in compliance with the connection's peak cell rate (PCR). The equivalent bandwidth EB allocated by CAC to a VBR connection lies generally between the SCR and the PCR values. As shown in the previous diagram, as the value of maximum burst size increases, the bandwidth requirement of the connection (in terms of the equivalent bandwidth (EB)) increases. For connections generating fairly long bursts at the source’s peak rate, the EB can get quite close to the connection’s PCR.

CAC for UBR traffic:

Since UBR traffic does not request any guarantees, there is no bandwidth related CAC for connections of the UBR traffic category. However, similar to the other types of connections, a UBR connection may be rejected for reasons such as lack of reachability to the destination, or lack of a VCI that can be allocated to the connection.

CAC can also permit a certain amount of VBR resource "overbooking" in order to increase statistical multiplex gain.

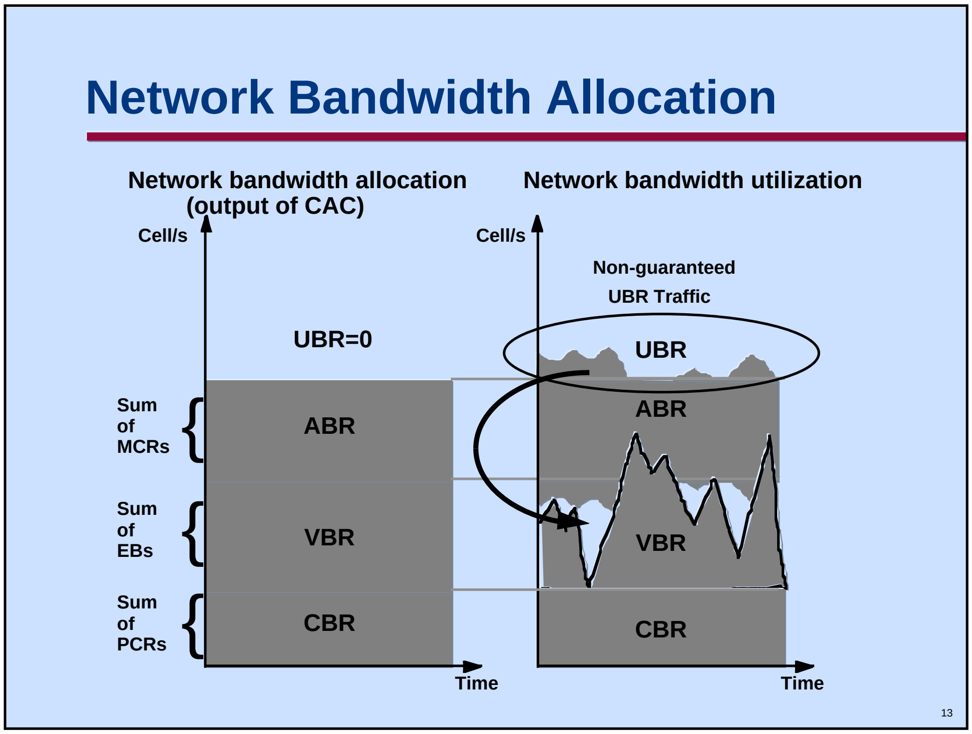

The above chart presents the bandwidth allocated by the CAC algorithm for the different types of connections at the switch interface (on the left), as well as the actual instantaneous utilization of this link carrying the traffic mix (on the right).

In both cases, CBR connections, constantly utilize the capacity allocated to them (sum of PCRs). For VBR connections, while the CAC algorithm executed at cell setup determines the equivalent baandwidth (EB) for each connection, the actual cell rate generated by the VBR sources varies with time, sometimes exceeding the EB allocation, and in other times requiring only much smaller capacity.

ABR connections, having the MCR allocated to each connection by the CAC algorithm, will attempt to utilize any bandwidth freed up by the VBR service, permitting ABR sources to send at a higher rate than the guaranteed MCR, while still maintaining the QOS guarantees .

UBR connections which have no bandwidth allocated to them at call setup, will also attempt to utilize any free bandwidth in the link. ABR being a more reliable service than UBR, due to the flow control algorithms employed at the ATM layer, takes precedence in utilizing link bandwidth freed up by non-active VBR sources.

UBR service usually supports a high degree of statistical multiplexing among sources. UBR connections include no notion of a per-VC allocated bandwidth resource. Transport of cells in UBR service is not guaranteed by mechanisms operating at the cell level. However, it is expected that network resources will be engineered for the UBR service in such a way as to make it usable for some set of applications.

Usage Parameter Control (UPC)

Function

- **To control the traffic offered by an ATM connection to ensure compliance with negotiated traffic contract**
- **ITU-T (in I.371) and ATM Forum standardized the generic cell rate algorithm (GCRA) enforcer**
- **Magellan has multiple provisionable enforcers, each based on the GCRA algorithm**

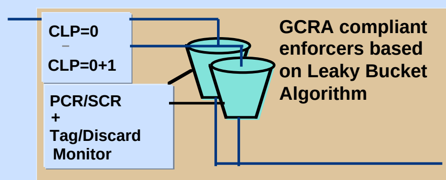

Magellan Usage Parameter Control (UPC)

14

Usage parameter control (UPC) is the set of actions taken by the network to monitor and control (i.e. police) offered traffic to protect network resources. The network has the responsibility to maintain the quality of service(QOS) granted to each connection. Therefore, the network must deny the creation of new connections which might impair its ability to maintain QOS for existing connections. In addition to this, the network must monitor and control the existing connections to assure that they do not violate their negotiated contract. These functions are carried out via the UPC.

If the traffic stream of an accepted connection (VCC or VPC) is non-compliant, for example, there are more cells than the contract allows, the ATM switch may either set the cell loss priority (CLP) bit to one (called “tagging”) or discard the non-conforming cell. A tagged cell may be transferred through the network, only if the current network capacity is sufficient. If not, the cell is discarded.

The Magellan UPC mechanisms implement multiple enforcers (commonly referred to as “leaky buckets”) to enforce the declared parameters of the established connections. In the Magellan UPC mechanisms, a "leaky bucket" can be configured to monitor the peak cell rate (PCR) or the sustainable cell rate (SCR) of a connection by specifying the appropriate leak rate and bucket size. Treatment of violating cells per enforcer can be provisioned to drop, tag, or simply monitor non-compliant cells. The programmability of UPC also allows the enforcement to be bypassed or to act on traffic with either CLP=0, CLP=1, or aggregate CLP=0 and CLP=1 (also referred to as CLP=0+1). For each connection, a combination of up to three leaky buckets may be specified along with the sequential order in which they should be processed.

