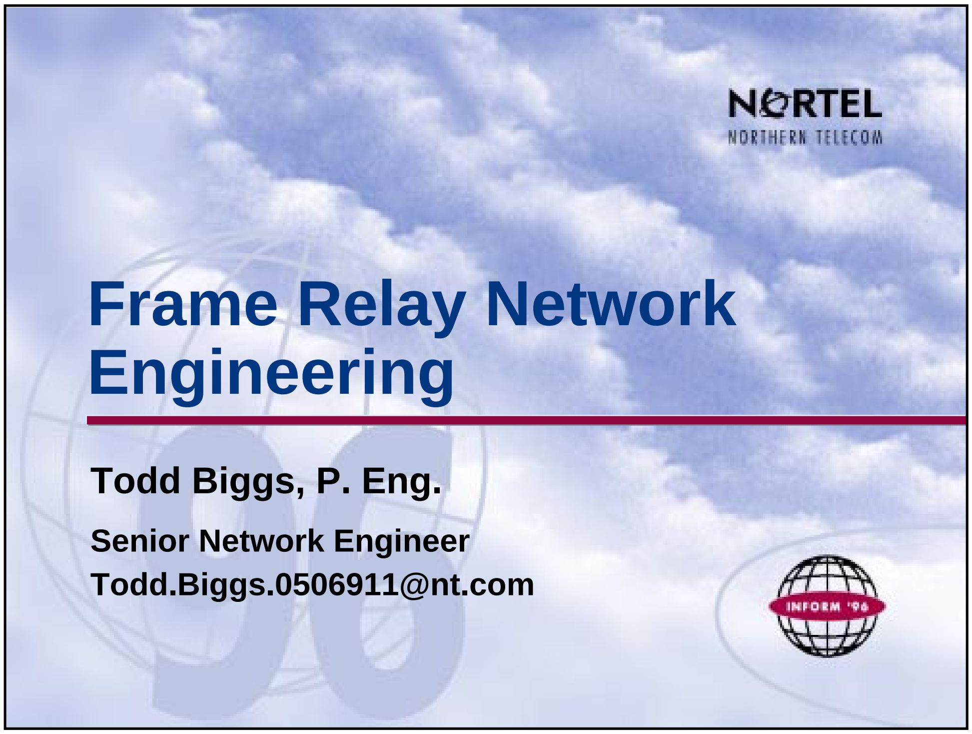

The objectives of this session are to review frame relay engineering basics and to present the role and impacts of congestion management, oversubscription and NNI implementation in network design.

The value of frame relay in today's networks is highlighted along with the parameters and considerations to properly engineer the access and backbone network. A basic foundation for designing frame relay networks is presented along with some advanced concepts. Particular attention will be given to congestion management, oversubscription, resiliency, and implementing NNI. Customers that are about to deploy frame relay networks will gain valuable insight into the design considerations and their impacts on networks. Customers with existing networks will acquire an understanding of features and techniques to better leverage their networks.

About the presenter:

Todd Biggs is a member of the Magellan Network Engineering Solutions group. His current responsibilities include designing and developing engineering guidelines for Magellan Passport and its features. Before joining Nortel, he held engineering positions in a satellite communications company and at international and domestic carriers. He has more than five years of experience designing networks for a variety of major private and government clients.

Agenda

- **The Frame Relay Explosion**
- The Network Design Process
- Frame Relay Network Design Parameters
- Frame Relay Access Design Parameters
- Magellan Congestion Management
- Frame Relay NNI
- Case Study
- Summary

2

The Frame Relay Explosion

- Increased peer-to-peer LAN traffic over the WAN
- Multiplexed mesh topologies cost effective
- Multiple transport options and bandwidth accounting options
- Frame relay is widely available and accepted
- Frame relay is a necessary step for multimedia, transition to ATM

3

Frame relay is an interface protocol that can provide substantial cost and performance benefits for bursty traffic such as LANs. Recent drivers include SNA as well as emerging multimedia applications.

Frame relay growth continues to be explosive and is not anticipated to slow down until at least 1998. The projected Compound Annual Growth Rate (CAGR) for 1994-1998 is 90%. U.S. frame relay service revenue passed $0.5 billion in 1995. European/Asian service revenue will pass $0.5 billion in 1996 representing 107% growth over 1995. (Source: Vertical Systems Group).

When frames are carried in native mode, it is more bandwidth efficient than ATM because it has less overhead. Many types of existing user equipment (e.g. routers) can be upgraded to frame relay without hardware changes. Frame relay is a viable and efficient WAN protocol for the following reasons:

- Increased peer-to-peer LAN traffic requires the availability of a high-speed, low-overhead service
- Relatively inexpensive way to interconnect multiple LANs when compared to leased lines, yet it provides good performance for LAN applications
- Resilient to failure when properly provisioned
- There is a real need for multiple priority transport options. Magellan frame relay has made practical use of service subscription and traffic class parameters
- Carriers can maintain account control and retain customers with a frame relay offering
- Frame relay is a must in today’s networks as an access vehicle to carry data over ATM. The frame relay adaptation to ATM will hasten the adoption and implementation of multimedia networks and applications. The reason is simple, no-one wants to forklift their existing infrastructure of terminal equipment

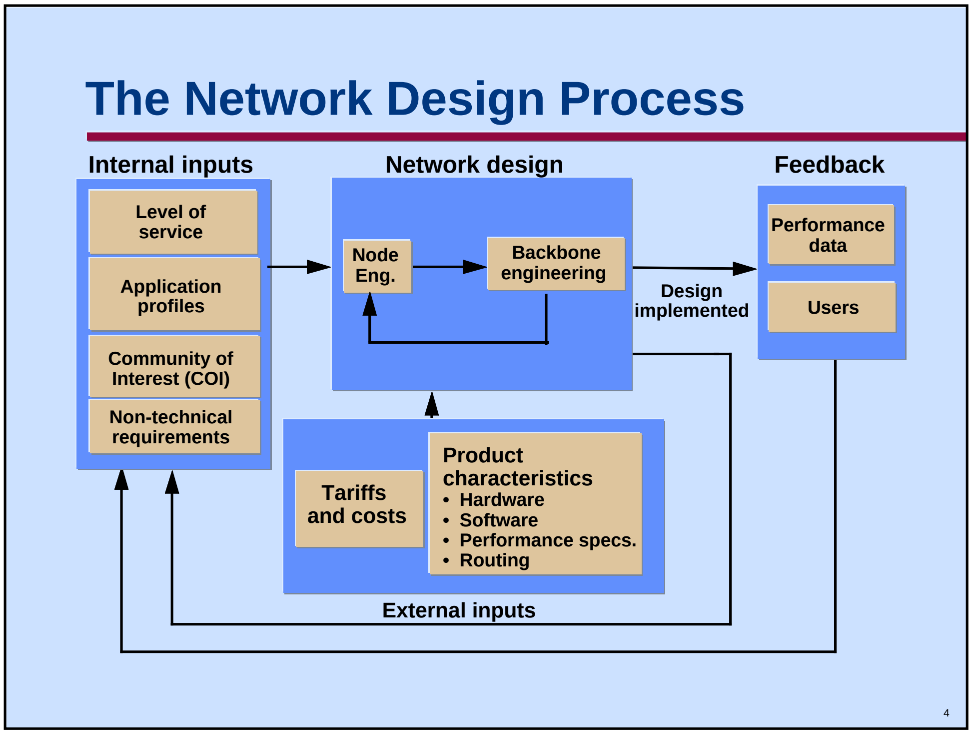

This slide captures, schematically, the network design process. This process applies equally well to new networks as well as to existing networks. Note, first of all, that it is a cycle with no end. The most successful network managers treat network engineering as an ongoing process, as opposed to something that is done once every couple of years.

New networks need to model the various access services and model their impact on the backbone (from a service provider perspective). Once the model is satisfactory, the network is implemented and the design is verified with on-switch performance numbers. Any modifications to the design are then made.

For existing networks, which have already experienced the initial modeling and verification cycle, 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).

Internal inputs are defined as inputs to the network design process which you as a network provider can determine and/or calculate. **External inputs** are inputs to the network design process which are given to the network provider or over which the network provider has little or no control.

The **network design** function can be defined as a series of calculations and decisions regarding network configuration based on internal and external inputs.

The **feedback** loop involves the gathering of performance data of the network to determine whether or not level of service objectives are being met. Based on this data, changes to assumptions/decisions about the inputs and the network design may be needed (Refer to Doug Bundgaard’s presentation “Using Statistics to Plan Your Network”).

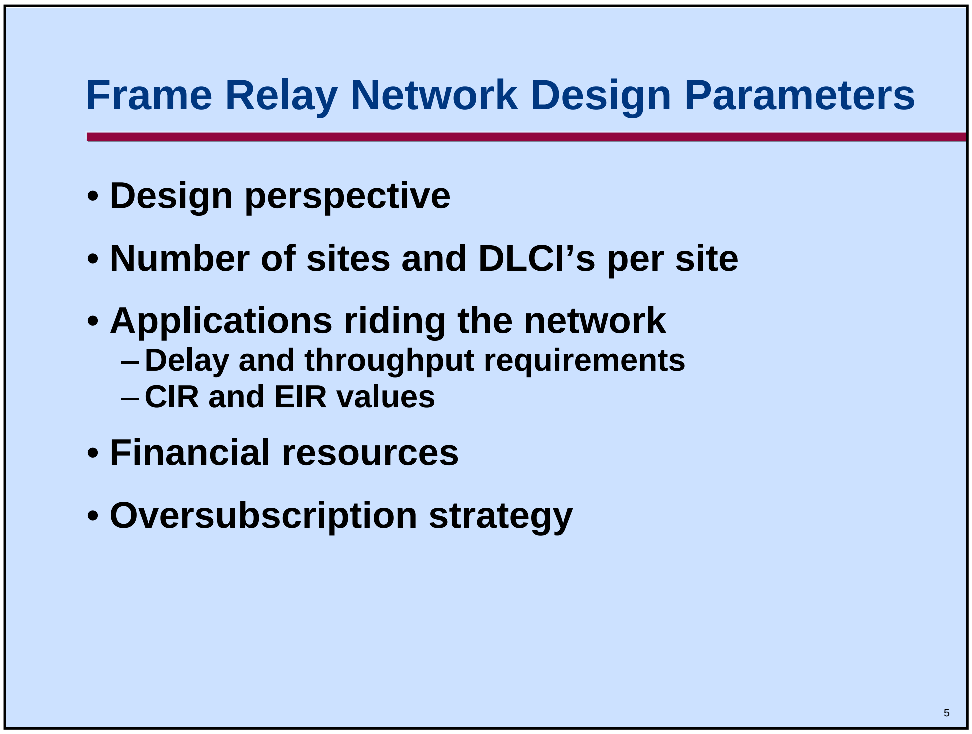

Depending on where you are looking at the frame relay service, you will see different solutions for the same problem. For instance, the network service provider would like to use the minimum bandwidth that would transport his customers' data (basically oversubscribe CIRs). In an enterprise network, the same professional does all network functions, and it is common not to use either CIR and EIR in frame relay service engineering. It is important, however, to perform calculations both as the user and as the service provider.

The number of DLCIs (or distinct frame relay connections) at each site is a known entity. It is important, however, to ensure that DLCI growth can be smoothly accommodated by the frame relay equipment.

