

Frame Relay to ATM Network and Service Interworking

Introduction

Interworking allows users to retain their existing Frame Relay services, and as their needs expand, migrate to the higher bandwidth capabilities provided by BPX ATM networks. Frame Relay to ATM Interworking enables frame relay traffic to be connected across high-speed ATM trunks using ATM standard Network and Service Interworking (Figure 14-1 and Figure 14-2).

Two types of Frame Relay to ATM interworking are supported, Network Interworking and Service Interworking. The Network Interworking function is performed by the AIT card on the IPX and by the BTM card on the IGX. The FRSM card on the AXIS supports both Network and Service Interworking.

Figure 14-1 Frame Relay to ATM Network Interworking

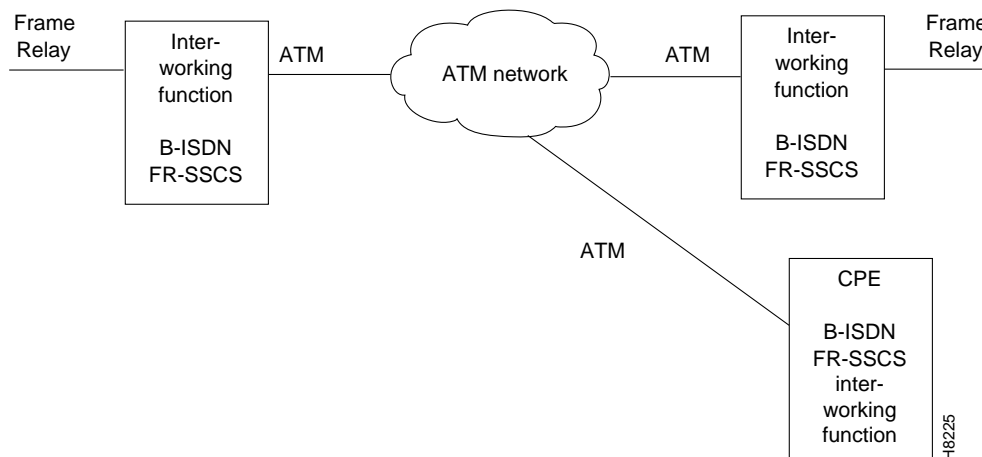

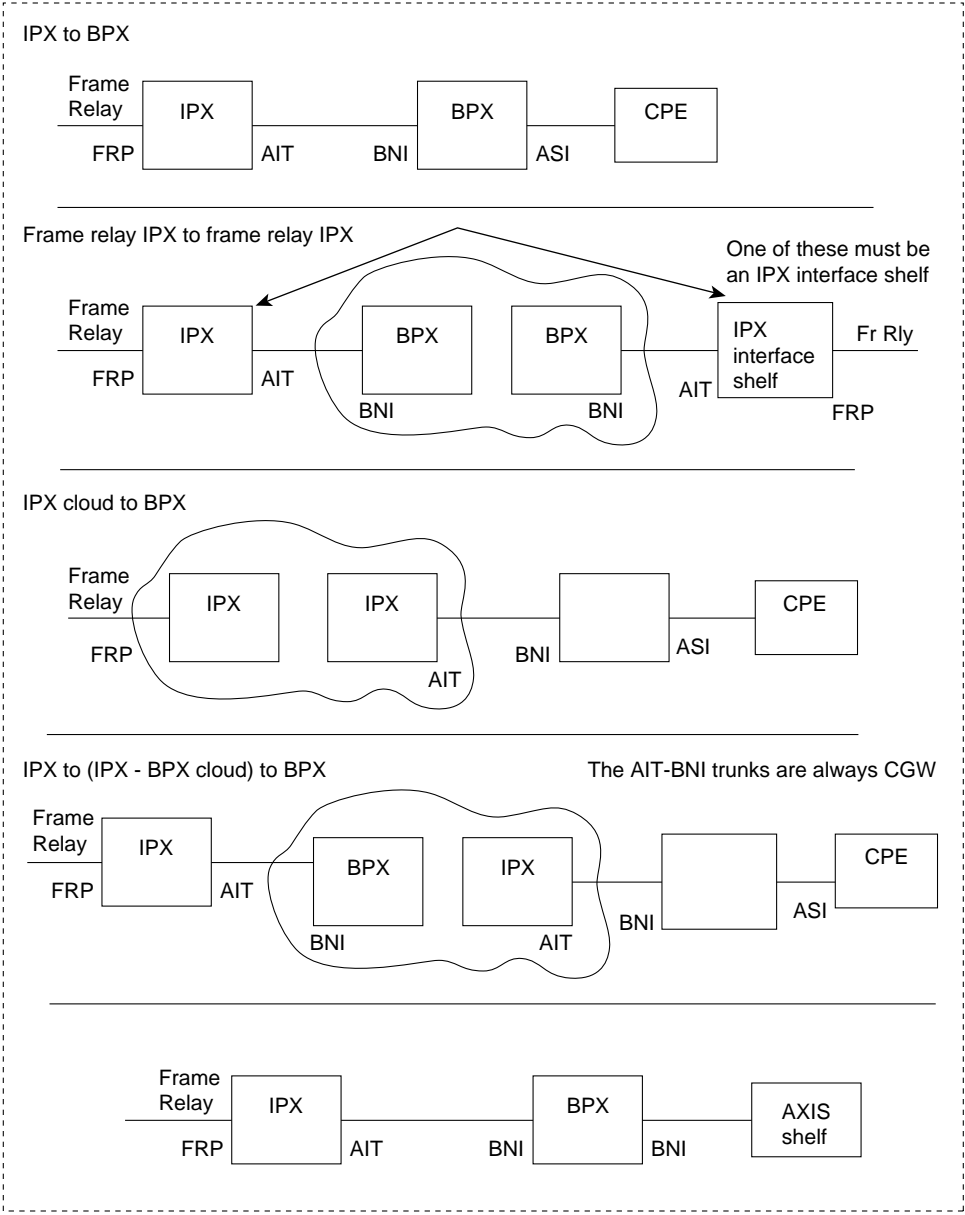


Figure 14-2 Frame Relay to ATM Service Interworking



See Figure 14-3 for some examples of ATM to Frame Relay Interworking.

