

Cisco Reader Comment Card

General Information

- 1 Years of networking experience _____ Years of experience with Cisco products _____
- 2 I have these network types: LAN Backbone WAN
 Other: _____
- 3 I have these Cisco products: Switches Routers
 Other: Specify model(s) _____
- 4 I perform these types of tasks: H/W Install and/or Maintenance S/W Config
 Network Management Other: _____
- 5 I use these types of documentation: H/W Install H/W Config S/W Config
 Command Reference Quick Reference Release Notes Online Help
 Other: _____
- 6 I access this information through: _____% Cisco Connection Online (CCO) _____% CD-ROM
_____% Printed docs _____% Other: _____
- 7 Which method do you prefer? _____
- 8 I use the following three product features the most:

Document Information

Document Title: Cisco WAN Switching System Overview

Part Number: 78-7041-01

S/W Release (if applicable): 9.1

On a scale of 1–5 (5 being the best) please let us know how we rate in the following areas:

_____ The document was written at my technical level of understanding.

_____ The information was accurate.

_____ The document was complete.

_____ The information I wanted was easy to find.

_____ The information was well organized.

_____ The information I found was useful to my job.

Please comment on our lowest score(s):

Mailing Information

Company Name

Date

Contact Name

Job Title

Mailing Address

City

State/Province

ZIP/Postal Code

Country

Phone (

Extension

Fax ()

E-mail

Can we contact you further concerning our documentation? Yes No


You can also send us your comments by e-mail to bug-doc@cisco.com, or fax your comments to us at (40 8)527-8089.

BUSINESS REPLY MAIL


FIRST-CLASS MAIL PERMIT NO. 4631 SAN JOSE CA

POSTAGE WILL BE PAID BY ADDRESSEE

ATTN DOCUMENT RESOURCE CONNECTION
CISCO SYSTEMS INC
170 WEST TASMAN DRIVE
SAN JOSE CA 95134-9883



NO POSTAGE
NECESSARY
IF MAILED
IN THE
UNITED STATES



Cisco WAN Switching System Overview

Release 9.1

July, 1999

Corporate Headquarters

Cisco Systems, Inc.
170 West Tasman Drive
San Jose, CA 95134-1706
USA

<http://www.cisco.com>

Tel: 408 526-4000
800 553-NETS (6387)

Fax: 408 526-4100

Customer Order Number: DOC-787041=
Text Part Number: 78-7041-01

THE SPECIFICATIONS AND INFORMATION REGARDING THE PRODUCTS IN THIS MANUAL ARE SUBJECT TO CHANGE WITHOUT NOTICE. ALL STATEMENTS, INFORMATION, AND RECOMMENDATIONS IN THIS MANUAL ARE BELIEVED TO BE ACCURATE BUT ARE PRESENTED WITHOUT WARRANTY OF ANY KIND, EXPRESS OR IMPLIED. USERS MUST TAKE FULL RESPONSIBILITY FOR THEIR APPLICATION OF ANY PRODUCTS

THE SOFTWARE LICENSE AND LIMITED WARRANTY FOR THE ACCOMPANYING PRODUCT ARE SET FORTH IN THE INFORMATION PACKET THAT SHIPPED WITH THE PRODUCT AND ARE INCORPORATED HEREIN BY THIS REFERENCE. IF YOU ARE UNABLE TO LOCATE THE SOFTWARE LICENSE OR LIMITED WARRANTY, CONTACT YOUR CISCO REPRESENTATIVE FOR A COPY.

The following information is for FCC compliance of Class A devices: This equipment has been tested and found to comply with the limits for a Class A digital device, pursuant to part 15 of the FCC rules. These limits are designed to provide reasonable protection against harmful interference when the equipment is operated in a commercial environment. This equipment generates, uses, and can radiate radio-frequency energy and, if not installed and used in accordance with the instruction manual, may cause harmful interference to radio communications. Operation of this equipment in a residential area is likely to cause harmful interference, in which case users will be required to correct the interference at their own expense.

The following information is for FCC compliance of Class B devices: The equipment described in this manual generates and may radiate radio-frequency energy. If it is not installed in accordance with Cisco's installation instructions, it may cause interference with radio and television reception. This equipment has been tested and found to comply with the limits for a Class B digital device in accordance with the specification in part 15 of the FCC rules. These specifications are designed to provide reasonable protection against such interference in a residential installation. However, there is no guarantee that interference will not occur in a particular installation.

Modifying the equipment without Cisco's written authorization may result in the equipment no longer complying with FCC requirements for Class A or Class B digital devices. In that event, your right to use the equipment may be limited by FCC regulations, and you may be required to correct any interference to radio or television communications at your own expense.

You can determine whether your equipment is causing interference by turning it off. If the interference stops, it was probably caused by the Cisco equipment or one of its peripheral devices. If the equipment causes interference to radio or television reception, try to correct the interference by using one or more of the following measures:

- Turn the television or radio antenna until the interference stops.
- Move the equipment to one side or the other of the television or radio.
- Move the equipment farther away from the television or radio.
- Plug the equipment into an outlet that is on a different circuit from the television or radio. (That is, make certain the equipment and the television or radio are on circuits controlled by different circuit breakers or fuses.)

Modifications to this product not authorized by Cisco Systems, Inc. could void the FCC approval and negate your authority to operate the product.

The Cisco implementation of TCP header compression is an adaptation of a program developed by the University of California, Berkeley (UCB) as part of UCB's public domain version of the UNIX operating system. All rights reserved. Copyright © 1981, Regents of the University of California.

NOTWITHSTANDING ANY OTHER WARRANTY HEREIN, ALL DOCUMENT FILES AND SOFTWARE OF THESE SUPPLIERS ARE PROVIDED "AS IS" WITH ALL FAULTS. CISCO AND THE ABOVE-NAMED SUPPLIERS DISCLAIM ALL WARRANTIES, EXPRESS OR IMPLIED, INCLUDING, WITHOUT LIMITATION, THOSE OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE AND NONINFRINGEMENT OR ARISING FROM A COURSE OF DEALING, USAGE, OR TRADE PRACTICE.

IN NO EVENT SHALL CISCO OR ITS SUPPLIERS BE LIABLE FOR ANY INDIRECT, SPECIAL, CONSEQUENTIAL, OR INCIDENTAL DAMAGES, INCLUDING, WITHOUT LIMITATION, LOST PROFITS OR LOSS OR DAMAGE TO DATA ARISING OUT OF THE USE OR INABILITY TO USE THIS MANUAL, EVEN IF CISCO OR ITS SUPPLIERS HAVE BEEN ADVISED OF THE POSSIBILITY OF SUCH DAMAGES.

Access Registrar, AccessPath, Any to Any, AtmDirector, CCDA, CCDE, CCDP, CCIE, CCNA, CCNP, CCSI, CD-PAC, the Cisco logo, Cisco Certified Internetwork Expert logo, *CiscoLink*, the Cisco Management Connection logo, the Cisco NetWorks logo, the Cisco Powered Network logo, Cisco System Capital, the Cisco Systems Capital logo, Cisco Systems Networking Academy, the Cisco Technologies logo, ConnectWay, ControlStream, Fast Step, FireRunner, GigaStack, IGX, JumpStart, Kernel Proxy, MGX, Natural Network Viewer, NetSonar, Network Registrar, New World Packet, PIX, Point and Click Internetworking, Policy Builder, Precept, RouteStream, Secure Script, ServiceWay, SlideCast, SMARTnet, StreamView, *The Cell*, TrafficDirector, TransPath, ViewRunner, VirtualStream, VisionWay, VlanDirector, Workgroup Director, and Workgroup Stack are trademarks; Changing the Way We Work, Live, Play, and Learn, Empowering the Internet Generation, The Internet Economy, and The New Internet Economy are service marks; and Asist, BPX, Catalyst, Cisco, Cisco IOS, the Cisco IOS logo, Cisco Systems, the Cisco Systems logo, the Cisco Systems Cisco Press logo Enterprise/Solver, EtherChannel, EtherSwitch, FastHub, FastLink, FastPAD, FastSwitch, IOS, IP/TV, IPX, LightStream, LightSwitch, MICA, NetRanger, Registrar, StrataView Plus, Stratm, TeleRouter, and VCO are registered trademarks of Cisco Systems, Inc. in the U.S. and certain other countries. All other trademarks mentioned in this document are the property of their respective owners. The use of the word partner does not imply a partnership relationship between Cisco and any of its resellers. (9906R)

Cisco WAN Switching System Overview, Release 9.1

Copyright © 2000, Cisco Systems, Inc.

All rights reserved. Printed in USA.



Prefac	xxv
Objectives	xxv
Audience	xx
Organization	xxvi
Related Documentation	xxviii
Conventions	xxix

CHAPTER 1

Cisco Wide Area ATM Networks	1-1
Introductio	1-1
Expanding Network Requiremen	t 1-2
ATM Networks	1-3
New with Release	8-3 .
Cisco BPX 8600 Series Broadband Swi	t 163
Cisco IGX 8400 Series Multiband Swit	c 1-3
Network	1-3
Cisco MGX 8220 Edge Concentra	t 164
StrataSphere Network Managemen	n 1-4
CiscoView Network Element Managemen	n 1-4
New with Release	8-4 .
Cisco IGX 8400 Series Multiband Swit	c 1-4
Network	1-5
Cisco MGX 8220 Edge Concentra	t 165
StrataSphere Network Managemen	n 1-5
New with Release	8-5 .
Continuing Features with Release 9.1, 8.5 a	n d 1-6 8 .
StrataSphere Network Managemen	n 1-6
Network	1-7
INS-DAS and INS-V	N 1-7
Cisco BPX 8600 Series Wideband Swit	c 1-7
Cisco IGX 8400 Series Multiband Swit	c 1-8
Cisco IPX Narrowband Switc	1-8
Cisco MGX 8220 Edge Concentrat	o 1-8

- Access Product **1-9**
- ATM Networks **1-9**
 - Enterprise Wide Area Network **1-9**
 - The Cisco IPX Narrowband Switch **1-10**
 - The Cisco IGX 8400 series Multiband Switch **1-10**
 - The Cisco BPX 8600 Series Broadband Switch with Cisco MGX 8220 Edge Concentrator **1-11**
 - The Cisco BPX 8600 Series Broadband Switch with Cisco MGX 8220 Edge Concentrator **1-11**
 - Service Provider Multi-Service Network **1-11**
- ATM NW Features **1-13**
 - Advanced Capabilities **1-13**
 - ABR with VSV **1-13**
 - Frame Relay to ATM Interworking **1-13**
 - Network Interworking **1-13**
 - Service Interworking **1-15**
 - Additional Information on Interworking **1-15**
 - Tiered Networks **1-15**
 - Routing Hubs and Interface Shelf **1-15**
 - Cisco IGX 8400 Series Multiband Switches Configured as Routing Hubs **1-16**
 - BPX Routing Hub **1-18**
 - IMA (Inverse Multiplexing ATM) **1-19**
 - Circuit Emulation Service (CES) **1-19**
 - Zero CIR for Frame Relay **1-19**
 - Virtual Trunking **1-19**
 - BXM Network and Service SONET and T3/E3 Interface **1-20**
 - BXM-622, 155, and T3/E3 Carrier **1-21**
 - Trunk Mod **1-21**
 - Service Mod **1-21**
 - BNI Network and ASI Service Interface **1-21**
 - BNI-155 Network Interface **1-21**
 - ASI-155 Service Interface **1-21**
- Traffic and Congestion Management **1-22**
 - FairShare **1-22**
 - OptiClass **1-23**
 - AutoRoute **1-23**
 - PNN **1-23**

Congestion Management, V S	/1-24/
Congestion Management, ForeSi g	1-24
ELMI	1-25
Cell Relay Networking, ATM and FastPacke	1-25
Asynchronous Transfer Mode (ATM)	1-26
PVC	1-27
SVCs	1-27
UNI/NNI Interface	1-27
Cisco WAN Switching FastPackets	1-27
ATM Product Family Overview	1-28
StrataSphere, Standards-Based Network Managem e	1-28
System Switch, ESP, Edge Concentrator, and Network Access Products Descri p t	1-29
The Cisco BPX 8600 Series Broadband Swi t	1-29
The E S	1-30
The Cisco MGX 8220 Edge Conce n t	1-31
The Cisco IGX 8400 Series Multiband Swi t	1-31
The Cisco IPX Narrowband Swit c	1-32
Network Access Product	1-33
FastPAD	1-33
Cisco 3810 Seri e	1-33
INS-VNS and INS-DA	1-33
System Software Description	1-34
Connections and Connection Rout i	1-35
Connection Routing Group	1-36
Network Synchronization	1-36
Network Availability	1-37
System Diagnostic	1-37
Alarm Reportin	1-37
Statistical Alarms	1-38
Failure Recovery	1-38
Standard	1-38

CHAPTER 2**BPX, IGX, and IPX Architectu r** 2-1

BPX Service Node (Broadband A T	2-1
General	2-1
Extended Services Processo	2-3

- Cisco MGX 8220 Shel 2-3
- BPX Architecture 2-3
- IGX (Multi-band ATM Switch) 2-4
 - General 2-4
 - IGX Architecture 2-5
- IPX (Narrowband ATM and FastPacket Swit c 2-5
 - General 2-5
 - IPX Architecture 2-6

CHAPTER 3

Cisco WAN Manager Network Management 3-1

- Cisco WAN Manage 3-1
 - Network Topolog 3-2
 - Network Performanc 3-2
 - Equipment Managem e 3-2
 - Connection Managem e 3-2
 - Alarm Reporting/Event L o 3-3
 - Software Updates 3-3

CHAPTER 4

Network Services Overvie 4-1

- Broadband ATM Trunk 4-1
- Narrowband FastPacket Tru n 4-2
- ATM Connections 4-3
- ATM Switched Virtual Circuits and the E S 4-3
- Intelligent Network Server DAS and VN 4-4
 - INS-DAS Dial-Up Frame Relay 4-5
 - Dial-Up Connectio 4-6
 - Dial-Backup Connection 4-7
 - INS-VNS 4-8
- Frame Relay Service 4-9
 - Basic Frame Relay Servic 4-11
 - Frame Relay SVCs 4-11
 - Frame Relay T1/E1 Por t 4-11
 - User-to-Network Interface (UNI) 4-12
 - Network-to-Network Interface (NNI) 4-12
 - Bundled Connecti o 4-13
 - Frame Forwarding 4-13

Frame Relay Congestion Avoidance	4-13
Credit Management	4-13
Congestion Notification (FECN/BECN)	4-14
Point-to-Point Data Connection	4-14
Synchronous Data Connection	4-14
DSOA Data Connection	4-15
Nx56K, Nx64K Connection	4-15
Data Frame Multiplexing (DFM)	4-15
Data Clocking Options	4-16
Synchronous Data Control Lead Option	4-16
Voice Connection	4-16
Voice SVCs	4-17
PCM Voice Connection	4-17
ADPCM Compression for Voice Connection	4-17
Voice Activity Detection Feature	4-17
Adaptive Voice	4-18
Instafax and Enhanced Instafax	4-18
Bandwidth Control Feature	4-18
Courtesy Downlink	4-18
Bandwidth Reservation	4-19
Graceful System Upgrade	4-19
Software Downloading	4-19
Firmware Download	4-20
System Software Restore	4-20
Network Synchronization	4-20
International Bridging Trunk	4-21
Reliability Feature	4-21
Distributed Network Intelligence	4-21
Redundant Node Powering	4-21
Redundant Card Set	4-22
Inter-Node Controller Communication	4-22
Configuration Save/Restore	4-22
Alarm Summary Relay Output	4-23
Network Management	4-23
NMS Interfaces	4-24

Node Hardware Administration 4-24
 Network Trunk Administration 4-24
 Embedded SNMP Agent 4-24
 Multiple User Access to Network 4-25
 Integration of FastPAD into Network 4-25

