

Networking

This chapter describes various types of networks or network applications that may be implemented with the StrataCom cell relay network switches, the IPX, the IGX, the BPX, and the associated equipment such as AXIS and FastPAD.

The IPX, IGX, and BPX, which are referred to as nodes, are the basic building blocks of StrataCom networks. Nodes are interconnected by transmission facilities called trunks or packet lines. User data inputs to the network at nodes through various user devices supplied by both StrataCom and OEMs.

Broadband ATM Networks

StrataCom ATM networking products are used to provide the advantages of Asynchronous Transfer Mode technology in a variety of network applications:

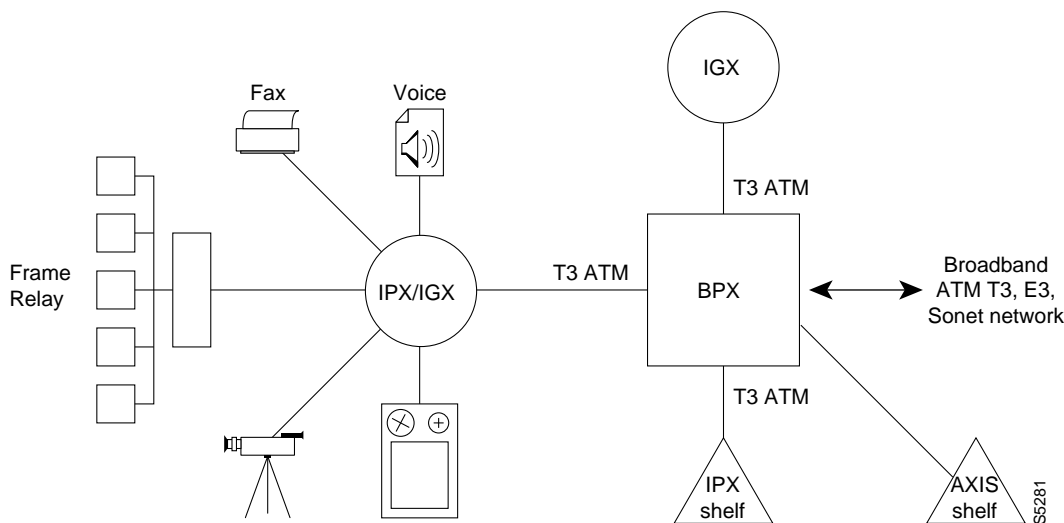
- Support of new, multi-media applications.
- Broadband high-speed backbones for large networks.
- Expansion of existing narrowband networks.
- ATM concentration for public service providers.
- Expansion of existing wide-area data networks.

Multi-Media Networking

One of the main advantages of ATM is its ability to integrate a wide variety of traffic types and provide a transport mechanism to equally support each of their unique requirements. ATM supports both constant bit rate and variable bit rate traffic equally well to open the door to multi-media networking.

ATM networks provide a suitable interface adaptation for each service type assuring the required amount of circuit bandwidth, the class of service, and the broadband aggregate capacity provided by the network ATM trunks.

This has already been successfully demonstrated over networks using the voice, data, video, and frame relay user interfaces currently available with the IPX, IGX, and the ATM connections to the BPX (Figure 5-1). Since the IPX has always been a cell-based node, it handles multimedia inputs with a maximum of efficiency. The IGX provides all the capabilities of the IPX but with additional multiband and multiservice capabilities.


Figure 5-1 Multi-media Network Application

High-Capacity Network Backbones

Another advantage of ATM technology is its scalability. The physical transmission facilities are separate from the various ATM data layers unlike other high-speed protocols such as FDDI. For example, an ATM backbone may be implemented initially with DS3 or E3 trunks operating over fiber or digital radio.

As network bandwidth requirements grow, the transmission facility may be upgraded to fiber operating at SONET rates of OC-3 or even OC-12. This will often require only a change-out of network trunk cards, not the whole node. The BPX currently supports T3, E3, and OC3 broadband trunks but the switching matrix and trunk slots provide the capacity to interface with network ports operating up to OC-12 rates.

Figure 5-2 illustrates a possible backbone configuration using T3, E3, or OC3 trunks arranged in a ring configuration for the ultimate in network robustness. Connections between nodes will be added to the network first on routes with the shortest path and the most available bandwidth. However, if a T3/E3/OC3 link should fail, the connections are automatically rerouted to reach their destination in the opposite direction assuming there is sufficient bandwidth available. In this example, as the network traffic load grows, a second or third T3/E3/OC3 ring can be added until such time that it becomes necessary to consider moving to even higher trunk capacities.

Figure 5-2 High-Capacity Backbone Configuration

Growth Path for Existing IPX/IGX Networks


Another application for the BPX is to provide a growth path for existing narrowband IPX networks. As network operators experience rapid growth in data communications requirements, existing IPX networks may become heavily meshed with tandem T1 connections approaching the capacity of the network.

BPX nodes may be added to existing IPX networks to greatly expand the network capacity and relieve congestion. The IPX or IGX can be configured either as routing nodes or as shelves. When configured as shelves in a tiered network configuration, they can provide additional ports without adding to the number of routing nodes. This simplifies the network routing, and at the same time frees up some of the IPX or IGX network ports. See Figure 5-3 for a typical example.

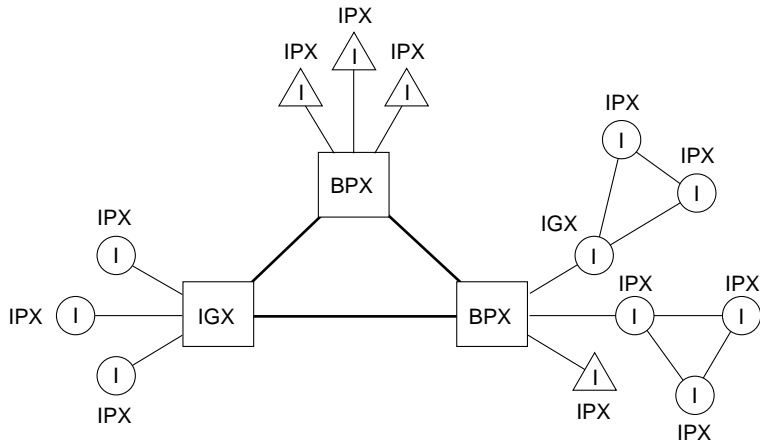
A BPX node operates as a peer with other IPX and IGX nodes in a mixed BPX/IGX/IPX network. From the network operator's perspective, the BPX node appears merely as a higher capacity IPX or IGX. The BPX can act as a regular node or as a "junction node" in structured networks (described later in this chapter).

IPX nodes are connected to a BPX node with the addition of an AIT card set. Similarly, IGX nodes are connected to a BPX with the addition of a BRM card set. The AIT card set operates at T3 or E3 rates and converts the IPX FastPacket format to ATM cell format. The AIT or BTM card set also provides frame relay to ATM interworking for the efficient transport of Frame Relay data across a ATM network. Normally, the BPX will be installed at sites where there are already several IPX or IGX nodes located, for example, at a network hub or large switching center and connect to these nodes with a short wire connection. The

Figure 5-3 Upgrading Existing IPX/IGX Networks



a. Heavily-meshed IPX/IGX network



b. BPX Add to IPX/IGX network

55172

Integrating a BPX backbone with its higher speed trunking and switching capacity expands the network capacity with minimal disruption of the existing network. The BPX provides expansion capabilities with full compatibility with the IPX and IGX and utilizes the same network management platform.

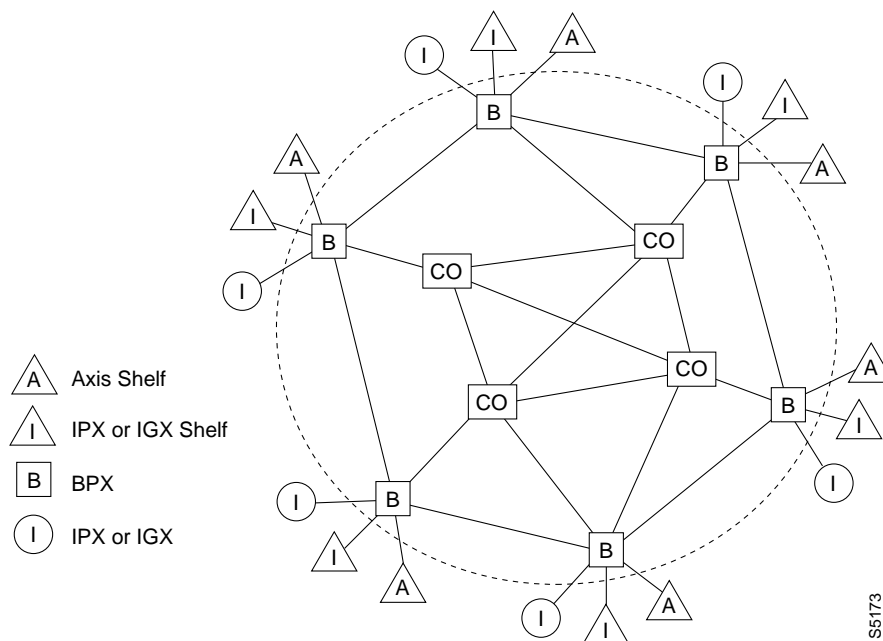
The actual bandwidth used on the trunks between a BPX node and an IPX or IGX node may be incrementally increased in T1 or E1 size steps. Therefore, existing IPX T1 and E1 trunks can be decommissioned one at a time and their cell traffic routed across the new BPX network backbone. The ATM interface to the IPX 8 economically serves small feeder locations and provides all of the same interfaces as the larger IPX 16/32.

Service Nodes in Public Networks

ATM can be used to economically provision a variety of services including frame relay, SMDS, compressed voice, on-demand video, and circuit emulation services. BPX nodes can serve as ATM service nodes for large central office digital switches in public switched networks (Figure 5-4). This allows the public service providers to offer ATM service to users who are not so concentrated as to justify the installation of a large ATM central office switch.

