

Connection Management

Introduction

This chapter is provided for users who wish to have an in-depth knowledge of the IPX and BPX connection management functions. Reading this chapter is not required in order to use the systems. It describes packet queuing and the various queue types. It also discusses circuit routing and rerouting, delay for various types of connections, and circuit bandwidth requirements and utilization.

IPX Packet Queuing

As previously discussed, packets may be created in an IPX node by any of the following cards: NPC, CDP, SDP, LDP, or FRP. Each of these cards creates one or more different types of packets, each of which is handled separately for purposes of packet queuing in the NTC cards.

Each NTC contains a routing table in RAM to determine which packets it should take from the system bus MUXBUS for transmission on its trunk. It checks the address on each MUXBUS packet against this table and, if a match is found, it reads the packet into one of its queues. The separate queues allow the NTC to set transmission priority for different packets depending on the type of information they carry.

Packets are removed from the system bus MUXBUS in the node and queued for transmission by a trunk card. Different types and models of trunk cards support different packet types. For example, the NTC Model B support only high priority, non-timestamped, timestamped, and voice packets. The NTC Model C support all six packet types.

In the NTC Model C or later, trunk cards, the queue service algorithm:

- Supports the two frame relay packet types (bursty data A and bursty data B).
- Gives high priority packets unrestricted bandwidth access.
- Guarantees that every packet type is provided some minimum bandwidth access regardless of the load of any queue.

This is accomplished using the Credit Manager as described in the section on frame relay. In these cards, high priority packets are not subject to the credit manager scheme. The other five packet types, however, are issued credits one by one at a rate equal to the configured load of that type of packets on the trunk. Furthermore, each packet type may not receive a credit if it has not used its previous credit.

As an example, assume that a trunk has the following load:

- 2000 voice packets per second.
- 1000 timestamped packets per second.
- 4000 non-timestamped packets per second.

Then, as frames go by, each of these packet types is eligible to accrue credits as indicated in Table 13-1.

Table 13-1 Credit Accumulation

Frame	1	2	3	4	5	6	7	8	9	10	11	12	13	14
Voice Credits	X				X				X				X	
Timestamped Credits			X								X			
Non-timestamped Credits		X		X		X		X		X		X		X

At every opportunity to send a packet, the NTC Model C runs the following queue service algorithm to determine which packet to send.

- Step 1** If there is a high priority packet, send it.
- Step 2** If there is no high priority packet, then examine each other queue in order of highest to lowest configured bandwidth. If a queue has a packet and a credit, send the packet.
- Step 3** If no queue has a credit, then use the following priority to send a packet.
- Non-timestamped data.
 - Timestamped data.
 - Voice and bursty data in a round-robin fashion.
- Step 4** If there is no packet to send, send a 4-byte idle packet.

This scheme allows every queue to use at least some minimum configured bandwidth. Any packets that exceed the configured bandwidth are handled in the order described, which gives a slight edge to non-timestamped, then timestamped data.

Delay in a Cell Network

The overall delay of information through the network includes:

- Packetization delay.
- Queuing delay.
- Transmission delay.
- Depacketization and null-timing delay.

Delay in Packet Frame Relay Networks

Bursty data packets are built in the FRP card as the data is received from the port. Therefore, the packetization delay is inversely proportional to the speed of the port. Essentially, the time to fill a packet is the time it takes to assemble 160 bits at the bit rate of the port. This delay is only relevant if the connection is not throttled in the FRP due to the credit manager scheme implemented there to prevent network congestion.

When designing an IPX frame relay network, the goals are to minimize delay, congestion, data loss and cost and to maximize bandwidth utilization. Minimizing delay and maximizing bandwidth utilization are usually conflicting goals.

Minimizing delay in the FRP card minimizes congestion and data loss at the source (or destination) point. However, care must be taken not to shift these problems to the network trunks. Note that the frame relay parameters that reduce delay always increase bandwidth utilization and vice versa. For instance, increasing MIR to reduce delay also increases the trunk bandwidth allocated. Or, reducing %utilization to reduce the trunk bandwidth allocated to frame relay can cause congestion in the NTC or AIT trunk cards, resulting in greater delay.

The following are a couple of general suggestions that can be applied when setting up an IPX frame relay network.

- At any given setting for port speed and connection MIR, the average queuing latency will be somewhat less in the FRP than in an access device. So, settings that move queuing delay from the access device to the FRP will generally lead to an improvement.
- Frames broken up into FastPackets move more quickly through the NTC/AiT trunk queues and trunks than they do through the connection virtual circuit queue (VC Q). So settings that move data more quickly onto the trunks are preferred.

Delay at the Source

Delay in the most common devices sending frames to the FRP (e.g. bridges, routers, etc.) follow the standard store-and-forward model for simple queues. This delay is changed by changing the port speed parameter (configured clock in **cnfrport** command). For a given amount of traffic, increasing port speed reduces delay in the access device.

However, when modifying port speed, delay in the access device must be considered in conjunction with delay in the FRP. Increasing port speed and holding MIN constant will increase delay in the FRP for that connection. While this produces a small overall reduction in delay, it also moves delay to the FRP where there is more control over it.