CHAPTER 5

Networking Architecture 5-1

Broadband ATM Network 5-1
 Multi-Media Network 5-2
 High-Capacity Network Backbone 5-2
 Growth Path for Existing Cisco IPX Narrowband/Cisco IGX 8400 Series Multiband Network 5-3
 Service Nodes in Public Network 5-5
 Expansion of Wide-Area Data Network 5-6
 T2 ATM Network Connection 5-8
 Cisco WAN Switching ATM Network Architectures 5-8
 ATM Service Interfaces, BXM, ASI, and Cisco MGX 8220 Edge Concentrator Shell 5-10
 BXM and ASI Interfaces 5-10
 Cisco MGX 8220 Edge Concentrator Shell 5-10
 Narrowband FastPacket Network 5-11
 T1 Networks 5-11
 E1 Networks 5-13
 Networks with an Eight Slot Cisco IPX Narrowband Switch 5-14
 Japanese J1/Y1 Network 5-15
 Frame Relay Networks 5-16
 Frame Relay Port Interface 5-17
 Port Concentrator 5-17
 Frame Relay Network Interfaces, Cisco IPX Narrowband Switch, Cisco IGX 8400 Series Multiband Switch, FastPAD, Cisco MGX 8220 Edge Concentrator 5-17
 User to Network Interface (UNI) 5-18
 Network to Network Interface (NNI) 5-19
 Frame Relay Addressing 5-19
 Networks with the FastPAD 5-20
 Simple FastPAD Network, Voice, Data, Frame Relay 5-20
 FastPAD Access to IGX Multiband or IPX Narrowband Frame Relay Network 5-21
 International Network 5-22
 Example of an International Network 5-22

E1/T1 Conversion	5-23
Setting International Node Parameters	5-24
Network Synchronization	5-24
Clock Sources	5-24
Clock Source Selection	5-25
Defining Clocks and Lines	5-25
Plesiochronous Network	5-26

CHAPTER 6**Tiered Network 6-1**

Routing Hubs and Interface Shelves	6-1
Cisco BPX 8600 Series Broadband Switches and Cisco IGX 8400 Series Multiband Switches Configured as Routing Hubs	6-1
Cisco IGX 8400 Series Multiband Routing Hubs in a Tiered Network	6-3
Tiered Network Implementation	6-4
General	6-5
Definitions	6-6
Upgrades	6-6
Co-locating Routing Hubs and Shelves	6-7
Network Management	6-7
Preferred Routing	6-7
Local and Remote Loopback	6-7
Testcon and Test	6-8
IGX Interface Shelf Description	6-8
Configuration and Management	6-8
Shelf Management	6-8
Alarm Management of Interface Shelf on the IGX Hub Node	6-9
Alarm Management on the IGX Interface Shelf	6-9
Port Management	6-9
Connection Management	6-9
Bandwidth Management	6-9
Bandwidth Efficiency	6-9
Statistics	6-10
User Interface Command	6-10
Shelf	6-10
Data Connection Command	6-10
Data Channel Command	6-10

- Voice Connection Command **6-10**
- Voice Channel Command **6-10**
- Cisco BPX 8600 Series Broadband Routing Hubs in a Tiered Network **6-11**
 - Tiered Network Implementation **6-12**
 - General **6-13**
 - Definitions **6-13**
 - Upgrades **6-13**
 - Co-locating Routing Hubs and Interface Shelves **6-14**
 - Network Management **6-14**
 - ForeSight **6-15**
 - Preferred Routing **6-15**
 - Local and Remote Loopback **6-15**
 - Tstcon and Tstde I **6-15**
 - IPX Interface Shelf Description **6-15**
 - Configuration and Management **6-16**
 - Interface Shelf Management **6-16**
 - Alarm Management of Interface Shelf on the Cisco BPX 8600 Series Broadband Switch Configured as a Hub Node **6-16**
 - Alarm Management on the IPX Interface Shelf **6-17**
 - Port Management **6-17**
 - Connection Management **6-17**
 - Bandwidth Management **6-17**
 - Statistics **6-17**
- Cisco WAN Manager NM **6-17**

CHAPTER 7

Network Maintenance 7-1

- Automatic Alarm Reporting **7-1**
- Network Troubleshooting **7-2**
 - Displaying Network Alarm **7-3**
 - Displaying Trunk and Line Status **7-5**
 - Displaying Node Status **7-6**
 - Displaying Connection Status **7-8**
 - Maintenance (Events) Log **7-9**
 - Display of Network Trunk Load **7-9**
- System Troubleshooting Tools **7-10**
 - User-initiated Test **7-10**

Loopback Tests	7-11
Connection Test	7-13
External Device Window	7-14
Network Statistics	7-14

CHAPTER 8**ATM and Broadband Trunk**

Asynchronous Transfer Mode	8-1
ATM Model	8-1
Physical Layer	8-2
ATM Layer	8-3
ATM Cell Headers	8-3
ATM Cell Addressing	8-6
ATM Adaptation Layer	8-8
PVCs vs. SVC	8-9
Cisco IGX 8400 Series Multiband and Cisco IPX Narrowband Trunk Interfaces	8-9
BAM	8-10
SAM	8-10
CAM	8-10
FastPacket Adaptation to ATM	8-11
Simple Gateway	8-11
ATM Cell Switching	8-12
Broadband (ATM) Trunk Format	8-15
DS3 PLCP Frame Structure	8-15
G.804 E3 Frame Structure	8-16
Virtual Trunks	8-17
Virtual Trunk Capacities	8-17
Virtual Trunk Traffic Classes	8-18
Virtual Trunk Addressing	8-19
Virtual Trunk Example	8-20
Full Mesh with Virtual Trunk	8-21
BNI One Stage Queueing	8-22
Virtual Trunk Statistics	8-22
BNI Virtual Trunk	8-22
Card Redundancy (Y-Redundancy)	8-23
Virtual Trunk Alarms	8-23
Trunk Specific Alarms	8-23

- BNI Virtual Trunk Alarm 8-23
- Trunk Port Alarms 8-23
- Feeder Trunk Support 8-23
- Connection Management 8-23
- Routing VPCs over Virtual Trunk 8-23
- Structured Networks Support 8-24
- Error Messages 8-24
- Command 8-24

CHAPTER 9

FastPackets and Narrowband Trunk 9-1

- FastPacket Format 9-1
 - High Priority Packet 9-2
 - PCM and ADPCM Voice Packet 9-3
 - Non-Timestamped Data Packet 9-3
 - Timestamped Data Packet 9-5
 - Frame Relay Data Packets 9-6
 - Idle Packet 9-6
 - Remote Alarm Packet 9-6
- Narrowband Trunk Format 9-7
 - Lines and Trunk 9-7
 - T1 Packet Trunk 9-8
 - Fractional T1 Packet Trunk 9-9
 - E1 Packet Trunk 9-10
 - Subrate Packet Trunk 9-12
 - Packet Framing (NTM and NTC) 9-13

CHAPTER 10

ATM Connections 10-

- ATM Connection Services 10-
- SVCs 10-2
- Traffic Management Overview 10-2
 - Standard ABR notes 10-3
 - VSVD Description 10-4
 - BXM Connections 10-4
 - ForeSight Congestion Control 10-5
- ATM Connection Requirements 10-
- Connection Routines 10-6

Addcon Command Synta	10-6
ATM Connection Configuration	10-7
CBR Connectio n	10-1
VBR and ATFR Connection	10-1
VBR Connectio n	10-13
ATFR Connections	10-1
ABR Notes	10-1
ABR and ATFST Connections	10-1
ABR Connectio n	10-16
ATFST Connections	10-1
UBR Connection	10-20
Traffic Policing Exampl e	10-2
Dual-Leaky Bucket (An Analogy)	10-21
CBR Traffic Policing Examp l	10-2
VBR Dual-Leaky Bucket Policing Examp l	10-2
Leaky Bucket	10-2
Leaky Bucket	10-2
Examples	10-2
ABR Connection Polici n	10-31
UBR Connection Polici n	10-31
Leaky Bucket	10-3
Leaky Bucket	10-3
LMI and ILMI Parameter	10-3

CHAPTER 11

SVCs, ATM and Frame Rela	11-1
ATM and Frame Relay SVC	11-1
PVCs and S V	10-2
PVC	11-2
SVCs	11-2
BPX Service Node Interface	11-3
Interim Inter-switch Protocol Routin	11-
PNN	11-3
Signaling Plan	11-
UNI Signaling Chan n	10-
NNI Signaling Cha n	11- e
Network Interworking Between Frame Relay and ATM	11-6

- Extended Services Processors **11-6**
 - ESP Interfaces **11-7**
 - Redundant ESP **11-8**
- Network Management **11-8**
- Resource Partitioning **11-8**

CHAPTER 12

Frame Relay Connection 12-1

- Cisco IPX Narrowband Switch, Cisco IGX 8400 Series Multiband Switch, and Cisco MGX 8220 Edge Concentrator **12-1**
 - Port Concentrator Shelf **12-1**
- Frame Relay Addressing **12-1**
- Frame Relay Signal Flow **12-6**
- Connection Types **12-7**
 - Normal Frame Relay Connections **12-7**
 - Bundled FR Connections **12-7**
 - Frame Forwarding Connections **12-9**
 - T1/E1 Frame Relay **12-1**
- Network Interfaces **12-1**
 - User to Network Interface (UNI) **12-11**
 - Network-Network Interfaces **12-1**
- Congestion Prevention and Notification **12-1**
 - Credit Manager **12-15**
 - Explicit Congestion Notification **12-1**
 - Source ECN **12-19**
 - Destination ECN **12-1**
 - ForeSight **12-1**
 - ForeSight Across Multiple Networks **12-2**
 - Discard Eligibility **12-22**
 - Cell Loss Priority **12-2**
 - Circuit Priority **12-2**
- Connection Parameters **12-2**
 - Configuring Connection Bandwidth **12-2**
 - Setting MIR, CIR, and PIR for Foresight Connections **12-2n**
 - Setting CIR **12-2**
 - Setting VC Q Delay **12-2**
 - Setting Cmax **12-28**

Setting ECN Queue Depth	12-28
Frame Relay Port Parameter	12-2
Setting Port Speed	12-3
Setting Port Queue Parameter	12-3
Port Concentrator Shelf Frame Relay Connections	12-3
Setting up Ports and Connections	12-3n
Port Commands	12-3
Connection Command	12-31
Maximum Number of Frame Relay Connections	12-3
Cisco WAN Manager Connections Management	12-3
Connection Command Sequence	12-32
Frame Forwarding Connection	12-3
Configuring Connection Bandwidth	12-3
Optimizing Traffic Routing and Bandwidth	12-3t
PCS Network Function	12-3
Frame Relay Interworking	12-3
Frame Forwarding	12-3
Maximum Throughput	12-3
Maximum Number of Connections	12-3
Delay	12-3
Frame Processing by the	12-3c
Frame Relay Format	12-3
Frame Queuing with the Port Concentrator	12-3
Signalling Protocol	12-3
ForeSight	12-3
Frame Forwarding	12-3
Bandwidth and Routing	12-3

CHAPTER 13**Frame Relay to ATM Network and Service Interworking 13-1**

Interworking	13-1
Service Interworking	13-
Networking Interworking	13-5
ATM Protocol Stack	13-8
AIT/BTM Interworking and the ATM Protocol Stack	13-1
AIT/BTM Control Mapping, Frames and Cells	13-12
Management, OAM Cells	13-11

- Functional Descriptio **13-1**
 - ATF Summary **13-1**
 - Feature **13-1**
 - Limitations **13-1**
 - Some ATF Connection Criteri **13-14**
 - Connection Managem e **13-1**
 - Port Managemen **13-1**
 - Structur **13-1**
 - Channel Statistic **13-1**
 - OAM Cell Supp o **13-16**
 - Diagnostic **13-17**
 - User Comman d **13-1**
 - Virtual Circuit Feature **13-1**
 - User Comman d **13-1**
 - User Comman d **13-1**
- Managemen **13-1**
 - Connection Managem e **13-1**
 - Routing **13-1**
 - Bandwidth Managemen **13-19**
 - User Interface **13-1**
 - Port Managemen **13-2**
 - Connection Managem e **13-2**
 - Signaling **13-2**
 - Alarms **13-2**

CHAPTER 14

Synchronous Data Conne c t i o n

- Data Connection Typ e **14-1**
 - Synchronous Data Connections (SD P **14-1**
 - Low-speed Data Connections (LD P **14-**
 - Channelized Data Connecti o **14-2**
 - Interface to DDS Networ **14-4**
- Data Block Diagram Signal Flo **14-5**
 - CDP Data Signal Flow **14-5**
 - SDP/LDP Synchronous Data Channel Signal Flo **14-6**
- Data Control Leads **14-8**
 - Non-Interleaved Control Le a **14-10**

Interleaved (Fast EIA) Control Lead	14-11
Partially Interleaved (Embedded) Control Lead	14-1
Control Lead Condition	14-1
Data Compression	14-1
Data Clocking	14-13
Synchronous Data Clock Source	14-13
Data Clock Configurations	14-1

CHAPTER 15**Voice Connections 15-1**

Voice Circuit Type	15-1
Voice Circuit Signal Flow	15-2
Voice Signal Flow with CDP Car	15-3
Voice Channel Signaling	15-5
Signaling Bit Flow with CD	15-6
Voice Processing Feature	15-
ADPCM Voice Compression	15-
Speech Detection	15-8
Modem Detection	15-8
Echo Cancellation	15-9
Voice and Signaling Condition	15-1
Level Adjustment	15-1

CHAPTER 16**Connection Management 16-1**

Packet Queuing	16-1
Delay in a Cell Network	16-2
Delay in Packet Frame Relay Networks	16-
Delay at the Source	16-3
Delay at the Sink	16-
Synchronous Data Connection Delay	16-4
Ways to Reduce Data Connection Delay	16-5
Voice Connection Delay	16-6
Maximum Hops—Voice Connection	16-
Routing and Rerouting	16-7
Load Mode	16-7
Load Model and Routing for Frame Relay	16-
Routing Algorithm	16-8

- Causes of Rerouting **16-9**
- Reroute Priority and Order **16-1**
- Courtesy Downin **16-1**
- System Message Traffic Routi **16-1**
- Bandwidth Allocatio **16-1**
 - Network Trunk Bandwidth **16-12**
 - Voice Compression Bandwidth Requireme n **16-13**
 - VAD and DFM Effect **16-1**
 - Data Channel Packet Generation Rate **16-1**
- Traffic Statistic **16-1**

APPENDIX A

Network Specifications A-1

- Broadband Trunk Interfa c ~~A-1~~
- Narrowband Trunk Interfac e **A-2**
- Narrowband Channel Interfa c ~~A-2~~
- ATM Trunk Interface (BXM-T3/E3 Cards) **A-3**
- ATM Trunk Interface (BXM-155 Cards) **A-3**
- ATM Trunk Interface (BXM-622 Cards) **A-5**
- ATM T3 Trunk Interface (BNI-T3, LM-3T3) **A-5**
- ATM E3 Trunk Interface (BNI-E3, LM-3E3) **A-7**
- ATM OC3 Trunk Interface (BNI-OC3, LM-OC3) **A-8**
- ATM Service Interface (BXM-T3/E3 Cards) **A-9**
- ATM Service Interface (BXM-155 Cards) **A-9**
- ATM Service Interface (BXM-622 Cards) **A-10**
- ATM Service Interface (ASI-1, LM-2T3) **A-10**
- ATM Service Interface (ASI-1, LM-2E3) **A-11**
- ATM Service Interface (ASI-2, LM-OC3) **A-11**
- Network Synchronization **A-12**
- Network Managemen **A-12**

<i>Figure 1-1</i>	An ATM Network Configuration	1-2
<i>Figure 1-2</i>	Example of an Enterprise Network Application	1-10
<i>Figure 1-3</i>	Example of a Service Provider Application	1-12
<i>Figure 1-4</i>	Frame Relay to ATM Network Interworking	1-14
<i>Figure 1-5</i>	Frame Relay to ATM Service Interworking	1-15
<i>Figure 1-6</i>	Tiered Network with Cisco BPX 8600 Series Broadband Switches and Cisco IGX 8400 Series Multiband Switches Configured as Routing Hubs	1-18
<i>Figure 1-7</i>	Virtual Trunking Example	1-20
<i>Figure 1-8</i>	Cisco BPX 8600 Series Broadband Switch Configuration	1-30
<i>Figure 1-9</i>	Cisco BPX 8600 Series Broadband Switch with Co-Located ESP and Cisco MGX 8220 Edge Concentrator	1-31
<i>Figure 1-10</i>	A Cisco IGX 8400 Series Multiband Switch Configuration	1-32
<i>Figure 1-11</i>	A Cisco IPX Narrowband Switch Configuration	1-33
<i>Figure 4-1</i>	ATM Switched Virtual Circuits	4-4
<i>Figure 4-2</i>	INS Dial-Up Application	4-6
<i>Figure 4-3</i>	INS Dial-Backup Application	4-7
<i>Figure 4-4</i>	Basic VNS SVC Call	4-8
<i>Figure 4-5</i>	VNS Areas	4-9
<i>Figure 5-1</i>	Multi-media Network Application	5-2
<i>Figure 5-2</i>	High-Capacity Backbone Configuration	5-3
<i>Figure 5-3</i>	Upgrading Existing Cisco IGX 8400 Series Multiband/IPX Networks	5-4
<i>Figure 5-4</i>	Service Node for Large C.O. Switches	5-6
<i>Figure 5-5</i>	ATM Network of the Future	5-7
<i>Figure 5-6</i>	T2 ATM Trunking	5-8
<i>Figure 5-7</i>	Cisco WAN Switching ATM Network Architectures	5-9
<i>Figure 5-8</i>	Direct ATM User-to-Network Connections	5-11
<i>Figure 5-9</i>	Typical T1 IPX Network	5-13
<i>Figure 5-10</i>	E1 IPX Network	5-14
<i>Figure 5-11</i>	Typical IPX Network with IPX 8 Remote Nodes	5-15
<i>Figure 5-12</i>	IPX 8 Used as a Feeder in Large Networks	5-15
<i>Figure 5-13</i>	Frame Relay PVCs Provide Full Mesh Connectivity	5-16
<i>Figure 5-14</i>	Frame Relay Network Interfaces	5-18
<i>Figure 5-15</i>	Frame Relay Network Connection	5-20