The applications using the network will have service requirements and impact CIR and EIR selection. A basic requirement in all telecommunications is how long we can or will wait for data transport. Delay requirements will vary with the environment and the transaction type. Interactive applications like SNA require low delay while other applications like file transfer can tolerate more delay. CIR and EIR depend on your perspective and the financial aspects.

Finally, an oversubscription strategy can be used to maximize the efficiency of the network.

Oversubscription

- Objective to minimize backbone bandwidth required to carry user data
- Requirements have statistical distribution
- Congestion and traffic management impacts
- Typically use values of 3.5 to 8

6

Frame relay service providers have the challenge of estimating the actual backbone bandwidth used, compared to what is subscribed. This leads to the practice of oversubscription. When this strategy is adopted, it is important to ensure that any frame relay equipment implementation is able to effectively handle congestion and traffic management.

The backbone bandwidth requirements are not necessarily an addition of the access traffic. Frame relay backbone traffic, similar to voice (Erlang estimations) and ATM (Bernoulli distribution), requires less bandwidth than the sum of the access speeds. The ratio of subscribed CIR to the allocated backbone bandwidth is known as oversubscription.

Consider the example of service provider ABC Tel offering a frame relay service. ABC charges on the basis of CIR subscribed, plus line access fees.

ABC needs to allocate backbone bandwidth for its customers. To benefit from the varying customer requirements, ABC assumes an over-subscription value based on the sum of CIR purchased by their customers. If one assumes the following:

- sum of subscribed CIR for all DLCIs between two nodes is 7 Mbit/s
- five is the over-subscription value used by ABC for their network design
- average frame size for all FRS traffic is 256 bytes, and 17 bytes overhead

then, ABC would assume a bandwidth utilization of {(7Mbit/s)*(1/5)*(256+17)/256= 1.49Mbit/s} approximately 1.5Mbit/s on the trunks along the data path of the traffic. Thus, a trunk size can be selected to accommodate this offered traffic.

Frame Relay Access Design Parameters

- Study applications
  - identify the required traffic classes
    - delay and throughput requirements
    - frame loss tolerance
  - determine the line speed
  - frame size is dependent on application
- Assign traffic profiles for each DLCI
  - decide on CIR and EIR
  - verify trunks have sufficient bandwidth for the frame size

7

When designing the access portion of a network, one must study the applications and identify the traffic classes that are required and assign the appropriate traffic profiles.

Identify the required traffic classes:

How much frame loss can the application tolerate? The protocols best suited for frame relay are those that can adjust to the available bandwidth by dynamically varying the window size.

Frame relay line speed:

A major advantage of frame relay is the ability to minimize serialization delays at network ingress and egress without increasing backbone network bandwidth.

If the Passport node is collocated with the router that it is to service, then always configure the maximum line speed possible for either the router or Passport. Use CIR to control the actual flow of data. If the Passport is not collocated, then line speed is dictated by economics and application requirements.

Note: It is recommended that a specific access line speed (plus overhead) not exceed trunk speed.

Frame size:

For Passport only networks, this number can be increased to 4,096 bytes and the Passport data path is optimized to use 4K byte frames. If a smaller subnet frame size is set, then additional execution of instructions are required by the processor.

Frame size impacts the percentage overhead and the total user bandwidth available. This may lead to increased processing requirements.

Assign Traffic Profiles per DLCI

- Mission critical traffic: CIR = peak
- Interactive traffic: CIR = averaged burst
- Background traffic:
  - CIR = 0
  - EIR = average bandwidth required
- File transfers:
  - CIR = Avg. file length (kbytes)/acceptable transfer time
  - EIR = (Peak - CIR)

8

In order to assign the appropriate profiles, the characteristics of the applications must be understood. One must determine the importance of the data, the transaction profile, and the level of burstiness. With these inputs, the engineer can select the CIR value.

Selecting the CIR requires balancing the cost of the service with the performance of the applications. The engineer wants the CIR low to save costs, but high enough to ensure the efficient transport of all data.

Another consideration in selecting the CIR is the equipment's interpretation of CIR. Switching equipment monitors the throughput of the applications to protect CIR to all users. Varying the period over which the CIR is monitored becomes increasingly important.

Some sample guidelines for the selection of CIR are given in this slide. Note that the above points are only examples. Many other factors can affect the CIR assignment. It is ultimately the engineer's decision and choice to set the CIR values. There are no hard and fast rules and, consequently, no absolute right or wrong answer.

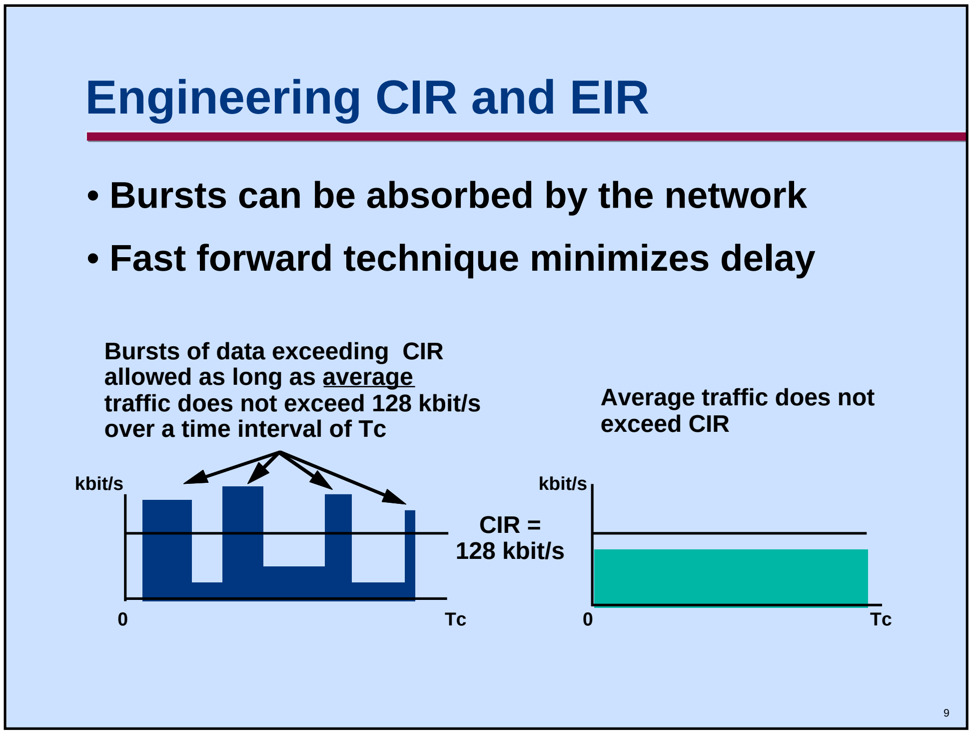

The committed information rate (CIR) and the excess information rate (EIR) are defined as follows:

CIR = Bc/Tc and EIR = Be/Tc

where Bc is the committed burst size, Be is the excess burst size and Tc is the measurement interval. Setting Tc =1 makes it easier to understand how the service is configured because CIR = Bc.

Set the CIR to be equal to the peak 15 minute traffic requirement of the end service. However, if this is not economically feasible, the lowest rate that can be tolerated by the application and/or users should be used.

For any particular DLCI, the CIR should never be greater than the trunk speed. If this guideline is not observed, then undue congestion may occur at the trunk FP.

CIR oversubscription, where the total of all CIRs is greater than trunk capacity, should be approached with caution. The amount of oversubscription should be based on the customer's traffic patterns, number of DLCIs, and other factors.

In addition, do not set the committed burst size Bc to be less than the maximum frame size or less than maximum window size of the protocol using the service. In the first case, setting Bc to a value less than the frame size would lead to rate enforcement of every frame. The second case occurs for an application riding over a poor protocol implementation (no selective re-transmits or dynamic window adjustment) that could feed back-to-back frames to the service. Rate enforcement would start before the last frame is transmitted. This means the whole window is sent again and the transfer fails again.

Magellan Congestion Management

- Overall strategy
- Traffic slasses and quality of service
- Magellan dynamic routing
- Loadsharing and loadspreading
- Failure recovery
- Impacts of oversubscription

Congestion management simplifies engineering

10

Inherent in the design of Magellan Passport is the philosophical goal to deliver data whenever possible and as fast as possible rather than discard the data or delay its transmission through the network. Passport has been architected for bursty traffic while providing high capacities and performance with superior scalability, key elements for a service provider to maximize revenue. Delivering all offered traffic in a fair manner requires a comprehensive traffic management system.

Magellan provides node level, network level and engineering applications to avoid and respond to congestion. Mechanisms include:

- System-wide implementation of MPS™ (Multiple Priority System) with multiple traffic classes for emission and discard priorities to differentiate traffic
- Rate enforcement to ensure equitable access to CIR and EIR subscriptions
- Fast forward feature to improve throughput
- Closed loop congestion feedback to move discards to the source where necessary. Implementation of the A-bit (warn and discard) to avoid carrying traffic unnecessarily with discards at the source for unavailable PVCs
- Signalling FECN/BECN to alert end devices
- Robust dynamic packet routing with a self-learning topology and support for delay and throughput classes of service to provide differentiated services. Passport overflow routing to divert traffic during peak conditions
- X-PLORER and X-AMINER tools for network planning and performance analysis to assist in proper network engineering

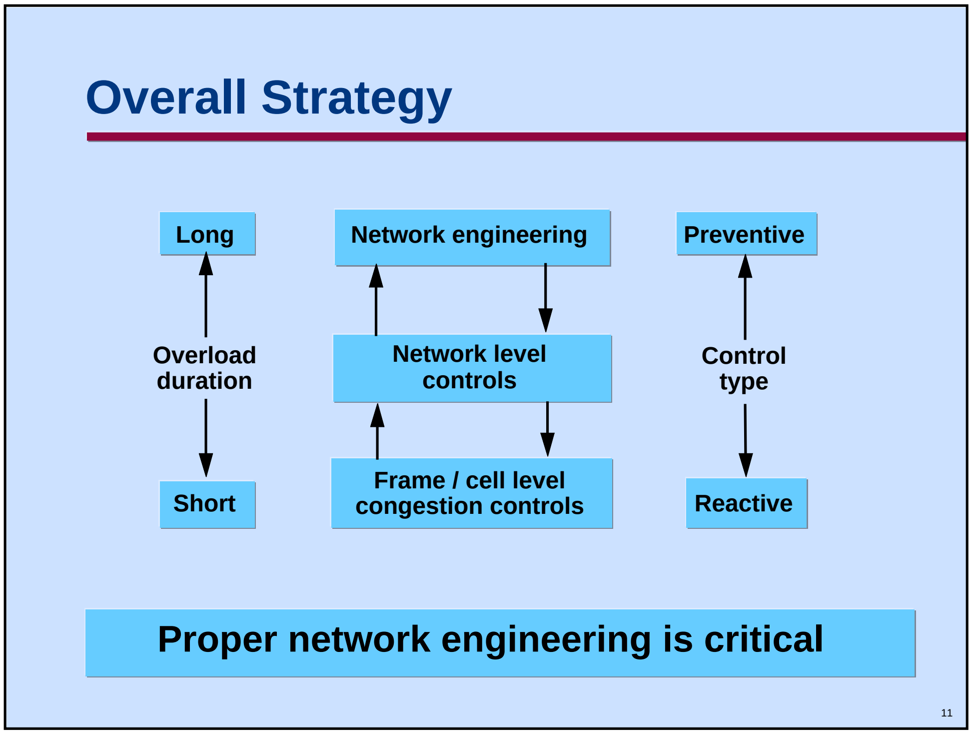

The Magellan congestion control strategy allows for both preventive and reactive controls. The objective is to minimize the occurrence of congestion and enable successful recovery from the various types of congestion (i.e. transient vs. sustained congestion). The three-level hierarchical control strategy provides effective controls:

- Network engineering mechanisms are preventive and include network dimensioning, setting the CIR for various connections, and assigning traffic priorities. It also involves an on-going process of monitoring network performance and making the required adjustments. The network generates detailed statistics, and the management system includes intelligent tools that assist with these tasks (see also “Closing the Loop on Planning and Analysis” by Doug Bundgaard).
- Network level controls require a global knowledge of the network and the co-operation of multiple networking nodes. They are most suitable for sustained kinds of congestion that last one round trip delay or longer. Some examples include routing features, overflow routing, and the closed-loop congestion control mechanism.
- Frame/cell level controls provide the last tool for fighting congestion and preserving the quality-of-service for the different services. Unlike previous mechanisms, these controls take effect immediately, and are suitable for recovering from transient congestion conditions. Some examples of frame/cell real-time controls are congestion notification, priority discarding as network queues start building up, diverting connectionless traffic away from congested links.