Delay in the FRP can be controlled by modifying MIR, Cmax, and VC Q depth. For a given amount of traffic, the greater value for MIR, the less the delay in the source FRP. Likewise, if MIR is less than the setting for port speed, the larger Cmax is, the smaller the delay. Increasing Cmax has the same effect on delay as increasing MIR. However, large Cmax can cause occasional congestion on the trunks. The value for VC Q depth sets the maximum allowable delay in the source FRP. But reducing this may result in discarded frames, which is generally unacceptable.

Delay in the Network

There are two primary sources of frame relay connection delay in the IPX network:

- intermediate node delay
- propagation delay

Intermediate node delay consists of processing, queuing, and transmit delay and is generally one to two milliseconds per hop even in a heavily loaded network. Even on a 10-hop connection, this delay would be under 20 milliseconds. This is assuming that care has been taken to prevent data loss that can significantly increase overall delay.

Propagation delay is generally small except in international networks. At roughly one millisecond per 100 miles, propagation delay on a 2600 mile connection (e.g. San Francisco to Boston) would be about 26 milliseconds.

If several of the connections on a trunk have large values for Cmax (on the order of 100 or more), then the possibility of short-term congestion arises. If all connections burst at once, the bursty data queue in the NTC or AIT will get very long. However, this condition should normally be of a short duration.

The connection utilization parameter, % utl, controls bandwidth allocation for frame relay connections on the trunks. Oversubscription, where the bandwidth allocated is significantly less than the connection MIR values, can allow trunks to become overloaded. This can result in congestion and data loss over network trunks resulting in significant increases in end-user delays.

Delay at the Sink

The sink FRP is the card that sends frames to the destination access device. Generally, delay in the sink FRP follows the same model as delay in the access device. However, if the sink access device is attached to a LAN, then increasing port speed can significantly reduce delay in the sink FRP without significantly increasing delay in the access device.

Another method for decreasing delay at the terminating FRP for selected connections is to assign a high priority to these connections. Frames for high priority connections are assembled in a separate output port queue from low priority connections. All frames in the high priority queue are transmitted before any frames are transmitted from the low priority queue.

Care should be taken when reducing the Port Queue Depth parameter in the **cnfrport** command as this could result in unnecessarily dropping frames if this queue should overflow. Where the sink FRP receives traffic from only one source and has a port speed that is greater or equal to the MIR, the queuing delay does not follow the normal model and is very small.

Synchronous Data Connection Delay

For all time-stamped and non-timestamped data packets, the number of information bytes in a packet varies from 4 to 21 bytes depending on the type and speed of the connection. The packetization delay for the two types of data packets can be calculated by using Table 13-2 or Table 13-3 or looked up in the tables at the end of this chapter.

There is a delay from the first information byte clocked into the packet buffer to the last. The lower the bit rate of the channel, the longer the packetizing delay would become. To keep this time low, packets are formed from as few as 4 bytes of information for low-speed channels. This, and the time necessary for the card's firmware to add address, priority, DFM and timestamp information to the buffer, constitutes packetization delay. The packet is then placed on the system bus.

In the SDP, non-timestamped packets are received for a particular channel, the header is discarded and the information placed in a flexible buffer. When the connection is first set up, the buffer is half-filled. This allows variations in transmission delay to be accommodated until the buffer overflows or underflows. It also allows for short-term variations in the clocks at the transmitting and receiving interfaces.

Timestamped packets are buffered in the receiving SDP until the timestamp has reached the maximum age set in the Configure System Parameter (cnfsysparm) command, then clocked out. Therefore all timestamped connections have a one-way delay approximately equal to the “maximum timestamped packet age” set in the Configure System Parameter command plus packetization and transmission delays.

EIA lead information (non-interleaved) and clock-speed information (isochronous connections) is sent in supervisory packets, SDP to SDP. These packets appear to the network like the data packets of the same connection. Therefore, their delay through the network should be the same as the data stream. However, because they are sampled asynchronously and packetized and depacketized through different paths of the SDP, their changes are time-shifted with respect to the data.

Ways to Reduce Data Connection Delays

Normally, data circuit delay is not a problem. For some user data devices transmitting over large networks, the data delay may appear to cause some minor problems. The following are several suggestions for reducing the network delay.

- Make sure connections take the shortest route (number of hops trades off against transmission distance).
- If many connections are sharing a trunk, there may be queuing delays in the transmitting trunk card. Reroute connections to balance the load through the network.
- Raise the baud rate or configure interleaved EIA.
- Reduce the “maximum timestamped packet age” (trade-off against dropped packets).
- Increase the “preage” value (trade-off against dropped packets).
- For EIA leads, increase the update rate or configure the connection with interleaved EIA. This is affected by “maximum timestamped packet age” for the network, and “time-stamp preaging” for the connection

Table 13-2 Calculation of Non-Timestamped Data Packet Overall Delay

Source	Delay (ms.)
Transmitting SDP packetizing	1–3
Transmission delay (terrestrial, per mile)	0.01
Transmission delay (satellite, per hop)	300
Miscellaneous dejitter delays (per hop)	0.25
Receiving SDP null timing buffer (per hop) ¹	2.5–5.0
Receiving SDP processing delays	4.0
Receiving SDP isochronous buffer delay ²	10.0
Minimum delay (one-hop, colocated nodes)	7.75

1. Includes trunk queuing delays

2. Isochronous connections only

Table 13-3 Calculation of Timestamped Data Packet Overall Delay

Source	Delay (ms.)
Transmitting SDP packetizing	3–33
Transmission delay (terrestrial, per mile)	0.01
Transmission delay (satellite, per hop)	300
Miscellaneous dejitter delays (per hop)	0.25
Receiving SDP null-timing buffer	40
Receiving SDP processing delay	4.0
Receiving SDP isochronous buffer delay ¹	10.0
Minimum delay (one hop, colocated nodes)	47.25

1. Isochronous connections only
2. Ages timestamp to 40 (default).

Voice Connection Delay

Non-VAD voice packets

For a voice channel without VAD (“p”, “d” or “a”), the packetization delay is constant. It is the time for 21 bytes or 42 bytes to be processed by the CDP or 2.625 msec for a “p” connection and 5.25 msec for an “a32” connection.