<i>Figure 5-16</i>	Cisco IPX Narrowband Network with FastPAD Access	5-21
<i>Figure 5-17</i>	IPX International Network	5-22
<i>Figure 5-18</i>	Example of IPX Network Clock Architecture	5-27
<i>Figure 5-19</i>	Example of a Pleisiochronous Network	5-27
<i>Figure 6-1</i>	Tiered Network with Cisco BPX 8600 Series Broadband and Cisco IGX 8400 Series Multiband Routing Hubs	6-2
<i>Figure 6-2</i>	Cisco IGX 8400 Series Multiband Shelves and Routing Hubs, Voice and Data Connections	6-3
<i>Figure 6-3</i>	Cisco IGX 8400 Series Multiband Shelves and Routing Hubs, Frame Relay Connections	6-4
<i>Figure 6-4</i>	Tiered Network with Cisco BPX 8600 Series Broadband Routing Hubs	6-12
<i>Figure 6-5</i>	Cisco WAN Manager Connection Manager	6-18
<i>Figure 7-1</i>	Automatic Alarm Reporting	7-2
<i>Figure 7-2</i>	Example of a Cisco WAN Manager Topology Map	7-3
<i>Figure 7-3</i>	Network Loopback Paths	7-13
<i>Figure 8-1</i>	B-ISDN Model	8-2
<i>Figure 8-2</i>	ATM Cell Format	8-3
<i>Figure 8-3</i>	UNI Header	8-4
<i>Figure 8-4</i>	NNI Header	8-4
<i>Figure 8-5</i>	STI Header	8-5
<i>Figure 8-6</i>	Virtual Paths and Virtual Channels	8-6
<i>Figure 8-7</i>	VP-only Switching	8-7
<i>Figure 8-8</i>	VP and VC Switching	8-7
<i>Figure 8-9</i>	Example of Adaptation Process	8-9
<i>Figure 8-10</i>	BAM, CAM, and SAM Configurations	8-11
<i>Figure 8-11</i>	Simple and Complex Gateway Formats	8-12
<i>Figure 8-12</i>	Operation of a Typical Crosspoint Switch Matrix	8-14
<i>Figure 8-13</i>	DS3 PLCP Frame Format	8-15
<i>Figure 8-14</i>	G.804 E3 Frame Format	8-16
<i>Figure 8-15</i>	Connection Identifier	8-20
<i>Figure 8-16</i>	Single Virtual Trunk Addition	8-21
<i>Figure 8-17</i>	Four Node Example of Virtual Trunking	8-22
<i>Figure 9-1</i>	General FastPacket Format	9-1
<i>Figure 9-2</i>	High Priority Packet Format	9-2
<i>Figure 9-3</i>	Voice Packet Format	9-3
<i>Figure 9-4</i>	Non-Timestamped Data Packet Format	9-3
<i>Figure 9-5</i>	Various Non-Timestamped Data Message Formats	9-4
<i>Figure 9-6</i>	Timestamped Data Packet Format	9-5

<i>Figure 9-7</i>	Frame Relay Data Packet Format	9-6
<i>Figure 9-8</i>	Lines and Trunks (sometimes called Circuit Lines and Packet Lines)	9-7
<i>Figure 9-9</i>	T1 Framing Requirement	9-8
<i>Figure 9-10</i>	T1 Frame Format	9-9
<i>Figure 9-11</i>	Fractional T1 Frame Format	9-10
<i>Figure 9-12</i>	E1 Frame Format	9-11
<i>Figure 9-13</i>	FastPackets with Various E1 Formats	9-12
<i>Figure 9-14</i>	Typical Subrate Trunk Setup	9-12
<i>Figure 10-1</i>	ATM Connections over a Cisco BPX 8600 Series Broadband Network	10-
<i>Figure 10-2</i>	ABR VSVD Flow Control Diagram	10-
<i>Figure 10-3</i>	CBR Connection Prompt Sequence	10-1
<i>Figure 10-4</i>	VBR Connection Prompt Sequence	10-1
<i>Figure 10-5</i>	ATFR Connection Prompt Sequence	10-1
<i>Figure 10-6</i>	ABR Standard Connection Prompt Sequence	10-1
<i>Figure 10-7</i>	Meaning of VSVD and Flow Control External Segments	10-1
<i>Figure 10-8</i>	ABR ForeSight Connection Prompt Sequence	10-1
<i>Figure 10-9</i>	ATFST Connection Prompt Sequence	10-2
<i>Figure 10-10</i>	UBR Connection Prompt Sequence	10-2
<i>Figure 10-11</i>	CBR Connection, UPC Overview	10-2
<i>Figure 10-12</i>	CBR.1 Connection with Bucket Compliant	10-24
<i>Figure 10-13</i>	CBR.1 Connection, with Bucket Discarding non-Compliant Cells	10-24
<i>Figure 10-14</i>	VBR Connection, UPC Overview	10-2
<i>Figure 10-15</i>	VBR Connection, Policing = 4, Leaky Bucket 1 Compliant	10-2
<i>Figure 10-16</i>	VBR Connection, Policing = 4, Leaky Bucket 1 Non-Compliant	10-2
<i>Figure 10-17</i>	VBR.2 Connection, Policing = 2, with Buckets 1 and 2 Compliant	10-28
<i>Figure 10-18</i>	VBR.2 Connection, Leaky Bucket 2 Discarding CLP (0) Cells	10-2
<i>Figure 10-19</i>	VBR.1 Connection, Policing = 1, with Buckets 1 and 2 Compliant	10-30
<i>Figure 10-20</i>	VBR.3 Connection, Policing = 3, with Bucket 2 non-compliant	10-3
<i>Figure 10-21</i>	UBR Connection, UPC Overview	10-33
<i>Figure 11-1</i>	Wide Area Network with BPX Service Nodes	11-2
<i>Figure 11-2</i>	BPX Service Node Network Signaling Plane	11-4
<i>Figure 11-3</i>	UNI Signaling Channels	11-5
<i>Figure 11-4</i>	ESP Signaling PVC	11-6
<i>Figure 11-5</i>	ESP Physical Interfaces	11-7
<i>Figure 12-1</i>	Example of Carrying Frame Relay Data with FastPackets	12-2

<i>Figure 12-2</i>	FastPackets for Frame Relay Data	12-4
<i>Figure 12-3</i>	Example of Frame Relay Addressing	12-5
<i>Figure 12-4</i>	Example of Frame Relay Global Addressing	12-6
<i>Figure 12-5</i>	Frame Relay Signal Flow	12-7
<i>Figure 12-6</i>	Basic Frame Relay Service	12-8
<i>Figure 12-7</i>	Bundled Connections	12-
<i>Figure 12-8</i>	Multiple and Single DSOs Forming a Logical Port	12-1
<i>Figure 12-9</i>	Frame Relay User-to-Network Interface	12-1
<i>Figure 12-10</i>	Example of NNI in a Multi-Network PVC	12-1
<i>Figure 12-11</i>	Cisco IPX Narrowband Switch Credit Manager Operation	12-1
<i>Figure 12-12</i>	Bandwidth and Credits	12-17
<i>Figure 12-13</i>	Explicit Congestion Notification Example	12-1
<i>Figure 12-14</i>	ForeSight Operation	12-2
<i>Figure 12-15</i>	ForeSight Operation Across Multiple Networks	12-2
<i>Figure 12-16</i>	PCS Frame Relay Access to an Cisco IPX Narrowband/Cisco IGX 8400 Series Multiband/Cisco BPX 8600 Series Broadband Network	12-3
<i>Figure 12-17</i>	Breakdown of PCS Frame Relay Data into FastPackets	12-35
<i>Figure 12-18</i>	PVC Ingress and Egress Queues	12-3
<i>Figure 13-1</i>	Frame Relay to ATM Network Interworking	13-3
<i>Figure 13-2</i>	Frame Relay to ATM Service Interworking	13-3
<i>Figure 13-3</i>	Frame Relay to ATM Interworking Examples with AIT Card on a Cisco IPX Narrowband Switch	13-4
<i>Figure 13-4</i>	Frame Relay to ATM Service Interworking Detail	13-5
<i>Figure 13-5</i>	Frame Relay to ATM NW Interworking Detail	13-6
<i>Figure 13-6</i>	ATF Connections, Simplified Example	13-
<i>Figure 13-7</i>	ATM Layers	13-9
<i>Figure 13-8</i>	Protocol Stack Operation	13-11
<i>Figure 14-1</i>	Examples of CDP Data Connections	14-
<i>Figure 14-2</i>	Typical DDS Network Configurations	14-
<i>Figure 14-3</i>	Data Circuit Flow with CDP	14-6
<i>Figure 14-4</i>	Synchronous Data Flow	14-7
<i>Figure 14-5</i>	Various non-Time Stamped Data Formats	14-9
<i>Figure 14-6</i>	Data Clocking Configurations	14-1
<i>Figure 15-1</i>	Generalized Cisco IPX Narrowband Switch/Cisco IGX 8400 Series Multiband Switch Voice Signal Flow	15-3
<i>Figure 15-2</i>	Voice Block Signal Flow with CDP	15-4
<i>Figure 15-3</i>	CDP-to-CDP Signalling	15-
<i>Figure 15-4</i>	Echo Effects Diagram	15-9



<i>Tabl e7-1</i>	System Troubleshooting Commands Available	7-11
<i>Tabl e7-2</i>	System Loopback Test	7-11
<i>Tabl e7-3</i>	Typical Statistics Collected	7-15
<i>Tabl e8-1</i>	SONET Data Rates	8-3
<i>Tabl e8-2</i>	STI Congestion Control Bits	8-6
<i>Tabl e8-3</i>	Classes of Traffic and Associated AAL Layers	8-8
<i>Tabl e8-4</i>	ATM Cell Addressing Modes	8-10
<i>Tabl e9-1</i>	Packet Types	9-2
<i>Tabl e10-1</i>	Standard ATM Traffic Classe	10-2
<i>Tabl e10-2</i>	Traffic Policing Definitio n	10-8
<i>Tabl e10-3</i>	Connection Parameters with Default Settings and Ran g e	10-9
<i>Tabl e10-4</i>	Connection Parameter Descript i o n	10-1
<i>Tabl e10-5</i>	CBR Policing Definitions	10-1
<i>Tabl e10-6</i>	VBR Policing Definitions	10-1
<i>Tabl e10-7</i>	jUBR Policing Definition	10-2
<i>Tabl e10-8</i>	LMI Parameters	10-3
<i>Tabl e10-9</i>	LMI Parameters	10-3
<i>Tabl e12-1</i>	StrataCom vs. Standard Frame Relay Parameters	12-2
<i>Tabl e12-2</i>	LMI Protocols Supporte	12-2
<i>Tabl e12-3</i>	LMI Protocol Parameters	12-30
<i>Tabl e14-1</i>	Data Control Lead Equivalenc	14-1
<i>Tabl e14-2</i>	LDP Embedded EIA Mode Features	14-1
<i>Tabl e15-1</i>	CDP/CVM Circuit Types	15-
<i>Tabl e15-2</i>	Loss Objectives for Common Carrier Trunks	15-1
<i>Tabl e15-3</i>	Quantization Distortion Pertaining to the Cisco IPX Narrowband Switch/Cisco IGX 8400 Series Multiband Switch	15-1
<i>Tabl e16-1</i>	Credit Accumulation	16-
<i>Tabl e16-2</i>	Calculation of Non-Timestamped Data Packet Overall Delay	16-5
<i>Tabl e16-3</i>	Calculation of Timestamped Data Packet Overall Delay	16-6
<i>Tabl e16-4</i>	Sources of Delay in Voice Connections	16-
<i>Tabl e16-5</i>	Typical Voice Connection Delays	16-7
<i>Tabl e16-6</i>	Priority for Rerouting	16-1

<i>Table 16-7</i>	Subrate Packet Line Bandwidth	16-1
<i>Table 16-8</i>	Cisco IPX Narrowband Switch Voice Grade	16-1
<i>Table 16-9</i>	Data Load Table with Standard EIA and No DFM	16-1
<i>Table 16-10</i>	Data Load Table with Standard EIA and DFM	16-17
<i>Table 16-11</i>	Data Load Table with Partially-Filled Packet and No DFM	16-1
<i>Table 16-12</i>	Data Load Table with Partially-Filled Packet and DFM	16-18
<i>Table 16-13</i>	Data Load Table with Fast EIA	16-1
<i>Table 16-14</i>	Data Load Table—Fast EIA with Partially-Filled Packet	16-20



Preface

This publication provides an overview of the operation of Cisco BPX 8600 series broadband switches, Cisco IGX 8400 series multiband switches, Cisco IPX narrowband switches, Cisco MGX 8220 edge concentrators, and associated ancillary Cisco WAN switching equipment.

Cisco documentation and additional literature are available in a CD-ROM package, which ships with your product. The Documentation CD-ROM, a member of the Cisco Connection Family, is updated monthly. Therefore, it might be more up to date than printed documentation. To order additional copies of the Documentation CD-ROM, contact your local sales representative or call customer service. The CD-ROM package is available as a single package or as an annual subscription. You can also access Cisco documentation on the World Wide Web at <http://www.cisco.com>, <http://www-china.cisco.com>, or <http://www-europe.cisco.com>.

Objectives

This publication is intended to provide system level descriptions of the operation of Cisco BPX 8600 series broadband switches, Cisco IGX 8400 series multiband switches, Cisco IPX narrowband switches, Cisco MGX 8220 edge concentrators, and associated ancillary Cisco WAN switching equipment. A high level overview is provided in Part 1 of the publication for those interested in a general knowledge of the equipment. Additional descriptions are provided in more detail in succeeding parts of the publication for those with a background and interest in specific features, including broad and narrowband trunks, ATM and Frame Relay traffic, and synchronous data and voice traffic.

Audience

This publication is intended for operators, network designers, system administrators, and others interested in an overview of the features, capabilities, and applications of the Cisco BPX 8600 series broadband switches, Cisco IGX 8400 series multiband switches, Cisco IPX narrowband switches, Cisco MGX 8220 edge concentrators, and associated ancillary Cisco WAN switching equipment in supporting multi-media traffic over a wide range of network configurations.

Organization

This publication is organized as follows:

Part 1, OVERVIEW

- Chapter 1 Cisco Wide Area ATM Networks** Provides an introduction to the Strategic Wide Area ATM networking products including an overview of the operation of Cisco BPX 8600 series broadband switches, Cisco IGX 8400 series multiband switches, Cisco IPX narrowband switches, Cisco MGX 8220 edge concentrators, and associated ancillary Cisco WAN switching equipment, including ESP, DAS, VNS, and StrataSphere NMS
- Chapter 2 BPX, IGX, and IPX Architectur** Describes the switches used to build cell-based wide area networks: the broadband BPX 8600 series, the multiband IGX 8400 series, the Cisco MGX 8220 edge concentrator, and the narrowband IPX.
- Chapter 3 Cisco WAN Manager Network Management** Describes the StrataView Plus NMS workstation advanced system applications and tools that provide integrated fault, performance, and configuration management functions unique to cell-based networks.
- Chapter 4 Network Services Overview** Provides an overview of the standard and optional network services (features) such as voice, data, ATM, and frame relay connections supported by networks utilizing broadband Cisco BPX 8600 series switches, multiband Cisco IGX 8400 series switches, Cisco MGX 8220 edge concentrators, and narrowband Cisco IPX switches.