Figure 14-3 Frame Relay to ATM Interworking Examples with AIT Card on IPX



S5239

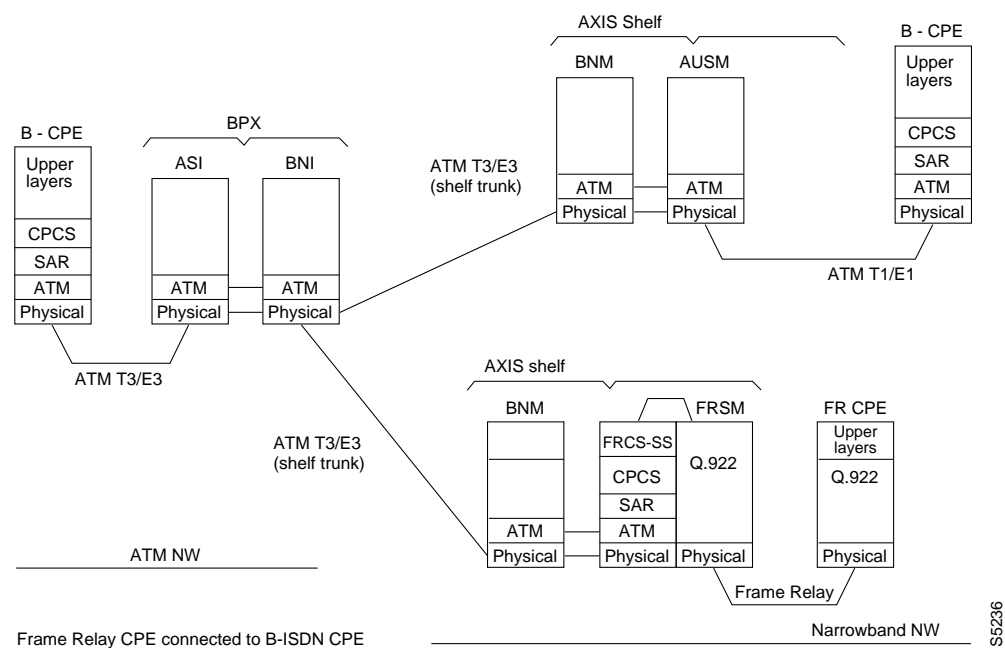
AIT Interworking Examples

Service Interworking

In Service Interworking, for example, for a connection between an ATM port and a frame relay port, unlike Network Interworking, the ATM device does not need to be aware that it is connected to an interworking function. The ATM device uses a standard service specific convergence sublayer, instead of using the Frame Relay FR-SSCS (Figure 14-4).

The frame relay service user does not implement any ATM specific procedures, and the ATM service user does not need to provide any frame relay specific functions. All translational (mapping functions) are performed by the intermediate IWF. The ATM endpoints may be any ATM UNI/NNI interface supported by the AXIS shelf, e.g., ASI, AUSM. Translation between the Frame Relay and ATM protocols is performed in accordance with RFC 1490 and RFC 1483.