As an ATM service node, the BPX with AXIS shelves, IPX and IGX shelves, and IPX and IGX nodes, provides ATM access and termination closer to the user and transports ATM cells to the C.O. ATM platforms. Each BPX node acts as a concentrator for the narrowband inputs from a number of local user devices. The IPX, IGX, and AXIS convert these inputs to cells that are forwarded to the BPX nodes over T3 or E3 trunks. Local connections can be switched back to the outer layer of IPX and IGX nodes, and longer-distance connections can be switched to the public network at a suitably equipped central office.

Figure 5-4 Service Node for Large C.O. Switches

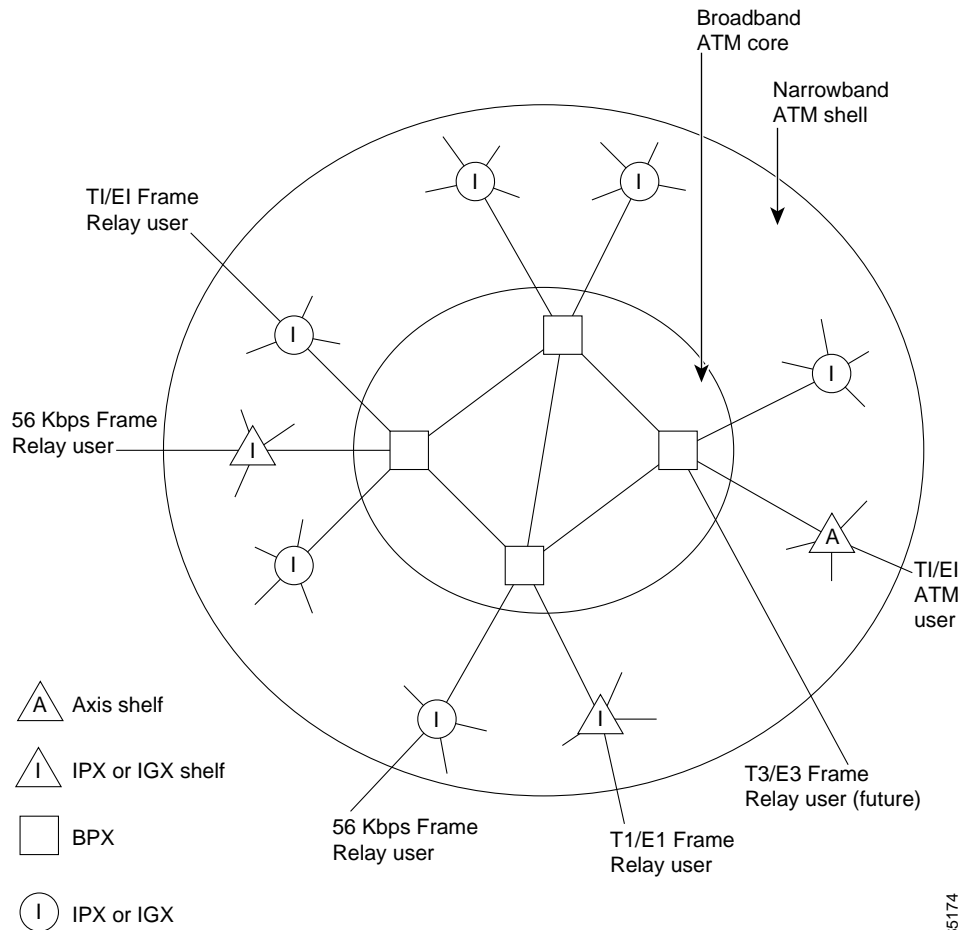


Expansion of Wide-Area Data Networks

Frame relay, linking LANs over wide areas, is a rapidly-growing offering of both public switched and private networks. Current user devices (routers, hubs, bridges, etc.) typically operate at rates of 56 or 64 Kbps and below on the network side. As the size of LANs grow and the operating rates of personal computers and workstations increase, new generations of LAN hubs, routers, etc. will operate at T1/E1 data rates and above offering many high-speed PVCs over a single port. The network bandwidth must also grow accordingly to provide the level of service initially provided.

It is envisioned that data networks will evolve to configurations employing a narrowband ATM shell ringing a broadband ATM core backbone network as illustrated in Figure 5-5. The narrowband shell, consisting primarily of IPX tributaries, IPX and IGX shelves, and IGX hubs will serve areas of low but increasing user density. This network configuration pushes the frame relay interface closer to the customer to reduce long-haul charges consolidating user data using the more economical hardware of IPX and IGX nodes.

Figure 5-5 ATM Network of the Future



S5174

For frame relay connections terminating at more remote destinations, the IPX and IGX nodes would feed this traffic to the higher-speed broadband core network utilizing the very high-speed tandem switching characteristic of the BPX nodes. These PVCs would likely traverse many hops. The higher speed of the broadband core network would offer reduced overall delay and increased trunk bandwidth avoiding the gridlock that could occur if this backbone traffic were routed over the narrowband shell.

For connections terminating at other local or intermediate destinations, the IPX nodes can communicate with adjacent IPX nodes using short-haul T1 or E1 facilities. Typical hop counts of 1 or 2 hops would be the norm for these connections.


To ease the transition from existing frame relay networks to networks based on ATM, StrataCom offers a wide range of LAN interconnect interfaces from low-speed frame relay e.g. 56 or 64 Kbps or less, T1/E1 interfaces with multiple virtual ports, to direct ATM data interfaces operating at T3 and E3 rates.

TT2 ATM Network Connections

A 6.312 Mbps T2 digital trunk interface is available from StrataCom. This is for markets where 6 Mbps trunks are more prevalent than T3 trunks e.g. Japanese markets. T2 trunks provide four times the bandwidth as T1 trunks. Since the IPX, IGX, or BPX network interface is either T3 or E3, a rate adapter is required.

This function is provided by StrataCom with the AT3-6M Interface Adapter. One unit is required at each node to be connected to the T2 network (see Figure 5-6). The nodes are software configured to restrict the bandwidth allocation on these trunks to avoid overloading the Interface Adapters and associated trunk.

Figure 5-6 T2 ATM Trunking



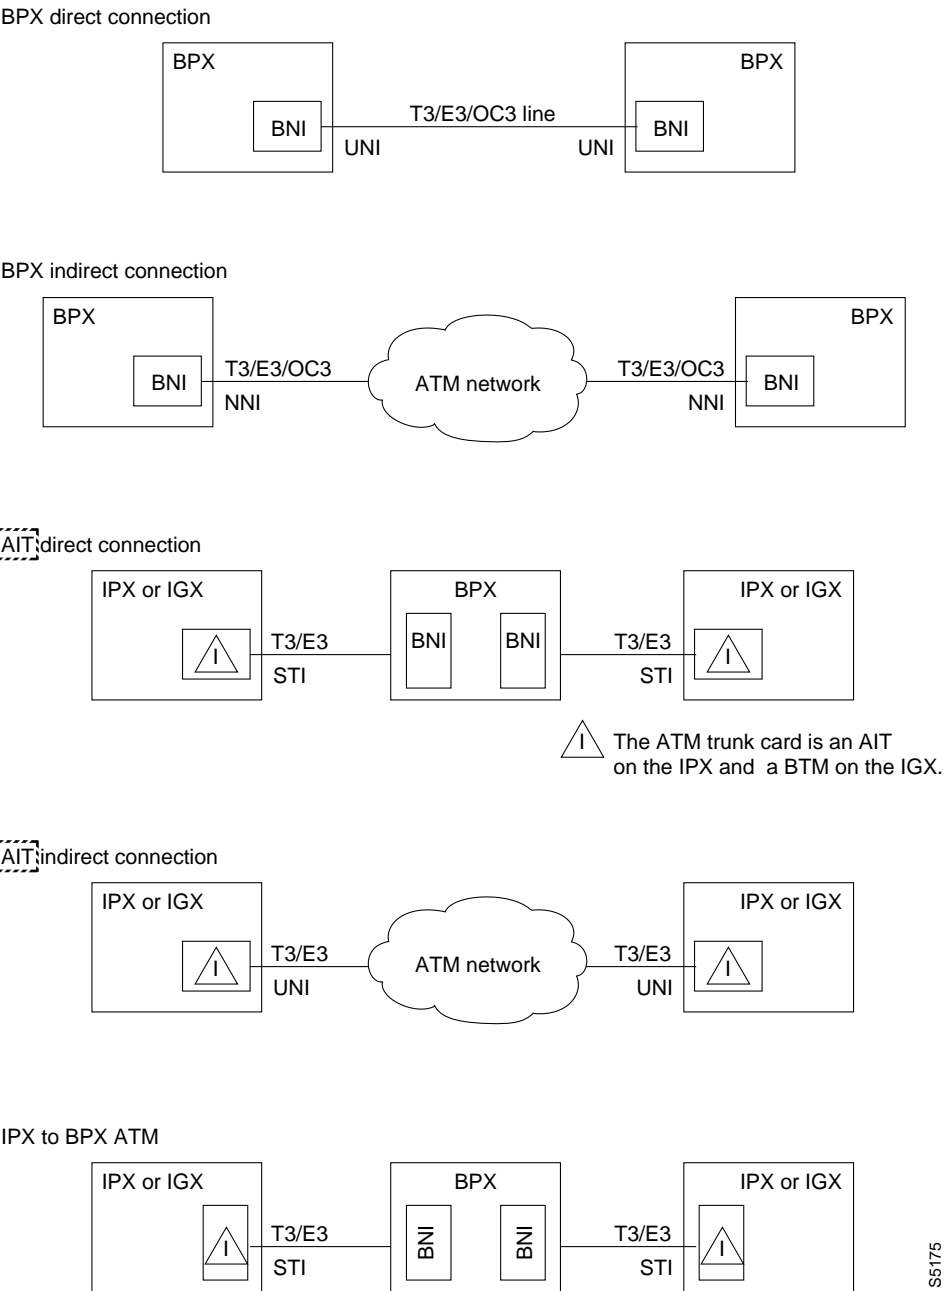
StrataCom ATM Network Architectures

Using the standardized UNI header allows the IPX, IGX, or BPX to be used in a generic ATM network (cloud) with a mix of various equipment. However, to take advantage of several features unique to StrataCom equipment, including ForeSight dynamic closed loop feedback congestion control, the STI header type, an extension of the standard UNI header, must be used. These five connection types are illustrated in Figure 5-7.