Part 2, NETWORKS

- Chapter 5 Networking Architecture** Describes various types of networks that may be implemented with the Cisco WAN switching cell relay network switches, broadband Cisco BPX 8600 series switches, multiband Cisco IGX 8400 series switches, narrowband Cisco IPX switches, and associated equipment such as the Cisco MGX 8220 edge concentrator and FastPAD
- Chapter 6 Tiered Networks** Describes the broadband Cisco BPX 8600 series tiered networks and multiband Cisco IGX 8400 series tiered networks that provide the capability of adding interface shelves/feeders (non-routing nodes) to a Cisco WAN switching routing network.
- Chapter 7 Network Maintenance** Describes some of the tools provided for detecting and identifying network and/or equipment problems that are available to the network operator.

Part 3, NARROW AND BROADBAND TRUNKS

- Chapter 8** **ATM and Broadband Trunks** Provided for users who wish to have an in-depth knowledge of the ATM and broadband trunks functions. It discusses ATM concepts and the various high-speed digital trunks that are used to carry ATM connections.
- Chapter 9** **FastPackets and Narrowband Trunks** Provided for users who wish to have an in-depth knowledge of FastPackets and the narrowband digital trunks used to carry them.

Part 4, ATM AND FRAME RELAY CONNECTIONS

- Chapter 10** **ATM Connections** Describes how ATM connection services are established over an ATM network by adding ATM and frame relay connections between ATM service interface ports in the network using ATM standard UNI 3.1 and Traffic Management 4.0. It describes BXM and ASI card operation and summarizes ATM connection parameter configuration.
- Chapter 11** **SVCs, ATM and Frame Relay** Provides a summary of switched virtual circuits with respect to the BPX Service Node (a Cisco BPX 8600 series broadband switch with a co-located Extended Services Processor). For additional information, refer to the *Cisco BPX Service Node Extended Processor Installation and Operations* document.
- Chapter 12** **Frame Relay Connections** Provided for users who wish to have an in-depth knowledge of network frame relay connections and related functions and also describes the Port Concentrator Shelf (PCS) which extends the port capacity of an FRP card on a Cisco IPX narrowband switch or of an FRM card on a Cisco IGX 8400 series multiband switch from 4 high-speed ports to 44 low-speed ports.
- Chapter 13** **Frame Relay to ATM Network and Service Interworking** Describes frame relay to ATM interworking which enables frame relay traffic to be connected across high-speed ATM trunks using ATM standard Network and Service Interworking

Part 5, DATA AND VOICE CONNECTIONS

- Chapter 14** **Synchronous Data Connections** Provided for users who wish to have an in-depth knowledge of Cisco IGX 8400 series multiband and Cisco IPX narrowband synchronous data connections and related functions. It describes the basic flow of information through the network for various data connection types. It also describes the data compression features, data clocking, and data channel conditioning available.
- Chapter 15** **Voice Connections** Provided for users who wish to have an in-depth knowledge of Cisco IGX 8400 series multiband and Cisco IPX narrowband voice connections and related functions. This chapter describes voice and signalling flow as well as digital processing topics such as voice compression, speech detection, modem detection, and echo cancelling.

Part 6, CONNECTION MANAGEMENT NOTES

Chapter 16 Connection Management Provided for users who wish to have an in-depth knowledge of Cisco IGX 8400 series multiband and Cisco IPX narrowband connection management functions. It describes packet queuing and the various queue types. It also discusses circuit routing and rerouting, delay for various types of connections, and circuit bandwidth requirements and utilization.

Part 7, REFERENCE

Appendix A Network Specifications Summarizes the network specifications.

Glossary Glossary of CiscoWAN switching-specific terms.

Related Documentation

The following Cisco WAN Switching publications contain additional information related to the operation of the Cisco BPX 8600 series broadband switches, the Cisco IGX 8400 series multiband switches, the Cisco IPX narrowband switches, the *Cisco MGX 8220 edge concentrators*, and associated equipment.:

- *Cisco WAN Manager Installation Guide* providing installation and configuration instructions for the Cisco WAN Manager network management system.
- *Cisco WAN Manager Operations Guide* providing procedures for using the CiscoWAN Manager network management system.
- *StrataSphere Network Design Tools* providing installation and configuration instructions for the StrataSphere Network Design Tools applications and procedures for modeling networks.
- Release 9.1 of the IGX 8400 Series, IPX, and BPX 8600 Series documentation, including:
 - *Cisco BPX 8600 Series Reference* providing a general description and technical details of the BPX 8600 series broadband switch.
 - *Cisco BPX 8600 Series Installation and Configuration* providing installation and configuration instructions for the BPX 8600 series broadband switches.
 - *Cisco IGX 8400 Series Reference* providing a general description and technical details of the IGX 8400 series multiband switches.
 - *Cisco IGX 8400 Series Installation* providing installation instructions for the IGX 8400 series multiband switches.
 - *Cisco WAN Switching Command Reference* providing detailed information on operating the BPX 8600 series broadband switches, IGX 8400 series multiband switches, and IPX narrowband switches through their command line interfaces.
 - *Cisco WAN Switching SuperUser Command Reference* providing detailed information on operating the BPX 8600 series broadband switches, IGX 8400 series multiband switches, and IPX narrowband switches using the command line interface commands requiring SuperUser access authorization.
 - *Cisco IPX Reference* providing a general description and technical details of the IPX narrowband switch.

- *Cisco IPX Installation* providing installation instructions for the IPX narrowband switch.
- Release 4.1 of the *Cisco MGX 8220 Edge Concentrator* documentation, including:
 - *Cisco MGX 8220 Edge Concentrator Installation and Configuration* providing installation and configuration instructions for the Cisco MGX 8220 edge concentrator.
 - *Cisco MGX 8220 Edge Concentrator Command Reference* providing detailed information for Cisco MGX 8220 edge concentrator command line usage.
- Release 2.0 of the *Cisco Extended Services Processor* documentation, including:
 - *Cisco BPX Service Node Extended Services Processor Installation and Operation* providing a general description and installation instructions for the Cisco Extended Services Processor (ESP).

Conventions

This publication uses the following conventions to convey instructions and information.

Command descriptions use these conventions:

- Commands and keywords are in **boldface**.
- Arguments for which you supply values are in *italics*.
- Elements in square brackets ([]) are optional.
- Alternative but required keywords are grouped in braces ({ }) and are separated by vertical bars (|).

Examples use these conventions:

- Terminal sessions and information the system displays are in `screen font`.
- Information you enter is in **boldface screen font**.
- Nonprinting characters, such as passwords, are in angle brackets (<>).
- Default responses to system prompts are in square brackets ([]).



Note

Means *reader take note*. Notes contain helpful suggestions or references to materials not contained in this manual.

PART 1

OVERVIEW

Cisco Wide Area ATM Networks

This chapter provides an introduction to the Cisco Wide Area ATM networking products including an overview of the Cisco IPX narrowband switch, the Cisco IGX 8400 series multiband switch, the Cisco BPX 8600 series wideband switch, ESP, the Cisco MGX 8220 edge concentrator, and CiscoWAN Switching access products, including the Cisco 3810, FastPAD, INS-DAS, INS-VNS, and the StrataSphere NMS.

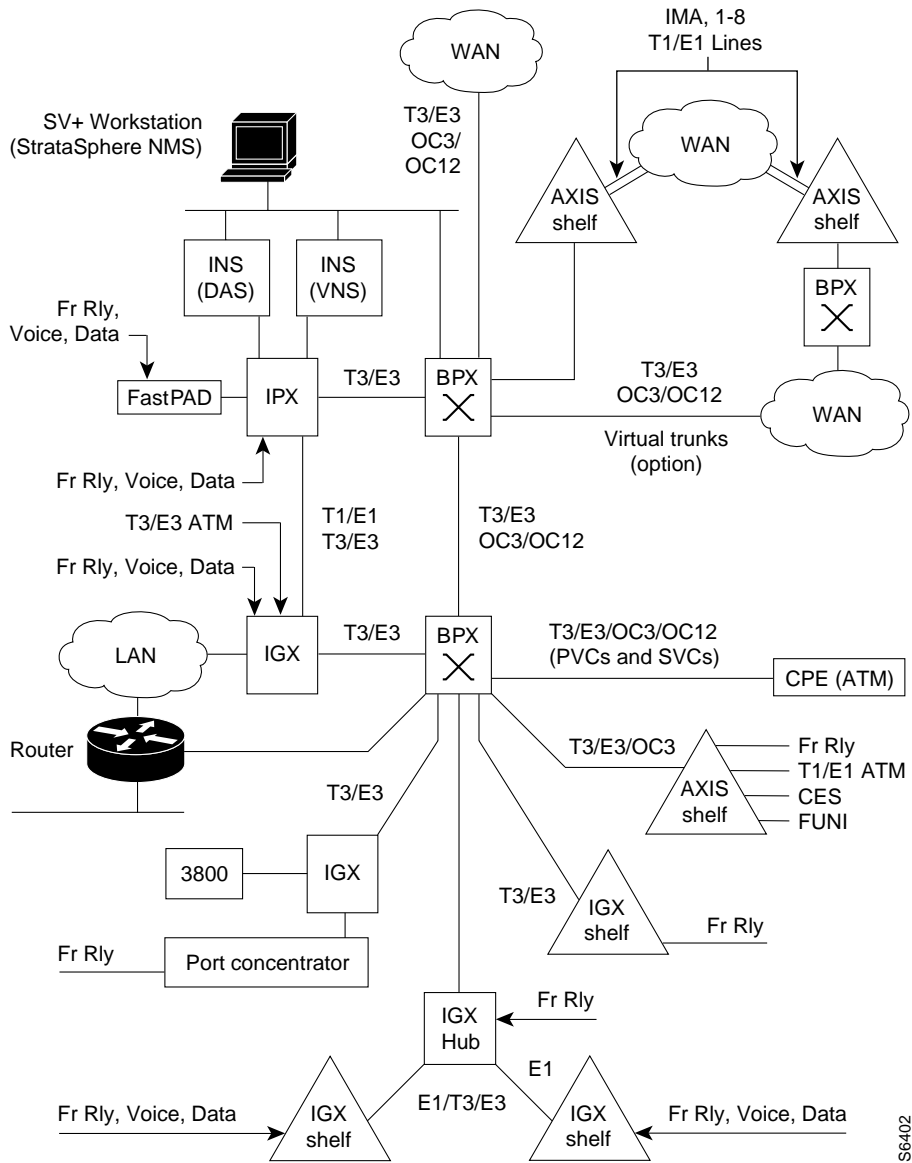
This chapter includes the following:

- Introduction
- New with Release 9.1
- New with Release 8.5
- New with Release 8.4
- Continuing Features with Release 9.1, 8.5 and 8.4
- ATM Networks
- ATM NW Features
- Traffic and Congestion Management
- Cell Relay Networking, ATM and FastPacket
- ATM Product Family Overview
- System Software Description
- Network Synchronization
- Network Availability
- Standards

Introduction

Cisco wide area ATM networks meet the expanding requirements of today's private enterprises and service providers. These ATM Wide Area Networks provide more bandwidth (up to OC-12/STM-4 rates of 622.08 Mbps), new services, reduced transaction costs, greater flexibility, scalability, service interworking, plus manageability and security for both enterprise and service provider networks (Figure 1-1).

Figur e1-1 An ATM Network Configuration



S6402

Expanding Network Requirements

Due to a number of events occurring in both enterprise and service provider networks; including on-going advances in computing power, more desktop interaction, the internet, more transactions, more visual content, an explosion of new applications, etc., greater demands have been placed on both private and public networks.

To maximize bandwidth utilization and flexibility, networks are moving from dedicated circuits with fixed bandwidth between devices to virtual networks. A virtual network comprises logical connections (virtual circuits) which dynamically share physical bandwidth capacity on an as needed basis with other logical connections (virtual circuits) or networks.

ATM Networks

These emerging virtual networks share bandwidth using the multiplexing technique called ATM asynchronous transfer mode, which allows networks to dynamically allocate capacity to connections on an as needed basis. ATM traffic is segmented into 53-byte cells for transmission of all types of traffic: voice, data, frame relay, video, ATM services, etc., at narrowband and broadband speeds.

New with Release 9.1

Cisco BPX 8600 Series Broadband Switch

- The Cisco BPX 8600 series broadband switch is configurable as a tag switching controller, supporting Cisco IP-based switches, routers and hubs.
- The BXM card on the Cisco BPX 8600 series broadband switch supports ATM connections with the UXM cards on Cisco IGX 8400 series multiband switch.
- The BME card on the Cisco BPX 8600 series broadband switch supports multicasting ATM connections.

Cisco IGX 8400 Series Multiband Switch

- The UXM card for the Cisco IGX 8400 series multiband switch supports a range of interfaces; including OC3, T3, E3, T1 or E1.
- The interfaces on a UXM card for the Cisco IGX 8400 series multiband switch can be configured as either ports (services) or trunks.
- The T1 and E1 interfaces on a UXM card for the Cisco IGX 8400 series multiband switch can be configured as IMA trunks.
- The UXM cards for the Cisco IGX 8400 series multiband switch support the following types of interworking:
 - interworking UXM port to FastPacket trunks
 - interworking UXM port to FastPacket Port cards (DAX connections)
 - interworking UXM trunk to FastPacket port cards
- The ATM trunks on a UXM card in the Cisco IGX 8400 series multiband switch support the industry standard UNI and NNI ATM cell formats.
- The ATM trunks on a UXM card in the Cisco IGX 8400 series multiband switch support configuration as virtual trunks

Network

- Support for Cisco IGX 8400 series multiband switches configured as hubs and associated interface shelves in tiered network
- Version interoperability between the 8.4, 8.5, and 9.1 releases, permitting individual node upgrades.

Cisco MGX 8220 Edge Concentrator

- MGX 8220 Release 4.1, which includes:
 - FRSM-8 with ELMI
 - AUSM-8 with IMA support
 - CESM 8 T1/E1
 - FRSM-HS1 (HSSI and X.21 interfaces)

StrataSphere Network Management

- StrataSphere NMS enhancements including additional management and provisioning capabilities including support of the UXM cards on the Cisco IGX 8400 series multiband switch and fab number support for all cards
- Support for the BME cards on the Cisco BPX 8600 series broadband switch
- Connection management support for the FRSM-HS1, CESM 8 T1/E1, FRSM-8 with ELMI, and AUSM-8 with IMA support on the Cisco MGX 8220 edge concentrator

CiscoView Network Element Management

- CiscoView network element management for the Cisco BPX 8600 series broadband switch, the Cisco IGX 8400 series multiband switch, and the Cisco MGX 8220 edge concentrator

New with Release 8.5

Cisco IGX 8400 Series Multiband Switch

- The Cisco IGX 8400 series multiband switch is configurable as a tiered network routing hub supporting Cisco IGX 8400 series multiband switches configured as interface shelves (feeders).
- A Cisco IGX 8400 series multiband switch configured as a routing hub supports both Cisco IGX 8400 series multiband switches configured as interface shelves and the Cisco 3810.
- A Cisco IGX 8400 series multiband switch configured as a routing hub supports up to four Cisco IGX 8400 series multiband switches configured as interface shelves.
- Cisco IGX 8400 series multiband switches configured as interface shelves connected to Cisco IGX 8400 series multiband switches configured as routing hubs support voice, data, and frame relay connections.
- Voice and data connections originating at a Cisco IGX 8400 series multiband switch configured as an interface shelf are routed across Cisco IGX 8400 series multiband intermediate switches and terminated on another Cisco IGX 8400 series multiband switch configured as an interface shelf.
- The Cisco IGX 8400 series multiband switch supports 30 trunks per routing node.
- The Cisco IGX 8400 series multiband switch supports 2750 virtual connections.
- Connection deroute delay.

- Connection routing groups by cell loading.

Network

- Support for Cisco IGX 8400 series multiband switches configured as hubs and associated interface shelves in tiered network
- The number of nodes supported in a network is increased to over 1100, of which 223 can be routed nodes.

Cisco MGX 8220 Edge Concentrator

- The Cisco MGX 8220 edge concentrator supports up to 12 StrataView Plus workstations.