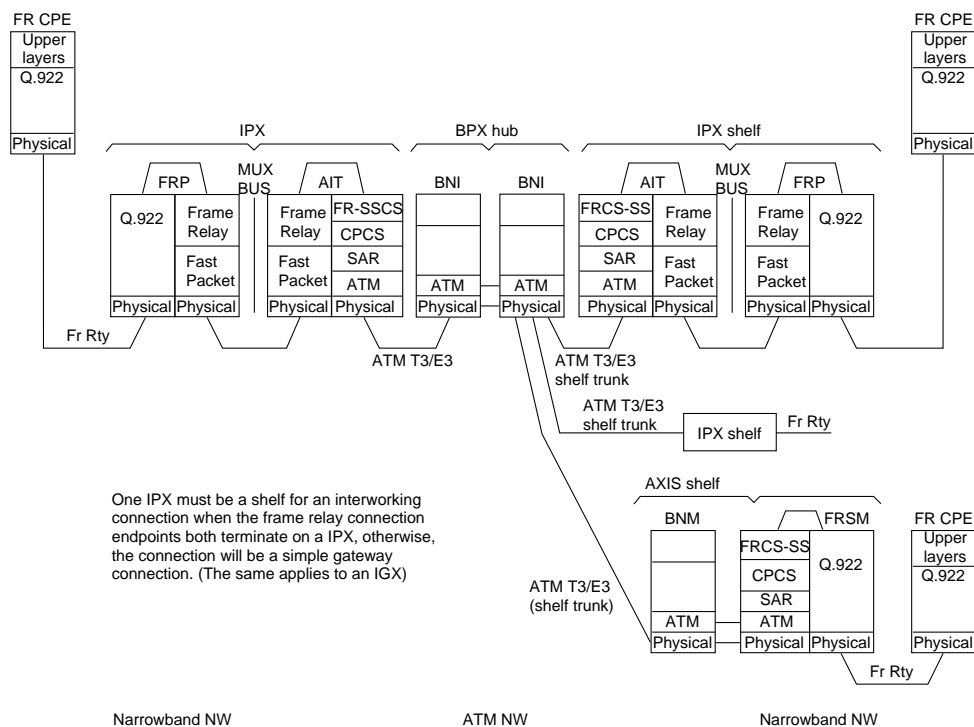
Figure 14-4 Frame Relay to ATM Service Interworking



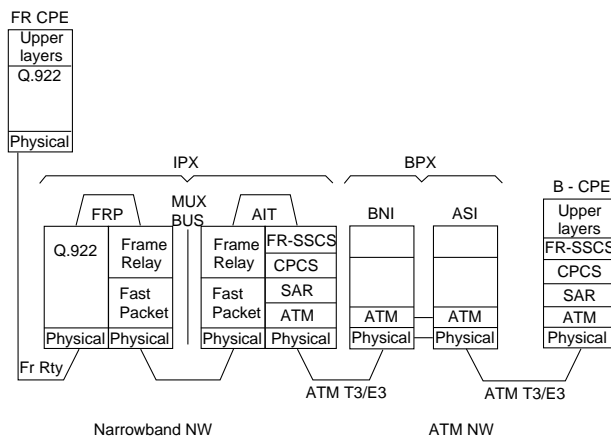
Networking Interworking

In Network Interworking, in most cases, the source and destination ports are frame relay ports, and the interworking function is performed at both ends of the connection as shown in Part A of Figure 14-5.

If a frame relay port is connected across an ATM network to an ATM device, network interworking requires that the ATM device recognize that it is connected to an interworking function (frame relay, in this case). The ATM device must then exercise the appropriate service specific convergence sublayer (SSCS), in this case the frame relay service specific convergence sublayer (FR-SSCS) as shown in Part B of Figure 14-5.

Figure 14-5 Frame Relay to ATM NW Interworking

Part A. Network interworking, frame relay to frame relay connected across an ATM network



Part B. Network interworking, frame relay to ATM CPE connected across an ATM network

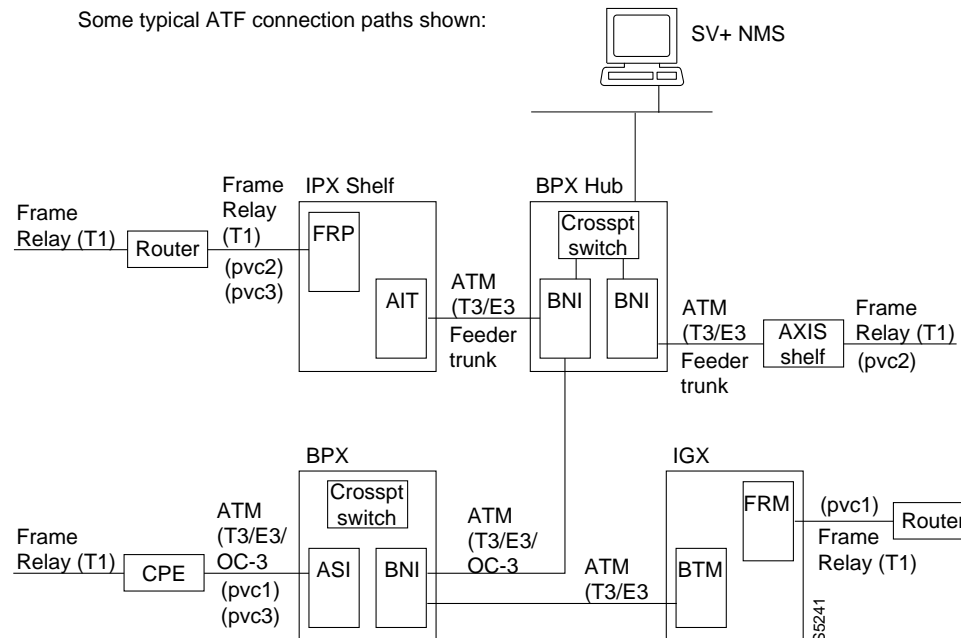
The frame relay to ATM networking interworking function is available as follows:

- IPX frame relay (shelf/feeder) to IPX frame relay (either routing node or shelf/feeder)
- AXIS frame relay to AXIS frame relay
- AXIS frame relay to IPX frame relay (either routing node or shelf/feeder)
- IPX frame relay (either routing node or shelf/feeder) to BPX or AXIS ATM port.
- AXIS frame relay to BPX or AXIS ATM port.
- In the items listed above, an IGX can be substituted for each instance of an IPX.

On the IPX, interworking is performed by the AIT card, and on the IGX by the BTM card. A simplified example of the connection paths is shown in Figure 14-6. In interworking, the AIT card receives FastPackets from the FRP, rebuilds the frames, and converts between frames and ATM cells. Data is removed from one package and placed in the other. Congestion information from the header is mapped to the new package. This processing by the AIT trunk card is called Complex Gateway. AIT trunk cards are required on every BPX to IPX hop in a Frame Relay to ATM connection's path.

Figure 14-6 ATF Connections, Simplified Example

Some typical ATF connection paths shown:



The cells within the frame are expected to possess the standard ATM Access Interface cell header. The traffic is assumed to have AAL-5 PDUs, and will not function properly otherwise (framing errors will result). Within the AAL-5 PDUs, the data must be packaged in standard frame relay frames, one frame per PDU (with respect to the AAL-5 layer).

The UPC and ForeSight algorithms are applied according to their configured values. The cell headers are converted into the proprietary StrataCom STI format before entering the network. The cells are delivered to their destination according to the configured route of the connection. Cells can be lost due to congestion.

Discard selection is based upon the standard CLP bit in the cells. When the routing path enters an IPX/IGX, an AIT/BTM card which supports Interworking traffic is required to convert the connection data from cells to frames (frames to fastpackets out onto MuxBus to FRP/cell bus to FRM), and visa versa. Additionally, the AAL-5 framing is removed upon conversion to frames, and added upon conversion to cells. At the destination (FRP), FastPackets are placed in the port queue and, when a complete frame has been assembled, the frame is played out the remote port in the original format (as provided in the frames delivered inside AAL-5 PDUs).

For each connection, only a single dlci can be played out for all traffic exiting the port, and is inserted into the frame headers. The standard LAPD framing format is played out the port on the FRP/FRM.

At the AIT/FRM card, several additional protocol mappings take place. First, the Interworking Unit acts as a pseudo endpoint for the purposes of ATM for all constructs which have no direct mapping into Frame Relay, such as loopbacks and FERF indications. Thus, end-to-end loopback OAM cells which come to AIT/FRM cards are returned to the ATM network without allowing them to proceed

into the Frame Relay network, which has no equivalent message construct. Further, AIS and supervisory cells and FastPackets (from the Frame Relay direction) are converted into their counterparts within the other network.

ATM Protocol Stack

A general view of the ATM protocol layers with respect to the Open Systems Interconnection model is shown in Figure 14-7. In this example, a large frame might be input into the top of the stacks. Each layer performs a specific function before passing it to the layer below. A protocol data unit (PDU) is the name of the data passed down from one layer to another and is the Service Data Unit (SDU) of the layer below it. For Frame Relay to ATM interworking, a specific convergent sublayer, Frame Relay Service Specific Convergent Sublayer, FR-SSCS is defined. This is also referred to as FR-CS, in shortened notation.

Figure 14-7 **ATM Layers**

AIT/BTM Interworking and the ATM Protocol Stack

ATM to Frame Relay interworking (ATF) performs various tasks including the following:

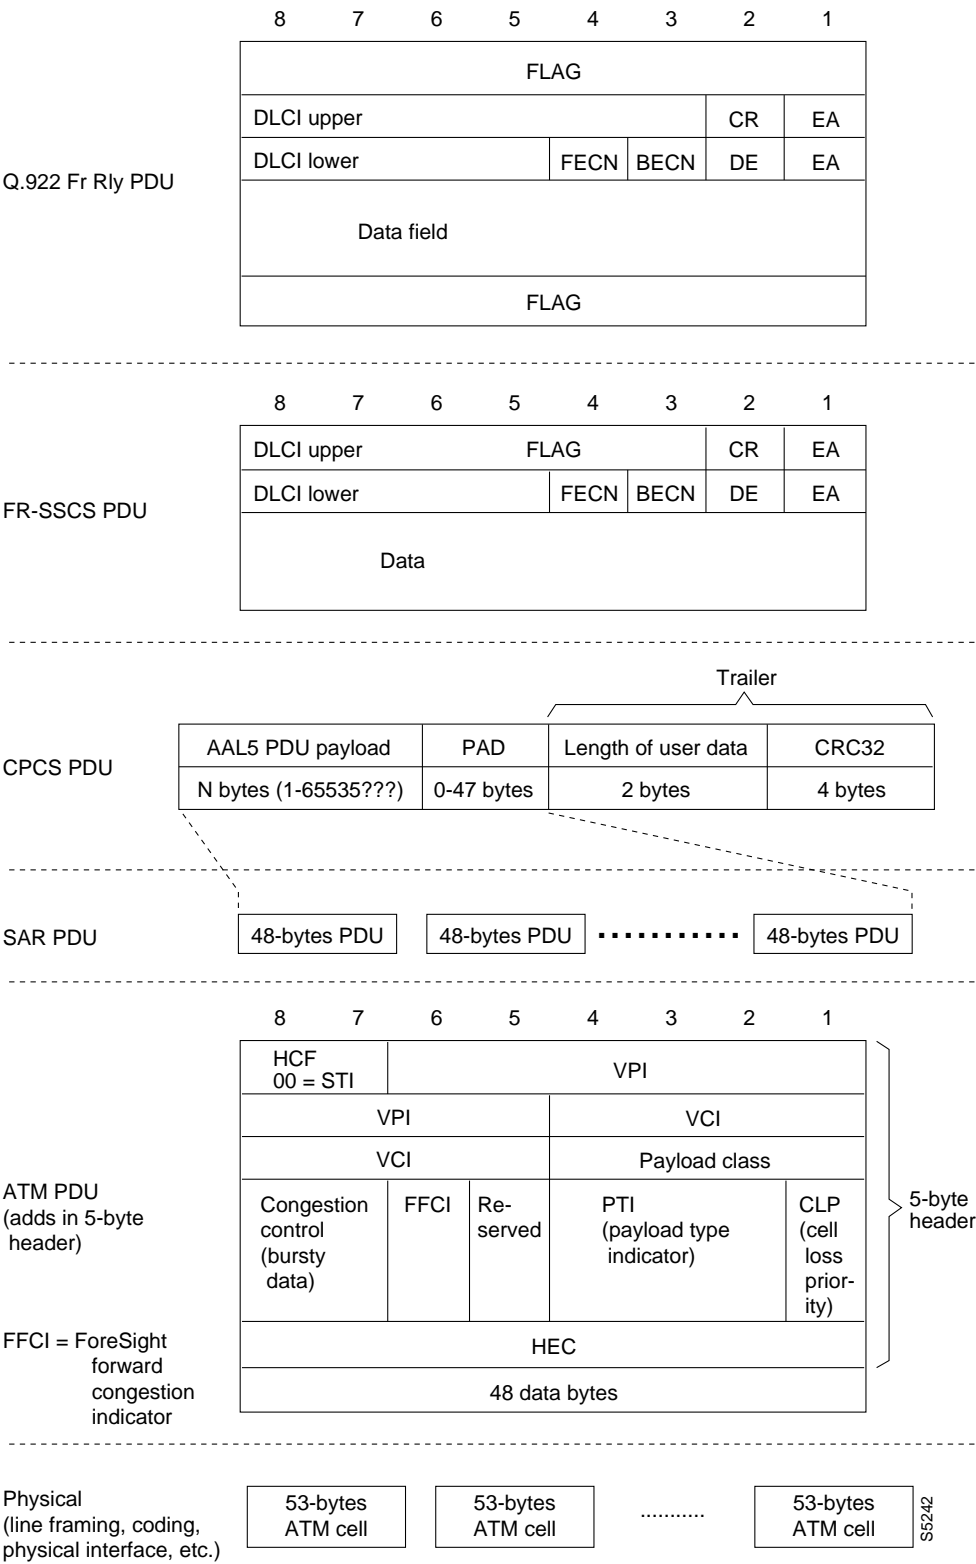
- Conversion of PDUs between the frame relay and ATM virtual circuits of the frame relay and ATM user devices.
- Conversion between frame relay traffic service and ATM quality of service parameters
- Mapping of management status, including connection, port, line, and trunk status and events

Figure 14-8 depicts the function of the protocol stack layers in the interworking between ATM and Frame Relay PDUs. Interworking by the AIT/BTM card in the IPX/IGX includes the following functions:

- Translating the ATM pvc identifier (vpi.vci) to the frame relay pvc identifier (dlci) and vice versa

- Mapping the Protocol Data Unit (PDU), which is essentially the data, between the Frame Relay Service Specific Convergence Sublayer (FR-SSCS) and the Frame Relay Q.922 core protocol, and vice versa.
- On the IPX, Incoming Frames are converted to FastPackets by the FRP card. The FastPackets are then routed to the AIT card via the IPX MUX bus and converted back into Frame Relay Q.922 frames by the AIT card. Then the AIT card interworking function executes four layers to convert the Frame PDU to ATM cells:
 - the FRCS layer (Frame Relay Service Specific Convergence Sublayer, FRSSCS, or FRCS for in shortened notation) which uses a PDU format identical to the Q.922 core (without CRC-16 or flags),
 - the CPCS layer (Common Part Convergence Sublayer) which appends a CS-PDU trailer to the FR-PDU to create a CS-PDU
 - the Segmentation and Reassembly layer (SAR) which segments the CS-PDU (Protocol Data Unit) into SAR-PDUs (48 byte data entities)
 - and the ATM layer which attaches an ATM header to each SAR-PDU to create an ATM-SDU (Service Data Unit). The same process is performed in the reverse order by the AIT card when transforming cells to frames.