This feature allows strategic flexibility to define multiple conformance rules for various current and future services and tariff requirements.

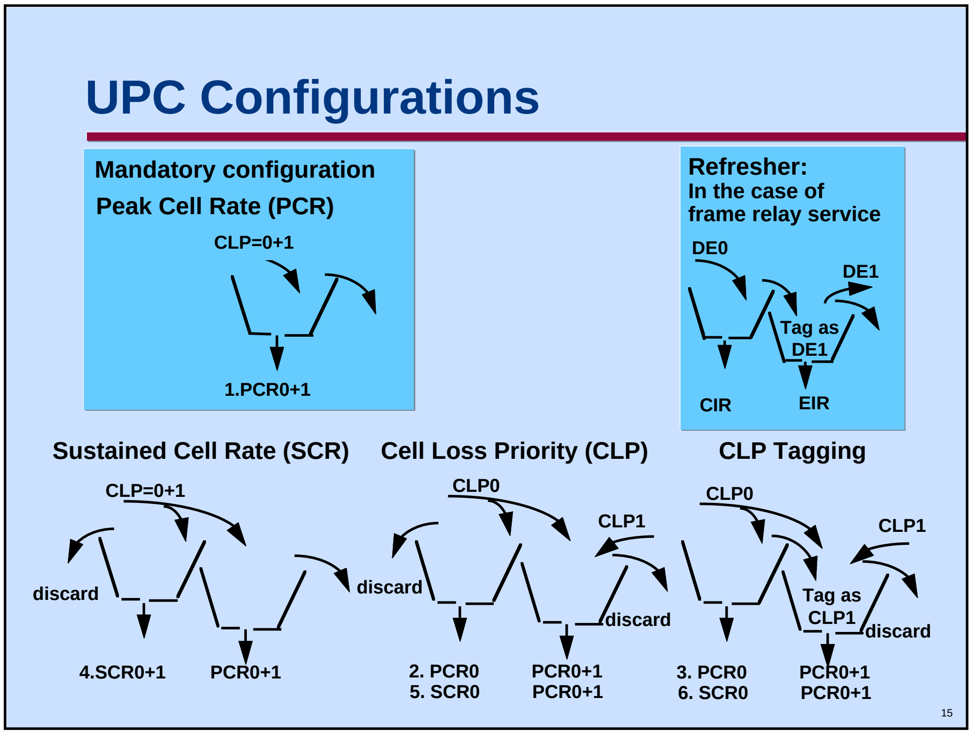

The following combinations of two leaky buckets and associated traffic parameters are specified by the ATM Forum for interoperability. These traffic parameters may be either explicitly specified, for example, in a signaling message; or implicitly specified, for example, by a default rule in the network.

1. PCR on CLP=0+1
2. PCR on CLP=0+1, and PCR on CLP=0 without tagging
3. PCR on CLP=0+1, and PCR on CLP=0 with tagging
4. PCR on CLP=0+1, and SCR+MBS on CLP=0+1 without tagging
5. PCR on CLP=0+1, and SCR+MBS on CLP=0+1 without tagging
6. PCR on CLP=0+1, and SCR+MBS on CLP=0+1 with tagging.

All of these configurations contain the peak cell rate (PCR) on the aggregate CLP=0+1 cell flow to achieve interoperability with the minimum requirement for ITU-T Recommendation I.371. Configuration 1 provides the check on PCR for the aggregate CLP=0+1 cell flow as the minimum, required configuration. All others are optional. Configuration 2 checks PCR conformance on the CLP=0 and aggregate CLP=0+1 cell flows separately. Configuration 3 is similar to configuration 2 but changes the CLP field from 0 to 1 for cells in the arriving CLP=0 flow that exceed the PCR. An example application for configurations 2 and 3 would be dual-level coded video where the critical information is coded as CLP=0, and the non critical information is coded as CLP=1. Configuration 4 checks the peak and sustainable rate on the aggregate CLP=0+1 flow, discarding any cells which violate either check. Configurations 5 and 6 are similar in form to configurations 2 and 3 as shown in the above chart; however the check is on peak and sustainable rate. The above chart also presents frame relay rate enforcement configuration, based on two leaky buckets for the enforcement of the committed information rate (CIR) and excess information rate (EIR) traffic parameters.

Example: UPC for Frame Relay/ ATM Service Interworking

Option 1

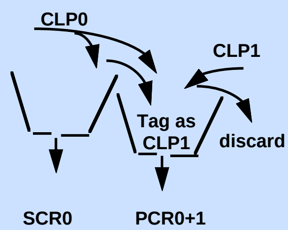

Traffic Parameter Mapping:
PCR $_{0+1}$ = CIR+EIR
SCR $_0$ = CIR
MBS $_0$ = Bc [1/(1-CIR/AR) +1]

Option 2

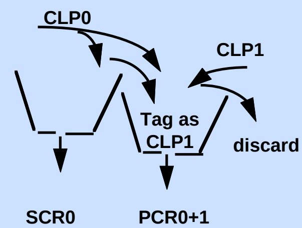

Traffic Parameter Mapping:
PCR $_{0+1}$ = Access Rate (AR)
SCR $_0$ = CIR
MBS $_0$ = Bc [1/(1-CIR/AR) +1]

16

The case of traffic parameter characterization for frame relay/ATM Service Interworking is considered, as an example for the application of the traffic enforcement function. In this case, 1-to-1 mapping between the frame relay connections and ATM connections takes place.

The above chart shows two approaches for the mapping of frame relay traffic parameters to ATM traffic parameters, and the corresponding UPC configurations.

In Option 1 the sum of the CIR and EIR of a frame relay connection is mapped to the PCR 0+1 of the ATM connection. In this case shaping of the traffic flow to ensure that the PCR does not exceed (CIR+EIR) is expected. The CIR is mapped to SCR for CLP=0 while the committed burst size is mapped to the MBS for CLP=0.

In Option 2 the frame relay access link rate (AR) is mapped to the PCR for CLP=0+1, while the CIR is mapped to SCR for CLP=0, and Bc is mapped to the maximum burst size MBS.