VAD voice packets

For a voice channel with VAD (“v” or “c”), there is a VAD software parameter, sample input delay (SID) that defines the size of a serial register in the CDP. This adjusts “front end clipping” but increases the end-to-end delay of the connection by the amount of buffer delay.

Transmission delay across a trunk is generally a function of the distance travelled. For a terrestrial trunk, signals travel an average of about 100 miles per millisecond, or 0.01 msec/mile. For a satellite trunk, the propagation delay of the signal to the satellite and back adds approximately 300 msec per satellite hop. Table 13-4 can be used to calculate delay for the four types of voice connections.

Table 13-4 Sources of Delay in Voice Connections

Delay Source	t & p	v	a16	a24	a32	c16	c24	c32
Circuit T1 transmitter dejitter	0.25	0.25	0.25	0.25	0.25	0.25	0.25	0.25
Transmitting CDP sample input delay	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5
Transmitting CDP packetization	2.6	2.6	10.5	7.0	5.25	10.5	7.0	5.25
Transmitting TXR queuing (per hop)	< 2.5	< 2.5	< 2.5	< 2.5	< 2.5	< 2.5	< 2.5	< 2.5
Transmission delay (terrestrial, per mile)	0.01	0.01	0.01	0.01	0.01	0.01	0.01	0.01
Transmission delay (satellite, per hop)	300	300	300	300	300	300	300	300
Miscellaneous dejitter delays (per hop)	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5

Table 13-5 illustrates some typical expected delays for various number of terrestrial hops. The calculations do not include any PBX or channel bank delay.

Table 13-5 Typical Voice Connection Delays

Expected Delay (ms.)	t or p	v	a	c
1-hop, colocated nodes ¹	9	30	14	35
2-hops, colocated nodes ¹	12	33	17	38
3-hops, colocated nodes ¹	15	36	20	41
4-hops, colocated nodes ¹	17	38	23	43
5-hops, colocated nodes ¹	20	41	25	45

1. For nodes that are not colocated nodes, add 0.01 ms./mile.

Maximum Hops—Voice Connections

With the NTC, as the number of hops for a connection increases, the possible fluctuation in delay increases also. This is because each NTC in the path may add delay up to the maximum allowed by the queuing parameters, depending on the other traffic passing over that trunk. The CDP has a large buffer so the buffer size does not limit the maximum number of hops for a voice connection.

Routing and Rerouting

This section discusses how the IPX determines the route each circuit takes when added to the network and the algorithms and considerations involved when the IPX must automatically reroute circuits because of a failure(s) detected in the network.

Load Model

The IPX maintains a load model of the network and uses it to make decisions for routing and failing to route connections. The inputs to this model are the number and type of all connections routed on each trunk, and the configured utilization figures for VAD and DFM connections (measured as a percent of nominal connection bandwidth).

These utilization figures are set by the network administrator. Defaults are 40 percent for VAD and 100 percent for DFM and frame relay. It is these figures that are used to determine how many connections may be routed over a trunk, and when no more bandwidth is available. This is without reference to the “real world” performance of VAD and DFM and frame relay.

The challenge of network optimization is to make these configured utilizations reflect reality, in order to gain the maximum possible E1 or T1 pair gain, and to predict and influence the performance of the network in extreme cases of trunk failures or unfavorable statistics.

Load Model and Routing for Frame Relay

Frame relay connections differ from others because they have a greater range of instantaneous packet rates. A connection with a minimum rate of 512 Kbps may generate no packets for a long time, then suddenly generate 10 or 20 packets in a row (depending on the value of Cmax) at the frame relay port speed.

Since the connection has accepted data and processed it through the FRP card very quickly, and since the delay across the connection depends directly on the queuing delay of the last packet in the frame, it is important to ensure there are no unnecessary bottlenecks in the network trunking.

When the first frame relay connection is routed over the trunk, the load model in software allocates the entire bursty data peak bandwidth. This is important for networks mixing frame relay with other traffic, as it ensures that when a frame relay burst reaches the IPX trunk card, the bandwidth available is at least the bursty data peak.

As more frame relay connections are routed over the same trunk, the statistical addition of the different sources allows them to share bandwidth more efficiently. Because of this, the user can allocate only a portion of the trunk bandwidth required for each new frame relay connection added (the default is 121%, which equates to 100% usage for user data and the remainder for the overhead of encapsulating the frame relay data into FastPackets). This oversubscription of bandwidth is can also be extended to the IPX MUXBUS bandwidth reserved for each FRP. This factor can be decreased for slots where there are many PVCs transmitting at lower rates (e.g. 56 Kbps and less).