A properly engineered network minimizes potential for discard while providing maximum performance and service levels to subscribers.

Traffic Classes

- Select traffic handling preferences
  - urgency → emission priority
  - importance → discard priority
  - reliability → overflow routing
- Traffic classes

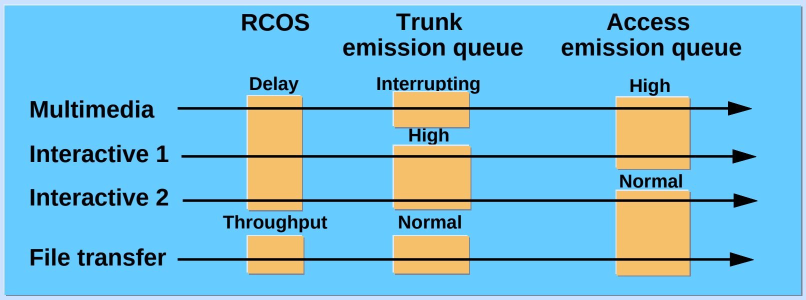

12

The traffic classes are subscriber options that allow the user to select the traffic handling preferences desired for different traffic types and utilize the most appropriate network resources. Multiple classes of service can be defined by setting a combination of urgency, importance and reliability. These parameters are individually provisionable on a per DLCI basis to match application needs and provide differentiated services.

Selection of a traffic class specifies both the trunk queuing priority and the FR service emission priority. The traffic class is also used by the packet forwarding and trunking system to make sure the network uses the appropriate type of transmission facility and preferred links within trunk groups, etc.

To define the **urgency** characteristic of traffic, Magellan offers an emission priority to differentiate delay sensitive traffic from normal traffic. The user can choose between two distinct traffic sensitivities: either delay or throughput. Selecting either class of service affects both the transmission path and the order of traffic transmission and is implemented network wide.

Importance is determined by multiple discard priorities. Upon congestion, the congestion management system intelligently determines which frames may potentially be discarded for various levels of congestion. Network discard priority is only examined within the subset and has no bearing on CIR/EIR rate enforcement which is performed at the access (if so subscribed).

Services can be assigned a **reliability** level of high (overflow onto alternate routes during congestion) or normal (no rerouting). The ability to overflow will better utilize existing bandwidth. Since facilities are purchased or provisioned in steps, there is almost always some extra capacity due to step size, in addition to room left for growth, etc.

Traffic Classes (notes)

- Notes continued

13

There are three egress queues at the trunking level which are normal priority, high priority, and interrupting priority queues. These queues are implemented in the hardware for high-performance operation. Interrupting priority is available to support a multimedia class of service for frame relay. Each egress queue has a series of thresholds which correspond to increasing levels of congestion. In this manner, one is able to discard traffic based on defined classes of service and the level of congestion which is being experienced. The closed loop congestion feedback mechanism, which is implemented as a system throughout the network, ensures fair access to available resources during periods of congestion.

In addition to the current delay and throughput traffic classes, a multimedia traffic class will be available. The multimedia traffic class is intended to minimize delay and delay variance for isochronous traffic carried across a frame relay network. The new traffic class, provisionable per DLCI, maps to the interrupting trunk priority across the backbone and the high emission priority across the destination frame relay link. This new traffic class delivers improved quality of service for multimedia applications over frame relay, by taking advantage of Passport's unique frame/cell trunking and Multiple Priority System.

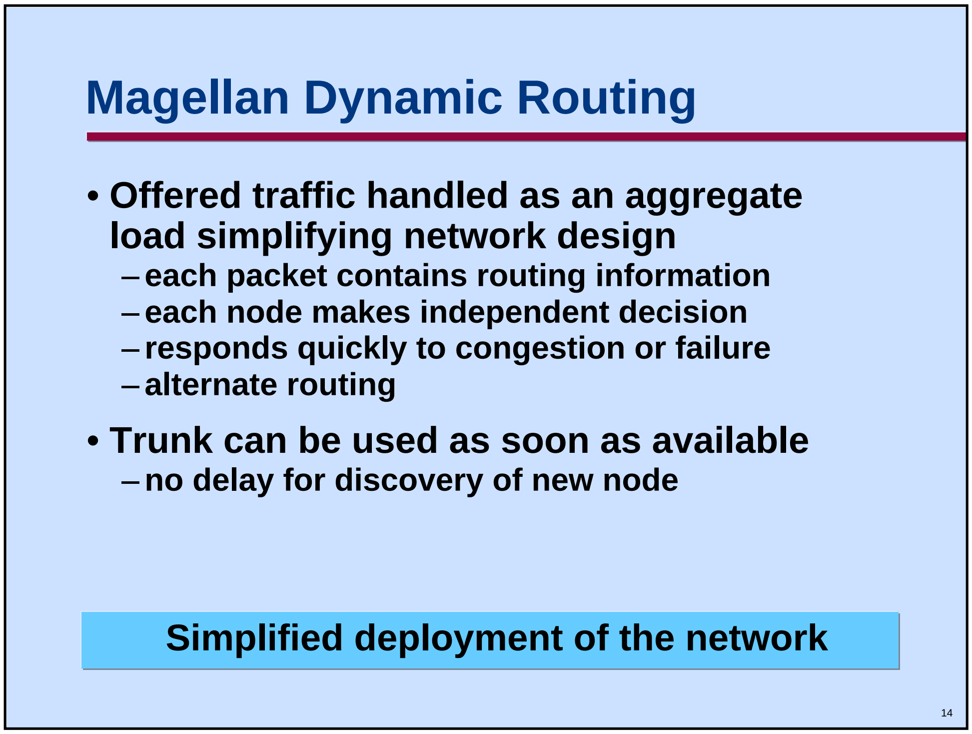

Frame relay uses Passport's dynamic packet routing system (DPRS) which is optimized for bursty data traffic such as frame relay and packet data. This provides the maximum resiliency and minimum overhead for frame based traffic. It ensures that frames are properly delivered while protecting CIR traffic without enforcing a restrictive mechanism of connection admission. In addition, loadsharing and overflow routing systems maximize use of available bandwidth in the network to deliver the data where possible.

Each packet contains a routing header and each node takes responsibility for the onward routing of the packet. This mechanism allows the rapid detection of potential congestion situations within the network and ensures that packets are rapidly re-routed over an alternative route to avoid potential problems such as congested trunks.

The system supports both delay sensitive and throughput sensitive traffic. Each traffic type has a set of metrics assigned to each trunk group (e.g. combined bandwidth and minimum delay). The path selection is based on using the minimum metric path. For the throughput-sensitive traffic, metrics are proportional to the bandwidth of the available trunk group. Metrics for the delay sensitive traffic are measured using the round trip of a 512 byte frame to reach the neighbouring nodes. Topology management for routing takes congestion into account in the updates.

Loadsharing and Loadspreading

- Loadspreading
  - spreads VC's over links
- Loadsharing
  - distributes frames in a VC over links
  - sharing helps prevent congestion

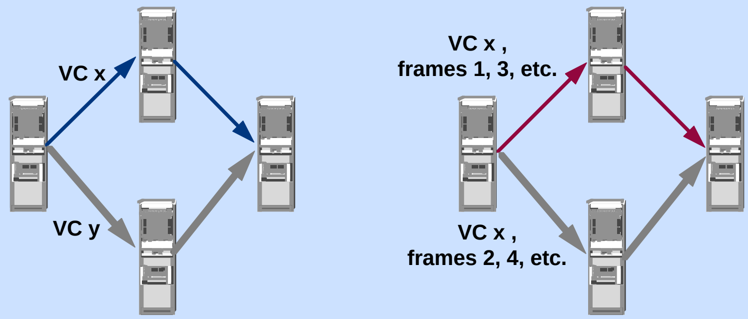

15

In loadspreading, all traffic for a given virtual circuit will follow the same route (i.e. use the same links) through the network. This preserves the order of the packets and minimizes delay. Different VC's will be spread across the available links. A separate overflow system allows traffic for a VC to be sent on an alternate link when the original link becomes congested.

In loadsharing, the traffic for a given VC will be distributed or shared across up to four links in up to two link groups. It is also sensitive to the capacity of each of the links as well as provisioned PORS traffic. Loadsharing effectively balances the overall load among the links. Loadsharing has many benefits including:

- Links of different speeds are effectively utilized
- Loadsharing the traffic can compensate imbalances due to PORS traffic
- All available bandwidth is effectively used
- Bandwidth aggregation - traffic from a single VC will be shared across all the links in a link group. This allows a single VC to have a bandwidth greater than that of a single link up to the aggregate bandwidth of the LinkGroup
- Overflow routing allows alternate routes to be used in the event of congestion on the preferred route. Access to the overflow routing feature may be provisioned on a per DLCI basis which provides greater control on an application basis as well as the ability for the carrier to offer this as a differentiated premium service to users.

Failure Recovery

- **Circuit re-establishment**
  - dynamic packet routing system
  - no need to provision alternate facilities
  - engineer trunk capacity to carry overflow in event of failure
- **UNI/NNI resiliency**

16

Magellan is designed to deliver data whenever possible as well as have the network survive in the event of catastrophic failures. The dynamic packet routing system (DPRS) is a key element in providing this superior ability. Passports DPRS makes independent decisions for each frame and at each hop for the route to the next hop. The decision is based on current network topology and loading. Each frame contains the full destination address. In the event of congestion or a failure, frames can be rerouted around the trouble area.

The routing system bases its routing decisions on the topology map, which means that its decisions are instantaneous and accurate. Instantaneous routing and rerouting is enabled by recalculating the best and next best routes for each traffic class of service after every routing update. This allows traffic to be switched to the next best route without having to go through a route discovery process. The routing system takes account of QoS requirements when choosing routes.

For access services, with Passport's sophisticated redundancy and routing features, many of the possible network failure scenarios do not cause noticeable disruption of traffic flow. For example, Passport's frame relay (DPRS routing) does not have to setup new calls as the best path and the second best path are always computed, and used to load-share traffic if the metrics are equal. As soon as the best path is unavailable, the frames are automatically and instantaneously routed to the next best path. When a better path becomes available, it will automatically become the best path (e.g. the failed trunk is returned to service). For DPRS traffic there is no dependence on the number of PVCs or SVCs using the failed link and setup times, because an alternative route is immediately chosen once the failed route is detected. Consequently Magellan networks are highly scalable. It has been designed for large networks and for rapid growth.