Note that in Option 1 explicit characterization of the access rate is not carried out, while Option 2 does not explicitly characterize the EIR rate. In both configurations the fact that the instantaneous rate of the frame relay stream can burst above CIR (up to the access rate AR) is taken into consideration when mapping Bc to MBS.

Traffic Shaping

- **Function**
  - to alter the traffic characteristics of a stream of cells on a VCC or VPC to achieve the desired characteristics
- **Benefits**
  - reduce source burstiness
  - simpler bandwidth allocation approaches
  - guarantees adherence to traffic contract for public UNI
  - compensation for cell delay variation introduced by network

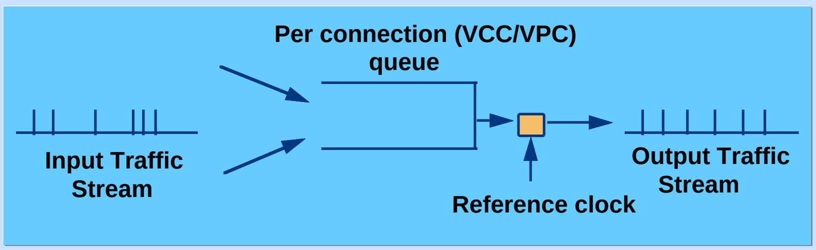

17

Traffic shaping is the management function performed at the user-network interface (UNI) of the ATM network. It ensures that the traffic generated by each source matches the contract negotiated between the user and the network during connection establishment.

Traffic shaping is a method for smoothing out traffic bursts, and is useful for ensuring conformance of transmitted traffic to subscribed traffic parameters, particularly the peak cell rate parameter. It regulates the emission interval of cells transmitted to the network by providing a per connection (per VPC or VCC) buffer and emitting the queued cells at the requested rate. Bandwidth allocation to connections supporting shaping is generally simpler, due to the fact that the cells are admitted into the network at a specified rate.

At the same time, shaping source traffic at the UNI will minimize the chances that a public network would discard user traffic due to non conformance to the agreed traffic characteristics.

Traffic shaping can also be applied on traffic received from the network as a mechanism to absorb the delay variation introduced by the queueing and buffering of cells along the path of the connection.

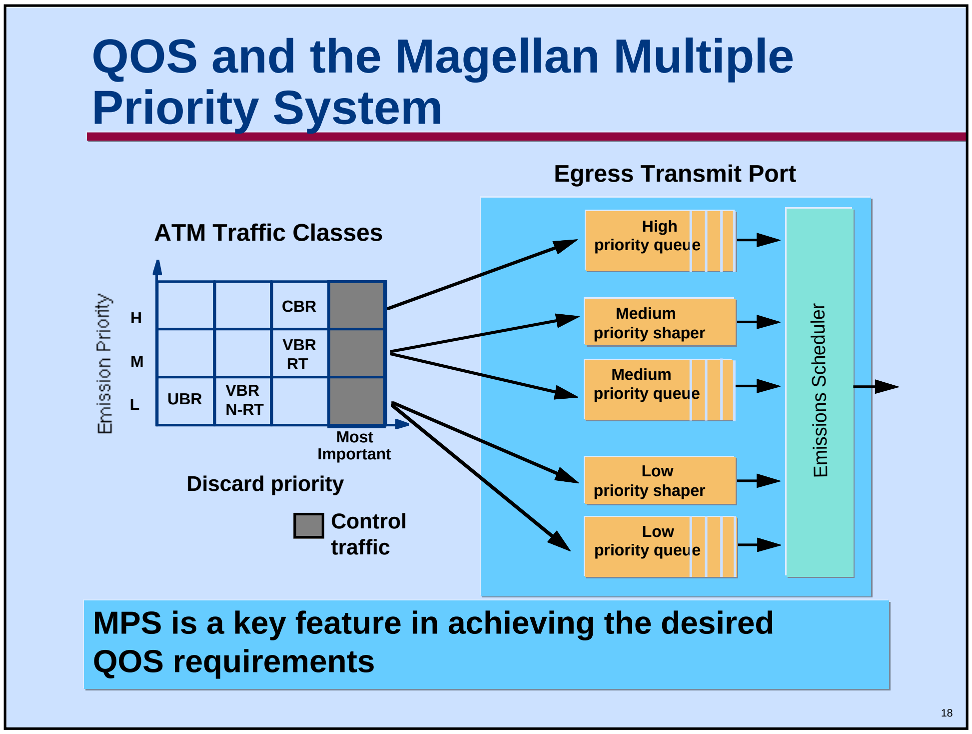

A key point in achieving the desired QOS requirements is the efficient support of multiple emission and discard priority, which allows for the differentiation between the requirements of real-time versus non real-time services. The above figure demonstrates the relationship between the available ATM QOS classes and the multiple emission and discard priorities, as part of the Magellan Multiple Priority System as implemented on Magellan Passport.

ATM services are separated according to their cell delay and cell delay variation requirements to the different emission priority queues, with CBR services having the highest emission priority, followed by real-time VBR, and non-real-time VBR given the lowest emission priority. Also, services requiring traffic shaping are provided with per-connection shaping queues that can be at either the medium and low emission priority. An emission scheduler is responsible for the orderly service of the different emission queues.

ATM services are also differentiated according to the importance of the traffic stream, by providing different discard levels per emission queue. This allows for selective cell discard within each queue, so that cells with less importance are sacrificed first, as congestion builds up. The highest discard priority level is reserved for the critical network and node control traffic. Three emission priorities and four discard priorities are supported.

The clear orthogonality between the emission and discard is a crucial element in supporting multiple QOS simultaneously, in an efficient manner.

Agenda

- ATM Traffic Management
- Bandwidth Allocation and Control
- **ATM Network Engineering: A Case Study**
  - network engineering cycle
  - model network access
  - model network backbone
  - design verification
- Summary

19

After examining key concepts in ATM traffic management affecting the design and engineering of ATM networks, and critical approaches for the allocation and control of bandwidth resources, a case study for the design and engineering of an ATM network is presented.

The typical network engineering cycle is discussed, stressing its iterative nature. The access services are then examined and modeled, based on the user requirements, and community of interest (COI) information. The network backbone is then designed and the ATM bandwidth requirements are specified. The proposed network design is verified to meet the required end-to-end QOS requirements, including robustness to single link failure scenario.