The routing algorithm (using frame relay optimization) allocates routes for new connections to minimize extra bandwidth allocation, and so tends to route frame relay connections over the same trunks that the earlier connection took. This results in good bandwidth efficiency.

A priority (high/low) can be assigned each ForeSight frame relay connection as it is added to the network. High priority connections are routed through a separate transmit queue in the FRP receiving the packets. The frames in the high priority queue are output before frames in the low priority queue. This reduces the queuing delay for these frames.

ForeSight is a closed-loop system that dynamically allocates trunk bandwidth based on the connection parameters set. If there is any excess bandwidth available after all the committed information rates have been satisfied, it will portion out the excess bandwidth based on each connection's CIR.

Routing Algorithm

Each node has, in a database, a representation of the network topology. This includes all trunks and their status, and all connections, their type and route. From this, the node calculates the load (packets/sec or cells/sec) on all trunks in the network.

When a connection is routed, the owner node determines the bandwidth requirements from lookup tables and the destination node and channel from the network database. If the connection has a preferred route (direct routing), it will attempt to comply if at all possible. The routing can also be specified by the operator to be restricted to a terrestrial route only or if a satellite route is acceptable.

The search for a circuit route is begun by first examining all trunks to adjacent nodes (in order of trunk number). If the route has enough bandwidth for the circuit (or bundle), and the terminating node is found to be the other end of the connection, the route has been found and the search is terminated. Otherwise, the search is continued.

When all single-hop routes have been examined but found lacking, the search is extended to nodes at a distance of two hops. The search radius is enlarged from the master node. Eventually, the search is successful, or completes without success, or the search times out without success. If a preferred route is specified and is unavailable for a connection that has directed routing, the connection will be marked immediately as failed.

When a search is successful, the route information (trunks and nodes on the path) is broadcast to all nodes on the chosen route so they can update their network topology models in their database.

The network does not continually look for new routes unless there are connections failed for lack of a route. If this is the case, the addition of trunks or deletion of connections is necessary. The network does not rearrange connections that are already routed to accommodate a connection that is not routed, even though the new connection may have a high priority Class of Service.

Likewise, if the statistical reserve on trunks is decreased, the network takes no action except to route any failed connections that can now be routed. However, if statistical reserve is increased, all connections in the network will be failed and rerouted as some previously used routes may no longer be available.

The IPX attempts to balance loads between trunks. This allows the adaptive voice feature to give better results, but affects all connections. The reroute algorithm finds all routes with the shortest hop count. It chooses the route where the current load on the most heavily loaded line of the route is a minimum. In order to force even balancing, the size of a routing bundle is restricted.

When a connection is first added to the network, software identifies the first route available in the usual way, finding the fewest hops given restrictions of trunk type (satellite/terrestrial) and current loading (there must be bandwidth available). It then finds all other routes of the same number of hops and chooses the route with the lowest loading factor.

Causes of Rerouting

The IPX does not move connections from existing routes unless one of the following conditions exists:

- The current route fails (i.e. trunk failure).
- Some other cause (i.e. card, circuit line) fails the connection, then clears.
- Parameters affecting permitted trunk loading (i.e. statistical reserve) are changed so as to reduce the bandwidth available for connections.
- The circuit has been “bumped” or preempted by a higher priority circuit where bandwidth is limited.

It is important to realize that the algorithm does not move working connections between trunks to balance load: the balancing occurs when a connection without a route is allocated one. A working connection is rerouted when its preferred route (when different from the current route) becomes available.

Reroute Priority and Order

For every connection, there is a master node (the owner). This node, where the connection was added is responsible for finding a route and rerouting the connection in the event of a failure. Master nodes act independently. If a trunk fails in a network, all nodes owning connections routed over that line recognize the failure since the information is broadcast to all nodes in the network. As each node recognizes the failure it attempts to reroute its connections without reference to the others.

For this reason, it is recommended that ownership of connections be concentrated in a small number of nodes. There will be fewer collisions in rerouting, and, since the class-of-service priority is followed within each node but not coordinated between nodes, performance will be more predictable and closer to that desired.

Within each node, the order of precedence for routing connections is determined by:

- **COS**—Connections with the lowest class-of-service (COS) are considered first. (Lowest number = highest priority.) There is a pause of 250 msec between each COS, to improve the network-wide effects of COS.
- **Bandwidth**—Connections with the highest bandwidth are considered first. When bandwidth on trunks is limited, it is easier to route a small bandwidth connection than a large one. Data connections vary in bandwidth, but the order is “p”, “d”, “a”, “v”, “c” for voice.
- **Bundle size**—Bundles are routed without being split, if possible. This makes the operation quicker, and fewer individual routes need to be found. If no route is available for the bundle, it will be split into single connections and the search repeated for each.

When a node has to find routes for a number of connections at the same time, it uses the rules above to determine the order in which it considers them. They are hierarchical. Bundle size will only be considered if there are a number of bundles of connections of the same type and COS. If a route cannot be found for a particular connection, the owning node will leave it failed and go on to the next. This is why the “largest first” rule is important. The network cannot reroute some connections to make room for others. Rerouting only occurs as the result of the failure of routed connections.