Figure 14-8 Protocol Stack Operation



AIT/BTM Control Mapping, Frames and Cells

In addition to performing DLCI to PVC/VCC conversion, the network interworking feature provided by the AIT card in the IPX or BTM in the IGX maps cell loss priority, congestion information, and management information between frame relay and ATM formats as follows:

CELL LOSS PRIORITY, Frame Relay to ATM Direction

Each frame relay to ATM network interworking connection can be configured as one of the following DE to CLP mapping choices:

- The DE bit in the frame relay frame is mapped to the CLP bit of every ATM cell generated by the segmentation process.

The following 2 choices are not available on IPX/IGX NIW (network interworking):

- CLP is always 0
- CLP is always 1

CELL LOSS PRIORITY, ATM to Frame Relay Direction

Each frame relay to ATM network interworking connection can be configured as one of the following CLP to DE mapping choices:

- If one or more ATM cells belonging to a frame has its CLP field set, the DE field of the frame relay frame will be set.

The following choice is not available:

- Choosing no mapping from CLP to DE

CONGESTION INDICATION, Frame Relay to ATM direction

- EFCI is always set to 0.

CONGESTION INDICATION, ATM to Frame Relay Direction

- If the EFCI field in the last ATM cell of a segmented frame is set, then FECN of the frame relay frame will be set.

For PVC Status Management

The AIT/BTM does convert OAM cells to OAM fastpackets, and vice-versa, including the AIS OAM. Also, “A-bit” status is now propagated via software messaging.

The ATM layer and frame relay PVC Status Management can operate independently. The PVC status from the ATM layer will be used when determining the status of the FR PVCs. However, no direct actions of mapping LMI A bit to OAM AIS will be performed.

Management, OAM Cells

OAM cell processing:

- F5 OAM loopback
- AIS
- FERF
- StrataCom Internal OAM

Functional Description

ATF Summary

Features

- Interworking: ATM to Frame Relay connections
- Connection Statistics
- Round Trip Delay measurements incorporated into the ForeSight algorithm
- Frame Based GCRA (FGCRA). This is an enhancement of the Generic Cell Rate Algorithm
- IBS (Initial Burst Size)
- cnfportq: 3 egress port queues are configurable CBR, VBR and VBR w/Foresight. (Queue Bin numbers and algorithm types are NOT user selectable.)
- BCM (Backward Congestion Messages)
- ILMI and associated configuration options and statistics
- Loopback functions: **tstdly**, **tstconseg**, **addrmtlp**, **addloclp**
- Selftest/ Background tests
- OAM flows: AIS, FERF, OAM loopback
- ASI/2 E3 support
- End-to-end status updates (per FR/ATM interworking)
- Annex G and associated configuration options and statistics
- ASI-1 as a clock source is supported.

Limitations

- Priority Bumping is not supported across the interface shelves, but is supported across the routing network.
- Statistical Line Alarms per Software Functional Specification (i.e., Bellcore standards).
- Programmable Opti Class: although 4 connection classes are supported: CBR, VBR, VBR with Foresight, ATF, and ATF with ForeSight. Configuration of egress port queues and BNI trunk queues for these connection classes is available.
- Port loopback “**tstport**”

- Test “**tstcon**” not supported at BPX endpoints; it is supported at IPX endpoints
- Gateway terminated inter-domain connections
- Via connections through IPX

Some ATF Connection Criteria

ATF connections are allowed between any combination of ATM and Frame Relay UNI and NNI ports. Virtual circuit connections are allowed. Virtual path connections are not.

ATF connections can be mastered by the IPX or BPX end.

ATF bundled connections and ATF point-to-point connections are not supported.

ATF connections use the frame relay trunk queues: bursty data A for non-ForeSight, bursty data B for ForeSight.

Bandwidth related parameters are defined using cells per second (cps) on the BPX, and bits per second (bps) on the IPX/IGX. On a given endpoint node, the bandwidth parms for both ends of the ATF connection are changed/displayed using this end's units. This saves the user from having to convert from cps to bps repeatedly.

ATF connections use the VBR egress queue on the ASI-1 card. ATF with ForeSight connections use the ABR egress queue.