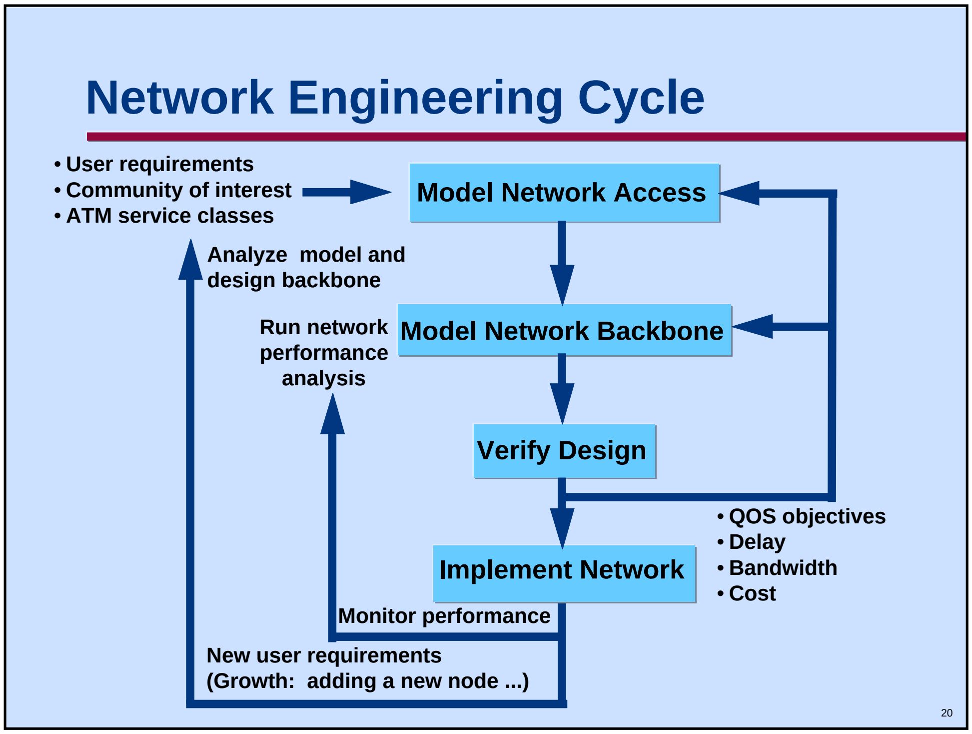

This slide captures, schematically, the network design process. The design process is iterative, and applies equally to new network design and to modification of existing network design.

New networks need to model the various access services and model their impact on the backbone (from a service provider or large enterprise customer perspective). Once the model is satisfactory, the network is implemented and the design must be verified with on-switch performance measurements. Any modifications to the design can then be made.

For existing networks, which have already experienced the growing pains of modeling and verification, a regular (routine) performance analysis and tuning of the network is required. Ongoing performance and design analysis is needed when any major change in the community of interest and/or level of service requirements occurs, such as adding a new node or servicing a new major contract.

Node Locations

The above chart presents the location of the network nodes. This network is based mainly in the US (New York, Chicago, Atlanta, Washington and San Francisco), with extensions in Europe, concentrated through a main network node in London, England.

New York and Washington act as hubs for the trans-atlantic traffic. Magellan Passport is the selected ATM platform to carry the user traffic.

Access Services: Voice Service Requirements

- Voice is treated as a rt-VBR service
- 64 kbit/s coding rate - no compression
- Speech activity detection (SAD) is ON
- 2% call blocking at the PBX
- Offered traffic of four DS1 at each node

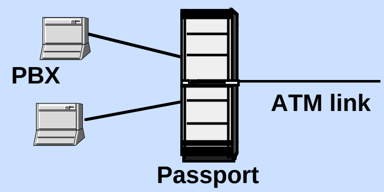

22

The first step in the network design process is to characterize the user requirements in terms of the available access services at the different network nodes, and their quality of service requirements.

In this case study, packetized voice traffic is carried in an uncompressed mode, with speech activity turned ON, in order to benefit from the silence periods in the voice stream to carry data traffic. As the offered traffic volume increases with time, compression of the traffic to 32 kbit/s or 24 kbit/s could be considered.

Call blocking at the PBX or central office feeding into Magellan passport is designed to be 2%. The blocked calls are then directed to the public voice network.

In this network design, voice is treated as an ATM VBR service to utilize the inherent burstiness in the voice traffic stream with SAD by data traffic, resulting in increased bandwidth efficiency.

Video Service Requirements

- Video service is generated as a CBR cell relay stream from the user terminal
- Multiple streams with an aggregate bandwidth requirement of 10 Mbit/s
- Video traffic is transported as a virtual path (VP) connection between source and destination (SF and London) across the Magellan network

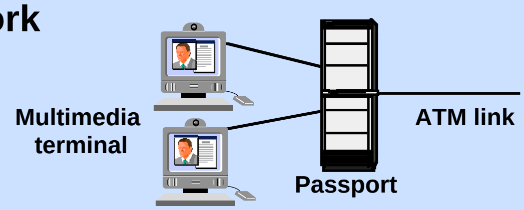

23

A critical user requirement is to carry video traffic resulting form video conferencing applications over the ATM network. The video terminal is equipped with an ATM adaptor card, generating constant bit rate traffic streams from the video conferencing applications.

The aggregate video traffic requirements is 10 Mbit/s of CBR traffic (sum of PCRs), carried by the network as a nailed up virtual path connection between the end nodes, San Francisco and London.

Data Service Requirements

- Frame relay traffic generated by frame relay routers
- Average frame size of 256 bytes
- CIR rate and community of interest (COI) information depend on user population

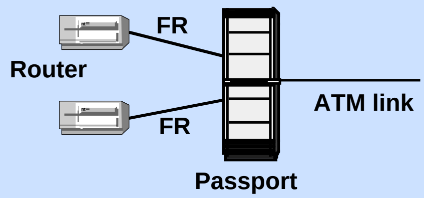

24

Data services for LAN interconnection are offered via routers carrying frame relay service. The average frame size is 256 bytes per frame.

Data services typically exhibit relaxed delay variation constraints, and hence can be classified as a non-real time VBR service.