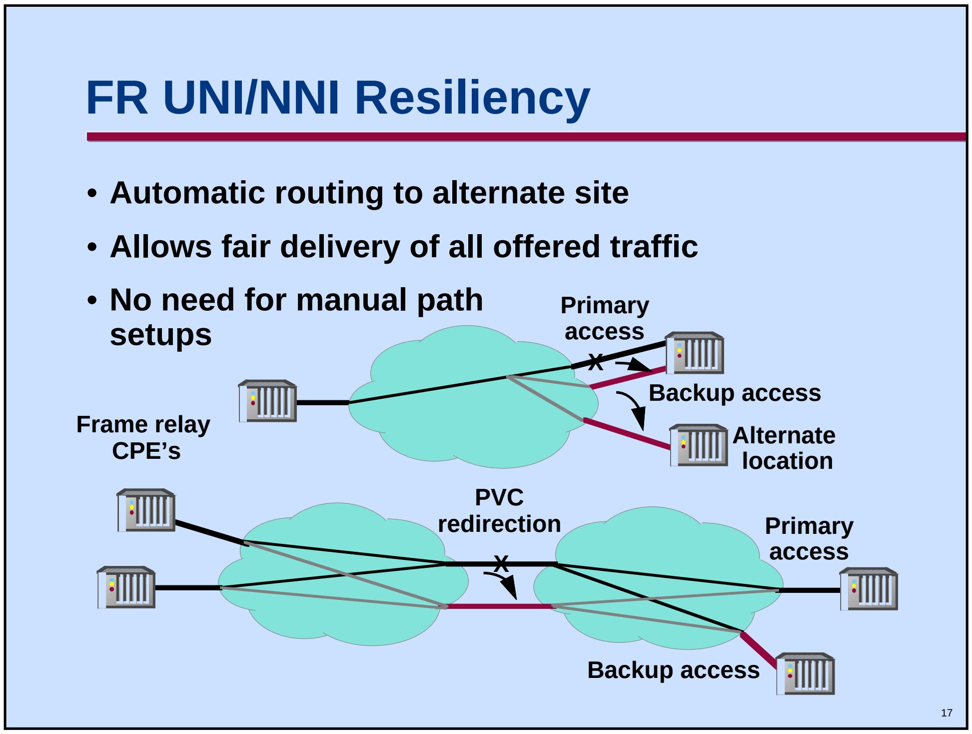

Passport supports resiliency at the NNI with PVC backup in the event of remote network failure, NNI facilities failure, or even failure of the NNI processor or node to redirect the calls to a backup NNI.

The resiliency feature is based on redirecting PVCs from a failed connection (e.g. an NNI where the facility is broken), to a designated backup. This can be activated manually or gated by operator activity. This feature does not require extensions to existing standards but would need a bi-lateral agreement, between the operators of the two connected networks, on the association of primary and backup NNIs.

Consider the example presented in this slide. One PVC with two remote accesses has the attraction of total transparency to the branch device because only a single PVC is involved to the central site(s). Magellan has considerable experience in this type of call redirection for both PVCs and SVCs. In the case of PVC, individual PVCs can be identified as candidates for redirection (by subscription). Triggering events for the Magellan redirection function are at the access level. This philosophy has evolved based on Magellan experience and lead customer requirements. Triggering events include DTE failure, access link failure, and network equipment failure or isolation. One useful application scenario is to have a number of client devices connected to a pair of centrally located servers which back each other up. Separating the central servers geographically provides even greater resiliency to disaster events.

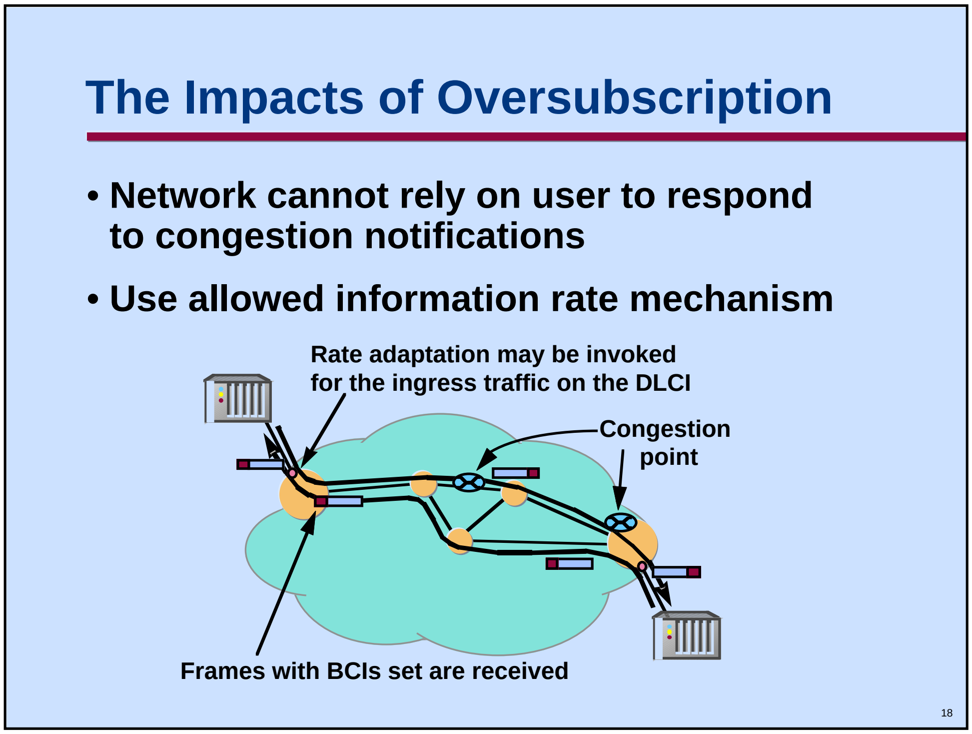

When an oversubscription strategy is adopted, it is possible that congestion may occur due to a statistical traffic burst. A mechanism is required to respond to and clear the congestion.

Passport employs a rate-based credit congestion control mechanism which allows data throughput to be accurately monitored without impacting the performance.

Rate adaptation - Allowed Information Rate (AIR):

The network cannot always rely on the user to respond quickly to congestion conditions in the network (e.g. devices that do not respond to BECNs). In these cases, the network can invoke a feature called rate adaptation - Allowed Information Rate (AIR). AIR provides a proactive congestion control mechanism by allowing the ingress access to make controlled discards when signs of congestion are detected.

Rate adaptation is an optional feature based on a closed-loop adaptive rate control scheme where the access automatically adapts its accepted load to better match network capacity. This is done by changing its EIR and/or CIR values during congestion. Under normal conditions the user is allowed to send data at a sustained rate equal to the CIR + EIR. The frame relay PVC service monitors the subnet for congestion conditions. If congestion is detected, the service will reduce the AIR. If the user does not reduce its load faster than the load reduction imposed by rate adaptation, then frames will be discarded by the network entry point. This feature reduces the needless use of valuable network resources required to transmit frames that are destined to be discarded anyway.

AIR is compatible with existing CCITT and ANSI standards.

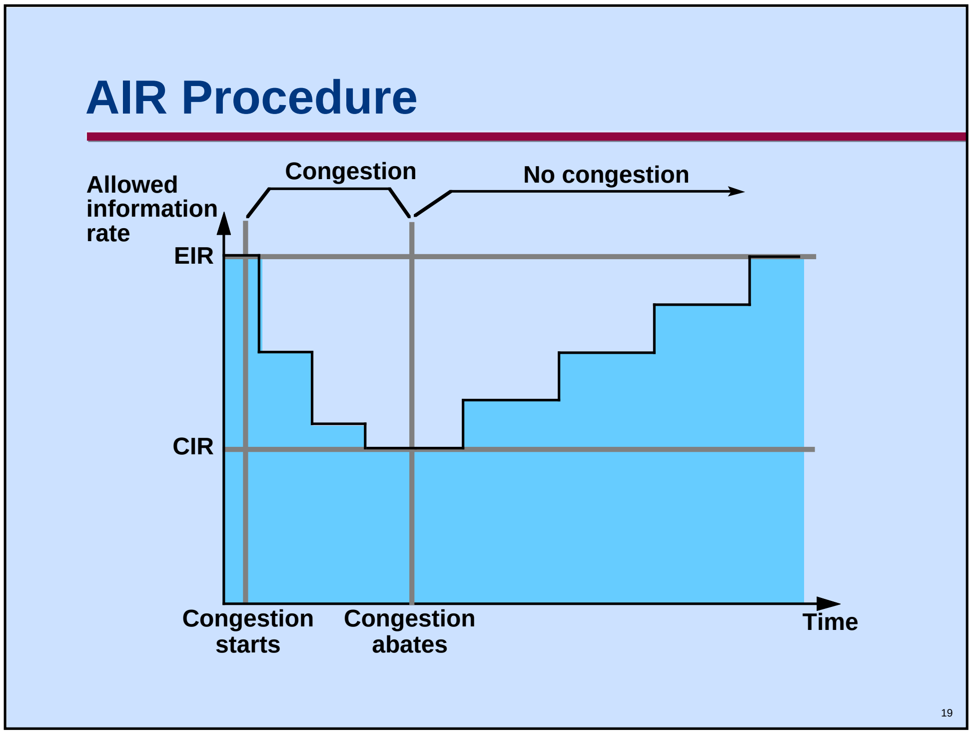

Recognizing the differences in service definitions from network provider to network provider, two rate adaptation schemes are supported:

• EIR adaptation only - When congestion is detected, EIR is reduced gradually to CIR until congestion abates, at which time EIR is restored gradually.

• EIR + CIR adaptation - When congestion is detected, EIR is reduced significantly, and if congestion persists, CIR is reduced gradually to a minimum value until congestion abates. CIR is then restored gradually followed by a gradual EIR restoration as well.

Under normal conditions the user is allowed to send data at the provisioned rates. The frame relay service monitors the subnet for feedback of congestion information. If congestion is detected, the service will reduce AIR by 25% of its current value, and will continue to monitor the subnet for further changes in congestion. Depending on the option, AIR will be either EIR or CIR. Based on this evaluation, it will decide whether the congestion condition has abated or not. If the congestion persists, the service will reduce by 25% of the current value of the AIR again. When the congestion condition has abated, the service will be restored at a linear rate in increments of 6% of provisioned CIR.