Connection Management

The following user commands are used to provision and modify ATF connections:

- **addcon**
- **cnfcls**
- **cnfcon**
- **delcon**
- **dspcls**
- **dspcon**
- **dspcons**

Port Management

The following features are added to the ASI-1 at the port level:

- An ASI-1 card can be configured to use the network-network interface (NNI) addressing format. This feature is only available on a per-card level, so changing one port to or from NNI changes the other one also, with appropriate warnings to the user.
- ILMI activation/configuration/statistics
- LMI Annex G activation/configuration/statistics
- Port egress queue configuration
- Backward congestion management

Structure

- NNI
- The NNI format supports a 12-bit VPI. A-bit status changes are passed to the remote end of the connection.
- ILMI
- The ILMI MIB and protocol was implemented in release 7.2. The additional support in consists of an activation and configuration interface, collection of statistics, and end-to-end status updates.
- LMI Annex G
- The LMI Annex G protocol was implemented in release 7.2. The additional support consists of an activation and configuration interface, collection of statistics, and end-to-end status updates.
- Port egress queue configuration
- Each of the pre-defined ASI-1 port egress queues can be configured by the user. These queues consist of CBR, VBR, and VBR with ForeSight (ABR). The configurable parameters are queue depth, EFCN threshold, and CLP thresholds.
- Backward congestion management
- Backward congestion management cells indicate congestion across the UNI or NNI. Transmission of these cells is enabled on a per-port basis. Software allows BCM to be configured on a UNI or NNI port for maximum flexibility should BCM over UNI be standards-defined.

The following user commands are used to configure ASI-1 port features:

- **cnfport**
- **cnfportq**

Channel Statistics

Statistics are supported on a per-channel basis. A range of traffic and error statistics are available. ASI-1 channel statistics are enabled by StrataView+ or by the BPX control terminal using the existing statistics mechanism. The existing collection intervals apply.

Channel statistics of the following general types are supported:

- Cells received/transmitted, dropped, tagged as non-compliant or congested
- Cell errors
- AAL-5 frame counts, errors

The following user commands are used to configure and display channel statistics

- **clrchstats**
- **cnfchstats**
- **dspchstats**
- **dspchstatcnf**
- **dspchstathist**

OAM Cell Support

OAM cells are detected and transmitted by the ASI-1 firmware. System software displays alarm indications detected by the firmware. Additionally, loopbacks between the ATM-UNI and the ATM-CPE can be established. ForeSight round-trip delay cells are generated by firmware upon software request.

System software deals with the following OAM cell flows:

- End-to-End AIS/FERF—software displays on a per-connection basis.
- External segment loopbacks—software initiates loopback of ATM-CPE via user command. The SAR creates the loopback OAM cell. External loopback cells received from the ATM-CPE are processed by the SAR.
- Internal ForeSight round trip delay—software commands the ASI-1 to measure the RTD excluding trunk queueing delay on each ForeSight connection. Software displays the result.
- Internal loopback round trip delay—software commands the ASI-1 to measure the RTD including trunk queueing delay on each ForeSight connection. Software displays the result.
- Internal Remote Endpoint Status—these cells are generated by one end of a connection due to remote network connection failure (A-bit = 0). The other end ASI-1 detects these cells and reports the connection status to software, which displays it.

The following user commands are associated with OAM cell status changes:

- **dspalms**
- **dspcon**
- **dspport**
- **tstconseg**
- **tstdly**

Diagnostics

- Loopbacks
- Local loopbacks loop data back to the local ATM-TE, via the local BPX. Remote loopbacks loop data back to the local ATM-TE, via the whole connection route up to and including the remote terminating card.
- Local and remote connection loopbacks, and local port loopbacks, are destructive.
- Card Tests
- The generic card selftest mechanism on the BPX is modified to include the ASI-1 card.
- The card background test that exists for the FRP card on the IPX is modified to work for the ASI-1 card.
- Connection Tests
- The tstcon command is not supported. The tstdly command is used for connection continuity testing. ASI-1 tstdly is non-destructive, as compared with the IPX tstdly.

User Commands

The following user commands are associated with diagnostics changes:

- **addloclp**
- **addrmtlp**
- **cnftstparm**
- **dellp**
- **dspalms**
- **dspcd**
- **dspcds**
- **tstdly**

Virtual Circuit Features

The following virtual circuit features are supported by the ASI-1:

- **Connection Groups**

Connection groups are supported for ASI-1 ATM and interworking connection types, allowing termination of up to 5000 (grouped) virtual circuits per BPX node. The connection grouping feature currently available on frame relay connections is expanded to include ASI-1 ATM and interworking connections.

- **FGCRA**

Frame-Based Generic Cell Rate Algorithm is an ASI-1 firmware feature that controls admission of cells to the network. It is configurable on a per-connection basis. It is a StrataCom enhancement of the ATM-UNI standard Generic Cell Rate Algorithm. System software allows configuration of FGCRA on a per-connection basis.