In general, the frame relay connections’ CIR rate depends on the user population, and so does the community of interest information.

In this example each city’s offered traffic CIR is considered to be 5 T1. To develop the offered traffic values and community of interest information, the cities’ population has been taken into consideration.

- Atlanta 2,737,000
- Chicago 6,216,000
- New York 16,198,000
- San Francisco 5,028,000
- Washington 3,374,000
- London, England 8,620,333

The population information has been extracted from *The Time Atlas of The World*, 9th complete edition.

Bandwidth Allocation for Frame Relay Service

Conservative

1. ATM trunk bandwidth $= \Sigma_{\text{over all access links}}$ Access Rate x overheadA(n)
2. ATM trunk bandwidth $= \Sigma_{\text{over all DLCI's}}$(CIR + EIR) x overheadB(n)
3. ATM trunk bandwidth = fraction x $\Sigma_{\text{over all DLCI's}}$(CIR + EIR) x overheadB(n)
4. ATM trunk bandwidth $= \Sigma_{\text{over all DLCI's}}$ CIR x overheadB(n)
5. ATM trunk bandwidth = fraction x $\Sigma_{\text{over all DLCI's}}$CIR x overheadB(n)

Risky

and

$$PCR_{0+1} = \text{bandwidth} / (53 \times 8) \text{ (cells/sec)}$$

where

n = number of user information bytes in a frame
overheadA(n)= frame relay access overhead x ATMoverhead
= n/(n+5) x ATMoverhead
overheadB(n)=ATMoverhead
fraction = real number less than 1 (inverse of over-subscription factor)

25

The frame relay traffic parameters can be used to determine the equivalent ATM traffic parameters for the case of N-to-1 mapping (over Magellan Passport's ATM logical trunking) as shown above.

Approach #1 is a conservative approach, in which case the traffic smoothing effect of statistical multiplexing of many connections is not assumed, and the worst case allocation is considered. The FR traffic is bounded by the sum of the access link rates, and the corresponding PCR for the ATM VCC is allocated, taking the access link overhead, and the AAL-5 overhead into consideration. In the formula provided above, the access link overhead takes into consideration the frame relay overhead of 5 bytes.

Approach #2 presents a more efficient engineering approach, where a sufficiently large number of frame relay PVCs share the access links, hence smoothing the resulting traffic stream due to statistical multiplexing. The PCR rate of the ATM VCC can be directly related to the sum of the individual FR connections' CIR and EIR rates, after taking the Passport routing/trunking overheads, as well as the AAL-5 overhead into consideration. The bandwidth allocation can be further reduced by assuming not all connections are running at their CIR rates at the same time (Approach #3).

Approach #3 will result in less bandwidth allocated for the FR over the ATM connection. This approach involves considerably more risk than approaches #1 and #2, as it relies on the statistical sharing of the access resources among the frame relay connections.

Approaches #4 and #5 do not explicitly allocate bandwidth to the EIR traffic stream of the frame relay connections, but focus on the bandwidth requirements of the CIR streams.

Depending on the network engineer’s knowledge of the user applications and traffic profile (including average versus peak utilization of access links), an appropriate engineering approach can be selected.

It should be stressed that approaches #3 and #5 involve over-subscribing the FR service over the ATM connections. This approach, although attractive because of its economical advantage, involves potential traffic performance risks that have to be carefully balanced by the network engineer.

In this network engineering example, bandwidth allocation for frame relay connections is based on approach #4. For 256 bytes frames, the overall overhead including Passport trunking/routing, ATM AAL-5 overhead (trailer and padding), as well as the ATM layer overhead (5 bytes per cell) is 62 bytes or 24.2% of the user information.

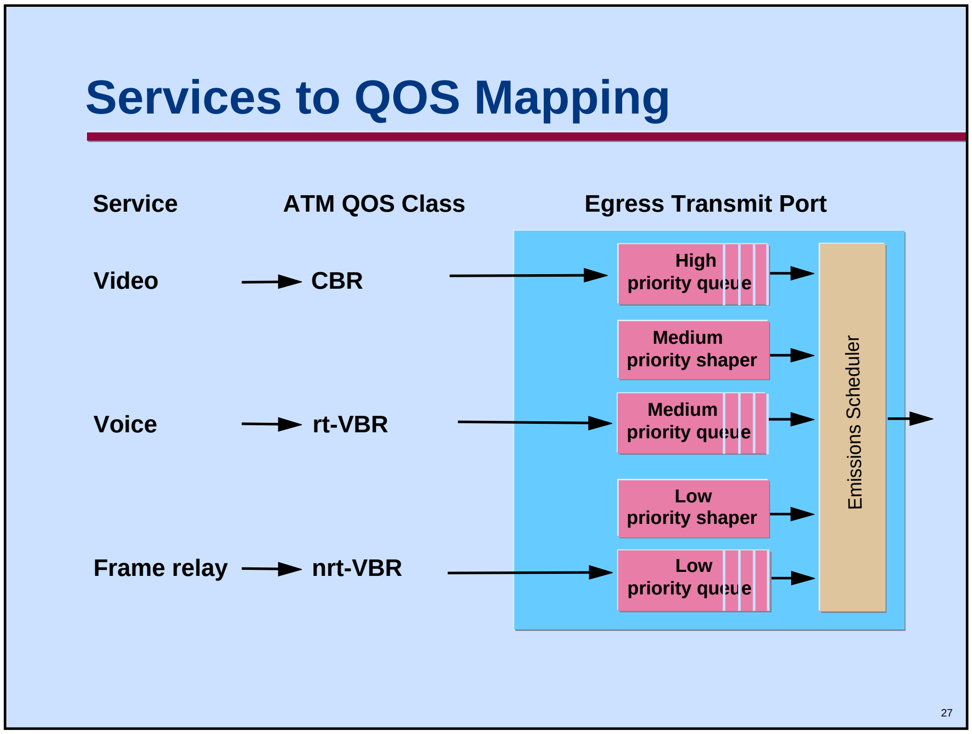

The above chart presents the selected mapping of user services (video, voice, and frame relay) to ATM QOS classes and from there to Magellan's MPS system as supported on the Passport ATM platform.