StrataSphere Network Management

- StrataSphere NMS enhancements including additional management and provisioning capabilities including support of Cisco IGX 8400 series multiband switch tiered network voice and data applications
- Support for 12 Cisco WAN Manager (formerly StrataView Plus) workstations
- Multi-network Cisco WAN Manager capability

New with Release 8.4

- The BXM cards provide a range of trunk and service interfaces and support ATM Forum Standards UNI 3.1 and ATM Traffic Management 4.0 including ABR connections with VS/VD congestion control. The BXM cards are implemented with Stratm technology which uses a family of custom Application Specific Integrated Circuits (ASICs) to provide high-density, high-speed operation. The three types of BXM cards are:
 - The BXM T3/E3 is available as an eight or twelve port card that provides T3/E3 interfaces at 44.376 or 34.368 Mbps rates, respectively. The BXM-T3/E3 can be configured for either trunk or access applications.
 - The BXM 155 is available as a four or eight port card that provides OC-3/STM-1 interfaces at 155.52 Mbps rates. The BXM-155 can be configured for either trunk or access applications.
 - The BXM 622 is available as a one or two port card that provides OC-12/STM-4 interfaces at 622.08 Mbps rates. The BXM-622 can be configured for either trunk or access applications.
- Enhanced network scaling:
 - 50/64 trunks per Cisco BPX 8600 series wideband switch equipped with BCC-32 or BCC-64, respectively
 - 72/144 lines per Cisco BPX 8600 series wideband switch equipped with BCC-32 or BCC-64, respectively
 - 223 routing nodes (with Cisco BPX 8600 series wideband switch or Cisco IGX 8400 series multiband switch)
 - trunk based loading

- BCC-3-64
- 7000 virtual connections (BCC-3-32)
- 12000 virtual connections (BCC-3-64)
- de-route delay timer
- connection routing groups by cell loading
- ATM and Frame Relay SVCs with Extended Services Processor

The Extended Services Processor (ESP) is an adjunct processor that is co-located with a Cisco BPX 8600 series wideband switch. The ESP provides the signaling and Private Network to Network Interface (PNNI) routing for ATM and frame relay switched virtual circuits (SVCs) via BXM cards in the Cisco BPX 8600 series wideband switch and AUSM and FRSM cards in the Cisco MGX 8220 edge concentrator.
- StrataSphere NMS enhancements including additional management and provisioning capabilities.
- BCC-3-64
- BCC-4 supporting 19.2 Gbps switching
- BXM cards support egress at up to 1600 Mbps and ingress at up to 800 Mbps
- Hot Standby Redundancy
- MGX 8220 Release 4.0, which includes:
 - BNM-155 interface to the BXM on the Cisco BPX 8600 series wideband switch
 - FRSM support for both SVC and PVC frame relay connections with ESP
 - AUSM support for both SVC and PVC ATM connections with ESP
 - FRSM-8 with ELMI
 - IMATM-B
 - AUSM-8
 - CESM/4T1E1
 - FRSM-HS1 (HSSI and X.21 interfaces)
 - SRM 3T3
- Access Products
 - Cisco 3800

Continuing Features with Release 9.1, 8.5 and 8.4

The following is a list of some of the continuing features with Release 9.1, 8.5 and 8.4:

StrataSphere Network Management

- StrataSphere frame relay connection and Cisco MGX 8220 equipment management by the Cisco WAN Manager Connection Manager and Equipment Manager.
- SNMP Enhancements for connection management and monitoring
- Support for Solaris 2.5.1

Network

- IMA (Inverse Multiplexing ATM)
- Frame relay to ATM network interworking (supported by the FRP on the Cisco IPX narrowband switch, the FRM and the UFM on the Cisco IGX 8400 series multiband switch, and the FRSM on the Cisco MGX 8220 edge concentrator)
- Frame relay to ATM service interworking (supported by the FRSM on the Cisco MGX 8220 edge concentrator and the UFM on the Cisco IGX 8400 series multiband switch)
- Tiered networks
- Automatic end-to-end routing of virtual connections (AutoRoute)
- Closed-loop, rate-based congestion management (using ABR with VS/VD or ForeSight)
- Effective management of quality of service (OptiClass)
- Per -VC queuing and per-VC scheduling (FairShare)

INS-DAS and INS-VNS

- INS-DAS: Dial Access Server provides frame relay dialup and dial backup services.
- INS-VNS: Voice Network Switching (VNS) provides voice switched virtual circuits for PBXs using QSIG and DPNSS signaling.

Cisco BPX 8600 Series Wideband Switch

- Virtual Trunking
- IMA (Inverse Multiplexing ATM)
- Enhanced Ingress buffers for ASI-155 and BNI-155 to 8K cells for Release 8.1 and up
- Cisco BPX 8600 series wideband switch OC3 network and service interfaces on the BNI and AS cards, respectively
- High-speed switching capacity
- Powerful crosspoint switching architecture
- 53-byte cell-based ATM transmission protocol
- Twelve 800 Mbps switch ports for network or access interfaces with BNI and ASI cards respectively.
- Three DS3 or E3 ATM network interface ports per card (BNI)
- Totally redundant common control and switch fabric
- Up to 20 million point-to-point cell connections per second between slots
- Switches individual connections rather than merely serving as a virtual path switch
- Easy integration into existing Cisco IPX narrowband and Cisco IGX 8400 series multiband networks
- Internal diagnostics and self-test routines on all cards and backplane, status indication on each card
- Collection of many ATM and other network statistics and transfer of the data collected to Cisco WAN Manager over high-speed Ethernet LAN interface

- Integration with the Cisco WAN Manager Network Management System to provide configuration, control, and maintenance
- Conformation to recommendations from all current ATM standards bodies: ATM Forum, ITU, ETSI, and ANSI
- Compliant with all applicable safety, emissions, and interface regulations. Meets requirements of NEBS for Central Office equipment

Cisco IGX 8400 Series Multiband Switch

- ALM-A
- ALM-B
- UFM-C, supporting ELMI
- UFM-U (HSSI, X.21, V.35), supports ELMI
- UVM, supporting CAS switching in conjunction with VNS
- Cisco IGX 8400 series multiband switches configured as interface shelves connected to Cisco BPX 8600 series wideband switches configured as routing hubs support frame relay connections.
- Port Concentrator Shelf provides low-cost frame relay service ports
- Supports FastPAD device connections including voice, data, and frame relay.
- Zero CIR (Uncommitted Service Option) for frame relay
- Cisco IGX 8400 series multiband switches configured as interface shelves interfaced to Cisco BPX 8600 series wideband switches configured as routing hubs support frame relay connections.
- Supports voice, data, multi-media, and frame relay connections.

Cisco IPX Narrowband Switch

- Cisco IPX narrowband switches configured as interface shelves connected to Cisco BPX 8600 series wideband switches configured as routing hubs support frame relay connections.
- Port Concentrator Shelf provides low-cost frame relay service ports.
- Supports FastPAD device connections including voice, data, and frame relay.
- Zero CIR (Uncommitted Service Option) for Frame Relay
- Supports voice, data, multi-media, and frame relay connections.

Cisco MGX 8220 Edge Concentrators

- IMA (Inverse Multiplexing ATM) for the Cisco BPX 8600 series wideband switch using T1/E1 interfaces
- CES T1/E1
- Cisco MGX 8220 T1/E1 frame relay and T1/E1 ATM service interfaces
- FUNI (Frame Based UNI overATM)

Access Products

FastPAD MM

FastPAD MP

ATM Networks

Cisco WAN switching networking systems support multiband ATM applications in private wide area networks and service provider service offerings, such as frame relay and native ATM. Cisco's WAN switching product family includes the Cisco IPX narrowband switch, the Cisco IGX 8400 series multiband switch, the Cisco BPX 8600 series wideband switch, the Cisco MGX 8220 edge concentrator, and the FastPAD, Cisco 3810, INS-DAS, INS-VNS, and StrataSphere products. These products integrate and transport a wide variety of information, including voice, data, frame relay, video, LAN traffic, image, and multimedia traffic ranging from narrowband to broadband ATM.

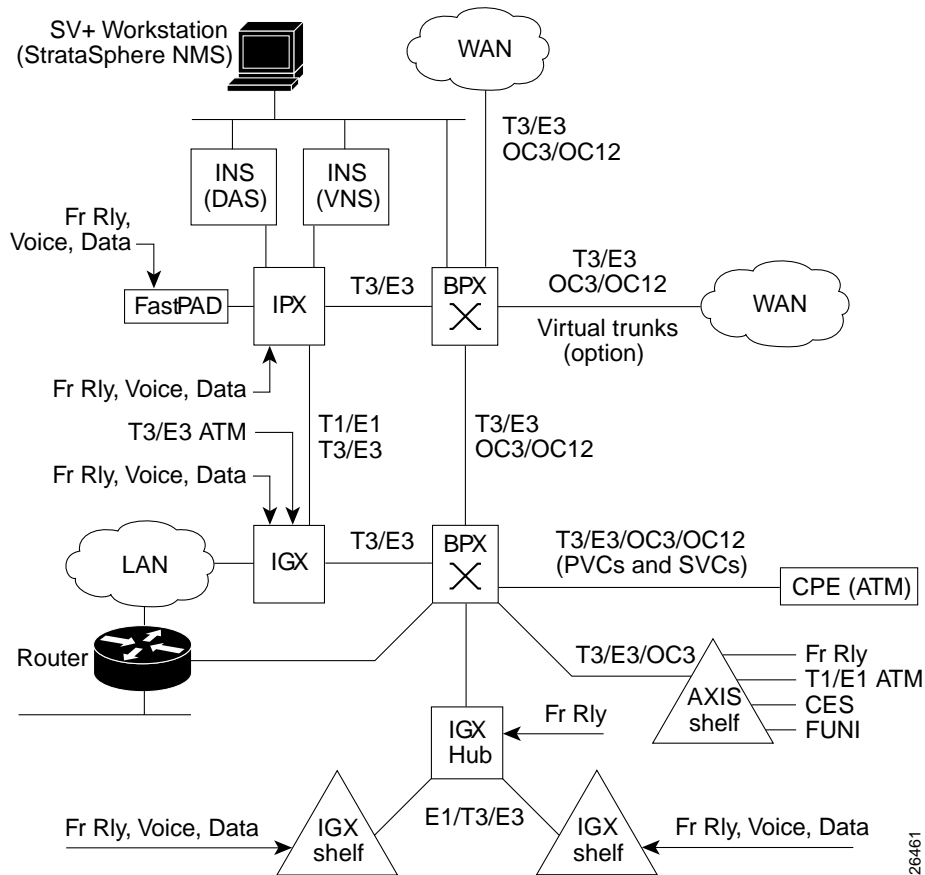
The Cisco IPX narrowband switches, the Cisco IGX 8400 series multiband switches, and the Cisco BPX 8600 series wideband switches are used to implement digital high-speed, wide area private and public networks (WANs) for interconnecting customer's local area networks (LANs). These cell relay networks are created by interconnecting Cisco WAN switching network switches with high-speed digital trunks provided by any of the number of public common carriers or private service providers.

Enterprise Wide Area Networks

Corporations, government agencies, universities, telecommunications service providers, and others with a need to link their communications facilities can use the Cisco IPX narrowband switch, the Cisco IGX 8400 series multiband switch, and the Cisco BPX 8600 series wideband switch as a basis on which to build their own private networks (Figure e1-2).

In many instances, the primary reason for implementing private WANs is to link far-flung LANs. With the additional bandwidth available and the flexibility of cell relay technology, a private user often can add voice circuits and even a video conferencing facility on the same network without adding trunks and with very little additional expense.

Figur e1-2 Example of an Enterprise Network Application



The Cisco IPX Narrowband Switch

The majority of private network locations have lower bandwidth requirements, fewer routes, small hubs, and a wide variety of service requirements. The Cisco IPX narrowband switch fits these applications by providing a wide offering of customer interfaces, several package sizes, and a scalable architecture. The Cisco IPX narrowband switch allows the user at each site to replace numerous low-speed dial-up and/or leased line circuits with a few high-speed T1 or E1 lines. The advantages often include faster response times, wider range of available services, more efficient utilization of bandwidth and the resulting cost savings. Plus, the private network often provides better management control, tighter security, and increased configuration flexibility under direct control of the end user.

The Cisco IGX 8400 series Multiband Switch

The Cisco IGX 8400 series multiband switch is a multiservice ATM networking switch that provides interfaces to support today's legacy and emerging broadband applications. Users have the advantage of ATM technology over narrowband and subrate T1 and E1, and broadband T3, and E3 trunks.

The Cisco IGX 8400 series multiband switch can be used as the basis for a leased-line Campus/MAN/WAN network, as an access device to high-speed digital services such as ATM, as a combination of both applications, and as a Value Added Network (VAN) service switch. Operating at 1.2 Gbps, the Cisco IGX 8400 series multiband switch seamlessly integrates with the Cisco IPX narrowband switch, the Cisco BPX 8600 series broadband switch, the Cisco MGX 8220 edge

concentrator, INS-DAS, INS-VNS, FastPAD access devices, and the Cisco 3810 to provide multiband solutions from the access interface to the core layer. The Cisco IGX 8400 series multiband switch can be configured as a routing hub and as an interface shelf. A Cisco IGX 8400 series multiband switch configured as an interface shelf connected to a Cisco IGX 8400 series multiband switch configured as a routing hub supports voice, data, and frame relay connections. The voice and data connections, in this case, are routed across Cisco IGX 8400 series intermediate nodes to another Cisco IGX 8400 series multiband switch configured as an interface shelf. A Cisco IGX 8400 series multiband switch configured as an interface shelf connected to a Cisco BPX 8600 series wideband switch configured as a routing hub supports frame relay connections across an ATM network.

The Cisco BPX 8600 Series Broadband Switch with Cisco MGX 8220 Edge Concentrators

Many network locations have increasing bandwidth requirements due to emerging applications. To meet these requirements, users can overlay their existing narrowband networks with a backbone of Cisco BPX 8600 series broadband switches to utilize the high-speed connectivity of the BPX operating at 19.2 Gbps with its T3/E3/OC3/OC12 network and service interfaces. The BPX service interfaces include BXM and ASI ports on the Cisco BPX 8600 series broadband switch and service ports on Cisco MGX 8220 edge concentrators. The Cisco MGX 8220 edge concentrators may be co-located in the same cabinet as the Cisco BPX 8600 series broadband switch, providing economical port concentration for T1/E1 Frame Relay, T1/E1 ATM, CES, and FUNI connections.

The Cisco BPX 8600 Series Broadband Switch with ESP

With a co-located Extended Services Processor (ESP), the Cisco BPX 8600 series broadband switch adds the capability to support ATM and frame relay switched virtual circuits (SVCs).

Service Provider Multi-Service Networks

The demand to provide LAN interconnections has driven most of the public service providers to consider ways to quickly react to this opportunity. Frame relay has proven to be a reliable, cost-effective, standards-based service for transmitting LAN data, which tends to be very bursty in nature. Typically, LANs access the network only at periodic intervals but when they do, they often require large amounts of bandwidth for short periods of time. It is not cost effective to provide sufficient bandwidth to every LAN connection on a full-time basis.

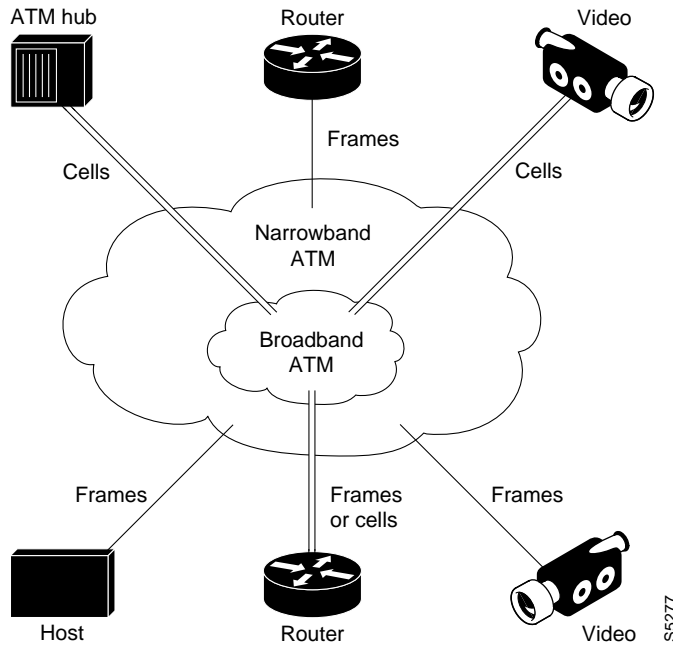
Since both the Cisco IPX narrowband switch and the Cisco IGX 8400 series multiband switch utilize cell-relay technology, there are significant service advantages when frame relay is implemented on them. Since cell network platforms only allocate bandwidth when there is demand, the unused bandwidth from idle frame relay connections can be used by active connections. This allows the active connections to “burst” or to send large amounts of data for a short interval above their committed information rate. Then as the connection goes idle, the bandwidth is utilized for yet another connection.

Another advantage of cell relay networks is the flexibility of offering Permanent Virtual Circuits (PVCs) to interconnect all LAN sites in a mesh topology, in contrast to using physical circuits that require a large investment in interface hardware and data circuits. Frame relay networks based on Cisco IPX narrowband switches, Cisco IGX 8400 series multiband switches and Cisco MGX 8220 edge concentrators offer minimal delay, maximum throughput, and avoid congestion.