When a group of connections is failed, a timer is started at the node owning the connections. COS 0 connections may be rerouted immediately, and there is a 250 millisecond delay before each subsequent COS may be rerouted. This is to allow COS to have a network-wide effect. Therefore, COS 8 connections will be rerouted after a pause of 2 seconds although there may still be COS 0 connections awaiting rerouting. The low COS gives a “head start” rather than absolute priority.

After the COS timer, priority is given to connections with the highest bandwidth (packets/second) of the group awaiting rerouting. This is because, as available bandwidth diminishes, it is more difficult to find routes for the higher bandwidth connections. The data block for each connection contains the packet/second requirement, so prioritizing is easy. The general rerouting priority order is given in Table 13-6.

When several similar connections have the same source and destination node, they can be routed as a bundle. This saves time, as only one route is found for several connections. Bundle size is the least important rerouting priority.

Table 13-6 Priority for Rerouting

Priority	Connection Type
1	high speed data connections (>64 Kbps)
2	"p" or "d" connections
3	"a" connections
4	"v" connections
5	"c" connections
6	low speed data connections (< 9.6 Kbps)

Priority Bumping/Courtesy Downing

The process of Priority Bumping controls several components. When bumping is invoked, it first looks for any locally owned connections that are currently not routable. If no connections are found, then this process is complete. When connections are found, the database that holds remote nodes' owned connections is searched to locate any other connections that need to be routed.

When remote nodes own connections that need to be routed with higher COS than the local node's connections, a 45 second wait period is started. This gives the remote nodes time to bump and route their connections before the local node attempts to bump. If no higher COS connections need to be routed, then the first node may proceed without waiting.

When the first node proceeds, it builds a list of up to N connections at this COS, including connections owned by other nodes. It then starts the background network load analysis process to find connections to deroute to make room for the connections on this list. It then starts a slow background timer to watch for excessively long processing.

If the timer expires before the background process completes, an error is logged, and the background process is aborted to be run again later. If the background process completes, then it returns a list of any connections to be derouted (bumped) to make room. This list of connections is used to send bump messages to each node that masters those connections. A broadcast message is sent to all nodes with the COS of the bumping connections to prevent lower COS connections from routing when bumping frees bandwidth.

The background process also returns a list of the local connections that could be routed. When any of the local connections cannot be routed, their connection state is changed to indicate they failed as a result of being bumped. This indicates that they should be analyzed for bump routing at a less frequent rate than other non-routable connections.

If there were connections owned locally that now should be routable, then a 15 second wait timer is set to allow the derouted connections time to send out updates. Finally at the end of this period of time, the rerouting process is started to allow it to route the connections.

System Message Traffic Routing

System messages are carried between node controller cards (NPC and BCC) in high priority packets called CC packets. The route used by any pair of controller cards to communicate is determined automatically by the network and is fixed as long as there are no changes to the network topology that affect the choice.

The criteria used to select a route between two controller cards are as follows.

- 1 The network selects the route with the fewest trunks that restrict controller traffic. A user may want to restrict a trunk that uses almost all of its bandwidth for customer traffic from carrying inter-node traffic to free up the bandwidth. This is done with the *Restrict CC Traffic* parameter in the Configure Trunk (**cnftrk**) command.
- 2 The network then considers the route with the fewest satellite trunks. A satellite trunk is entered in the *Link type* option of the Configure Trunk (**cnftrk**) command. The network has no way of determining whether a trunk actually uses a satellite.
- 3 The network then selects the route with the biggest "choke point." The network determines for every route, a "choke point", which is the trunk that has the least total bandwidth capacity. The network then selects the route with the least restrictive choke point.
- 4 The network then selects the route with the least total number of hops.
- 5 If there are still choices available, internode communication will travel over the lowest numbered trunk (of the choices being considered) on the node that has the lowest number of the two nodes.

Every packet or cell that is sent between node controllers is acknowledged by the recipient. The maximum time that a controller will wait for an acknowledgment is 1.7 seconds. If no acknowledgment is received in time, the node will retransmit the packet/cell and wait another 1.7 seconds.

The maximum number of attempts, 5 or 7, depending on whether there are satellite trunks in the communication path between the nodes or not. If acknowledgment is received after the maximum allowed attempts, the far node is declared unreachable. This represents a **communication break** condition.

Bandwidth Allocation

One of the benefits of the IPX network is the compression of voice (VAD) and data (DFM) connections to allow cost savings through pair-gain. However, these features both depend on statistical properties of the data offered to the system. Therefore, their exact level of effectiveness is not easily predicted. VAD may result in a 0 percent to 70 percent bandwidth savings, for instance, whereas the effectiveness of ADPCM, (50 percent savings for 32 Kbps ADPCM), is predictable and unchanging.

Since the total traffic capacity of an IPX network is somewhat difficult to predict, StrataCom has developed a software Network Modeling Tool (NMT). This allows the user to analyze his proposed network to determine if there will be sufficient capacity available. For further information on the NMT, refer to the *Network Modeling Tool User's Guide*.

Network Trunk Bandwidth

The system calculates the available bandwidth of each network trunk as follows:

T3 Framed:

- The available T3 trunk bandwidth is 44.736 Mbps minus framing and overhead.
- Each PLCP frame carries 53 eight-bit cells that occur at an 8 KHz rate for a total of 40.704 Mbps of user data, equivalent to 96,000 cells/sec.
- Each cell can transport up to two FastPackets for a rate of 192,000 FastPackets/sec.
- T3 trunks terminating on IPX nodes are limited to 80,000 FastPackets/sec. This allocation is half of the IPX MUXBUS capacity.