There are five types of ATM network connections supported:

- Direct connection of two or more BPX nodes via T3/E3/OC3 lines. Provides broadband networking for connections originating on IPX, IGX, and BPX service interfaces. Uses STI cell header to provide ForeSight.
- The connection of BPXs to a generic, standards-based, ATM network. This allows the BPX to be a part of a non-StrataCom network. Uses UNI cell header.
- Direct connection of two IPX (or IGX) nodes via a T3/E3 line. This provides a point-to-point broadband trunk between IPXs. Uses StrataCom Interface (STI) cell header to provide ForeSight.
- Indirect connection of IPXs (or IGXs) via an ATM network. This provides a broadband interface into a non-StrataCom ATM network. Uses User Network Interface (UNI) cell header.
- Indirect connection of two IPX (or IGX) nodes via a BPX network. This provides a switched broadband trunk between IPXs. Uses STI cell header to provide ForeSight.

Figure 5-7 StrataCom ATM Network Architectures



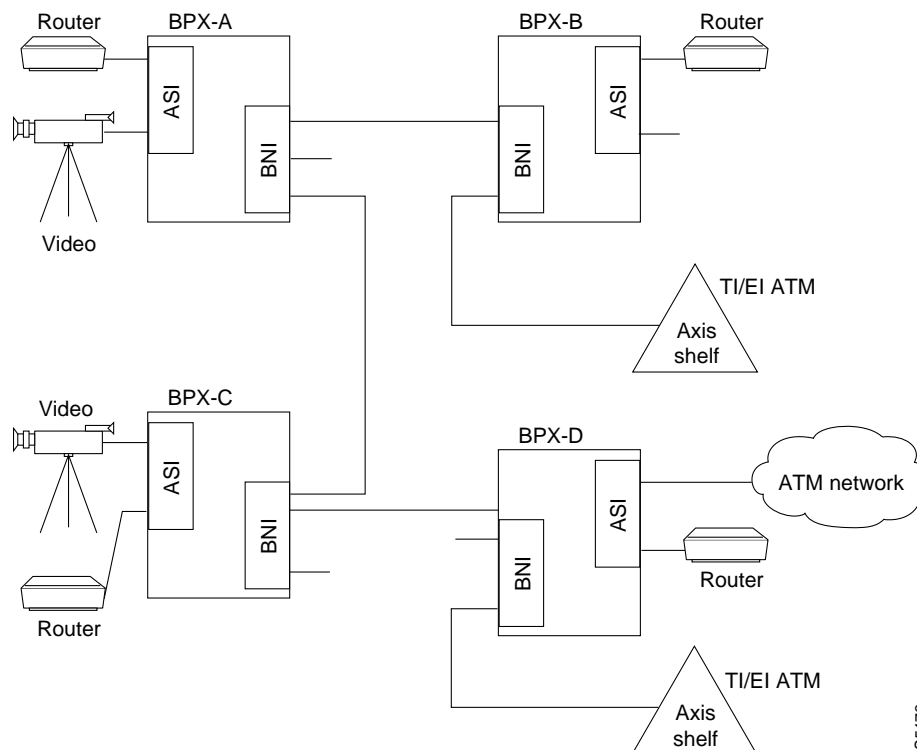
S5175

ATM Service Interfaces, ASI and AXIS Shelves

ATM service (connections) can be provisioned directly on a StrataCom ATM network through the use of an ATM Service Interface (ASI-1 or ASI-155) card set. The ATM Service Interface provides direct connection of high-speed user devices (such as routers or hubs) with ATM standard interfaces to StrataCom ATM wide-area networks.

Each ASI-1 card set, installed in the BPX, provides two T3 or DS3 ATM service interfaces (ports). Each ASI-155 card set, installed in the BPX, provides two OC3/STM-1 ATM service interfaces (ports). Either of the two ports can be configured as a User-to-Network (UNI) interface for directly interfacing user ATM devices to the ATM network as shown in Figure 5-8. Or the port may be configured as a Network-to-Network Interface for interface between two dissimilar ATM networks (Figure 5-8)Figure 5-8.

Figure 5-8 Direct ATM User-to-Network Connections



ATM connections are provisioned as permanent virtual circuits and operate in much the same manner as frame relay circuits. As indicated previously, ATM connections can be provisioned to support multiple services, suited for applications with different service requirements. Currently, these include:

- Circuit Emulation Services (CES) for applications carrying constant bit rate data such as encoded voice and video (compressed or uncompressed).
- Bursty Data Services (BDS) for such variable bit rate applications as LAN-to-LAN connectivity (e.g., frame relay) and variable bit rate video.

ATM connections are identified by a two-part address field consisting of a Virtual Path Identifier (VPI) and a Virtual Circuit Identifier (VCI) similar to the manner in which frame relay connections are identified. In this release, the number of PVCs available is 1000 per ASI-1 card, 2000 per

ASI-155 card, and 1000 total per node, or 5000 grouped connections per BPX node. Connections are automatically routed (AutoRoute) throughout a BPX network and terminate on an ASI card set at the remote BPX node. In this release, the connections may not route through IPX nodes.

StrataCom ATM networks have a flexible architecture and configuration interface that allows the user to define service parameters that combine features of both Constant Bit Rate and Variable Bit Rate to tailor their connection types to new and unique applications that will arise in the future. Up to 32 different classes of service can be predefined for various connection types.

ATM connections using the ASI (ATM Service Interface) will include user access controls, such as Usage Parameter Control (UPC) and input queue configuration on a per-PVC basis to assure equitable access to the network bandwidth. UPC operates using the Generic Cell Rate Algorithm (GCRA) defined by the ATM standards bodies or a unique Frame-based Generic Cell Rate Algorithm (F-GCRA) developed by StrataCom. Various input parameters, such as Cell Loss Priority and Forward Congestion Indicator thresholds can be set by the user.

Advanced network congestion control features, such as Selective Cell Discard and Enhanced Forward Congestion Indicator (EFCI) are provided on all connections. ForeSight, the congestion control mechanism used so successfully on frame relay connections for congestion control and bandwidth optimization, has been adapted to ATM connections and is available as an option.

AXIS

The AXIS shelf provides an interface to the BPX for T1/E1 frame relay, T1/E1 ATM, CES, and FUNI connections. With interworking, the frame relay connections may be terminated at another AXIS, an ASI, an IPX, or IGX. The AXIS shelf T1/E1 ATM connections may be routed to another AXIS or to an ASI.

Narrowband FastPacket Networks

The IPX and IGX provide bandwidth management capabilities for narrowband networks to:

- Use T1, E1, or subrate digital trunk bandwidth efficiently.
- Integrate voice and data on these digital transmission facilities.
- Provide digital access and cross connections (DAC).
- Provide automatic network rerouting and equipment redundancy.
- Give network managers complete control over all aspects of network operation.

T1 Networks

A simplified example of an IPX 16/32 T1 network is shown in Figure 5-9. The various node, consisting of IPX 16/32 T1 network switching bandwidth managers, are interconnected with T1 packet trunks that carry packets of data at 1.544 Mbps. The data input to the network may consist of a mixture of voice transmission from D4 channel banks or digital PABXs, various low-speed point-to-point data circuits, or frame relay data. The following network interfaces are supported by the IPX and IGX:

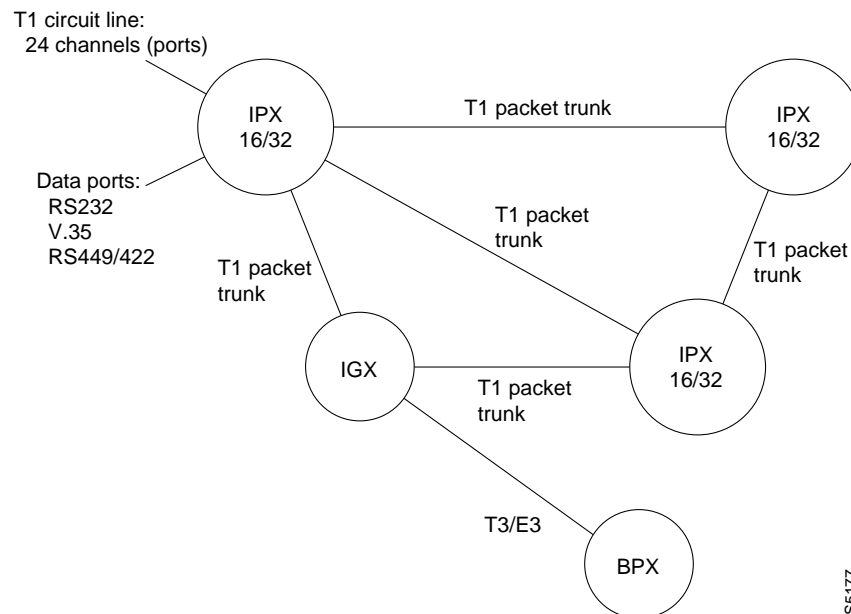
Comment: In the following, the FRM, HDM, and CVM cards perform the same functions in the IGX as the FRP, SDI, and CDP do in the IPX.

- Unchannelized 1.544 Mbps interface as specified in AT&T Pub 62411 for the network trunk interface.

- The V.35, X.21, and T1/E1 frame relay port interfaces are implemented using a FRP with either a FRI-V.35, FRI-X.21, or FRI-T1/E1 back card.
- The balanced RS422 synchronous data port interface is implemented using an SDI-RS-449 back card and adapter cables.
- The V.35 synchronous data port interface is implemented on an SDI-V.35 back card.
- The unbalanced RS232 C or D data port interface is implemented on a SDI-RS232 (synchronous) or LDI-RS232 back card (asynchronous).
- DS0A and DS0B data connections are implemented on a CDP in data mode with BC-T1 back card.
- 24-channel encoded voice interface is implemented on a CDP with BC-T1 back card. The port side signalling protocols that the CDP handles are:
 - Robbed-bit D4 with A, B, C, and D signaling channels.
 - ESF with B8ZS signaling.
 - 64 Kbps clear channel.