Current frame relay networks offer LAN circuit interconnection at rates from 56 Kbps to 2 Mbps. As frame relay traffic increases and customers demand more bandwidth for advanced applications, Cisco IPX narrowband and Cisco IGX 8400 series multiband FastPacket architectures give service providers a clear upgrade path to broadband ATM capabilities. Broadband networks, utilizing high-speed trunks

and ATM cell switching, can be overlaid on a narrowband FastPacket network by adding a backbone of Cisco BPX 8600 series broadband switches. An existing network can be upgraded by adding a high-speed network ATM backbone utilizing Cisco BPX 8600 series wideband switches with gigabit switching as indicated in Figure 1-3.

Figur e1-3 Example of a Service Provider Application



The Cisco BPX 8600 series broadband switch provides ATM UNI services from the same platform using BXM and ASI cards, and also interfaces to Cisco IGX 8400 series multiband switches and Cisco IPX narrowband switches configured as routing nodes and to Cisco MGX 8220 edge concentrators and Cisco IPX narrowband switches and Cisco IGX 8400 series multiband switches configured as interface shelves to provide cost-effective multi-media services.

Frame relay ports can also be provided directly on Cisco BPX 8600 series broadband switches using Cisco MGX 8220 edge concentrators or Cisco IPX narrowband switches configured as interface shelves for maximum port density.

This multiband ATM service model provides both narrow and broadband interfaces from a common cell switching infrastructure. Connection management, congestion control, and network management are extended seamlessly across the entire network.

ATM NW Features

Advanced Capabilities

The Cisco WAN switching ATM networks include sophisticated system software for management and control of the network. Cisco WAN Switching system software is fully distributed in each switch to provide the fastest response time to network provisioning and network problems. Advanced capabilities include ABR with VSVD, Frame Relay to ATM Network and Service Interworking, Tiered Network operation, Virtual Trunking, Sonet/SDH Interfaces, AutoRoute, Opti-Class, FairShare, and ForeSight.

ABR with VSVD

The BXM series of cards, BXM T3/E3, BXM-155, and BXM-622 using high density application specific integrated circuit (ASIC) technology, provide advanced ATM networking features including ABR with VSVD supporting explicit rate (ER) feedback and congestion indication (CI) options. The BXM cards provide a high degree of scalability, providing 8 or 12 T3 or E3 ports per slot with the BXM-T3/E3, 4 or 8 OC3 ports with the BXM-155, and 1 or 2 OC12 ports with the BXM-622.

Frame Relay to ATM Interworking

Interworking allows users to retain their existing 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

Two types of frame relay to ATM interworking are supported, Network Interworking (Figure 1-4) and Service Interworking (Figure 1-5). The Network Interworking function is performed by the AIT card on the Cisco IPX narrowband switch, the BTM card on the Cisco IGX 8400 series multiband switch, and the FRSM card on the Cisco MGX 8220 edge concentrator. The FRSM cards on the Cisco MGX 8220 edge concentrator and the UFM cards on the Cisco IGX 8400 series multiband switch also support Service Interworking.

The frame relay to ATM network and service interworking functions are available as follows:

Network Interworking

Part A of Figure 1-4 shows typical frame relay to network interworking. In this example, a frame relay connection is transported across an ATM network, and the interworking function is performed by both ends of the ATM network. The following are typical configurations:

- Cisco IGX 8400 series multiband or Cisco IPX narrowband frame relay (shelf/feeder) to Cisco IGX 8400 series multiband or Cisco IPX narrowband frame relay (either routing node or shelf/feeder)
- Cisco MGX 8220 edge concentrator frame relay to Cisco MGX 8220 edge concentrator frame relay
- Cisco MGX 8220 edge concentrator frame relay to Cisco IGX 8400 series multiband or Cisco IPX narrowband frame relay (either routing node or shelf/feeder)

Part B of Figure 1-4 shows a form of network interworking where the interworking function is performed by only one end of the ATM network, and the CPE connected to the other end of the network must itself perform the appropriate service specific convergence sublayer function. The following are example configurations:

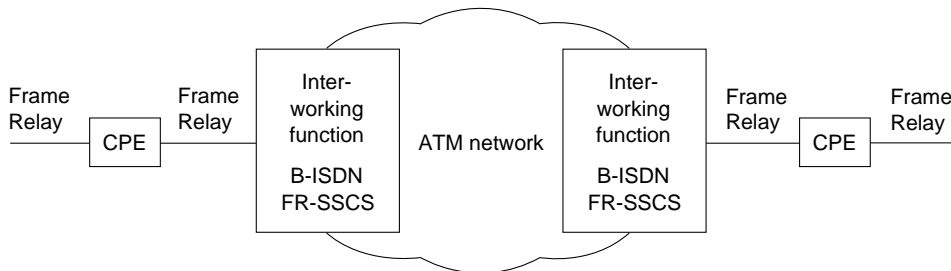
- Cisco IGX 8400 series multiband or Cisco IPX narrowband frame relay (either routing node or shelf/feeder) to Cisco BPX 8600 series broadband switch or Cisco MGX 8220 edge concentrator ATM port.
- Cisco MGX 8220 edge concentrator frame relay to Cisco BPX 8600 series broadband switch or Cisco MGX 8220 edge concentrator ATM port.

Network Interworking is supported by the FRP on the Cisco IPX narrowband switch, the FRM, UFM-C, and UFM-U on the Cisco IGX 8400 series multiband switch, and the FRSM on the Cisco MGX 8220 edge concentrator. The Frame Relay Service Specific Convergence Sublayer (FR-SSCS) of AAL5 is used to provide protocol conversion and mapping.

Figur e1-4 Frame Relay to ATM Network Interworking


Part A

Network interworking connection from CPE Frame Relay port to CPE Frame Relay port across an ATM Network with the interworking function performed by both ends of the network.



Part B

Network interworking connection from CPE Frame Relay port to CPE ATM port across an ATM network, where the network performs an interworking function only at the Frame Relay end of the network. The CPE receiving and transmitting ATM cells at its ATM port is responsible for exercising the applicable service specific convergence sublayer, in this case, (FR-SSCS).




Service Interworking

Figure 1-5 shows a typical example of service interworking. Service interworking is supported by the FRSM on the Cisco MGX 8220 edge concentrator and the UFM-C and UFM-U on the Cisco IGX 8400 series multiband switch. Translation between the Frame Relay and ATM protocols is performed in accordance with RFC 1490 and RFC 1483. The following is a typical configuration for service interworking:

- Cisco MGX 8220 edge concentrator Frame Relay (FRSM card) to Cisco BPX 8600 series broadband switch or Cisco MGX 8220 edge concentrator ATM port
- Cisco IGX 8400 series multiband Frame Relay (FRM-U or FRM-C card) to Cisco BPX 8600 series broadband switch or Cisco MGX 8220 edge concentrator ATM port

Figur e1-5 *Frame Relay to ATM Service Interworking*



Additional Information on Interworkin

For additional information about interworking, refer to Chapter 13, “Frame Relay to ATM Network and Service Interworking”.

Tiered Networks

Networks may be configured as flat (all nodes perform routing and communicate fully with one another), or they may be configured as tiered. In a tiered network interface shelves are connected to routing hubs, where the interface shelves are configured as non-routing nodes.



Note

A routing hub is a standard routing node that is also connected to interface shelves via feeder trunks.

With Release 8.5 and beyond, tiered networks support voice and data connections as well as frame relay connections. With this addition, a tiered network can now provide a multi-service capability (frame relay, circuit data, voice, and ATM). By allowing CPE connections to connect to a non-routing node (interface shelf), a tiered network is able to grow in size beyond that which would be possible with only routing nodes comprising the network.

Routing Hubs and Interface Shelve

In a tiered network, interface shelves at the access layer (edge) of the network are connected to routing nodes via feeder trunks (Figure 1-6). Those routing nodes with attached interface shelves are referred to as routing hubs. The interface shelves, sometimes referred to as feeders, are non-routing nodes. The routing hubs route the interface shelf connections across the core layer of the network.

The interface shelves do not need to maintain network topology nor connection routing information. This task is left to their routing hubs. This architecture provides an expanded network consisting of a number of non-routing nodes (interface shelves) at the edge of the network that are connected to the network by their routing hubs.

For detailed information about tiered networks, refer to Chapter 6, “Tiered Networks.”

Cisco IGX 8400 Series Multiband Switches Configured as Routing Hubs

Voice, data, and frame relay connections originating on Cisco IGX 8400 series multiband switches configured as interface shelves (feeders) are routed across the routing network via their associated Cisco IGX 8400 series multiband switches configured as routing hubs. For voice and data connections originating on Cisco IGX 8400 series multiband switches configured as interface shelves, the intermediate routing nodes must be Cisco IGX 8400 series multiband switches. A frame relay connection originating at a Cisco IGX 8400 series multiband switch configured as an interface shelf may terminate on a Cisco MGX 8220 edge concentrator, a Cisco IPX narrowband switch, or a Cisco IGX 8400 series multiband switch configured as an interface shelf, or a Cisco IPX narrowband switch or a Cisco IGX 8400 series multiband switch configured as a routing node.

The following applies to IGX routing hubs and interface shelves:

- A Cisco IGX 8400 series multiband switch configured as a routing hub supports up to four Cisco IGX 8400 series multiband switches configured as interface shelves.
- A Cisco IGX 8400 series multiband switch configured as an interface shelf can have only one feeder trunk to the routing network.
- A Cisco IGX 8400 series multiband switch configured as an interface shelf is the only type of interface shelf that can connect to Cisco IGX 8400 series multiband switches configured as routing hubs.
- No direct trunking between Cisco IGX 8400 series multiband switches configured as interface shelves is supported.
- No routing trunk is supported between the routing network and interface shelves.
- The feeder trunks between Cisco IGX 8400 series multiband switches configured as routing hubs and Cisco IGX 8400 series multiband switches configured as interface shelves are connected to a BTM-E1, or T3/E3 backcards on each end of the trunk.
- Telnet is supported to an interface shelf; the vt command is not.
- Remote printing by the interface shelf via a print command from the routing network is not supported.

The following applies to voice and data connections over Cisco IGX 8400 series multiband switches configured as interface shelves:

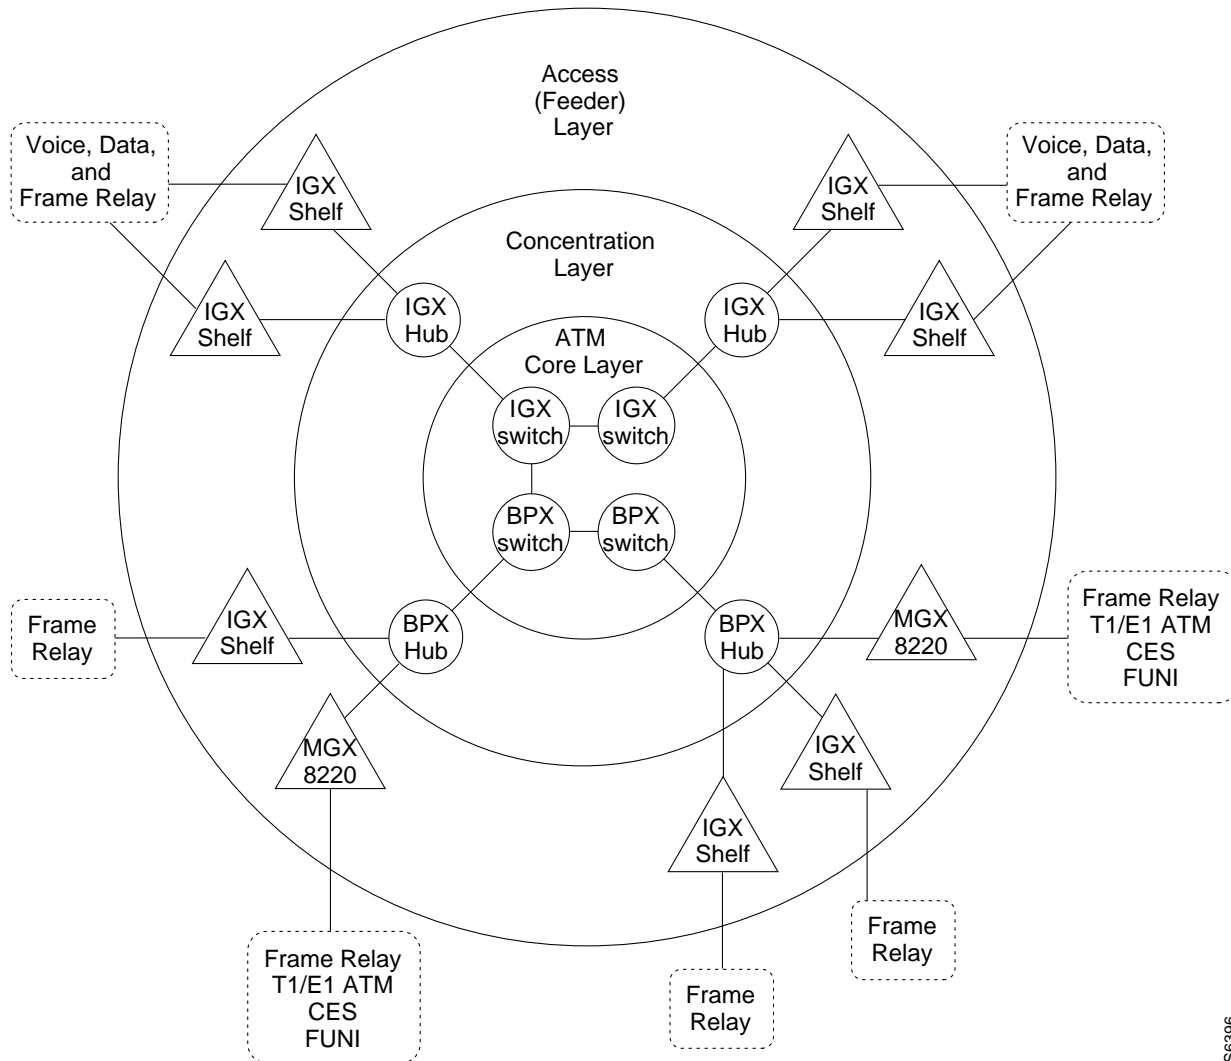
- Multi-service connection management to a Cisco IGX 8400 series multiband switch configured as an interface shelf is provided by Cisco WAN Manager.
- 3-segment connections are supported, that is: originating Cisco IGX 8400 series multiband interface shelf data or voice card to a Cisco IGX 8400 series multiband switch configured as a routing hub, across Cisco IGX 8400 series multiband intermediate switches, as applicable, to a Cisco IGX 8400 series multiband switch configured as a routing hub, to a data or voice card on a terminating Cisco IGX 8400 series multiband switch configured as an interface shelf.
- 2-segment voice or data connections are not supported, (data or voice card on a terminating Cisco IGX 8400 series multiband switch configured as an interface shelf to a Cisco IGX 8400 series multiband switch configured as a routing hub).

- Routing of voice and data connections through the middle segment of the three segment connection is done via Cisco IGX 8400 series multiband routing switches using CBR mode and simple gateway over the IGX trunks.
- Connection statistics are supported at user endpoints only.
- Adaptive voice is not supported.

The following applies to frame relay connections originating at Cisco IGX 8400 series multiband switches configured as interface shelves connected to Cisco IGX 8400 series multiband switches configured as routing hubs:

- Multi-service connection management to a Cisco IGX 8400 series multiband switch configured as an interface shelf is provided by Cisco WAN Manager.
- 3-segment connections are supported, that is: originating at a frame relay card in a Cisco IGX 8400 series multiband switch configured as an interface shelf to a Cisco IGX 8400 series multiband switch configured as a routing hub, across Cisco IGX 8400 series multiband intermediate switches, as applicable, to a Cisco IGX 8400 series multiband switch configured as a routing hub, to a data or voice card on a terminating Cisco IGX 8400 series multiband switch configured as an interface shelf.
- 2-segment frame relay connections are supported, (voice connections on a Cisco IGX 8400 series multiband switch configured as an interface shelf to a Cisco IGX 8400 series multiband switch configured as a routing hub, to any routing node or routing hub, to a terminating frame relay card on a Cisco IPX narrowband switch, on a Cisco IGX 8400 series multiband switch, on a Cisco MGX 8220 edge concentrator, etc.).
- Routing of frame relay connections through the routing nodes is done using ATFR mode and simple or complex gateway, as applicable to the trunk cards.
- Connection statistics are supported at user endpoints only.
- Connection management is provided by CiscoWAN Manager.