E3 Framed:

- The available E3 trunk bandwidth is 34.368 Mbps minus framing and overhead.
- Each ITU-T G.804 frame carries 530 payload octets that occur at an 8 KHz rate for a total of 33.920 Mbps of user data.
- The G.804 E3 frame can transmit 10 ATM cells per frame at 8000 frames/sec. or 80,000 cells per second.

T1 Framed:

- The available trunk bandwidth is 1.544 Mbps – 8 Kbps framing = 1.536 Mbps.
- Framed T1 lines can carry 1.536 Mbps / 192 bits/packet = 8,000 packets/sec.

E1 Framed:

- The available trunk bandwidth is 2.048 Mbps – 64 Kbps framing = 1.984 Mbps.
- Framed E1 lines can carry 1.984 Mbps / 192 = 10,333 packets/sec.

E1 Unframed:

- The available bandwidth is 2.048 Mbps.
- Unframed E1 lines carry 2.048 Mbps / 192 = 10,666 packets/sec.

Subrate:

Depends on the number of DS0's available in the subrate trunk. See Table 13-7.

Table 13-7 Subrate Packet Line Bandwidth

DS0s	BW	DS0s	BW	DS0s	BW	DS0s	BW
1	n/a	9	3000	17	5666	25	8333
2	n/a	10	3333	18	6000	26	8666
3	n/a	11	3666	19	6333	27	9000
4	1333	12	4000	20	6666	28	9333
5	1666	13	4333	21	7000	29	9666
6	2000	14	4666	22	7333	30	10000
7	2333	15	5000	23	7666	31	10333
8	2666	16	5333	24	8000	32	10666

Note It is recommended that a subrate trunk be configured with at least four DS0s to provide sufficient statistical reserve for inter-node communications traffic.

A packet slice on the TDM bus is 1000 packets/sec, therefore an E1 trunk requires 11 packet slices of TDM bandwidth for a total of 11,000 packets/sec per E1 trunk. The T1 trunk requires 8 packet slices for a total of 8,000 packets/sec. per T1 trunk.

The total bandwidth available on the IPX backplane MUXBUS, excluding NPC-reserved bandwidth, is 30.72 Mbps. This corresponds to 30.72 Mbps / 192 = 160,000 packets/sec. Therefore, the maximum number of E1 trunks in a node is 160,000 / 11,000 = 14. The maximum number of T1 trunks per node is 160,000 / 8,000 = 20. However, software limits this to 16 trunks.

Each 64 Kbps time slot, or DS0, provides 1/3 X 1000 or approximately 333 packets per second of available bandwidth on a trunk. Table 13-7 shows the packet bandwidth available on a subrate trunk as a function of the number of DS0s regardless of the trunk type.

Voice Compression Bandwidth Requirements

The bandwidth required on a trunk to carry the information on a DS0 circuit depends on which one of the five compression types is selected for the circuit as indicated in Table 13-8. The equivalent bit rate after compression is also listed in this table.

Compression is an effective means of reducing the network bandwidth requirements but does degrade the quality of the voice circuit. Note, however, that any circuit that may at times have a fast modem or FAX connection will automatically revert to a “p” connection during the transmission with attendant increase in bandwidth required.

Table 13-8 IPX Voice Grades

Type	Equivalent Bit Rate	Required BW
p	64 Kbps	381 pkts/sec.
t	64 Kbps	381 pkts/sec.
v	32 Kbps	191 pkts/sec.
a32	32 Kbps	191 pkts/sec.
a24	24 Kbps	143 pkts/sec.
a16	16 Kbps	95 pkts/sec.
a16(z)	16 Kbps	95 pkts/sec.
c32 ¹	16 Kbps	95 pkts/sec.
c24 ¹	12 Kbps	72 pkts/sec.
c16 ¹	8 Kbps	48 pkts/sec.
c16(z) ¹	8 Kbps	48 pkts/sec.

1. Assumes 50% VAD.

VAD and DFM Effects

Voice activity detection takes place in the CDP card before speech is transmitted over a trunk. If speech is present, packets are sent. If speech is not present, no packets are sent and the trunk bandwidth may be used by other connections.

Similarly, DFM allows packets whose contents can be predicted by the receiving card, those containing repetitive patterns, to be suppressed. It should be noted that a DFM packet uses one more byte for control information (sequence byte) than the packet for a corresponding non-DFM connection. If the data cannot be compressed to less than 90 percent utilization by DFM, bandwidth savings will be made by disabling DFM for that connection.

Before the development at StrataCom of statistical tools, VAD was assumed to save 60 percent of nominal bandwidth. Experience has shown this to be a good estimate in most cases. But if a connection is not off-hook all the time (less than 36 ccs/hr) this estimate may be too high. Likewise, if there is high background noise on the circuit, this estimate may be too low.

With new statistical tools provided by StrataView Plus NMS, this utilization can now be measured on an active network. Voice and data compression can be treated in similar ways. For effective traffic studies, it is necessary to configure utilization figures for voice in the same way as data and this section treats both forms of compression similarly.

Data Channel Packet Generation Rates

The synchronous data channels use widely varying amounts of trunk bandwidth depending on whether they use timestamped data packets or not and how the control lead information is carried. Refer to Table 13-9 through Table 13-13 for bandwidth requirements or calculate as follows.