As video service is generated as a constant bit rate traffic stream with tight cell delay variation constraints, it is given the CBR QOS class and directed to the high emission priority ATM queue on the egress port. Voice with SAD has cell delay variation and cell loss constraint that are more relaxed than those for the CBR stream. Hence voice traffic is given the rt-VBR service class and directed to the medium emission priority queue.

Frame relay service is considerably tolerant of delay variation, and is hence directed to the low emission priority queue on the egress port.

In general, the mapping of the user services to an ATM QOS class is flexible in order to accommodate the specific user's networking requirements.

Community of Interest - Voice

- **Community of interest (COI) information for voice service is as follows (number of channels)**

	Atlan	Chicag	NY	SF	Wash	Lond
Atlan	36	12	12	12	12	12
Chicag	12	36	12	12	12	12
NY	12	12	36	12	12	12
SF	12	12	12	36	12	12
Wash	12	12	12	12	36	12
Lond	12	12	12	12	12	36

- **Around 30% of offered voice traffic is local, while 70% is external**

28

For the case of voice service, the offered voice traffic at each network node is 4T1 (i.e. 96 channels). Of the offered voice traffic, 30% is local, while 70% is external, i.e. distributed among the other nodes in the network. For the sake of this case study, the external voice traffic is assumed to be divided equally among the destination nodes in the network.

The above table shows the COI table for voice traffic.

Community of Interest – Frame Relay

- Community of Interest (COI) information for frame relay service is as follows (CIR in kbit/s)

	Atlan	Chicag	NY	SF	Wash	Lond
Atlan	2560	768	2048	640	512	1152
Chicag	415.1	2560	2214	692	553.5	1245.4
NY	568.9	1137.8	2560	948.15	758.5	1706.7
SF	404.2	808.4	2155.8	2560	539	1212.6
Wash	393.8	787.7	2100.5	656.4	2560	1181.5
Lond	349.1	698.2	1861.8	581.8	465.5	2560

- 30% of offered frame relay traffic is local, while 70% is external

29

For the case of data (frame relay) service, the offered frame relay traffic CIR at each network node is 5T1. Of the offered frame relay traffic 30% is local, while 70% is external, i.e. distributed among the other nodes in the network. The distribution of the external frame relay traffic is selected to be proportional to the destination population;

i.e. the traffic from Atlanta to Chicago =

external traffic generated by Atlanta x (population of Chicago/overall population excluding Atlanta).

The above table shows the COI table for the frame relay traffic. Note that this COI distribution will result in asymmetric bandwidth requirements between networks nodes.

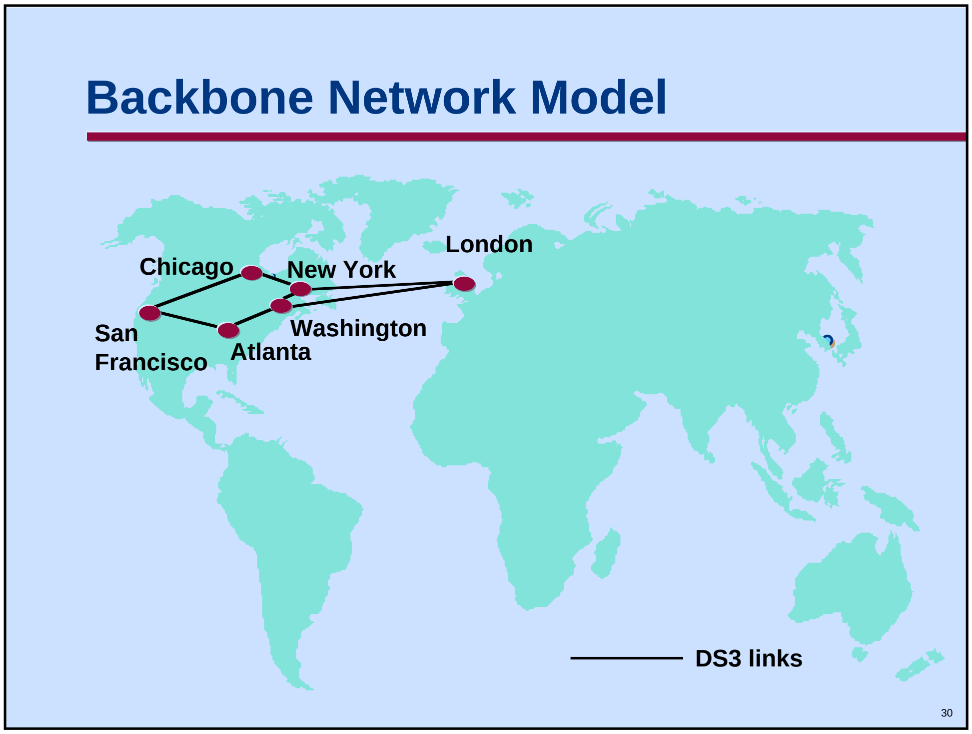

The above chart presents the network topology. In this case, each city is connected to its two closest neighbors. Such a design results in a ring topology connecting the network nodes in the US.

London is connected to both Washington and New York for reasons of reliability and survivability.

Due to tariff considerations, the available link capacity is DS3 for trunking between nodes in the US as well as between the US and London.