Figur e1-6 Tiered Network with Cisco BPX 8600 Series Broadband Switches and Cisco IGX 8400 Series Multiband Switches Configured as Routing Hubs



S6396

BPX Routing Hub

T1/E1 Frame Relay connections originating at Cisco IPX narrowband switches and Cisco IGX 8400 series multiband switches configured as interface shelves and T1/E1 Frame Relay, T1/E1 ATM, CES, and FUNI connections originating at Cisco MGX 8220 edge concentrators configured as interface shelves are routed across the routing network via their associated Cisco BPX 8600 series wideband switches configured as routing hubs.

The following requirements apply to Cisco BPX 8600 series wideband switches configured as routing hubs and their associated interface shelves:

- Only one feeder trunk is supported between a routing hub and an interface shelf.
- No direct trunking between interface shelves is supported.
- No routing trunk is supported between the routing network and interface shelves.

- The feeder trunks between Cisco BPX 8600 series wideband switches configured as routing hubs and Cisco IPX narrowband switches or Cisco IGX 8400 series multiband switches configured as interface shelves are either T3 or E3.
- The feeder trunks between Cisco BPX 8600 series wideband switches configured as routing hubs and Cisco MGX 8220 Edge Concentrators configured as interface shelves are T3, E3, or OC3-c/STM-1.
- Frame Relay Connection management to Cisco IPX narrowband switches and Cisco IGX 8400 series multiband switches configured as interface shelves are provided by StrataView Plus
- Frame Relay and ATM connection management to a Cisco MGX 8220 edge concentrator configured as an interface shelf is provided by StrataView Plus
- Telnet is supported to an interface shelf; the vt command is not.
- Remote printing by the interface shelf via a print command from the routing network is not supported.

IMA (Inverse Multiplexing ATM)

Where greater bandwidths are not needed, the Inverse Multiplexing ATM (IMA) feature provides a low cost trunk between two Cisco BPX 8600 series wideband switches. The IMA feature allows Cisco BPX 8600 series wideband switches to be connected to one another over from 1 to 8 T1 or E1 trunks provided by an IMATM module on a Cisco MGX 8220 edge concentrator. A BNI or BXM port on each Cisco BPX 8600 series wideband switch is directly connected to an IMATM or an AUSM 8 module in a Cisco MGX 8220 edge concentrator by a T3 or E3 trunk. The AIMNM modules are then linked together by from 1 to 8 T1 or E1 trunks. Refer to the *Cisco MGX 8220 Edge Concentrator Installation and Configuration* and the *Cisco MGX 8220 Edge Concentrator Command Reference* documents for further information.

Circuit Emulation Service (CES)

The Cisco MGX 8220 edge concentrator supports CES over T1/E1 lines with either an 8 port or 4 port Circuit Emulation Service Module. Data is transmitted and received over the network in AAL-1 cell format.

Zero CIR for Frame Relay

This feature allows users to take advantage of lower cost uncommitted frame relay service. The feature applies to frame relay connections that originate and terminate on FRP (Cisco IPX narrowband switch) or FRM (Cisco IGX 8400 series multiband switch) ports. It does not apply to frame relay to ATM connections. The CIR on the connection is configured to zero, and the MIR must be non-zero.

Virtual Trunking


Virtual trunking provides the ability to define multiple trunks within a single physical trunk interface. Virtual trunking benefits include the following:

- Reduced cost by configuring the virtual trunks supplied by the service provider for only as much bandwidth as needed instead of at full T3, E3, or OC3 bandwidths.

- Utilization of the full mesh capability of the service provider to reduce the number of leased lines needed between nodes in the CiscoWAN switch subnetworks.
- Or, choice of keeping existing leased lines between nodes, but using virtual trunks for backup.
- Ability to connect BNI trunk interfaces to a public network using standard ATM UNI cell format
- Virtual trunking can be provisioned via either a Public ATM Cloud or a Cisco WAN Switching ATM cloud.

A virtual trunk may be defined as a “trunk over a public ATM service”. The trunk really doesn’t exist as a physical line in the network. Rather, an additional level of reference, called a **virtual trunk number**, is used to differentiate the virtual trunks found within a physical trunk port. Figure 1-7 shows four Cisco WAN Switching sub-networks, each connected to a Public ATM Network via a physical line. The Public ATM Network is shown linking all four of these subnetworks to every other one with a full meshed network of virtual trunks. In this example, each physical line is configured with three virtual trunks.

Figur e1-7 Virtual Trunking Example



For further information on Virtual Trunking, refer to Chapter 8, ATM and Broadband Trunks. Also, refer to the *Cisco WAN Switching Command Reference document*.

BXM Network and Service SONET and T3/E3 Interfaces

To meet the need for high performance backbone networking, the BXM-622, BXM-155, and BXM T3/E3 cards provide OC12/STM-4 (622.08 Mbps), OC3/STM-1 (155.52 Mbps), and T3/E3 (44.736/34.368 Mbps) interfaces, respectively. The BXM cards may be user configured to function in either of two modes:

- in trunk mode to provide trunk interfaces (node/network connection)
- in access (port) mode to provide access service ports to customer CPE.

The BXM-622, BXM-155, and BXM-T3/E3 are designed to support the ATM UNI 3.1 and ATM Forum TM 4.0 standards.

BXM-622, 155, and T3/E3 Card

The BXM-622, BXM-155, and BXM-T3/E3 cards support the full range of ATM service types per ATM Forum TM 4.0.

To support the BXM-622 front cards, there are SMF and SMFLR backcards in either one or two port versions as required. To support the BXM-155-XX front cards there are MMF, SMF, and SMFLR back cards available in either 4 or 8 port configurations. To support the BXM-T3/E3 there is the Cisco BPX 8600 series broadband switch-T3/E3-8/12 back card available in either 8 or 12 port versions. Any of the 12 general purpose Cisco BPX 8600 series broadband switch slots can be used to hold these cards. The same backcards are used whether the BXM ports are configured as trunks or lines.

Trunk Mod

When configured for trunk mode, the BXM cards provide high-speed ATM interconnections between Cisco BPX 8600 series broadband switches and networks. The large cell buffering capability provided by the BXM cards ensures highly reliable ATM trunk connections.

Service Mod

The BXM-622, BXM-155, and BXM-T3/E3 are designed to support all the following service classes per ATM Forum TM 4.0: Constant Bit Rate (CBR), Variable Bit Rate (VBR), Available Bit Rate (ABR with VS/VD, ABR without VS/VD, and ABR using Foresight), and Unspecified Bit Rate (UBR). ABR with VS/VD supports explicit rate marking and Congestion Indication (CI) control.

BNI Network and ASI Service Interfaces

The BNI-155 and ASI-155 provide OC-3 network and service interfaces, respectively, but provide a more limited range of service types and configurations than the BXM-155 card. The BNI T3 and BNI E3 provide T3 and E3 network interfaces respectively. The ASI T3 and E3 provide T3 and E3 service interfaces, respectively.

BNI-155 Network Interfac

The BNI-155 operates at the standard OC-3/STM-1 (155.52 Mbps) rate to provide high-speed ATM trunking between Cisco BPX 8600 series broadband switches. The BNI-155 supports up to 12 Classes of Service (CoS) over Cisco BPX 8600 series broadband switch networks.

The physical interface options include multi-mode fiber (MMF), single-mode fiber intermediate reach (SMF-IR) and single-mode fiber long reach (SMF-LR) for optical terminations.

ASI-155 Service Interfac

The ASI-155 provide broadbands connectivity between the Cisco BPX 8600 series broadband switch and ATM CPE. The ASI-155 is a two port OC-3/STM-1 (155.520Mbps) ATM service interface card that can be plugged into any of the 12 general purpose card slots in the Cisco BPX 8600 series broadband switch. The ASI supports a more limited range of service types than the BXM cards.

The physical interface options include multi-mode fiber (MMF), single-mode fiber intermediate reach (SMF-IR) and single-mode fiber long reach (SMF-LR) for optical terminations.

Traffic and Congestion Management

The Cisco BPX 8600 series broadband switch provides ATM standard traffic and congestion management per ATM Forum TM 4.0 using BXM cards.

The Traffic Control functions include:

- Usage Parameter Control (UPC)
- Traffic Shaping
- Connection Management Control:
- Selective Cell Discarding
- Explicit Forward Congestion Indication (EFCI)

In addition to these standard functions, the Cisco BPX 8600 series broadband switch provides advanced traffic and congestion management features including:

- Support for the full range of ATM service types per ATM Forum TM 4.0 by the BXM-T3/E3, BXM-155, and BXM-622 cards on the Cisco BPX 8600 series broadband switch.
- FairShare, dedicated queue and rate controlled servers for each VPC/VCC at the network ingress.
- OptiClass guarantees Quality of Service (QoS) for individual connections by providing up to 16 queues with independent service algorithms for each trunk in the network.
- AutoRoute, end-to-end connection management that automatically selects the optimum connection path based upon the state of the network and assures fast automatic alternate routing in the event of intermediate trunk or node failures.
- PNNI, a standards based routing protocol for ATM and Frame Relay switched virtual circuits (SVCs).
- FBTC (Frame Based Traffic Control) for AAL5 connections, including early and partial frame discard.
- ForeSight, an end-to-end closed loop rate based congestion control algorithm that dynamically adjusts the service rate of VC queues based on network congestion feedback.
- ABR Standard with VSVD congestion control using RM cells and supported by BXM cards on the Cisco BPX 8600 series broadband switch.

FairShareTM

FairShare provides per-VC queueing and per-VC scheduling. FairShare provides fairness between connections and firewalls between connections. Firewalls prevent a single non-compliant connection from affecting the QoS of compliant connections. The non-compliant connection simply overflows its own buffer.

The cells received by a port are not automatically transmitted by that port out to the network trunks at the port access rate. Each VC is assigned its own ingress queue that buffers the connection at the entry to the network. With ABR with VSVD or with ForeSight, the service rate can be adjusted up and down depending on network congestion.

Network queues buffer the data at the trunk interfaces throughout the network according to the connections class of service. Service classes are defined by standards-based QoS. Classes can consist of the four broad service classes defined in the ATM standards as well as multiple sub-classes to each of the four general classes. Classes can range from constant bit rate services with minimal cell delay variation to variable bit rates with less stringent cell delay.

When cells are received from the network for transmission out a port, egress queues at that port provide additional buffering based on the service class of the connection.

OptiClass™

OptiClass provides a simple but effective means of managing the quality of service defined for various types of traffic. It permits network operators to segregate traffic to provide more control over the way that network capacity is divided among users. This is especially important when there are multiple user services on one network.

Rather than limiting the user to the four broad classes of service initially defined by the ATM standards committees, OptiClass can provide up to 16 classes of service (service subclasses) that can be further defined by the user and assigned to connections. Some of the COS parameters that may be assigned include:

- Minimum bandwidth guarantee per subclass to assure that one type of traffic will not be preempted by another.
- Maximum bandwidth ceiling to limit the percentage of the total network bandwidth that any one class can utilize.
- Queue depths to limit the delay.
- Discard threshold per subclass.

These class of service parameters are based on the standards-based Quality of Service parameters and are software programmable by the user. The Cisco BPX 8600 series broadband switch provides separate queues for each traffic class.

AutoRoute

With AutoRoute, connections in Cisco WAN Switching cell relay networks are added if there is sufficient bandwidth across the network and are automatically routed when they are added. The user only needs to enter the endpoints of the connection at one end of the connection and the Cisco IPX narrowband switch, the Cisco IGX 8400 series multiband switch, and the Cisco BPX 8600 series broadband switch software automatically sets up a route based on a sophisticated routing algorithm. This feature is called *AutoRoute*. It is a standard feature on all Cisco WAN Switching nodes.

System software automatically sets up the most direct route after considering the network topology and status, the amount of spare bandwidth on each trunk, as well as any routing restrictions entered by the user (e.g. avoid satellite links). This avoids having to manually enter a routing table at each node in the network. AutoRoute simplifies adding connections, speeds rerouting around network failures, and provides higher connection reliability.

PNNI

The Private Network to Network Interface (PNNI) protocol provides a standards-based dynamic routing protocol for ATM and frame relay switched virtual circuits (SVCs). PNNI is an ATM-Forum-defined interface and routing protocol which is responsive to changes in network resources, availability, and will scale to large networks. PNNI is available on the Cisco BPX 8600 series broadband switch when an Extended Services Processor (ESP) is installed. For further information about PNNI and the ESP, refer to the *Cisco WAN Switching BPX Service Node Extended Services Processor Installation and Operation* document.

Congestion Management, VS/VD

The Cisco BPX 8600 series broadband switch/Cisco IGX 8400 series multiband switch/Cisco IPX narrowband switch networks provide a choice of two dynamic rate based congestion control methods, ABR with VS/VD and ForeSight. This section describes Standard ABR with VSVD.



Note

ABR with VSVD is an optional feature that must be purchased and enabled on a single node for the entire network.

When an ATM connection is configured for Standard ABR with VSVD per ATM Forum TM 4.0, RM (Resource Management) cells are used to carry congestion control feedback information back to the connection's source from the connection's destination.

The ABR sources periodically interleave RM cells into the data they are transmitting. These RM cells are called forward RM cells because they travel in the same direction as the data. At the destination these cells are turned around and sent back to the source as Backward RM cells.

The RM cells contain fields to increase or decrease the rate (the CI and NI fields) or set it at a particular value (the explicit rate ER field). The intervening switches may adjust these fields according to network conditions. When the source receives an RM cell it must adjust its rate in response to the setting of these fields.

When spare capacity exists within the network, ABR with VSVD permits the extra bandwidth to be allocated to active virtual circuits.

Congestion Management, ForeSight

The Cisco BPX 8600 series broadband switch/Cisco IGX 8400 series multiband switch/Cisco IPX narrowband switch networks provide a choice of two dynamic rate based congestion control methods, ABR with VS/VD and ForeSight. This section describes ForeSight.



Note

ForeSight is an optional feature that must be purchased and enabled on a single node for the entire network.

ForeSight may be used for congestion control across Cisco BPX 8600 series broadband switches/Cisco IGX 8400 series multiband switches/Cisco IPX narrowband switches for connections that have one or both end points terminating on other than BXM cards, for example ASI cards. The ForeSight feature is a dynamic closed-loop, rate-based, congestion management feature that yields bandwidth savings compared to non-ForeSight equipped trunks when transmitting bursty data across cell-based networks.

ForeSight permits users to burst above their committed information rate for extended periods of time when there is unused network bandwidth available. This enables users to maximize the use of network bandwidth while offering superior congestion avoidance by actively monitoring the state of shared trunks carrying frame relay traffic within the network.

ForeSight monitors each path in the forward direction to detect any point where congestion may occur and returns the information back to the entry to the network. When spare capacity exists with the network, ForeSight permits the extra bandwidth to be allocated to active virtual circuits. Each PVC is treated fairly by allocating the extra bandwidth based on each PVC's committed bandwidth parameter.

Conversely, if the network reaches full utilization, ForeSight detects this and quickly acts to reduce the extra bandwidth allocated to the active PVCs. ForeSight reacts quickly to network loading in order to prevent dropped packets. Periodically, each node automatically measures the delay experienced along a frame relay PVC. This delay factor is used in calculating the ForeSight algorithm.

With basic frame relay service, only a single rate parameter can be specified for each PVC. With ForeSight, the virtual circuit rate can be specified based on a minimum, maximum, and initial transmission rate for more flexibility in defining the frame relay circuits.

ForeSight provides effective congestion management for PVC's traversing broadband ATM as well. ForeSight operates at the cell-relay level that lies below the frame relay services provided by the Cisco IPX narrowband switch. With the queue sizes utilized in the Cisco BPX 8600 series broadband switch, the bandwidth savings is approximately the same as experienced with lower speed trunks. When the cost of these lines is considered, the savings offered by ForeSight can be significant.

ELMI

ELMI is an enhancement to LMI. ELMI adds capabilities that are not currently supported in LMI so that network switches, e.g., Cisco BPX 8600 series broadband switches, Cisco IGX 8400 series multiband switches, etc., can inform a user (routers, bridges, etc.) about network parameters such as various quality of service (QoS) parameters. Depending on the implementation, these might be such parameters as Committed Information Rate (CIR), Committed Burst Size (Bc), Excess Burst Size (Be), maximum Frame Size, etc. Currently, ELMI support includes the UFM-U and UFM-C cards on the Cisco IGX 8400 series multiband switch.

Cell Relay Networking, ATM and FastPacket

Cell relay technology is also referred to as cell switching, FastPacket, or asynchronous transfer mode (ATM). Cell switching is just another name for cell relay technology. The Cisco WAN switching FastPacket technology uses fixed 24-byte length cells. ATM uses a standards based 53-byte cell. All these terms describe a switching and multiplexing technique in which user data is placed into fixed length cells that are routed to their destination without regard to content.