The T1 packet trunks contain the standard framing bit for channel synchronization but the remaining T1 frame data is not channelized. Nodes connect to other nodes using a minimum of one packet trunk and in larger networks often interconnect in ring or star configurations to provide alternate routing paths for maximum protection against outages that may result from a failed T1 line.

Figure 5-9 Typical T1 IPX Network




E1 Networks

A simplified example of an IPX E1 network is shown in Figure 5-10. Each IPX node is connected to the user's PBXs and channel banks via E1 circuit-switched lines. Data connections are made via various data port interface types. The following network interfaces are supported by the IPX and IGX E1:

Comment: In the following, the FRM, HDM, and CVM cards perform the same functions in the IGX as the FRP, SDI, and CDP do in the IPX.

- The 2.048 Mbps standard interface as specified in G.703. This interface is available on the circuit line and trunk side of the IPX.
- The X.21 frame relay port interface is implemented using a FRP with FRI-X.21 or FRI-E1 back card.
- The V.11/X.21 synchronous data port interface is implemented using an SDI back card (RS-449) and adapter cables.
- The V.24 data port interface is implemented on an SDI-RS232 (synchronous) or LDI-RS232 back card (asynchronous).
- The V.35 synchronous data port interface is implemented on an SDI-V.35 back card.
- 30-channel encoded voice interface is implemented on a CDP with BC-E1 back card. The port side signalling protocols that the CDP handles are:
 - Channel Associative Signalling (CAS).
 - ISDN Q.931 common channeling signalling pass through.
 - Digital Private Network Signalling System No 1 (DPNSS) pass through.

Figure 5-10 E1 IPX Network



Networks with IPX 8

The IPX 8 operates seamlessly in existing and new networks with StrataCom's existing family of cell relay switches. It also provides both T1 and E1 packet line and circuit line interfaces for use in international networking. The IPX 8 can be used in large networks where there is a mixture of IPX 8, IPX 16, and IPX 32 nodes and IGX nodes (Figure 5-11). It can be used as a feeder (Figure 5-12) to a larger node, or it can be configured as an IPX shelf and connected to a BPX, as can an IGX. Or, it can be used in small networks by itself. Typical applications include:

- Connection of a remote branch office via a single trunk to a regional office on an IPX backbone network.
- Connection of a redundant loop of branch offices to a regional office on an IPX backbone network and one or two T1 ports. Redundant systems provide automatic switchover on equipment or network failure.

Figure 5-11 Typical IPX Network with IPX 8 Remote Nodes

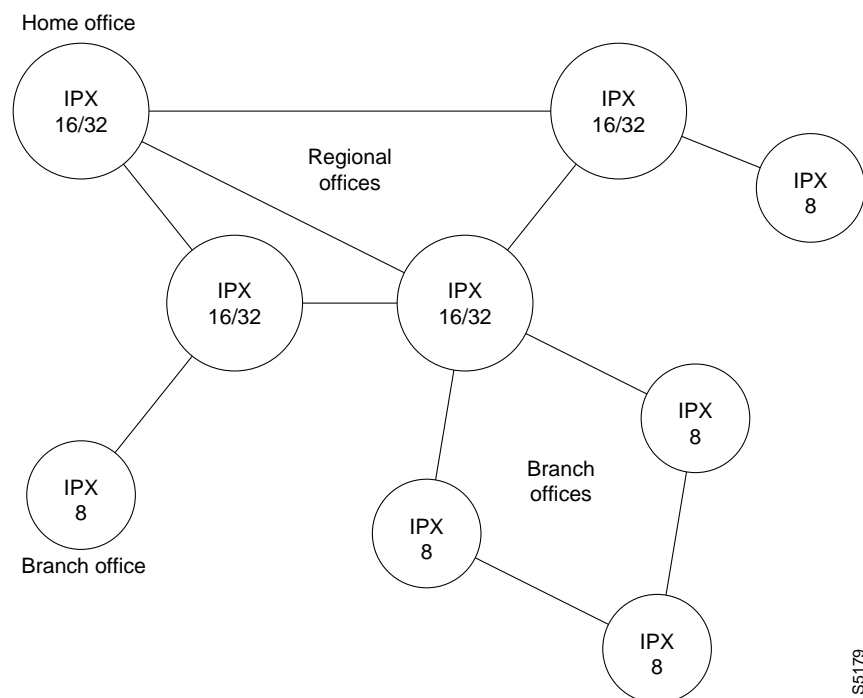



Figure 5-12 IPX 8 Used as a Feeder in Large Networks



Japanese J1/Y1 Networks

The BC-J1 and BC-Y1 back cards provide IPX nodes with the capability to interface to J1 circuit lines and Y1 network trunks in Japan. Y1 trunks are used in Japan and are similar to T1 trunks described earlier except they use a Coded Mark Inversion (CMI) line code instead of Alternate Mark Inversion used by the T1 lines. This assures a line signal with sufficient one's density for the line repeaters.

The BC-Y1 back card is used in combination with an NTC (NTM in an IGX) front card to provide an interface for a trunk in an IPX or IGX network using a Japanese standard "Y", 1.544 Mbps circuit (referred to hereafter as a Y1 circuit).

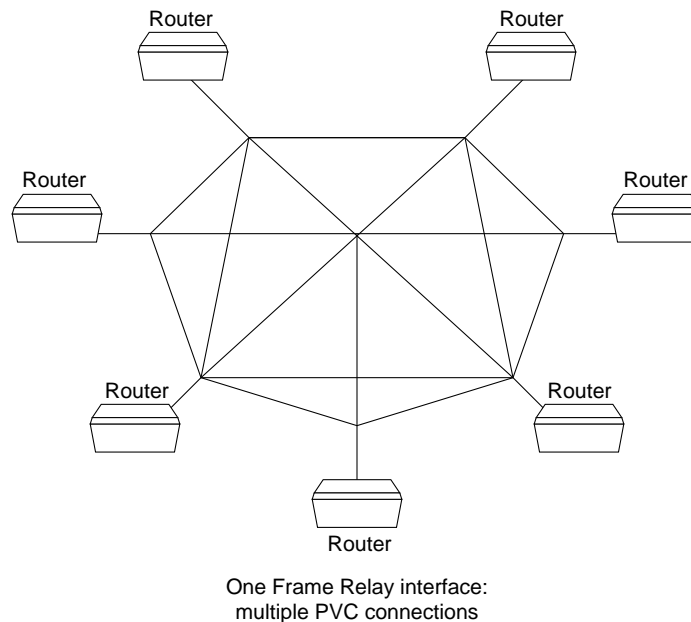
Likewise, the J1 circuit line resembles a channelized E1 circuit line except that it runs at the 1.544 Mbps. T1 rate and also uses the CMI line code. The BC-J1 back card is used in combination with a CDP (CVM in IGX) front card to provide an interface for a circuit line to a PBX using a Japanese standard TTC-JJ-20, 2.048 Mbps circuit (J1 circuit). In all other aspects, these networks are similar to the T1 and E1 networks just described.

Frame Relay Networks

Frame relay networks have grown rapidly in popularity for their ability to provide full mesh connectivity for LAN-to-LAN interconnections. Instead of transferring bits between fixed locations like a normal synchronous data circuit, a frame relay data circuit acts like a wide area LAN (WAN). A sending device places an addressed frame into the network and it arrives at its destination.

To be totally effective, every LAN should have one or more connections to every other LAN. To do this with dedicated circuits, can require a large number of physical circuits for networks with more than a just a few locations. For example, Figure 5-13 Figure 5-13 illustrates a network interconnecting seven LANs. This will require 21 circuits for full mesh connectivity. In general, a network interconnecting "n" LANs will require $n(n-1)/2$ circuits for full mesh connectivity.

Figure 5-13 Frame Relay PVCs Provide Full Mesh Connectivity



55181

StrataCom Frame cell relay networks utilize Permanent Virtual Circuits (PVCs) to carry frames of user data. With a PVC instead of a dedicated circuit, the user perceives a connection that is permanently in-place but one that doesn't utilize any network bandwidth until there is actually data to be transmitted. Frame relay PVCs are routed through the network over packet or ATM trunks using either NTC or AIT cards in the IPX.

Frame Relay Port Interfaces

Frame relay derives its cost advantages by passing many data circuits between the user device, e.g. router, to a single, high-speed, port on the network node. These ports on a StrataCom node can operate at speeds from 56 Kbps up to 2 Mbps.

Frame relay ports are added to an IPX node using a Frame Relay PAD (FRP) front card and associated Frame Relay Interface (FRI) back card or to an IGX node using the Frame Relay Module (FRM) and FRI back card.

On an AXIS, the frame relay interface is provided by the FRSM card. The IPX and IGX card sets provide four or more UNI ports for interface to frame relay routers, hubs, or bridges as indicated below:

- V.35, four ports per card.
- X.21, four ports per card.
- T1 or E1, single DS0 or multiple DS0 rates, 6 or 24 DS0 ports per card.

In the IPX or IGX, each frame relay card can carry up to 252 PVCs. Each node can accommodate up to 1024 PVCs, which would require up to four frame relay card sets (FRP/FRI) to achieve maximum capacity.

Port Concentrator

The Port Concentrator provides up to 44 low-cost frame relay ports and connects to an IPX via an FRP-2 card or to an IGX via an FRM-2 card. Connection management using the Port Concentrator ports is the same as that for frame relay connections to the FRP and FRM cards.