Features:

• Excess Information Rate (EIR) is throttled gradually under congestion

• EIR is gradually re-attained when congestion alleviates

• Geometric decrease, arithmetic increase minimizes oscillation

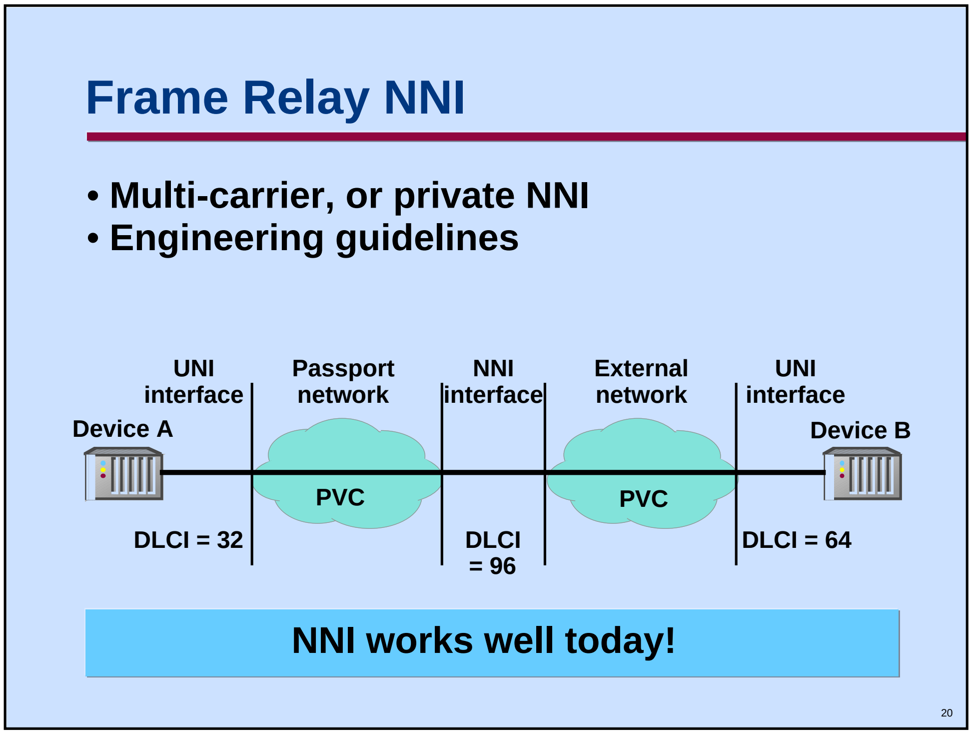

The Passport NNI provides a communication interface between a Passport frame relay network and other frame relay networks. In addition to transfer of user data between different networks, the NNI is required to properly receive, process and propagate the network status signalling information from a global perspective. The end user must have an accurate picture of the the network connection that may span several different frame relay networks.

This slide illustrates the connection between two frame relay users over a hybrid frame relay network. Device A is uses a Passport frame relay UNI service to access a Passport frame relay network. The Passport frame relay network uses the Passport NNI service to connect with an external frame relay network. Device B uses a non-Passport frame relay UNI service to connect to the external network. The NNI service ensures multi-vendor compatibility with those networks that conform to the frame relay Forum Implementation Agreement FRF.2 specification.

Passport provides a very high performance, resilient NNI with a variety of high-speed interfaces. This offers high performance and flexible interworking between both Magellan customers and multi-vendor/multi-carrier environments.

Magellan frame relay NNI is fully functional today and has been used extensively in trials both in laboratory environments, where operation has been examined and exercised in detail, and also for the carriage of live traffic.

Frame Relay NNI (notes only)

- NNI is operational today
- Standards compliance
  - Nortel editorship of X.76
  - Annex D of T1.617 (ANSI)
  - Annex A of Q.933 (ITU-T)
- Other activity
  - regular contributor to frame relay Forum Technical Committee
  - chair of Frame Relay Forum Market Development and Education committee

21

With appropriate engineering, the current NNI is fully ready to assume service between networks carrying live data. Nortel is strongly committed to participation and contribution to emerging standards for data communications in general, and NNI enhancement in particular. Nortel holds the editorship of recommendation X.76 in the ITU-T and is a regular contributor to the frame relay Forum Technical Committee in the area of NNI enhancement. Also, Nortel holds the chair position in the Frame Relay Forum Market Development and Education committee where service provider issues such as network interconnection arrangements and issues are frequently raised.

Magellan Passport frame relay NNI fully complies with the following standards:

- Annex D of Standard T1.617 (ANSI) with two byte frame addressing, note asynchronous status reporting is fully supported.
- Annex A of Recommendation Q.933 (ITU-T) with two byte frame addressing, note asynchronous status reporting is fully supported.
- Implementation Agreement FRF.2 (Frame Relay Forum).

Magellan frame relay NNI has been connected to frame relay NNI implementations from Stratacom, Cascade, Newbridge, Cisco, Bay (Wellfleet) and N.E.T. The A-Bit signalling has proven robust and all operational alarms and messages are functional and useful following minor enhancements suggested during early trials.

NNI Engineering Guidelines

- Coordination with UNIs
- Where available, use co-located platforms to minimize facilities costs and maximize availability
- Use natural interconnection points at high speed
- Initially engineer sum of CIRs < NNI access rate, then develop oversubscription strategy for each NNI
- Where possible, provision alternate NNI

22

Often, the real issue for multicarrier NNI is not so much technical, but the possible differences in the tariffs, management, and administration of the service. Users buy frame relay for the high performance and end-to-end management. With appropriate engineering of the NNI, these drivers are maintained. Proper provisioning requires a coordinated strategy across both the UNIs and the NNIs that connect them.

Use links between co-located platforms for minimum facilities costs and maximum availability. Use natural interconnection points at backbone/access demarcation points. As traffic increases and economics balance traffic efficiency over fanout, deploy additional NNIs and higher speeds to satisfy bandwidth and community of interest.

Where possible, high-speed NNI links should be used. This will minimize capital costs and provide minimal latency. High-speed NNI allows improved scalability and reduced administration with less physical connections. NNI deployed at 8 Mbit/s using V.11 is available now. Passport can easily support NNI at E3 and DS-3. Alternatively use HSSI with appropriate CSU/DSU.

Where physically and economically feasible, an alternate NNI link should be provisioned. This allows the network to take advantage of the NNI resiliency feature. This feature is not dependent on standards development and will work with other vendors as no proprietary extensions are required. This can be automatic or set to be manual.

When implementing the NNI, the sum of CIRs should initially be less than the NNI access rate. This is consistent with the FRF.2 recommendation "to limit the sum of the subscribed CIRs (egress from the network) of all PVCs on a given NNI to be less than the NNI access rate." Based on the QoS required and the traffic profiles, one can monitor the NNI to determine if overbooking is possible based on peak traffic profiles.

NNI Engineering Guidelines

- **Rate enforcement**
  - traffic shape may have been altered
- **Rate adaptation**
  - faster reaction to potential congestion

23

Rate enforcement involves setting the CIR, Be and Bc to dictate traffic flow and bandwidth usage. Rate enforcement should only be controlled by the networks directly connected to the users.

If rate enforcement is already applied at the UNI level, it should not be necessary for it to be applied at the NNI. Exercise caution if rate enforcement is applied at the NNI. There may be occasions where an external network has altered the original traffic characteristics. The NNI is a natural concentration point. Random bursts from several independent sources may arrive at the NNI close together (i.e. in a clumped fashion). This can have the effect of creating larger instantaneous bursts at that point. In this case, a CIR/Be/Bc violation can be created when in fact no violation has occurred at the ingress. Although the average CIR over time is acceptable, a shorter duration may indicate a violation. It may be advisable to increase Bc and Tc to account for the change in traffic shape.

One must decide whether to perform rate adaptation at the network ingress and/or at the NNI by considering the tradeoffs. If adaptation is performed at both the access and the NNI, there is a shorter feedback loop and faster reaction to congestion in the second network. This is at the cost of policing the user twice, which may increase the chance of discard (minimal). If adaptation is done at the access only, the user is policed once and the chance of false discards is reduced. In general, rate adaptation should be performed at both points.

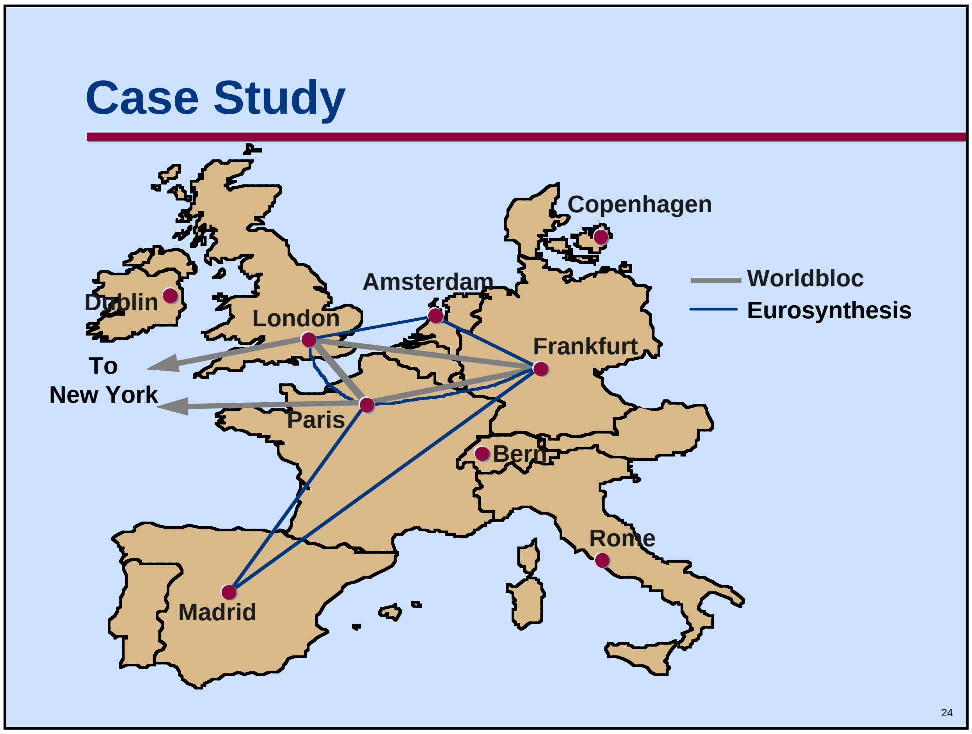

This case study will demonstrate a strategy for the implementation of NNI between two networks.

Worldbloc is a consortium of carriers providing services around the world. They offer a frame relay service based on Passport. Transatlantic facilities are provided through points of presence in New York, London, Paris and Frankfurt.

Eurosythesis is a relatively young service provider in Europe. They offer a frame relay service on a different platform between London, Paris, Amsterdam, Frankfurt, and Madrid.

Recently, Eurosynthesis decided to offer an international frame relay service. They have entered into an agreement to have Worldbloc provide the overseas connection. Currently, engineering groups are in the process of determining the impact of the new traffic and where to provision the NNI or NNIs.