- Timestamped data packets, 8/8 coded, not fast EIA, have 160 data bits per packet.
- Timestamped data packets, 7/8 coded, not fast EIA, have 140 data bits per packet.
- Non-timestamped data packets, 8/8 coded, not fast EIA, have 168 data bits per packet.
- Non-timestamped data packets, 7/8 coded, not fast EIA, have 147 data bits per packet.
- Fast EIA packets, 8/8 coded, have 80 data bits per packet.
- Fast EIA packets, 7/8 coded, have 70 data bits per packet.
- DFM packets, 8/8 coded have 152 data bits per packet.
- DFM packets, 7/8 coded have 133 data bits per packet.

Exceptions are the low-speed connections listed in Table 13-11, Table 13-12, and Table 13-13, where partially-filled packets are used to reduce packetization delay. Divide the bit rate of the connection by the number of user bits per packet and the result is the number of packets/second.

For DFM connections, the actual packet generation rate will depend upon the actual utilization. The load model uses the user-configured utilization to calculate the expected number of packets/second. Add between 0 and 20 packets/second for EIA updates (an isochronous clock implies 20 packets/second in the direction the clock is propagated).

Traffic Statistics

The IPX provides a number of statistical tools to assist in traffic studies. The object of such a study is to collect enough information so that an accurate figure for configured utilization may be chosen for each connection. The display of IPX statistics requires a StrataView Plus workstation connection to the IPX network. The StrataView Plus collects all of the operating statistics for a network and stores it in its database (usually on its own hard disk). Refer to the *StrataView Plus Operations Manual* for details of statistics displays and examples.

Table 13-9 Data Load Table with Standard EIA and No DFM

Bit Rate	7/8 Coding			8/8 Coding		
	Kbps	Pkt/sec	Byte/pkt	Delay, ms	Pkt/sec	Byte/pkt
1.2	43	4	24	38	4	27
1.8	65	4	16	57	4	18
2.4	35	10	29	30	10	33
3.2	46	10	22	40	10	25
3.6	52	10	20	45	10	22
4.8	35	20	29	30	20	33
6.4	46	20	22	40	20	25
7.2	52	20	20	45	20	22
8	58	20	18	50	20	20
9.6	69	20	15	60	20	17
12	86	20	12	75	20	14
12.8	92	20	11	80	20	13
14.4	103	20	10	90	20	11
16	115	20	9	100	20	10
16.8	120	20	9	105	20	10
19.2	138	20	8	120	20	9
24	172	20	6	150	20	7
28.8	206	20	5	180	20	6
32	229	20	5	200	20	5
38.4	275	20	4	240	20	5
48	343 ¹	20 ¹	3 ¹	300	20	4
56	381	21	3	350	20	3
57.6	392	21	3	360 ¹	20 ¹	3 ¹
64	436	21	3	381	21	3
72	490	21	3	429	21	3
76.8	523	21	2	458	21	3
84	572	21	2	500	21	2
96	654	21	2	572	21	2
112	762	21	2	667	21	2
115.2	784	21	2	686	21	2
128	871	21	2	762	21	2
144	980	21	2	858	21	2
168	1143	21	1	1000	21	1
192	1307	21	1	1143	21	1
224	1524	21	1	1334	21	1
230.4	1568	21	1	1372	21	1
256	1742	21	1	1524	21	1

Bit Rate	7/8 Coding			8/8 Coding		
	Kbps	Pkt/sec	Byte/pkt	Delay, ms	Pkt/sec	Byte/pkt
288	1960	21	1	1715	21	1
336	2286	21	1	2000	21	1
384	2613	21	1	2286	21	1
448	3048	21	1	2667	21	1
512	3483	21	1	3048	21	1
672	4572	21	1	4000	21	1
768	5225	21	1	4572	21	1
772	5252	21	1	4596	21	1
896	6096	21	1	5334	21	1
1024	6966	21	1	6096	21	1
1152	7837	21	1	6858	21	1
1344	9144	21	1	8000	21	1

1. Connections below this rate generate time-stamped data packets. Connections above this rate generate non-time-stamped data packets.

Table 13-10 Data Load Table with Standard EIA and DFM

Bit Rate Kbps	7/8 Coding			8/8 Coding		
	Pkt/sec	Byte/pkt	Delay, ms	Pkt/sec	Byte/pkt	Delay, ms
1.2	58	3	18	50	3	20
1.8	86	3	12	75	3	14
2.4	39	9	27	34	9	30
3.2	51	9	20	45	9	23
3.6	58	9	18	50	9	20
4.8	37	19	28	32	19	32
6.4	49	19	21	43	19	24
7.2	55	19	19	48	19	22
8	61	19	17	53	19	19
9.6	73	19	14	64	19	16
12	91	19	12	79	19	13
12.8	97	19	11	85	19	12
14.4	109	19	10	95	19	11
16	121	19	9	106	19	10
16.8	127	19	8	111	19	10
19.2	145	19	7	127	19	8
24	181	19	6	158	19	7
28.8	217	19	5	190	19	6
32	241	19	5	211	19	5
38.4	289	19	4	253	19	4
48	361	19	4	316	19	4
56	422	19	3	369	19	3
57.6	434	19	3	379	19	3
64	482	19	3	422	19	3
72	542	19	2	474	19	3
76.8	578	19	2	506	19	2
84	632	19	2	553	19	2
96	722	19	2	632	19	2
112	843	19	2	737	19	2
115.2	867	19	2	758	19	2
128	963	19	2	842	19	2

*All of the connections below 56 Kbps generate time-stamped data packets.