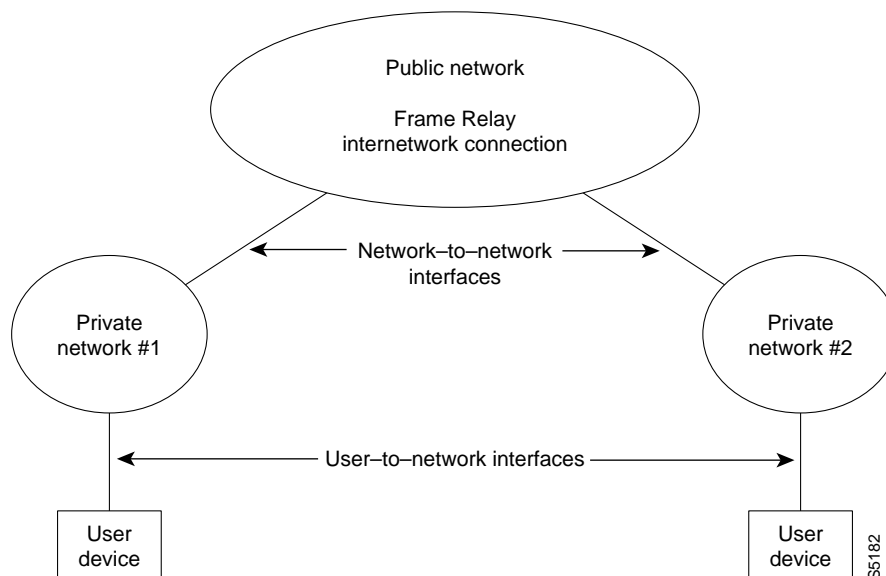
Frame Relay Network Interfaces, IPX, IGX, FastPAD, AXIS

There are two types of network interfaces protocols that are supported by StrataCom frame relay ports:

- User-to-Network Interface (UNI).
- Network-to-Network Interface (NNI).

Figure 5-14 illustrates these two frame relay interfaces. The User-to-Network Interface is defined as the port where a user device, such as a router, interfaces with a StrataCom wide area network carrying the frame relay traffic. Currently, the IPX and IGX support UNI via the following protocols: ITU-T Q.933 Annex A, ANSI T1.617 Annex D and StrataCom LMI, which offers additional features over the other protocols. Refer to the *AXIS Reference Manual* and the *FastPAD Reference Manual* for further information of the protocols they support.

Figure 5-14 Frame Relay Network Interfaces



A Network-to-Network Interface is a port that forms a boundary between two independent wide area networks, e.g. a StrataCom network and another network and may or may not consist of StrataCom equipment. No user device is connected, only another network port.

User to Network Interface (UNI)

The User-to-Network Interface for frame relay permanent virtual circuits (PVC) is a defined set of protocols and procedures. Currently, the IPX supports UNI via the following protocols: StrataCom LMI, ITU-T Q.933 Annex A, and ANSI T1.617 Annex D.

LMI transmits on a logical connection between the IPX and IGX and the user device (router) separate from the data path using DLCI 1023. This connection is a special PVC, carrying messages between the IPX or IGX and the user device. The messages transmitted via the LMI protocol provide the following information to the user device:

- Keepalive/Administration Configuration/Flow Control.
- Network notification of the active and available PVCs.
- Network notification of the removal or failure of a PVC.
- Real time monitoring of the status of the physical and logical link between the network and each user device.
- Network notification of a change in PVC status.
- Notification of the minimum bandwidth allocated by the network for each virtual circuit.
- Notification of the priority of each virtual circuit.
- XON/XOFF type flow control mechanism to prevent buffer overflow.

Some user devices can obtain the network configuration dynamically using LMI messages. With these devices, the Network Administrator assigns Data Link Connection Identifiers (DLCIs) for both ends of each connection in the network and the user device interrogates the frame relay port to determine the DLCI assignment. If the user device does not have this feature, then the Network Administrator must manually configure the user device to use the DLCIs programmed into the IPX/IGX network.

Network to Network Interface (NNI)

Currently, IPX and IGX frame relay ports report the status, active or failed, of each frame relay connection within the network to the user devices at both ends of the circuit. The status is reported using the frame relay User-Network Interface (UNI) protocol and consists of the status of the far end user device as well as any failure or congestion on the overall connection.

When a circuit spans multiple networks, it is desirable to maintain this status reporting for the entire length of the PVC. This necessitates obtaining the circuit status from the foreign network and including it with the status reported by the local network. The Frame Relay Network-to-Network Interface (NNI) feature is used when extending connections between a StrataCom frame relay network and another independent, StrataCom or non-StrataCom frame relay network.

At each port where one network connects to another network, Frame Relay NNI passes the circuit status for internetwork connections between the two networks. NNI status messages indicate not only the status of the remote user device, but also the combined status of each network link traversed by the multi-network PVC.

Within a StrataCom network, the status of every frame relay PVC is known by every node in the network since it is distributed network-wide by system software communicating with each node. The status of each PVC is sent to the user device over the UNI as discussed in *paragraph 3.3.3*. There are three possible status to report:

- Active-connection OK.
- Failed connection.
- Disconnected-DLCI removed from service.


However, there may be PVCs that extend beyond the boundaries of a StrataCom network and connect to a user in a totally separate network. The frame relay Network-to-Network Interface (NNI) is used to forward the connection status through to the User-to-Network Interface (UNI) ports within the different networks. A frame relay port is defined as a UNI or an NNI port with software configuration. In order for this feature to be completely implemented, the interconnected ports in both networks must support NNI.

Frame Relay Addressing

The frame relay port uses a number, the Data Link Connection Identifier (DLCI), to identify each bidirectional frame relay connection at the local interface between the user device and the frame relay network. This number is assigned by the network operator when the connection is added to the network.

An example is shown in Figure 5-15. In this illustration, there are three possible destinations for this data frame, A, B, or C, and each needs a unique DLCI to identify the PVC to these destinations.

Figure 5-15 Frame Relay Network Connection



Note that the DLCI has only a local significance. The frame relay port and the user equipment (router) need to coordinate the use of DLCIs but the DLCI does not matter to anything else in the network. Consequently, DLCI numbers only have to be unique at each frame relay port to the network. The same numbers may be reused at other ports in the network although this numbering scheme may be somewhat more difficult to administer.

Networks with the FastPAD

The FastPAD is used for small networks or for branch sites where it may be difficult to justify the expense of a fully-equipped IPX or IGX node. It is a unique device that resembles a small PABX and router combined into one package. It can be used by itself, connecting to a public network, or it can be integrated into a StrataCom frame relay network.

Each FastPAD provides up to eight analog voice telephone circuits and up to six data ports. It supports direct interface to telephone with both pulse or tone dialing and a variable rate voice compression for efficient bandwidth utilization. Customer data rates supported range from 2.4 to 19.2 Kbps synchronous or asynchronous for five of the data channels to 64 Kbps synchronous for the sixth channel.

Simple FastPAD Network, Voice, Data, Frame Relay

The FastPAD takes the various frame relay, voice, and data inputs and multiplexes them into frames for transmission over a composite link. The resulting frames are encapsulated in a standard frame relay format header with a unique Data Link Connection Identifier (DLCI) assigned to each FastPAD. These frames are then applied to any standard frame relay network for forwarding to another FastPAD at the destination. DLCI bundling of multiple logical connections within single PVCs reduces the number of PVCs required within the network. With a FastPAD network in place, a user automatically has the advantage of a private telephone network without the expense of providing a PBX or key system at each site.

Each FastPAD is assigned a unique dialing code in the network when it is initially configured. Thus, in small site applications, where there are no PABXs, telephones can be directly connected and the network will provide the call switching between sites equipped with FastPADs.

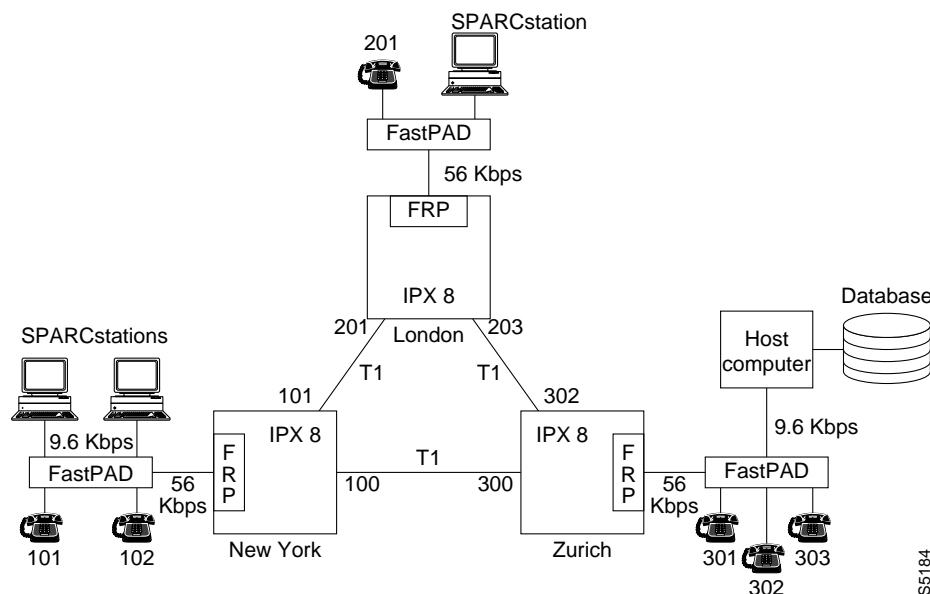
Similarly, a FastPAD network provides a cost-effective way of linking a number of small LANs, some of which may consist of only one or two terminals, without having to purchase a separate router at each location. The FastPAD acts very much like a router itself.

FastPAD Access to IPX Frame Relay Network

The FastPAD Access Multiplexer provides an economical access device for sites where only a few of voice and data circuits are required but where it is desired to gain access to a network that perhaps links many widely separated locations. Since the composite link of the FastPAD is directly compatible with the FRP/FRI interface of an IPX, or the FRM/FRI interface of an IGX, Permanent Virtual Circuits (PVCs) can be established between IPXs to connect many FastPADs using a frame relay network.

An example of a simple network using FastPADs across an IPX or IGX frame relay network is illustrated in Figure 5-16. This might serve to link a corporate headquarters in Zurich where a host computer and MIS database is located with two field offices, one in London and the other in New York. The field offices require several phone connections and several low-speed database access circuits. Speed dialing between the various offices is a plus.