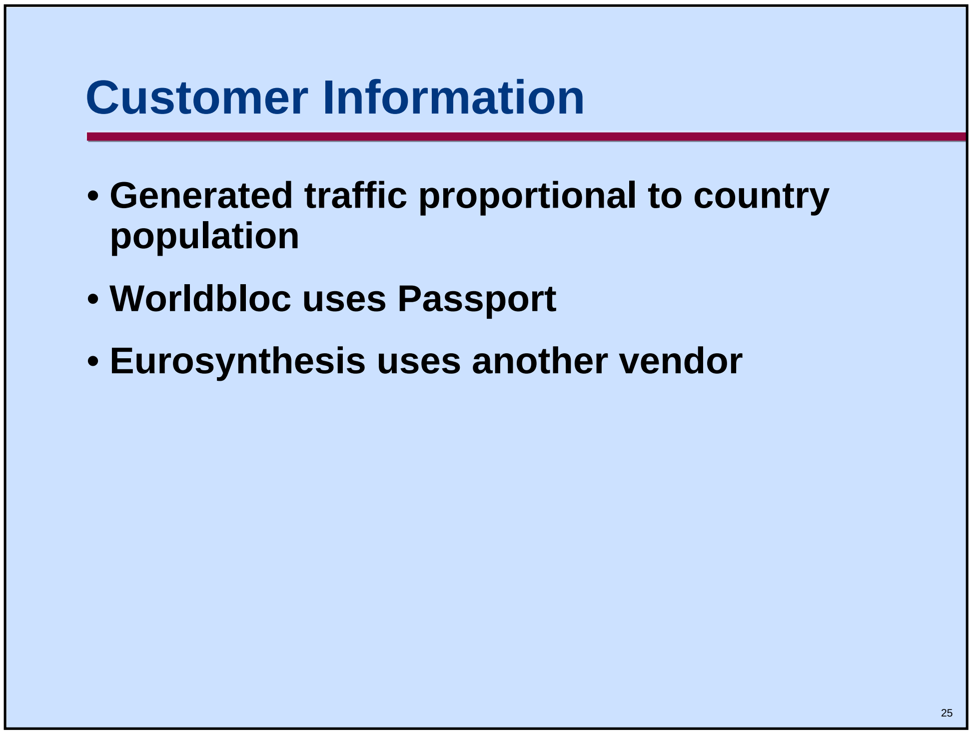

To develop the offered traffic values, an aggregate network load is chosen and allotted according to the population of the country. The current Eurosynthesis network carries a total offered load equal to two E1's of traffic. The Worldbloc network currently carries an aggregate load of two DS1's to New York. The new international service will introduce a new offered traffic load of 30% of Eurosynthesis' current load.

Country populations are as follows:

• France	56,556,000
• United States	248,709,873
• United Kingdom	55,514,500
• Netherlands	15,019,000
• Spain	38,991,000
• Germany	78,500,000

All stated figures are from “The Times Atlas of the World,” 9th complete edition.

Traffic flows are given on the following slide for information purposes.

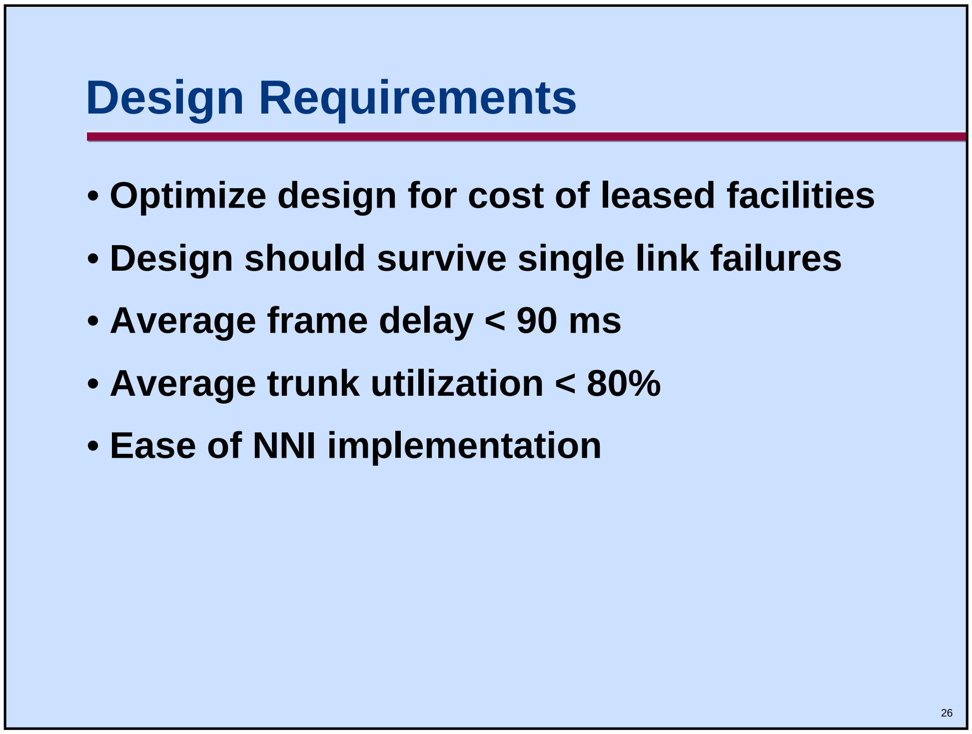

In selecting the location of the NNI interface, we must adhere to a set of design requirements while following our guidelines.

Costs are a typical driver behind many networks. In this study, network costs will be optimized. Changing the topology of an established network can be expensive but must still be considered if evidence leads to this conclusion.

The networks must be able to withstand a single link failure. This is a common requirement, especially for a service provider. Loss of connectivity can result in immediate financial losses, in addition to lost credibility and the resulting impact on future sales. Since both entities in our study are service providers, this is a critical requirement.

The average frame delay of the overseas traffic is not to exceed 90 ms. This value was selected to account for the distances and propagation delay involved in an overseas connection. This translates into a maximum of 30 ms for the Eurosynthesis network, and 60 ms for the Worldbloc network.

Trunk utilization should be kept below 80%. This is a recommended design guideline so that the network can handle traffic fluctuations. Trunk utilizations of 90% or higher will be viewed as a network failure. This is due to the fact that frame delay begins to increase asymptotically, leading to a violation of our delay requirement.

Finally, the location of the NNI interface must not impose any unusual requirements on either carrier (e.g. provision a facility between two distant locations).

Design Assumptions

- Average frame size is 256 bytes
- Standard deviation is 128 bytes
- Backbone trunking is terrestrial
- Bandwidth is available in N*128 kbit/s increments
- Nodal delay is 0.18 ms
- Prices are based on simulated tariffs

27

The design assumptions considered above are standard for frame relay service providers.

- An average frame size of 256 bytes with a standard deviation of 128 bytes is acceptable
- The trunking fabric is terrestrial
- Routers constitute the frame relay access into the network
- Bandwidth is in N*128 Kbit/s when available, the majority of service providers decrease their bandwidth resolution (offering) as bandwidth requirements exceed 50% of country specific ISDN primary rate interface
- With Passport used as the service providing equipment, the measure nodal delay is 0.18 ms (that number may be ignored for all practical reasons)
- The prices used and quoted are simulated service provider tariffs

Eurosynthesis Current Network

Network Summary

	Normal	Failure
Average Frame Relay Data Field (bytes)	256	256
Standard Deviation of Data Field (bytes)	128	128
Average Mesh Path Delay Per Frame (ms)	5.6	7.5
Average Frame Relay Access Speed (kbps)	256	256
Average Access + CIR Delay Per Frame (ms)	16.6	16.6
Average Total Delay Per Frame (ms)	22.2	24.2
Trunk Cost (Pounds/Month)	45,480	45,480
Data Traffic Carried (Mbit/s)	3.969	3.969
Average Trunk Utilization (Total)	24.1%	30.6%

Bandwidth Summary

				Normal		Failure	
Source	Dest	km	Available BW (kbps)	Used Tot BW	% Tot Util	Used Tot BW	% Tot Util
Amsterdam	Frankfurt	361	1920	518	26.1	1053	54.8
Amsterdam	London	311	1920	565	28.5	962	50.1
Frankfurt	Amsterdam	361	1920	609	30.7	1144	59.6
Frankfurt	Madrid	1158	1920	353	17.8	546	28.4
Frankfurt	Paris	482	1920	450	22.7	792	41.3
London	Amsterdam	311	1920	539	27.1	935	48.7
London	Paris	362	1920	535	27.0	0	0.0
Madrid	Frankfurt	1158	1920	299	15.1	476	24.8
Madrid	Paris	925	1920	359	18.1	181	9.5
Paris	Frankfurt	482	1920	398	20.0	755	39.3
Paris	London	362	1920	535	26.9	0	0.0
Paris	Madrid	925	1920	390	19.7	198	10.3

Failed major link

28

The Eurosynthesis network has been designed and recently expanded with the expectation of future growth, including the introduction of the international service. Thus, the E1 facilities between the nodes currently have a large unused capacity. Even in the event of a link failure, the average utilization is 30.6%, and no link is over 60%. The network is well positioned to accept additional traffic load and still survive a single link failure.

The current Eurosynthesis network is used to provide a benchmark and a reference point for the study of the NNI interface location.

Worldbloc Current Network

Network Summary

	Normal	Failure
Average Frame Relay Data Field (bytes)	256	256
Standard Deviation of Data Field (bytes)	128	128
Average Mesh Path Delay Per Frame (ms)	30.2	34.2
Average Frame Relay Access Speed (kbps)	256	256
Average Access + CIR Delay Per Frame (ms)	16.6	16.6
Average Total Delay Per Frame (ms)	46.9	50.8
Trunk Cost (Pounds/Month)	149,364	149,364
Data Traffic Carried (Mbit/s)	3.072	3.072
Average Trunk Utilization (Total)	23.5%	28.4%

Bandwidth Summary

				Normal		Failure	
Source	Dest	km	Available BW (kbps)	Used Tot BW	% Tot Util	Used Tot BW	% Tot Util
Frankfurt	London	601	1920	0	0.0	657	34.2
Frankfurt	Paris	482	1920	657	34.2	0	0.0
London	Frankfurt	601	1920	0	0.0	657	34.2
London	New York NY	5569	1920	465	24.2	1595	83.1
London	Paris	362	1920	0	0.0	473	24.6
New York NY	London	5569	1920	465	24.2	1595	83.1
New York NY	Paris	5539	1920	1130	58.9	0	0.0
Paris	Frankfurt	482	1920	657	34.2	0	0.0
Paris	London	362	1920	0	0.0	473	24.6
Paris	New York NY	5539	1920	1130	58.9	0	0.0

Failed major link

29

The Worldbloc network has been designed with the expectation of future growth. Thus, the E1 facilities between the nodes currently have a large unused capacity. A full mesh design has been used to provide single link survivability. This results in some links being idle under normal operation. As can be seen from the tables, these links are required to carry traffic in the event of a link failure. As well, the network is able to survive a failure of one of the overseas connections. Even in the event of a link failure, the average utilization is 28.4%, and the internal network is below 35%. The internal network is well positioned to accept additional traffic load and still survive a single link failure. Future growth will probably require the expansion of the overseas links.

The current Worldbloc network is used to provide a benchmark and a reference point for the study of the NNI interface location.

Where Locate NNI Interface

- How will additional traffic affect Worldbloc network
- How will NNI point affect flows in Eurosynthesis network
- Access closest to traffic source, or gateway
- Impact of link failure

30

The starting design is the current network topologies.

Many questions need to be asked when determining the location of the NNI interface. Different issues will be important depending on the environment and the perspective of the designer. A single company multi-platform environment can have collocated resources. However, a multi-carrier environment will likely have to provide a leased line between switches. A private enterprise NNI is much like providing a UNI connection.

This slide provides some of the questions that must be addressed in this case study. Throughout the study, one must consider the engineering guidelines in order to determine the optimal design.

The common points of presence will be examined for their respective impacts on each service provider network. Selection of the final NNI location will be based on the costs incurred by each provider, and adherence to the design requirements.

Note that as the locations are considered, the average frame delay is the sum of the delays in each network. That is, the overseas end-to-end delay is equal to the delay in the Eurosynthesis network, plus the delay in the Worldbloc network.