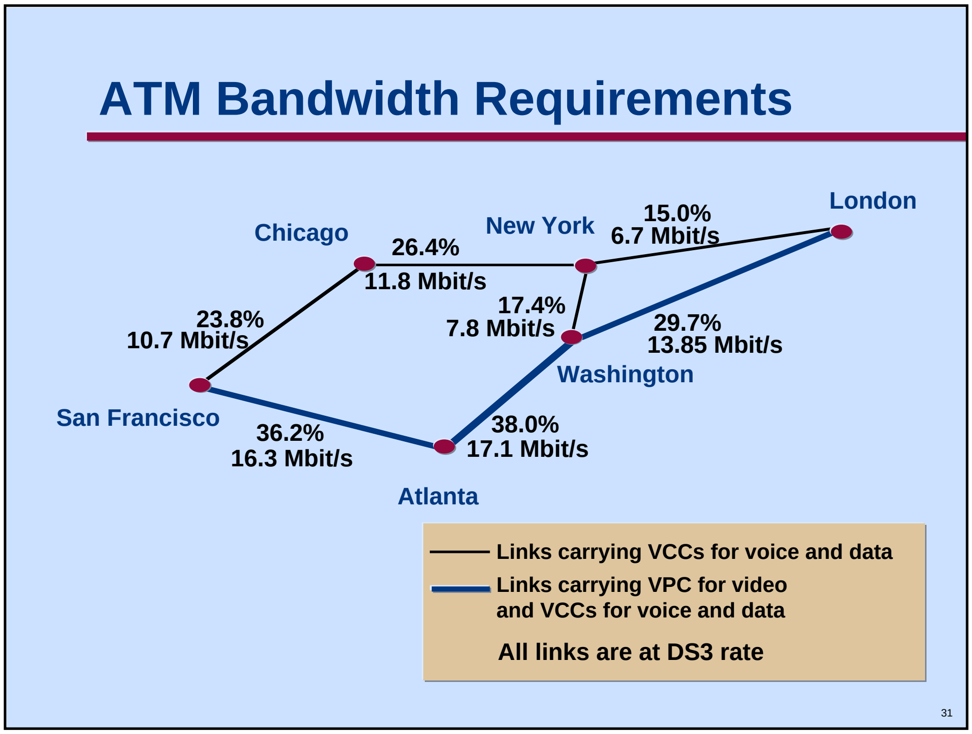

After determining the user traffic requirements, the QOS requirements and the COI information for the different services, we are ready to evaluate the ATM bandwidth requirements to carry the offered services.

The first step here is to examine the bandwidth requirement for the virtual path (VPC) connection, in particular because it is carrying a CBR QOS traffic stream. The bandwidth requirement of 10 Mbit/s of the VPC connection is reserved along the selected path of the connection (San Francisco-Atlanta-Washington-London).

Then the rest of the offered traffic from voice and data services over ATM is overlayed on the network topology and the resulting bandwidth requirements are as shown in the above chart.

Note that voice and (teleconferencing) video services will result in symmetric bi-directional bandwidth requirements. On the other hand, frame relay will result in asymmetric bi-directional bandwidth requirements . The figures shown on the network chart correspond to the higher bandwidth requirement of both directions, as this is the determining factor in the suitability of the proposed DS3 link capacity.

The above design shows that the lower half of the ring topology is the more heavily used path, with the supported VPC connection contributing significantly to this utilization. At the same time, the link utilization does not exceed 50% on any of the links, leaving room for the extra bandwidth in the ring to carry over additional bandwidth rerouted due to link/node failure or increase in offered traffic demand.

Design Verification

City 1	City 2	mi	Data Total Delay ms	Data Queueing Delay ms	Voice Total Delay ms
Atlanta GA	Chicago IL	3991	31.0	0.09	30.92
Atlanta GA	London	4207	32.7	0.09	32.58
Atlanta GA	New York NY	746	6.0	0.09	5.96
Atlanta GA	San Francisco CA	2136	16.6	0.08	16.55
Atlanta GA	Washington DC	542	4.4	0.08	4.28
Chicago IL	Atlanta GA	3991	31.0	0.08	30.92
Chicago IL	London	4172	32.4	0.07	32.31
Chicago IL	New York NY	712	5.6	0.07	5.58
Chicago IL	San Francisco CA	1855	14.4	0.06	14.38
Chicago IL	Washington DC	916	7.3	0.07	7.26
London	Atlanta GA	4207	32.6	0.08	32.58
London	Chicago IL	4172	32.4	0.07	32.31
London	New York NY	3460	26.8	0.06	26.73
London	San Francisco CA	6027	46.7	0.07	46.69
London	Washington DC	3665	28.4	0.08	28.30
New York NY	Atlanta GA	746	6.0	0.08	5.96
New York NY	Chicago IL	712	5.6	0.06	5.58
New York NY	London	3460	26.8	0.06	26.73
New York NY	San Francisco CA	2566	20.0	0.07	19.96
New York NY	Washington DC	204	1.7	0.06	1.68
San Francisco CA	Atlanta GA	2136	16.6	0.08	16.55
San Francisco CA	Chicago IL	1855	14.4	0.07	14.38
San Francisco CA	London	6027	46.8	0.08	46.69
San Francisco CA	New York NY	2566	20.0	0.08	19.96
San Francisco CA	Washington DC	2678	20.9	0.09	20.83
Washington DC	Atlanta GA	542	4.3	0.08	4.28
Washington DC	Chicago IL	916	7.3	0.07	7.26
Washington DC	London	3665	28.4	0.08	28.30
Washington DC	New York NY	204	1.7	0.06	1.68
Washington DC	San Francisco CA	2678	20.9	0.09	20.83

32

After evaluating the bandwidth requirements of the offered traffic on the proposed network topology, the resulting design has to be verified to meet the expected quality of service requirements for the different services carried by the ATM network.

Using an internal tool, the resulting average data (frame relay) mesh delay is 19.6 msec. The above table shows that the average path delay does not exceed 47 msec, and it is dominated by the propagation delay component. The paths that introduce the highest delay, as highlighted in the above chart, are Chigaco-London, and London-San Francisco. Due to the relatively moderate traffic loading on the DS3 links, the delay introduced by queueing in the network nodes is around 0.1 msec.

For voice traffic, the resulting average mesh delay is 19.5 msec with the average path delay not exceeding 47 msec. The voice average delay is also dominated by the propagation delay component.

As voice traffic is carried as a rt-VBR service, with certain cell delay variation constraints (around 30 msec of end-to-end CDV), it is essential to verify that the network design meets these CVD requirements. The end-to-end delay variation for the offered voice traffic on the longest path (between SF and London) is below 0.8 msec at the 10^{-6} percentile of the end-to-end delay distribution. A multi-node network model has been used, based on GI/D/1 queueing model for each node, to obtain the end-to-end delay distribution, taking into consideration the first two moments of the voice cell inter-arrival time.