Figure 5-16 IPX Network with FastPAD Access



In the example of Figure 5-16, the FastPAD voice slots have been assigned dialing codes of 101 and 102. At London, the dialing code is 201, and in Zurich they are 301, 302 and 303. If New York wants to call London, they can just dial 201. The FastPAD keeps an internal routing table of all the dialing codes and destinations.

Since each port attached to the FastPAD has two destinations to reach, it requires two DLCIs. At New York, an arbitrary DLCI of 100 is used to route the PVC to Zurich while 101 routes a second PVC to London. The DLCIs are in reality assigned indirectly by the FastPAD when the user assigns a node number and slot number to each unit.

Since the FastPAD connects to the IPX via an FRP, the FastPAD is considered to be a frame relay user device attached to a frame relay network device. When linking a FastPAD to an IPX, the FastPAD should be equipped with V.35 interface cards/cables and the frame relay port should use the FRI-V35 backcard. The link speeds supported between the FastPAD and the IPX frame relay port include 56, 64, 112 or 128 Kbps. Refer to the *FastPAD User's Guide* for additional configuration information.

International Networks

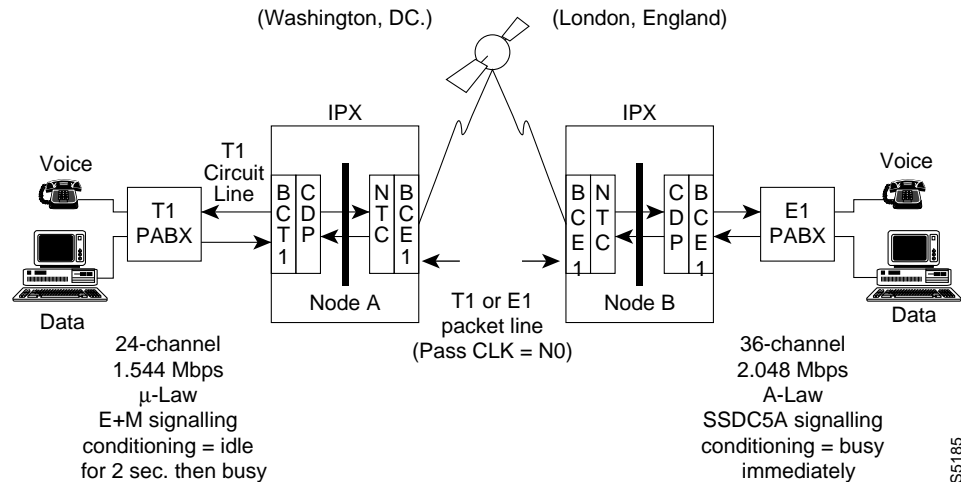
StrataCom cell-relay networks can easily be configured to provide direct communication links between domestic and international locations each using the differing transmission standards, e.g., T1 and CEPT E1.

For example, in Figure 5-17, an IPX network is used to extend voice and data circuits in a T1 network with a node in Washington D.C. to connect via satellite T1 or E1 packet trunks to a node in London, England serving a separate CEPT E1 network. Not only is the bit rate (and the frame format) of the two networks different, but the VF encoding law, signalling formats, and signalling conditioning requirements are different.

Comment: In the following, the FRM, HDM, and CVM cards perform the same functions in the IGX as the FRP, SDI, and CDP do in the IPX.

Building a seamless international network has always been facilitated by a cell relay network but the introduction of the CDP card makes protocol conversion even easier than before. All of the voice and signalling parameters indicated in Figure 5-17 can be programmed by the network operator within each CDP.

Figure 5-17 IPX International Network



Example of a StrataCom International Network

In this example, the CDP at node A, terminates the T1 circuit line to the PABX in Washington with a corresponding BC-T1 back card. It is programmed for the voice parameters required by the Washington PABX in the transmit direction and for the parameters required by the London PABX in the receive direction. The CDP in Node A makes the appropriate conversions, adds gain or loss as required for level adjustment, packetizes the channels (up to 24), and sends them to the NTC, with its attendant back card (E1 in this example but could be T1 depending on the satellite circuits used).

In London, the Node B NTC and associated BC-E1 terminates the E1 line, receives the packets, and forwards those packets with the proper destination header to the CDP connected to the E1 PABX as shown. The Node B CDP is programmed for the requirements of the London PABX with no conversion (all the conversion was performed in Node A in Washington). This CDP assembles 30-channel E1 frames from the packet data received, and when a frame is full, sends it out to the PABX.

Note, the packet trunk connecting nodes A and B has to be either E1 or T1 with the same back card at each NTC. However, the two circuit lines connecting the PABXs can be T1 at one end of the circuit and E1 at the other. A distinct requirement is that the total number of DS0 channels assigned on this virtual circuit be the same at both ends. If 24 circuits are assigned, based on the T1 out of the Washington PABX, then the London PABX will have six channels that can be assigned to other circuits ($30 \text{ E1} - 24 \text{ T1} = 6$).

These two nodes will usually be part of two separate networks as far as clocking is concerned, each operating at their geographic area's basic network rate (1.544 Mbps in the US network and 2.048 Mbps in the UK network). Since the packet trunk connecting Node A and Node B is over satellite it is not permitted to pass clock. Therefore, the two networks will be isolated from each other as far as network timing, an example of pleisiochronous networks (discussed more fully in Pleisiochronous Network). IPX and IGX nodes operate with no loss of data in pleisiochronous networks using frame slips during idle frames to accommodate the differences in the two network clocks.

E1/T1 Conversion

Comment: In the following, the FRM, HDM, and CVM cards perform the same functions in the IGX as the FRP, SDI, and CDP do in the IPX.

International networks often transition across T1 and E1 environments. As such, they require circuits that allow voice and signalling communication between T1 nodes and CEPT E1 nodes. IPX/IGX networks handle the code conversions, varying data rates, and different signalling formats through software conversion and inherent transparency of the packet network.

The IPX Channelized Data Pad (CDP) card provides communication between T1 and CEPT (E1) connections. This card provides data rates of from 512 Kbps to 2.4 Kbps. These connections can terminate with standard RS-232, RS-449, V.35 interfaces. This card also provides echo cancellation and compression of 8:1 per channel with 16 Kbps ADPCM and 2:1 Voice Activity Detection.

IPX transition functionality is provided by the CDP on both the CEPT (E1) IPX node and the T1 node. This function allows a voice port on a channel bank or PBX in T1 to connect to a voice port on a channel bank or PBX on the CEPT (E1), and assumes a connection is required by a voice protocol in a CEPT (E1) system that is A-law encoded and a voice port in the T1 system that is μ -law encoded.

A T1 system obtains ADPCM voice connections on the CDP card. Signalling conversion is provided between 4W E&M and SSDC5A (British signalling convention).

Note This function does not support rotary dialing, V.25 modem connections, or 64 Kbps voice.

Voice communication paths require that the PCM samples sent between the T1 and E1 IPX nodes be converted from μ -law to A-law. In order for this conversion to take place, a CDP must be installed in the T1 IPX node, and a CDP must also be installed in the CEPT IPX node.

Setting International Node Parameters

The Configure Line (cnfln) command includes parameters such as line coding (HDB3 or AMI, encoding type (μ -law or A-law), and clocking type (loop clock or not) that can be programmed on individual or groups of channels. Since the circuit line at each end of an international hop is separate, the parameters are programmed separately and do not have to be the same. The IPX and IGX make the conversion automatically.

The Configure Interface Type for Voice Channel (cnfvchtyp) command is used to set the channel parameters for the CDP (in IPX) or CVM (in IGX) used in an international (or any other node). There are currently 23 standard channel or trunk interface types already defined (for example 2WE&M & 4WE&M, FX, DP, DX, RD, R1 & R2). These channel definitions are used primarily to define what to do with various VF circuit parameters during circuit failure.

Each type has preset A and B signalling conditioning values as well as the substitute PCM voice sample sent to the attached equipment (often an idle code) in case the connection fails. If none of the predefined conditions fits the needs of the international node, the user can specify the signalling conditioning and voice code to be transmitted by using a template.

The Configure Transmit Signalling (cnfxmstsig) command allows the user to pass, block, or convert any or all of the A, B, C and D signalling bits in the transmit direction. A similar command is available for Receive Signalling. The Configure Channel Gain (cnfchgn) command specifies the loss or gain (–8 dB to +6 dB) to add to the channel to set the end-to-end transmission levels for each voice circuit. There are similar commands for configuring and converting various data channel signalling lead and alarm output patterns.

Network Synchronization

Each node contains a highly stable internal clock source. But because all network nodes are interconnected with trunks, it is important that the various node clocks be synchronized to one timing source. If not, the data rate on some trunks will be faster than the node timing. On other trunks, the data rate may be slower. Either case will eventually cause loss of frame synchronization in the cards terminating the trunks.

Fortunately, because the nodes resynchronize by dropping idle cells/packets, there usually is no loss of user data. But if there are few idle cells/packets, as on a fully utilized trunk, the resynchronizing may have to drop cells/packets containing valid data just to remain in frame.

Another problem that occurs when the data rate on the incoming trunk does not match the node timing, is that the packet buffers will eventually overflow or underflow. This results in either lost data or underutilization of the buffer. Network-wide clock synchronization eliminates these problems.

Clock Sources

Each node in the network has available to it four sources of timing:

- Internal—using the nodes internal oscillator for clock source.
- External—using an input from an external device for clock source.
- TRK—using the recovered receive timing from one of the network trunks.
- CLN—using the recovered receive timing from one of the circuit lines.

To obtain clock from a Common Carrier, the packet trunks defined as clock sources must pass through a Digital Access and Crossconnect System (DACS) or similar device that retimes the data stream. A trunk must be activated and clear of alarms before it can be used as a network clock source. In addition, a packet trunk must be configured not to pass clock before it can be defined as a clock source. Each source of clock for a node is ranked according to its source stability as follows:

- Primary—highest stability. Use this, above all, if available.
- Secondary—less stable. Use only if primary source is unavailable.
- Tertiary—least stable. Use only if secondary source is unavailable.