- **IBS**

Initial Burst Size is an ATM bandwidth parameter that is used by firmware to allow short initial bursts, similar to the Cmax mechanism on the IPX. It is configurable on a per-connection basis

- **Full VPI/VCI addressing range**

The entire range of VPI and VCI on both UNI and NNI interfaces is supported. For ATM-UNI, 8 bits of VPI and 16 bits of VCI are supported. For ATM-NNI, 12 bits of VPI and 16 bits of VCI are supported. In either case, VPC connections only pass through the lower 12 bits of the VCI field.

- **Connection Classes**

ATM and interworking connection classes are defined with appropriate bandwidth parameter defaults. These classes only apply at addcon time. They are templates to ease the user's task of configuring the large number of bandwidth parameters that exist per connection.

User Commands

The following user commands are associated with virtual circuit feature changes:

- **addcon**
- **addcongrp**
- **cnfcon**
- **cnfatmcls**
- **delcon**
- **delcongrp**
- **dspatmcls**
- **dspcongrps**
- **grpcon**

ATM-E3

The following areas are modified to support the ASI-1 E3 interface:

- Card management
- Line management
- Statistics
- Alarms
- Displays/logging

User Commands

The following user commands are modified to support ASI-1 E3:

- **cnfln**
- **cnflnstats**
- **dspcd**
- **dspcds**
- **dsplncnf**
- **dsplns**
- **dsplnstatcnf**
- **dsplnstathist**
- **dspyred**
- **prtyred**

Management

Connection Management

Interworking connections may be added from either the BPX, the IPX, the IGX, or the AXIS. Intra- and inter-domain interworking connections are supported.

Connection configuration parameters are endpoint-specific. Thus, the ATM-only parameters are only configurable on the BPX end. The IPX does not know about these parameters, so they cannot be configured or displayed at the IPX end. Parameter units are endpoint-specific also. Units on the BPX are cells per second, units on the IPX are bits per second.

Bundled interworking connections are not supported.

Virtual path interworking connections are not supported.

Routing

Interworking connections use the complex gateway feature of the AIT trunk card to repackage data from frames to ATM cells, and vice-versa. All BPX-IPX hops these connections route over must provide the complex gateway function. IPX-IPX hops (frame relay connections) can be any trunk card type. This requirement simplifies the routing mechanism when dealing with structured networks, as software does not know the type of trunks in remote domains.

Bandwidth Management

Bandwidth calculations for interworking connections assume a large frame size, which minimizes the loading inefficiency of packets vs. cells. In other words, the translation between packets and cells assumes 100 percent efficiency, so the conversion is simply based on 20 payload bytes per fastpacket vs. 48 payload bytes per ATM cell.

This mechanism keeps the fastpacket/cell conversion consistent with the bits per second/cells per second conversion. Thus, conversion of endpoint rates to trunk loading is straightforward.

User Interface

ATM connection classes are added for convenience. Classes can be configured as interworking or regular ATM. The **cnfcls** command is used to configure a class. The class is specified as part of the **'addcon'** command. ATM connection classes are maintained on all BPX nodes. IPX nodes do not know about these classes.

A special ATM class is defined as the default interworking class. When an interworking connection is added from the frame relay end, the ATM-only parameters for this connection are taken from this default class.

Network-wide ForeSight parameters are supported for the frame relay end of interworking connections. The **'cnffstparm'** command is used to configure these parameters. Since the ATM end of interworking connections has per-virtual circuit ForeSight parameter configurability, the network-wide ForeSight parameters do not apply.

Note that the default ATM ForeSight parameters will match the default frame relay ForeSight parameters, with appropriate units conversion.

Port Management

The 'cnfport' command supports the following new features:

- An ASI-1 card can be configured to be UNI or NNI.
- An ASI-1 UNI or NNI port can be configured to transmit Backwards Congestion Messages (BCM) to indicate congestion to the foreign ATM network.
- An ASI-1 UNI or NNI port can be configured for LMI, ILMI, or no local management.

The 'cnfportq' command supports configuration of queue depth, EFCN threshold, and CLP thresholds for all port egress queues (CBR, VBR, VBR w/ForeSight).

Connection Management

The NNI cell format has 12 bits for the VPI, so 'addcon' allows specification of VPI 0-4095 on NNI ports.

Signaling

System software supports the following LMI/ILMI signaling actions:

- Internal network failure: software informs LMI/ILMI to set A bit = 0 for failed connections. Software informs ASI-1 to transmit AIS to port for failed connections.
- Port failure/LMI Comm Failure: software informs remote nodes terminating all affected connections. Remote node BCC informs LMI/ILMI to set A bit = 0, and ASI-1 to transmit AIS.
- LMI A = 0: software polls ILMI agent periodically for A-bit status. Status changes are reflected in the 'dspcon' screen.

Alarms

LMI communication failure on an ASI-1 causes declaration of a minor alarm. The 'dspport' screen shows the failure, as does the 'dspalms' screen.

A-bit = 0 on an NNI port causes declaration of a minor alarm. The 'dspcon', 'dspcons', and 'dspalms' screens show this failure.