The nodal cell loss ratio for the voice traffic is below 10^{-6}, and for the CBR traffic stream the nodal cell loss ratio is below 10^{-10}

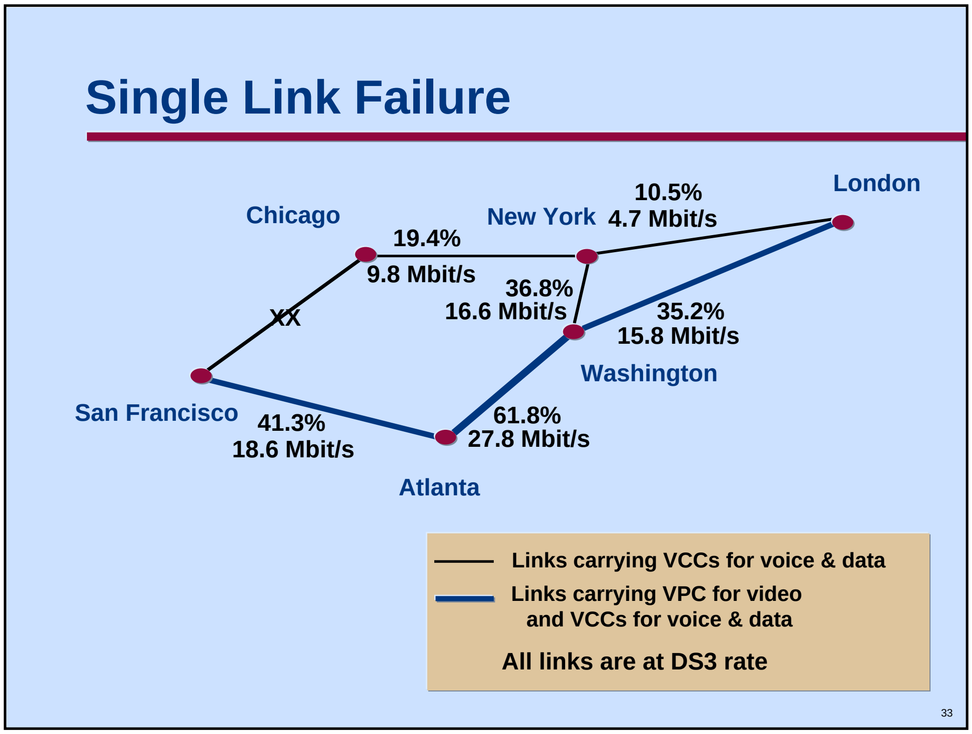

The network design should also be verified to accommodate single link failure scenario.

To illustrate this point, the link between San Francisco and Chicago has been disabled, resulting in re-routing of the network traffic around the failed link. The link utilization on the Atlanta-Washington link has increased to around 62%, which is still a relatively conservative utilization. At the same time, the utilization of Chicago-New York link has dropped slightly becuase of the diversion of the traffic originating from San Francisco.

The average path delay on the San Francisco-London path has increased to around 49 msec, which is still within acceptable bounds.

Agenda

- ATM Traffic Management
- Bandwidth Allocation and Control
- ATM Network Engineering: A Case Study
- **Summary**

34

The following chart summarizes the key points covered in this presentation.

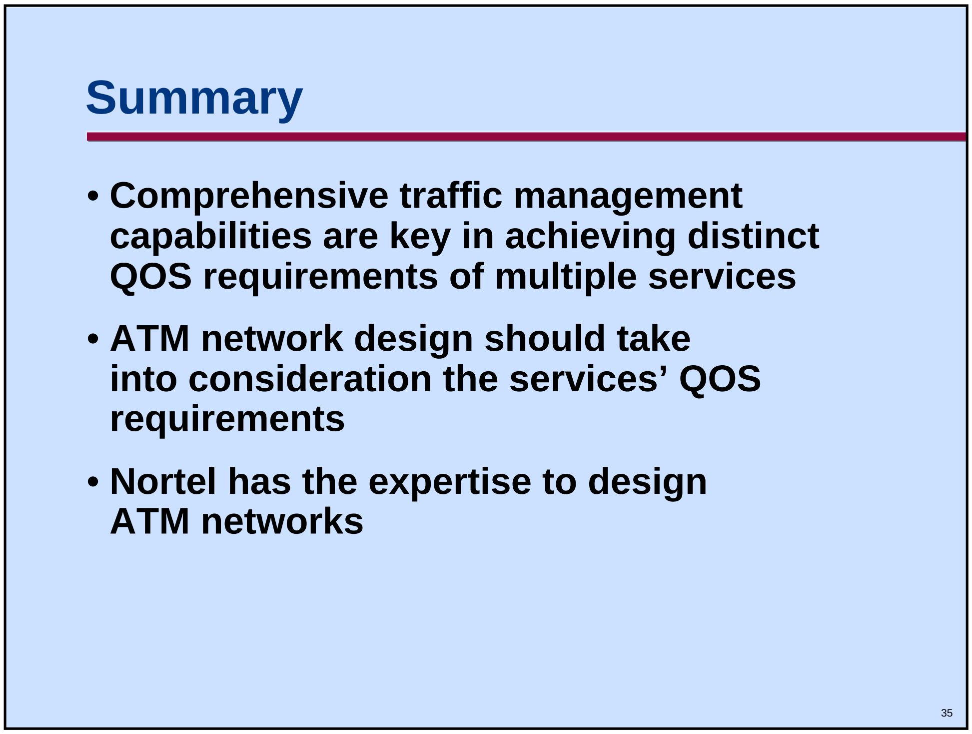

Comprehensive traffic management capabilities are key in achieving distinct QOS requirements of multiple services supported by a common ATM network.

Proper engineering of ATM networks aims at meeting the different quality of service (QOS) requirements of various service classes such as real-time services (i.e. packetized voice, video), and non-real time services (i.e. file transfer, and various data applications), while efficiently and cost-effectively utilizing the network resources.

Nortel has the expertise and tools to work jointly with our customers in designing efficient and reliable ATM networks that meet the end-users' requirements.

Appendix A: Terminology

ABR	available bit rate
ATM	asynchronous transfer mode
BT	burst tolerance
CAC	connection admission control
CBR	constant bit rate
CDV	cell delay variation
CDVT	cell delay variation tolerance
CER	cell error rate
CLR	cell loss ratio
CLP	cell loss priority
COI	community of interest
GOS	grade of service
ITU-T	International Telecommunication Union - Telecommunications
MBS	maximum burst size
MPS	Multiple Priority System
PCR	peak cell rate
QOS	quality of service
SAD	speech activity detection
SCR	sustained cell rate
UBR	unspecified bit rate
UNI	user network interface
UPC	usage parameter control
VBR	variable bit rate
rt-VBR	real-time VBR
nrt-VBR	non real-time VBR
VCC	virtual circuit connection
VCI	virtual connection identifier
VPC	virtual path connection
VPI	virtual path identifier