These sources are defined in the node database using the Configure Clock Source (**cnfclksrc**) command. They can be displayed at any time using the Display Clock Sources (**dspclksrc**) command. Once defined, the location and type of the network clock source is broadcast to all nodes in the network. Each node maintains a list of all available clock sources for the network. This list is preserved during a power failure or controller card switch on the node.

Clock Source Selection

StrataCom cell relay networks use a fault-tolerant network synchronization scheme similar to that used by telcom digital switching offices. The node clocking scheme ensures that all nodes in the network are automatically synchronized to the nearest, highest level clock available. This synchronization remains in effect even after line failures, line repairs, joining of sub networks and all other network topology changes. Each node selects the clock source in the following hierarchy:

- 1 If a clocking hierarchy is defined, it automatically selects the nearest working, primary source if it is available.
- 2 Where there are two packet or circuit line sources to choose from, it selects the lower numbered line to use.
- 3 If a primary source is unavailable, it selects the nearest secondary source.
- 4 If a secondary source is unavailable, it selects the nearest tertiary source.
- 5 If none of the defined sources are available, as may be the case in a catastrophic network failure, the IPX, IGX, or BPX reverts to its internal oscillator.
- 6 If no clocking hierarchy is defined in the network, the IPX arbitrarily selects the internal clock from one of the nodes as the active network source.

The clock source selected is monitored by internal circuits in the node. If it is lost or if it drifts out of range, it is marked as failed and the node controller looks for another source. The node's active clock source and the path to that clock source can be displayed by the operator using the Display Current Clock (**dspcurclk**) command.

Defining Clocks and Lines

A network's clock source and line characteristics are configured as part of the node installation process. Thereafter, clock sources would be redefined when networks are reconfigured or line status is changed. Engineering considerations by which to assess and define clock sources include:

- Stratum level of each clock source.
- Reliability of each source.
- Network configuration (topology, for example backbone, ring star, mesh).
- Availability of multiple clock sources in a pleisiochronous network.

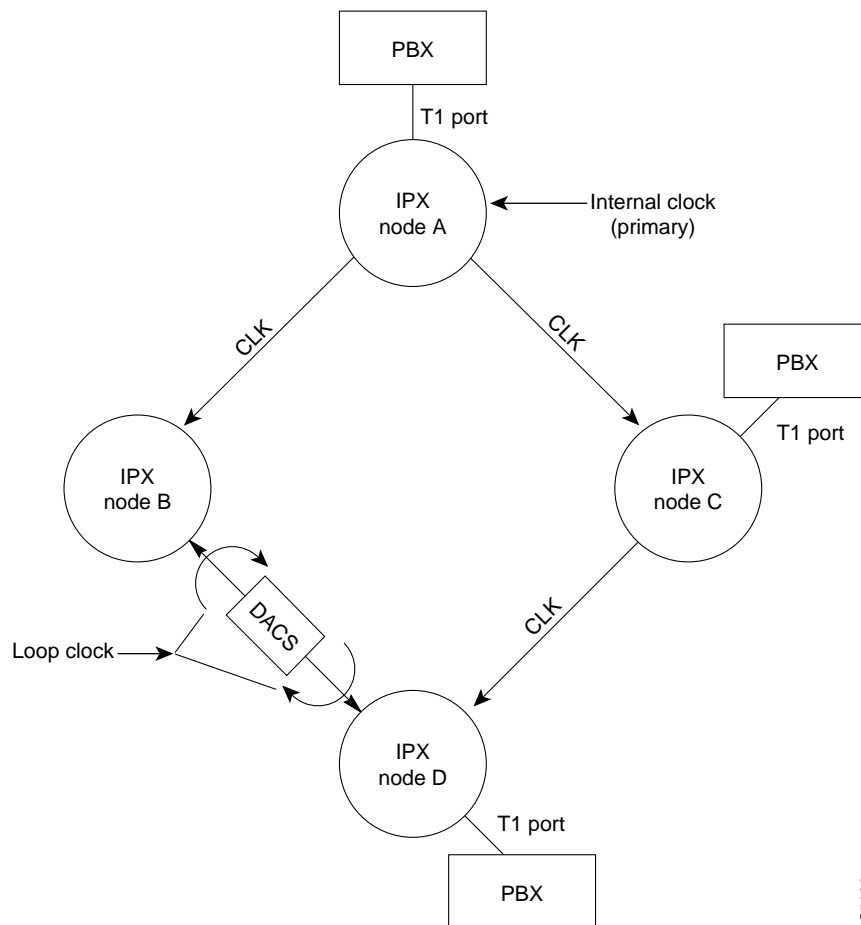
Trunks may be classmarked as being able to pass clock or not, or to loop clock or not, using the Configure Trunk (**cnftrk**) command. A trunk passes clock if the clocking transmitted at one end is recovered at the receive end and may be used to clock the node. If a trunk or circuit line comes from a source that should be isolated from the network, specifying looped clock allows the data to be clocked into and out of the CDP or NTC using this timing source but it does not propagate through the node and into the network. The difference is that the node clock is not retransmitted down the trunk or circuit line.

Normal T3, T1 and E1 spans, by default, pass clock. Satellite lines, on the other hand, do not normally pass clock. The Display Trunk (**dsprtrk**) command can be used to show which lines are marked to pass or loop clock.

Circuit lines can also be defined as clock sources in the same manner as trunks. If a small network consists of a star of several circuit line routes and one of them is a master at a headquarters site, for example, its clock source can be used to synchronize the network and possibly all the other routers permitting them to pass data in the most efficient manner between themselves.

Figure 5-18 illustrates a typical network clock architecture. The primary source is the internal oscillator in IPX Node A. The clock propagates to Nodes B and C over trunks A-B and A-C, which should be defined as passing clock (default). It is propagated to Node D over trunk C-D only.

Figure 5-18 Example of IPX Network Clock Architecture



In this example, a digital central office (DCO) using a DACS is located in the trunk B-D hop and this source of timing is to be isolated from the IPX network. To do this trunk B-D is marked to loop clock. The timing from the DCO is used only to time the data on trunk B-D and is kept out of the network. As an alternative, the IPX network could be set to clock off the DACS. In this case, the trunks between node B and node D would be set to pass clock.

SS186

The PABXs off Nodes A, B and C are, for this example, all digital PABXs with direct T1 interfaces. These should all be set up for loop timing so they will be synchronous with the network. Their associated circuit lines should be marked as not passing clock unless one or more of the PABXs are used for the secondary or tertiary source of timing for this network.


Pleisiochronous Network

A pleisiochronous network is one in which there are two or more independent, active clock sources. This may result when the network trunks are provided by two or more OCCs and pass through their digital switching offices. Or it may be the result of a clock source assigned as primary in two geographic regions.

The Common Carriers each have their own master clocks and, to date, they are not synchronized to each other. For this reason, it is best to avoid, if possible, mixing packet trunks from different carriers. If this is unavoidable, consider using the CDP or NTC cards. These cards can turn clock around back towards the intervening digital central office(s) to isolate the undesirable timing source.

Figure 5-19 shows an example of a network that may operate as a pleisiochronous IPX network. Nodes 1, 2, and 3 are in Region A while nodes 4, 5, and 6 are in a separate Region B. Each region has a node whose internal clock is used as a primary clock source. If the E1 packet trunk connecting node 1 to node 4 is defined as not passing clock, the two regions will be synchronous unto themselves but not to each other. This is called **pleisiochronous operation**. With this configuration, the E1 frames on the node 1 to node 4 packet trunk may need to be slipped occasionally to account for the two independent network clocks.

Figure 5-19 Example of a Pleisiochronous Network



Flat, Tiered, and Structured Networks

As the size of networks continued to expand, a time is reached when the size and complexity of the networks demand a new approach to network configuration. This new approach is Structured Networks and offers an entirely new process for configuring large, widespread data/voice networks, encompassing many IPX, IGX, and BPX nodes. This network configuration provides additional growth potential to permit current networks to expand to accommodate many more nodes, yet retain the features of smaller networks such as efficient bandwidth utilization, high-speed rerouting, and centralized management.

Structured vs. Flat vs. Tiered Networks

IPX networks may be configured into two types of network structures, flat or structured. A flat network, as described in this document, is defined as a network where all nodes are treated equally and reachable with no hierarchy attached to the routing of the virtual connections. In a flat network, every node stores the configuration details of every other node and network messages are exchanged between every node.

Flat Network

A flat network is the most straightforward configuration to implement and is generally suitable until the number of nodes in the network approaches a certain limit. The size of a flat network is limited by addressing and processing constraints. For networks based on IPX nodes using the current system software and hardware, a network of more than 48 nodes must be configured into a structured network.

Tiered Network

Either a structured or flat network may be also be configured as a tiered network by adding interface shelves, e.g., AXIS, IPX, or IGX shelf interfaces to a BPX. The IPX and IGX are configured as non-routing nodes and then added to the BPX as interface shelves.

Structured Network

A structured network, on the other hand, is a network that is partitioned in some well-defined way to simplify its operation. The most common approach to structuring is based on communities of interest or geographic boundaries but can be based on any other criteria that the user desires.

The Structured Network feature provides expansion for existing networks with minimal effort or expense. Once the network is configured as a Structured Network, it currently can expand until it reaches a total of 384 nodes in the network. Advantages inherent in structured networks over large, flat networks include:

- Permits the overall network size to expand considerably without the exponential increase in network control traffic.
- Automatic routing and rerouting decisions for local traffic are limited to primarily within the domain to maintain short response time for reroutes around a trunk failure. Long-haul connections (across several domains) are still automatically routed and rerouted.
- It is easier to manage and administer a number of smaller subnetworks (domains) than a large single network.
- Network management may, if desired, be divisionalized, separating the area responsibility between domains (regions).