NNI in London - Impact on Eurosynthesis

Network Summary

	Normal	Failure
Average Frame Relay Data Field (bytes)	256	256
Standard Deviation of Data Field (bytes)	128	128
Average Mesh Path Delay Per Frame (ms)	6.1	162.3
Average Frame Relay Access Speed (kbps)	256	256
Average Access + CIR Delay Per Frame (ms)	16.6	16.6
Average Total Delay Per Frame (ms)	22.8	178.9
Trunk Cost (Pounds/Month)	45,480	45,480
Data Traffic Carried (Mbit/s)	5.809	5.809
Average Trunk Utilization (Total)	37.5%	40.0%

Bandwidth Summary

				Normal		Failure	
			Available	Used	%	Used	%
Source	Dest	km	BW (kbps)	Tot BW	Tot Util	Tot BW	Tot Util
Amsterdam	Frankfurt	361	1920	914	47.6	329	17.1
Amsterdam	London	311	1920	1038	54.1	0	0.0
Frankfurt	Amsterdam	361	1920	1006	52.4	394	20.5
Frankfurt	Madrid	1158	1920	353	18.4	353	18.4
Frankfurt	Paris	482	1920	450	23.5	1488	77.5
London	Amsterdam	311	1920	1011	52.7	0	0.0
London	Paris	362	1920	1018	53.0	1890	98.5
Madrid	Frankfurt	1158	1920	299	15.6	299	15.6
Madrid	Paris	925	1920	556	28.9	556	28.9
Paris	Frankfurt	482	1920	398	20.7	1409	73.4
Paris	London	362	1920	1017	53.0	1917	99.8
Paris	Madrid	925	1920	587	30.6	587	30.6

Failed major link

31

The tables presented in this slide indicate the impact the NNI interface in London would have on the Eurosynthesis network. Notice that, under normal conditions, the network can accommodate the increase of 30% in the offered traffic and performs well. The average utilization is 37.5% with an average frame delay of 22.8 ms.

However, in the event of a link failure, the current network is unable to handle the rerouted traffic. A failure in the London-Amsterdam trunk requires traffic to be rerouted across the Frankfurt-Paris and Paris-London trunks. As the tables show, the Paris-London trunks is almost 100% utilized. This causes frames to encounter queuing and other delays resulting in an average frame delay of 178.9 ms.

Eurosynthesis could alleviate this problem by providing additional facilities between Paris and London. It may also be beneficial to add capacity to the Frankfurt-Paris link. More detailed analysis will be postponed until the other locations are studied.

Summary:

Initial capability: Excellent

Failure handling: Network failure

Potential costs: Additional facilities Paris-London

NNI in London - Impact on Worldbloc

Network Summary

Network Summary	Normal	Failure
Average Frame Relay Data Field (bytes)	256	256
Standard Deviation of Data Field (bytes)	128	128
Average Mesh Path Delay Per Frame (ms)	32.4	Fail
Average Frame Relay Access Speed (kbps)	256	256
Average Access + CIR Delay Per Frame (ms)	16.6	16.6
Average Total Delay Per Frame (ms)	49.1	Fail
Trunk Cost (Pounds/Month)	149,364	149,364
Data Traffic Carried (Mbit/s)	5.453	5.453
Average Trunk Utilization (Total)	36.3%	54.0%

Bandwidth Summary

				Normal		Failure	
			Available	Used	%	Used	%
Source	Dest	km	BW (kbps)	Tot BW	Tot Util	Tot BW	Tot Util
Frankfurt	London	601	1920	0	0.0	0	0.0
Frankfurt	Paris	482	1920	657	34.2	657	34.2
London	Frankfurt	601	1920	0	0.0	0	0.0
London	New York NY	5569	1920	1700	88.6	0	0.0
London	Paris	362	1920	0	0.0	1700	88.6
New York NY	London	5569	1920	1700	88.6	0	0.0
New York NY	Paris	5539	1920	1130	58.9	2830	147.4
Paris	Frankfurt	482	1920	657	34.2	657	34.2
Paris	London	362	1920	0	0.0	1700	88.6
Paris	New York NY	5539	1920	1130	58.9	2830	147.4

Failed major link

32

This slide shows the impact the NNI interface in London would have on the Worldbloc network. The first point to consider is the additional traffic that the network must carry. All of the new traffic from the Eurosynthesis network is offered to Worldbloc at the NNI interface location. This means that the overseas trunks must carry 1,190.4 kbit/s of new traffic in addition to the normal offered load. In the London case, the offered traffic is 1,700 kbit/s.

Notice that, under normal conditions, the network can just accommodate the new traffic load. The average utilization is 36.3% with an average frame delay of 49.1 ms. The London-New York trunk is heavily utilized at 88.6%. This is less than optimal and should be alleviated.

Also, in the event of a link failure, the current network is unable to handle the rerouted traffic. A failure in the London-New York trunk requires traffic to be rerouted across the Paris-London trunks, and onto the Paris-New York trunk. The Paris-London trunk is 88.6% utilized. The Paris-New York trunk cannot accept the offered load since it is greater than the capacity. Thus, the network will not survive a single link failure.

Obviously, the Worldbloc network must add capacity to the overseas trunks to handle the new contract. The design issue is to minimize this additional cost. Examining the new load, the offered traffic for Paris is 2,366 kbit/s. This will obviously necessitate the provisioning of additional facilities from Paris to New York. Also, Frankfurt will introduce a load of 1,893 kbit/s to the internal network , resulting in 2,366 kbit/s of traffic on the Paris-New York trunk. The Frankfurt case will be examined later.

For the remainder of the study, we will assume that Worldbloc doubles the overseas capacity.

NNI in London - Impact on Worldbloc with New Capacity

Network Summary

	Normal	Failure
Average Frame Relay Data Field (bytes)	256	256
Standard Deviation of Data Field (bytes)	128	128
Average Mesh Path Delay Per Frame (ms)	28.5	34.2
Average Frame Relay Access Speed (kbps)	256	256
Average Access + CIR Delay Per Frame (ms)	16.6	16.6
Average Total Delay Per Frame (ms)	45.1	50.8
Trunk Cost (Pounds/Month)	214,743	214,743
Data Traffic Carried (Mbit/s)	5.453	5.453
Average Trunk Utilization (Total)	21.6%	39.3%

Bandwidth Summary

				Normal		Failure	
			Available	Used	%	Used	%
Source	Dest	km	BW (kbps)	Tot BW	Tot Util	Tot BW	Tot Util
Frankfurt	London	601	1920	0	0.0	0	0.0
Frankfurt	Paris	482	1920	657	34.2	657	34.2
London	Frankfurt	601	1920	0	0.0	0	0.0
London	New York NY	5569	1920	1700	44.3	0	0.0
London	Paris	362	1920	0	0.0	1700	88.6
New York NY	London	5569	1920	1700	44.3	0	0.0
New York NY	Paris	5539	1920	1130	29.4	2830	73.7
Paris	Frankfurt	482	1920	657	34.2	657	34.2
Paris	London	362	1920	0	0.0	1700	88.6
Paris	New York NY	5539	1920	1130	29.4	2830	73.7

Failed major link

33

This slide shows the impact the NNI interface in London would have on the Worldbloc network with the additional overseas capacity.

Under normal conditions, the network can easily accommodate the new traffic load. The average utilization is 21.6% with an average frame delay of 45.1 ms.

In the event of a London-New York link failure, the current network is able to handle the rerouted traffic. A failure in the London-New York trunk requires traffic to be rerouted across the Paris-London trunks, and onto the Paris-New York trunk. As the tables show, the Paris-London trunks is 88.6% utilized; this may be considered acceptable since it is under a failure condition. The Paris-New York trunk can now handle the extra load since the capacity has been increased.

Summary:

Initial capability: Excellent

Failure handling: Acceptable

Potential costs: None, may consider additional facilities Paris-London

NNI in Paris - Impact on Eurosynthesis

Network Summary

	Normal	Failure
Average Frame Relay Data Field (bytes)	256	256
Standard Deviation of Data Field (bytes)	128	128
Average Mesh Path Delay Per Frame (ms)	5.4	9.3
Average Frame Relay Access Speed (kbps)	256	256
Average Access + CIR Delay Per Frame (ms)	16.6	16.6
Average Total Delay Per Frame (ms)	22.0	26.0
Trunk Cost (Pounds/Month)	45,480	45,480
Data Traffic Carried (Mbit/s)	5.799	5.799
Average Trunk Utilization (Total)	33.0%	44.3%

Bandwidth Summary

				Normal		Failure		
			Available	Used	%	Used	%	
Source	Dest	km	BW (kbps)	Tot BW	Tot Util	Tot BW	Tot Util	
Amsterdam	Frankfurt	361	1920	518	27.0	1409	73.4	
Amsterdam	London	311	1920	641	33.4	1242	64.7	
Frankfurt	Amsterdam	361	1920	609	31.7	1501	78.2	
Frankfurt	Madrid	1158	1920	353	18.4	546	28.4	
Frankfurt	Paris	482	1920	847	44.1	1545	80.5	
London	Amsterdam	311	1920	614	32.0	1216	63.3	
London	Paris	362	1920	891	46.4	0	0.0	← Failed major link
Madrid	Frankfurt	1158	1920	299	15.6	476	24.8	
Madrid	Paris	925	1920	556	28.9	379	19.7	
Paris	Frankfurt	482	1920	794	41.4	1508	78.5	
Paris	London	362	1920	891	46.4	0	0.0	← Failed major link
Paris	Madrid	925	1920	587	30.6	395	20.5	

34

The impact of the NNI interface in Paris is now examined. Under normal conditions, the Eurosynthesis network can accommodate the additional offered traffic and performs well. The average utilization is 33.0% with an average frame delay of 22.0 ms.

The network can also withstand a London-Paris link failure. The average frame delay is still acceptable when combined with that of the Worldbloc network on the next slide. The only considerations are the utilizations on the Frankfurt-Paris and the Frankfurt-Amsterdam trunks. These are close to our guideline to remain with 80%. This is acceptable since this is under a link failure condition. However, it is not strictly optimal.

Summary:

Initial capability: Excellent

Failure handling: Good

Potential costs: None

NNI in Paris - Impact on Worldbloc

Network Summary

	Normal	Failure
Average Frame Relay Data Field (bytes)	256	256
Standard Deviation of Data Field (bytes)	128	128
Average Mesh Path Delay Per Frame (ms)	28.6	34.6
Average Frame Relay Access Speed (kbps)	256	256
Average Access + CIR Delay Per Frame (ms)	16.6	16.6
Average Total Delay Per Frame (ms)	45.3	51.3
Trunk Cost (Pounds/Month)	214,743	214,743
Data Traffic Carried (Mbit/s)	5.453	5.453
Average Trunk Utilization (Total)	21.6%	39.4%

Bandwidth Summary

				Normal		Failure	
			Available	Used	%	Used	%
Source	Dest	km	BW (kbps)	Tot BW	Tot Util	Tot BW	Tot Util
Frankfurt	London	601	1920	0	0.0	657	34.2
Frankfurt	Paris	482	1920	657	34.2	0	0.0
London	Frankfurt	601	1920	0	0.0	657	34.2
London	New York NY	5569	3840	465	12.1	2830	73.7
London	Paris	362	1920	0	0.0	1709	89.0
New York NY	London	5569	3840	465	12.1	2830	73.7
New York NY	Paris	5539	3840	2366	61.6	0	0.0
Paris	Frankfurt	482	1920	657	34.2	0	0.0
Paris	London	362	1920	0	0.0	1709	89.0
Paris	New York NY	5539	3840	2366	61.6	0	0.0

Failed major link

35

This slide considers the impact of the Paris NNI interface on the Worldbloc network. As determined in the London impact study, the overseas trunk capacity has been doubled.