Table 13-11 Data Load Table with Partially-Filled Packet and No DFM

Bit Rate	7/8 Coding			8/8 Coding		
Kbps	Pkt/sec	Byte/pkt	Delay, ms	Pkt/sec	Byte/pkt	Delay, ms
2.4/4	86	4	12	75	4	14
3.2/4	115	4	9	100	4	10
3.6/4	129	4	8	113	4	9
4.8/10	69	10	15	60	10	17
4.8/4	172	4	6	150	4	7
6.4/10	92	10	11	80	10	13
6.4/4	229	4	5	200	4	5
7.2/10	103	10	10	90	10	12
7.2/4	258	4	4	225	4	5
8/10	115	10	9	100	10	10
9.6/10	138	10	8	120	10	9
12/10	172	10	6	150	10	7
12.8/10	183	10	6	160	10	7
14.4/10	206	10	5	180	10	6

*All of the above connections generate time-stamped data packets.

Table 13-12 Data Load Table with Partially-Filled Packet and DFM

Bit Rate	7/8 Coding			8/8 Coding		
Kbps	Pkt/sec	Byte/pkt	Delay, ms	Pkt/sec	Byte/pkt	Delay, ms
2.4/4	115	3	9	100	3	10
3.2/4	153	3	7	134	3	8
3.6/4	172	3	6	150	3	7
4.8/10	77	9	14	67	9	15
4.8/4	229	3	4	200	3	5
6.4/10	102	9	10	89	9	12
6.4/4	305	3	4	267	3	4
7.2/10	115	9	9	100	9	10
7.2/4	343	3	3	300	3	4
8/10	127	9	9	112	9	9
9.6/10	153	9	7	134	9	8
12/10	191	9	6	167	9	6
12.8/10	204	9	5	178	9	6
14.4/10	229	9	5	200	9	5

*All of the above connections generate time-stamped data packets.

Table 13-13 Data Load Table with Fast EIA

Bit Rate Kbps	7/8 Coding			8/8 Coding		
	Pkt/sec	Byte/pkt	Delay, ms	Pkt/sec	Byte/pkt	Delay, ms
1.2f	35	5	29	30	5	33
1.8f	52	5	20	45	5	22
2.4f	35	10	29	30	10	33
3.2f	46	10	22	40	10	25
3.6f	52	10	20	45	10	22
4.8f	69	10	15	60	10	17
6.4f	92	10	11	80	10	13
7.2f	103	10	10	90	10	11
8f	115	10	9	100	10	10
9.6f	138	10	8	120	10	9
12f	172	10	6	150	10	7
12.8f	183	10	6	160	10	7
14.4f	206	10	5	180	10	6
16f	229	10	5	200	10	5
16.8f	240	10	5	210	10	5
19.2f	275	10	4	240	10	5
24f	343 *	10 *	3 *	300	10	4
28.8f	412	10	3	360 *	10 *	3 *
32f	458	10	3	400	10	3
38.4f	549	10	2	480	10	3
48f	686	10	2	600	10	2
56f	800	10	2	700	10	2
57.6f	823	10	2	720	10	2
64f	915	10	2	800	10	2
72f	1029	10	1	900	10	2
76.8f	1098	10	1	960	10	2
84f	1200	10	1	1050	10	1
96f	1372	10	1	1200	10	1
112f	1600	10	1	1400	10	1
115.2f	1646	10	1	1440	10	1
128f	1829	10	1	1600	10	1
144f	2058	10	1	1800	10	1
168f	2400	10	1	2100	10	1
192f	2743	10	1	2400	10	1
224f	3200	10	1	2800	10	1
230.4f	3292	10	1	2880	10	1
256f	3658	10	1	3200	10	1

Bit Rate	7/8 Coding			8/8 Coding		
	Kbps	Pkt/sec	Byte/pkt	Delay, ms	Pkt/sec	Byte/pkt
288f	4115	10	1	3600	10	1
336f	4800	10	1	4200	10	1
384f	5486	10	1	4800	10	1
448f	6400	10	1	5600	10	1
512f	7315	10	1	6400	10	1

*Connections below this rate generate time-stamped data packets. Connections above this rate generate non-time-stamped data packets.

Table 13-14 Data Load Table—Fast EIA with Partially-Filled Packet

Bit Rate	7/8 Coding			8/8 Coding		
	Kbps	Pkt/sec	Byte/pkt	Delay, ms	Pkt/sec	Byte/pkt
1.2f/2	86	2	12	75	2	14
1.8f/2	129	2	8	113	2	9
2.4f/5	69	5	15	60	5	17
2.4f/2	172	2	6	150	2	7
3.2f/5	92	5	11	80	5	13
3.2f/2	229	2	5	200	2	5
3.6f/5	103	5	10	90	5	12
3.6f/2	258	2	4	225	2	5
4.8f/5	138	5	8	120	5	9
6.4f/5	183	5	6	160	5	7
7.2f/5	206	5	5	180	5	6

*All of the above connections generate time-stamped data packets.