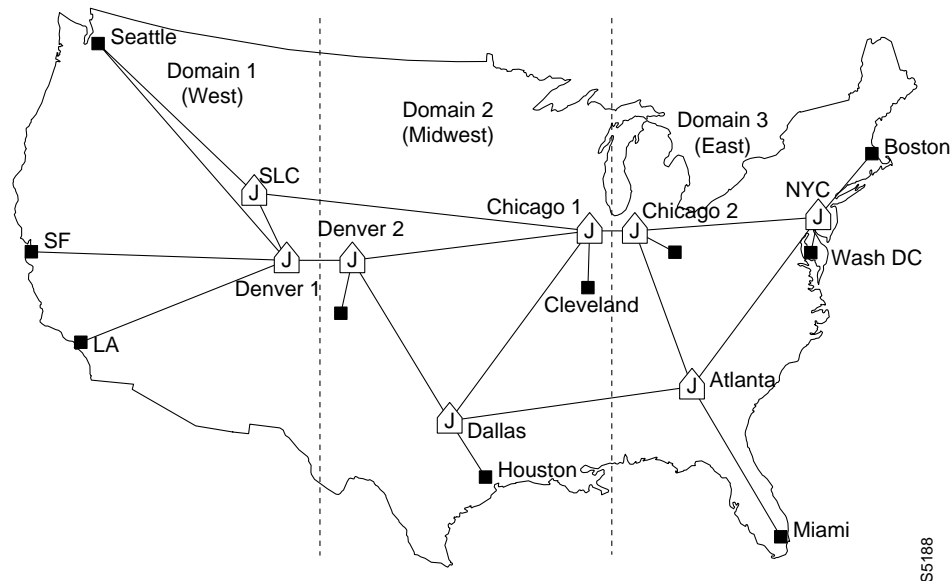
Domains

A structured network is a partitioning of a network into groups of nodes referred to as *domains*. These domains have all the same characteristics and behavior as a flat network. Often the nodes in a domain have a geographic or functional association with all the other nodes within the domain as illustrated in Figure 5-20. However, there are no absolute rules or restrictions used in determining which nodes may be included in a domain.

Local traffic flows throughout the domain just as if it were a small, flat network. There is full-connectivity between all the nodes within a domain. However, connections that are routed to destinations in other domains must follow specific routing paths.

Likewise, every node has a complete database representation of the layout and status of its domain that is kept current by system messages that flow between node processors. However, a domain has only limited knowledge of all other domains in the network as only a select, few system messages pass across domain boundaries. These contain much of the activity of system updates and reporting within the domain.

Figure 5-20 Dividing a US Network into 3 Domains



Topology of Structured Networks

A domain consists of a grouping of IPX nodes. One or two nodes in each domain are designated specifically to inter-network with adjacent domains and are called *junction nodes*. All connections that cross domain boundaries must route through a junction node in each of the domains. Junction nodes have routing tables that are network-wide and can do route translation when necessary.


All junction nodes in the network are linked with the other junction nodes in the network over junction trunks. The junction trunks must terminate on a junction node at each end and not pass through a local node. For example, in Figure 5-20, all junction nodes, from Denver, through Chicago, to New York City are connected by junction trunks with no intervening nodes. They do not, however, have to be connected in a star or mesh configuration. The subnetwork of junction nodes and junction trunks is referred to as the *junction network*.

The other nodes within a domain that are not junction nodes are referred to as *local nodes*. Local nodes have routing tables only for the domain. Local nodes connect to each other and to the junction nodes using normal packet trunks; there are no special considerations.

Refer to Figure 5-21 illustrating a simplified structured network with three domains, domain 1, domain 2, and domain 3. Domain 1 is formed by one junction node, J1, and three local nodes A, B, and C. Likewise, domain 2 consists of junction node, J2, and local nodes D, E, and F. The junction network is formed by nodes J1, J2, and J3 linked by their associated junction trunks.

For clarity when illustrating a structured network, the junction nodes are represented by small polygons while local nodes are represented as a small square. A dotted line represents the boundary of each domain. Domains are numbered and given a name. In node screen displays domains are indicated to be part of a domain by preceding the node name with the prefix “Dn” where n is a number, followed by a name. For example, domain 1, in this example might be referred to as D1.alpha.


Figure 5-21 Simplified Structured Network




A network may be configured with up to eight junction nodes. Currently, with each domain equipped with one junction node, a structured network may be equipped with up to 384 nodes (8 domains with one junction x 63 nodes per domain). In some network configurations, each domain may be equipped with two junction nodes. In these cases, the network is limited to 192 nodes (4 domains with two junctions x 63 nodes per domain).

A domain with two junction nodes is often desirable for networks where domains are bounded geographically on two sides by adjacent domains (e.g. domain 2 in Figure 5-22)Figure 5-22. In this network configuration, it is necessary to connect junction nodes, J2 and J3, directly even though a path exists through local nodes C and D. Interdomain traffic between domain 1 and domain 3 must pass directly through junction nodes with no intermediate local nodes.

Figure 5-22 Structured Network with Non Redundant Junction Nodes



Two junction nodes can also provide an increased measure of network robustness as they provide alternate routes in or out of a domain in the unlikely event that a junction node should completely fail (Figure 5-23Figure 5-23). For instance, if junction node, J1, should fail, inter-domain traffic could route through J2 to domain 2. If J1 were the only junction node, domain 1 would be isolated for the duration of the failure.

Figure 5-23 Structured Network with Redundant Junction Nodes

Connections in Structured Networks

Structured network connections are categorized as either intra-domain and inter-domain.

Intra-domain connections, originating and terminating within the same domain, have the much the same characteristics as any flat network connection. They are owned by the node where the connection was added and that node determines the overall connection route.

Inter-domain connections, which cross one or more domain boundaries, are a bit more restricted. In general, inter-domain connections are like intra-domain connections.

- They must be added or deleted from a local node.
- They cannot originate from or terminate on junction nodes.
- They must pass through a junction node on each side of a domain boundary.
- They are routed automatically, just like an intra-domain connection but the routing is determined partially by the originating local node and partially by the junction nodes that the connection passes through.
- They must pass through junction nodes only when traversing a domain where they do not terminate.
- They are restricted in their use of certain features. These features require knowledge of the available bandwidth along the whole path that is not known by all nodes.

The major difference between these two connection types is their routing. A network operator does not need to specify whether a connection is intra-domain or inter-domain, it is merely determined by the specified connection endpoints when the connection is added to the network. Both connections are automatically routed by the IPX/IGX/BPX network, no manual routing is required. Likewise, any rerouting around network trunk failures is completely automatic.

Design Considerations

There are certain factors that must be considered when converting an existing flat network to a structured network. These factors are listed below:

- A domain in a structured network supports all IPX, IGX, and BPX features and functions found

Comment: Intra-domain signifies within the domain, and inter-domain signifies between domains.

in a flat network.

- Certain network-wide administrative features are not supported across domains. User ID, password, user access privilege level, event logging, statistics gathering, maintenance printing, and network clock sources are all maintained separately for each domain.
- Voice connections across domain boundaries must use a CDP (IPX) or CVM (IGX) card at each end.
- Voice Activity Detection (VAD) for voice connections and Data Frame Multiplexing (DFM) for data connections are supported across domain boundaries.
- Adaptive Voice, the automatic enabling and disabling of VAD depends on the knowledge of available bandwidth throughout the network and is not supported across domain boundaries. It is supported on intra-domain voice connections.
 - Routing features like Priority Bumping and Preferred Routing function only on connections within a domain.
 - Automatic routing and rerouting of intra-domain connections between local nodes in the same domain must be contained within domains. An intra-domain connection that originates and terminates within the same domain can not be routed across a domain boundary even to route around a failure. Otherwise it is considered an inter-domain connection.
 - Connections that originate in one domain and route across an intermediate domain to terminate in a third domain (a via-domain connection) may not pass through local nodes in the intermediate node.
 - An unlimited number of intra-domain connections may be routed between two junction nodes in the same domain. However, there is a limit to the number of via-domain connections (typically 50) that may be routed in this manner.
 - Junction trunks do not support the up connection (**upcon**) and down connection (**dncon**) commands. These commands must be entered from a local node at one end or the other of the connection.
 - The test connection (**tstcon**) command operates within the domain as it does with a flat network. For inter-domain connections, it will report a pass/fail for the connection but cannot isolate the failure to a card.
 - System parameters are defined (using the **cnfsysparm** command) and updated for each domain separately and cannot be defined network-wide.

The performance of a network divided into domains improves when the nodes are clustered so that intra-domain traffic exceeds inter-domain traffic. Once the decision is made to divide an existing flat network into domains, the first task is identifying the communities of interest in the network with the intent of turning them into domains.

When selecting groups of nodes to include in the various domains consider the following:

- Inter-domain traffic is more costly than intra-domain traffic. Inter-domain connections utilize more trunks and bandwidth. When dividing the network into domains, try to maximize the amount of intra-domain traffic versus the amount of inter-domain traffic. Nodes with many connections or requiring high throughput between them, should be included in the same domain.
- The number of junction nodes in the network increases with the number of domains in the network. To reduce cost and improve rerouting time, effort should be made to minimize the number of domains.
- Look at the geographical layout of the network. If administration and maintenance are to be performed on a domain-by-domain basis, consider any territorial or regional boundaries that may already exist.
- Look at any existing tariff boundaries that may make logical domain boundaries.

- Once domains have been identified, select the junction nodes as follows. The nodes with the least number of connections are good candidates to become junction nodes. This helps to minimize the number of inter-domain connections.

BPX in Structured Networks

The BPX node may be used in Structured Networks along with the IPX and IGX. The BPX provides greatly increased bandwidth capacity and higher switching speeds. An IPX replaced by a BPX can be reassigned to serve as a local node at the same site.

In very large networks, it may be desirable to use the BPX for all junction nodes.

With junction domains equipped with BPXs, the maximum size of the network increases to 32 junction nodes maximum. But, more importantly, the backbone junction network provides much more bandwidth for interdomain trunking with the use of T3 or E3 trunks.

Another advantage to using BPX junction nodes is that it converts the network backbone to an ATM network. Extending the ATM service to the users is facilitated with User-Network Interface provided by the BPX Access Interfaces and existing IPX and IGX local node accesses.