Under normal conditions, the network can easily accommodate the new traffic load. The average utilization is 21.6% with an average frame delay of 45.3 ms.

A Paris-New York link failure is handled fairly well. The combined average frame delay (26.0+51.3=77.3 ms) is within our 90ms limit. However, the utilization on the London-Paris trunk is not within our guideline of 80% and is almost 90%. The network will survive a single link failure, but some service degradation may be experienced on the London-Paris trunk if the traffic fluctuates.

This could be remedied by providing additional facilities on the London-Paris trunk. This analysis will be postponed until all cases are reviewed.

Summary:

Initial capability: Excellent

Failure handling: Adequate, some possible degradation

Potential costs: None, although may consider additional facilities London-Paris

NNI in Frankfurt - Impact on Eurosynthesis

Network Summary

Network Summary	Normal	Failure
Average Frame Relay Data Field (bytes)	256	256
Standard Deviation of Data Field (bytes)	128	128
Average Mesh Path Delay Per Frame (ms)	5.8	8.9
Average Frame Relay Access Speed (kbps)	256	256
Average Access + CIR Delay Per Frame (ms)	16.6	16.6
Average Total Delay Per Frame (ms)	22.4	25.5
Trunk Cost (Pounds/Month)	45,480	45,480
Data Traffic Carried (Mbit/s)	5.585	5.585
Average Trunk Utilization (Total)	33.8%	37.3%

Bandwidth Summary

				Normal		Failure	
			Available	Used	%	Used	%
Source	Dest	km	BW (kbps)	Tot BW	Tot Util	Tot BW	Tot Util
Amsterdam	Frankfurt	361	1920	874	45.5	0	0.0
Amsterdam	London	311	1920	846	44.1	329	17.1
Frankfurt	Amsterdam	361	1920	966	50.3	0	0.0
Frankfurt	Madrid	1158	1920	550	28.7	507	26.4
Frankfurt	Paris	482	1920	736	38.3	1654	86.1
London	Amsterdam	311	1920	819	42.7	394	20.5
London	Paris	362	1920	535	27.9	1409	73.4
Madrid	Frankfurt	1158	1920	496	25.8	448	23.3
Madrid	Paris	925	1920	359	18.7	407	21.2
Paris	Frankfurt	482	1920	683	35.6	1515	78.9
Paris	London	362	1920	535	27.8	1501	78.2
Paris	Madrid	925	1920	390	20.3	433	22.6

Failed major link

36

The impact of the NNI interface in Frankfurt is now examined. Under normal conditions, the Eurosynthesis network can accommodate the additional offered traffic and performs well. The average utilization is 33.8% with an average frame delay of 22.4 ms.

Also, since the majority of the new overseas traffic originates in Frankfurt, this location minimizes the amount of traffic Eurosynthesis must carry on its facilities. This leaves more room for future growth.

The network can basically withstand an Amsterdam-Frankfurt link failure although the utilization of the Frankfurt-Paris trunk is high. The average frame delay is still acceptable when combined with that of the Worldbloc network on the next slide. The high utilization is acceptable since this is under a link failure condition. However, it is not strictly optimal.

This could be remedied by providing small additional facilities on this trunk. This analysis will be postponed until all cases are reviewed.

Summary:

Initial capability: Excellent

Failure handling: Acceptable

Potential costs: None, although may consider additional facilities Frankfurt-Paris

NNI in Frankfurt - Impact on Worldbloc

Network Summary

	Normal	Failure
Average Frame Relay Data Field (bytes)	256	256
Standard Deviation of Data Field (bytes)	128	128
Average Mesh Path Delay Per Frame (ms)	62.2	63.8
Average Frame Relay Access Speed (kbps)	256	256
Average Access + CIR Delay Per Frame (ms)	16.6	16.6
Average Total Delay Per Frame (ms)	78.9	80.4
Trunk Cost (Pounds/Month)	214,743	214,743
Data Traffic Carried (Mbit/s)	5.453	5.453
Average Trunk Utilization (Total)	34.5%	39.4%

Bandwidth Summary

				Normal		Failure	
			Available	Used	%	Used	%
Source	Dest	km	BW (kbps)	Tot BW	Tot Util	Tot BW	Tot Util
Frankfurt	London	601	1920	0	0.0	1893	98.6
Frankfurt	Paris	482	1920	1893	98.6	0	0.0
London	Frankfurt	601	1920	0	0.0	1893	98.6
London	New York NY	5569	1920	465	12.1	2830	73.7
London	Paris	362	1920	0	0.0	473	24.6
New York NY	London	5569	1920	465	12.1	2830	73.7
New York NY	Paris	5539	1920	2366	61.6	0	0.0
Paris	Frankfurt	482	1920	1893	98.6	0	0.0
Paris	London	362	1920	0	0.0	473	24.6
Paris	New York NY	5539	1920	2366	61.6	0	0.0

Failed major link

37

The impact of the Frankfurt NNI interface on Worldbloc is now examined. Under normal conditions, the network cannot accommodate the additional offered traffic. As the Frankfurt location reduces the load on the Eurosynthesis network, it increases the load on the local Worldbloc network. As noted earlier, the NNI introduces a load of 1,893 kbit/s on the Frankfurt-Paris trunk. This yields an excessive utilization of 98.6%. This is considered a network failure.

A similar overutilization condition is seen on the Frankfurt-London trunk if the Paris-New York trunk fails. Worldbloc would be required to add capacity to the Frankfurt-Paris and Frankfurt-London trunks in addition to doubling the overseas capacity. This is a large cost and is not an optimal condition.

Summary:

Initial capability: Network failure

Failure handling: Network failure

Potential costs: Large cost for additional facilities Frankfurt-Paris and Frankfurt-London

Location Summary

- Compromise between the providers

	London		Paris		Frankfurt	
	Eurosyn.	Worldbloc	Eurosyn.	Worldbloc	Eurosyn.	Worldbloc
Initial Capability	excellent	excellent	excellent	excellent	excellent	failure
Failure handling	failure	acceptable	good	adequate	acceptable	failure
Potential costs	one trunk	none	none	none	none	2 trunks

38

Eurosynthesis would like to have the interface in Frankfurt. It reduces the load on their network and would only require small additional facilities for the Frankfurt-Paris trunk to provide better single link failure survival. Paris is another good choice. Normal and failure conditions are handled approximately the same as the Frankfurt case. The extra traffic load reduces the free capacity for growth. London is the last choice since extra capacity must be added to provide single link failure survival. Paris and Frankfurt can essentially handle the failure without additional facilities, although some utilizations are high.

Worldbloc would suffer the greatest costs if the NNI interface is in Frankfurt. The current network cannot handle the additional load on the local network. Thus, great costs would be incurred to provide the additional facilities. Paris and London are both acceptable to Worldbloc. The scenarios behave almost identically under normal and failure conditions. In both cases, the Paris-London trunk should be expanded to handle the single link failure case, although it is not strictly required.

In conclusion, the best compromise location for the NNI interface is Paris.

Multiple NNI interfaces

- Provide enhanced reliability and survivability
- NNI resiliency at local node or distant node
- May provide more even distribution of offered traffic load on the networks

39

Provisioning more than one NNI interface point has many advantages.

Additional NNI interfaces enhance the reliability and survivability of the networks. Obviously, if only one NNI interface exists, a failure would result in loss of connectivity which is unacceptable. With multiple NNI interfaces, networks can take advantage of NNI resiliency features. Also, traffic could be balanced across the NNI interfaces resulting in a more even distribution of traffic loads across both networks.

These advantages are extended if the additional NNI interfaces are in different cities. This helps protect against node failure and provides a higher level of protection to the networks.

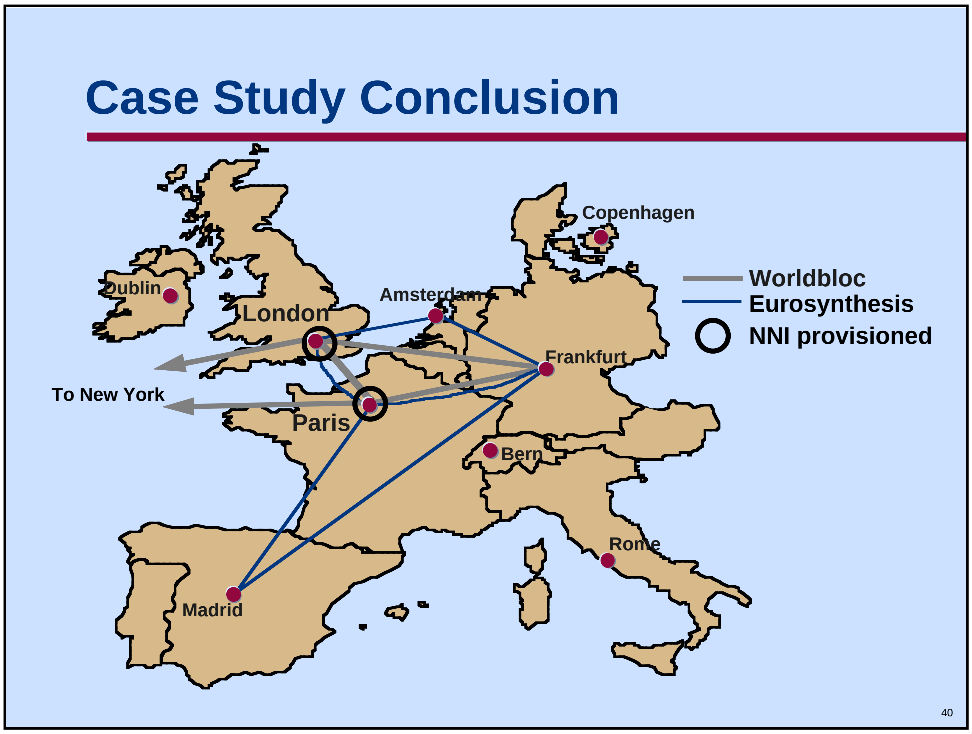

The analysis has indicated that Paris is the best location for a NNI. In order to meet the criteria of single link failure survivability, multiple high-speed NNI's should be provisioned.

In addition, Eurosynthesis and Worldbloc can take advantage of a more even distribution of traffic load if a NNI is also located in London. Eurosynthesis will incur a small incremental charge for facilities between Paris and London to protect against a failure of the NNI in London. However, this is offset by the added redundancy, node diversity and growth opportunity.

In this case study, the providers should locate multiple high-speed NNI interfaces in London and Paris.

Summary

- Magellan products simplify task of network engineering
  - dynamic routing simplifies design
  - differentiate traffic with multiple classes
  - increase efficiency with loadspreading and loadsharing
  - exploit benefits of oversubscription
- NNI is operational today
- Frame relay engineering is evolving
  - focused on facilitating your growth and developing your business opportunities

41

The engineering of frame relay networks is continuing to evolve as the service grows. New applications are being developed and the integration of frame relay with other technologies is becoming more important. For Nortel, this evolution includes partnering with our customers to deliver solutions to the business challenges of today and tomorrow. Nortel is focussed on facilitating your growth, and on developing opportunities for your business.