- ATM is the standards based 53-byte cell switching technology.
- FastPacket is Cisco's proprietary narrowband 24-byte cell switching technology.
- Cell Relay or Cell Switching both refer to switching packages (cells) of fixed length.

Cell relay communications networks use high-speed digital trunks to link network nodes which provide customer access and network routing functions. Cell relay networks are characterized by very high throughput, short delays, and very low error rates. These networks provide highly reliable transport services to the user without the overhead associated with extensive error control implementations.

There are currently three basic methods employed for transmitting data over digital trunks: the classic time division multiplexing techniques, frame switching, and cell relay.

Cell relay networks utilize small, fixed-length data packets, called cells, that contain an address identifying the network connection and a payload. The use of a common packet format for the transport of all network traffic results in simplified routing and multiplexing.

Unlike Time Division Multiplexing technology used in previous systems, cell relay technology only uses network bandwidth when there is information to send. Connections are established in the network configuration but do not generate cells when idle or there is no data to be sent. These connection types are referred to as Permanent Virtual Circuits (PVCs). Once set up, they are permanent but they are virtual as they do not use any network bandwidth unless there is information to transmit.

When PVCs need to transmit information, the data is segmented into cells and the cells are then assigned to the proper trunk that has bandwidth available. As a result, cell relay networks use about half the bandwidth of TDM networks for most voice and data applications. The net result is a statistical sharing of the network trunk facilities that effectively increases the amount of traffic a network can accommodate.

Cell relay networks are especially useful for LAN to LAN interconnects. LAN data tends to occur in bursts with periods of inactivity in between the bursts. Cell relay connections provide bandwidth on demand for these bursty data applications and can dynamically allocate unused bandwidth from idle connections to active connections.

Cells typically consist of a short header with a destination address and a payload for carrying the user data. Cell length can be either fixed or variable depending on the system type. Cell buffers are employed to temporarily store data to allow for processing and routing.

Asynchronous Transfer Mode (ATM)

Asynchronous Transfer Mode is a well-defined standard for broadband, cell-switching networks. It offers the ability to intermix various types of traffic and dynamic bandwidth allocation to maximize the utilization of network bandwidth. Traffic type may be intermixed from cell to cell and all cells have relatively small variability in the end-to-end delay.

ATM is a connection-oriented network protocol using small, fixed-length cells for carrying data. The small cell size minimizes the delay in building the cell and other queuing delays in transmitting the information across the network. The fixed-length cell size simplifies processing, supports higher switching speeds, and minimizes the uncertainty in delays experienced by variable cell/frame length of some common LAN protocols.

ATM traffic is carried in fixed length (53-byte) cells at high speeds (typically DS3 or E3, OC3/STM-1, OC12/STM-4 and above). The cell size was chosen as a compromise between a small cell with short delay, better for voice quality, and a larger cell size with a better ratio of data to overhead, best for data transmission.

The asynchronous aspect of ATM refers to the fact that data is transmitted only when there is actual information to be sent unlike synchronous transfer modes, such as TDM, where data is continuously being sent, even when it is an idle code. This leads to better utilization of network bandwidth.

The ATM protocol provides a clearly defined delineation between the transport layer and the application layer. Since the ATM protocol is independent of the transmission speed of the connection, it simplifies the network data processing requirements and facilitates scaling of the transmission facilities to accommodate the needs of each individual user while providing an economical growth path as demands on the network increase.

Unlike most LAN protocols, ATM is connection-oriented. Before data transfer can occur, an end-to-end connection must be established. Once a connection is defined, ATM cells are self-routing in that each cell contains a header with an address that indicates the destination. This saves processing time at each intermediate node since the routing is pre-determined.

Each ATM cell header contains two address fields, a Virtual Path Identifier (VPI) and a Virtual Connection Identifier (VCI). These two fields serve to identify each connection across a single link in the network. ATM switches use either the VPI alone or the VPI and VCI fields to switch cells from the input port to an output port at each network node.

PVCs

Permanent Virtual Circuits are connections which, after being added to the network, remain relatively static. PVCs are generally defined by the network operator. While no bandwidth is allocated on a link to a PVC when setting up the connection, the operator informs the network about the characteristics of the desired circuit and sufficient network bandwidth is reserved.

SVCs

Switched Virtual Circuits (SVCs), on the other hand, are dynamic in that they are established by the user on-demand. Enhancements of the ATM standards include provisions for SVCs. SVCs can provide dial-up capabilities controlled by the user. The Cisco BPX 8600 series broadband switch with the ESP installed provides SVC capability for ATM and frame relay connections.

UNI/NNI Interface

Two network interfaces are defined in ATM standards, a User to Network Interface (UNI) and a Network to Network Interface (NNI). The UNI is any interface between a user device and an ATM network. This interface terminates on an ATM switch. The UNI is used to send messages from the network to the user device on the status of the circuit and rate control information to prevent network congestion. NNI is used at the boundary between two different ATM networks e.g. a private network and a public network. Information passing across a NNI is related to circuit routing and status of the circuit in the adjacent network.

Cisco WAN Switching FastPackets

Unlike early X.25 networks, which used low bit-rate facilities, the Cisco IPX narrowband switch FastPacket networks were designed to take advantage of the widely-available, high-speed T1 and E1 networks. The fixed length, 24-byte Cisco IPX narrowband switch FastPacket cell length was specifically designed to fit into the 192-bit payload of a standard T1 frame.

Because a short, fixed-length cell is used to carry information, delay through the network is held to a minimum permitting delay-sensitive applications, such as digitized voice and SNA to be successfully carried by FastPacket networks in addition to LAN traffic.

FastPacket networks depended on digital transmission facilities that transmit with very low error rate. Taking advantage of this, only minimal error checking is performed and only at the FastPacket destination rather than at all intermediate network nodes. Using a simplified protocol for transmitting data, FastPacket networks are able to utilize hardware-based switch fabrics, resulting in very high switching speeds.

As a result, FastPacket networks have very high throughput and low delays. They can be used for all kinds of communication traffic: voice, synchronous data and video, as well as the low speed data applications that are being serviced by conventional packet networks to date. FastPacket transmits all information across the digital trunk in a single packet format, including voice, data, video, and signaling, and all packets are transported through the network using common switching, queuing, and transmission techniques, no matter what the connection type or its bandwidth requirements.

FastPacket networks clearly demonstrate the advantages of cell-relay technology. Currently available wide area information networks are being built using digital trunks with bandwidths up to approximately 2 Mbps. But rapid growth of Local Area Networks, linking communities of personal computers and workstations with distributed databases, has placed increased bandwidth requirements on these wide area networks.

The migration path is to broadband networks, employing higher speed trunks, from T3 to OC3 rates and beyond, to satisfy this increased demand for bandwidth.

As bandwidth demands increase in a network, the Cisco IPX narrowband switch and the Cisco IGX 8400 series multiband switch may be connected to ATM networks, and, by converting between similar cell protocols (FastPackets and ATM cells), take advantage of the high speed, flexibility, and scalability offered by ATM.

ATM Product Family Overview

The Cisco WAN switching product family includes: the Cisco BPX 8600 series broadband switch which can include the Extended Services Processor (ESP) for SVC switching of ATM and frame relay connections, the Cisco IPX narrowband switches, and the Cisco IGX 8400 series multiband switches, Cisco MGX 8220 edge concentrators configured as interface shelves, Cisco IGX 8400 series multiband switches configured as interface shelves, and Cisco IPX narrowband switches configured as interface shelves, StrataSphere network management products, and Access Products including the Cisco MGX 8220 edge concentrator, the Cisco 3810, and FastPAD.

Cisco's WAN switching systems are flexible, modular, cell-based platforms that support network requirements ranging from a few voice and data connections, frame relay connections, on up to a multi-service ATM network with thousands of users.

A common software architecture and baseline ensures full interoperability within the Cisco WAN switching product line including the Cisco IPX narrowband switch, the Cisco IGX 8400 series multiband switch, and the Cisco BPX 8600 series broadband switch.

StrataSphere, Standards-Based Network Management

Conforming to the Network Management Forum's advanced management framework for integrated service management and press automation, StrataSphere is a standards-based multi-protocol management architecture. StrataSphere combines embedded management intelligence distributed throughout the network elements (for fast implementation) along with centrally located NMS workstation advanced system applications and tools to provide integrated fault, performance, and configuration management functions unique to cell-based networks.

- Cisco WAN Manager
- StrataSphere BILLder
- StrataSphere Modeler
- StrataSphere Adaptor
- StrataSphere Agent

StrataSphere automates key network management processes such as service provisioning, billing, statistics collection, and network modeling and optimization. StrataSphere provides a high-volume standards-based usage billing solution for emerging services, such as ATM, as well as cell-based frame

relay services, and supports the high level of statistics collection—up to one million statistics per hour per billing station required in a next-generation ATM product. Additional information on StrataSphere Network Management is provided in Chape r3, CiscoWAN Manager Network Management.

System Switch, ESP, Edge Concentrator, and Network Access Products Description

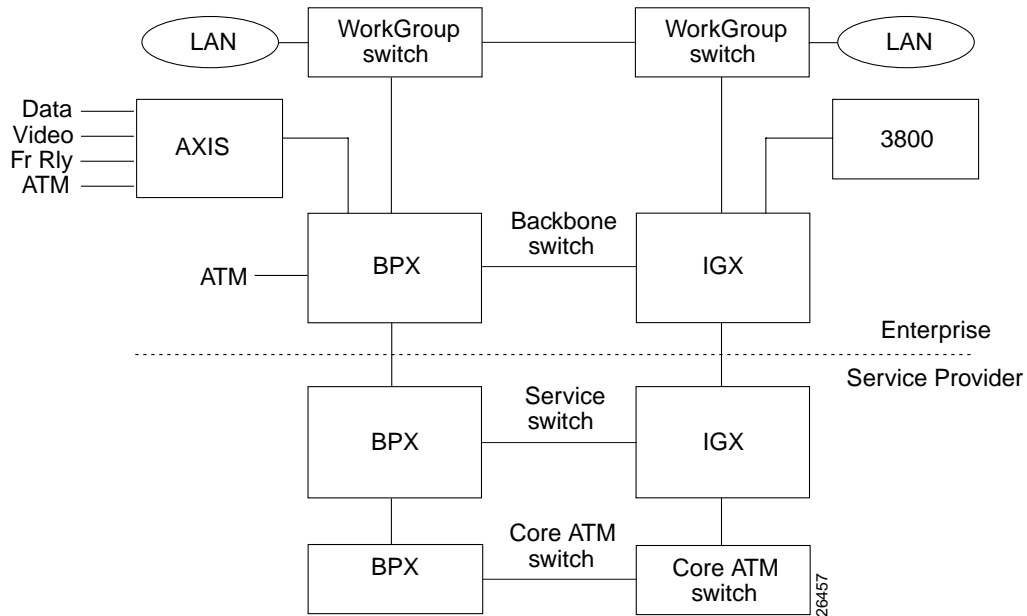
The Cisco BPX 8600 Series Broadband Switc

The Cisco BPX 8600 series broadband switch is a standards based high-capacity, 19.2 Gbps, broadband ATM switch that with the co-located Cisco MGX 8220 edge concentrator and ESP provides backbone ATM switching and delivers a wide range of user services (Figure 1-9). Fully integrated with the Cisco IPX narrowband switch and the Cisco IGX 8400 series multiband switch, the Cisco BPX 8600 series broadband switch is a scalable, standards-compliant unit. Using a multi-shelf architecture, the Cisco BPX 8600 series broadband switch supports both narrowband and broadband user services. The modular, multi-shelf architecture enables users to incrementally expand the capacity of the system as needed.

The Cisco MGX 8220 edge concentrator configured as an interface shelf supports a wide range of narrowband interfaces. It converts all non-ATM traffic into 53-byte ATM cells and concentrates thi traffic for high speed switching by the Cisco BPX 8600 series broadband switch.

Similarly, the Cisco IPX narrowband switches or Cisco IGX 8400 series multiband switches may be configured as shelves and connected to a Cisco BPX 8600 series broadband switch configured as a routing hub to provide a low-cost service input for frame relay to ATM interworking for the Cisco BPX 8600 series broadband switch. The Cisco IPX narrowband switches and Cisco IGX 8400 series multiband switches configured as interface shelves concentrate this traffic over an ATM trun connected to the Cisco BPX 8600 series broadband switch.

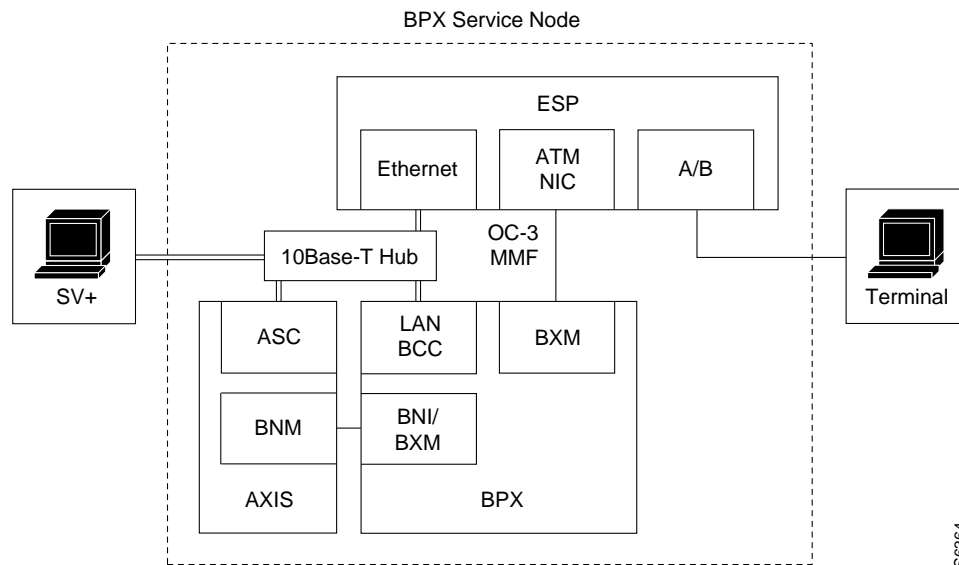
Figur e1-8 Cisco BPX 8600 Series Broadband Switch Configuration



The ESP

The Extended Services Processor (ESP) is an adjunct shelf co-located with the Cisco BPX 8600 series broadband switch (Figure 1-9). Typically, a Cisco MGX 8220 Edge Concentrator is also co-located with the Cisco BPX 8600 series broadband switch to provide ATM and frame relay switched virtual circuits (SVCs). For further information about the ESP, refer to the *Cisco WAN Switching BPX Service Node Extended Services Processor Installation and Operation* document.

Figur e1-9 Cisco BPX 8600 Series Broadband Switch with Co-Located ESP and Cisco MGX 8220 Edge Concentrator



The Cisco MGX 8220 Edge Concentrator

The Cisco MGX 8220 edge concentrator is a standards-based ATM interface shelf that provides a low-cost service interface to multi-service networks and is usually co-located with a Cisco BPX 8600 series broadband switch (Figure 1-9). The Cisco MGX 8220 edge concentrator provides a broad range of narrowband user interfaces. Release 4 of the Cisco MGX 8220 edge concentrator provides T1/E1 and subrate frame relay, FUNI (Frame Based UNI over ATM), T1/E1 ATM, T1/E1 Circuit Emulation Service (CES), frame relay to ATM network and service interworking for traffic over the ATM network via the Cisco BPX 8600 series broadband switch, HSSI and X.21 interfaces and SRM-3T3 enhancements. The Cisco MGX 8220 edge concentrator allows users to concentrate large numbers of PVC connections over high-speed ATM trunks. For further information, refer to the *Cisco MGX 8220 Edge Concentrator Installation and Configuration* and the *Cisco MGX 8220 Edge Concentrator Command Reference* documents. In conjunction with the Cisco BPX 8600 series broadband switch and the ESP, the Cisco MGX 8220 edge concentrator also supports ATM and frame relay switched virtual circuits (SVCs).

The Cisco IGX 8400 Series Multiband Switc

The Cisco IGX 8400 series multiband switch is a standards based, 1.2 Gbps, highly-scalable ATM switch that provides interfaces to support current legacy and emerging broadband applications (Figure 1-10). The Cisco IGX 8400 series multiband switch is designed for use in Wide Area Networks (WANs) public or private, using subrate, fractional T1, E1, T3, or E3 transmission facilities. The Cisco IGX 8400 series multiband switch supports multiservice traffic including voice, data, Frame Relay, and ATM T3/E3 connections. The Cisco IGX 8400 series multiband switch currently supports CBR and VBR ATM connection types. The Frame Relay connection interfaces include T1/E1 channelized and unchannelized, HSI, X.21, and V.35.

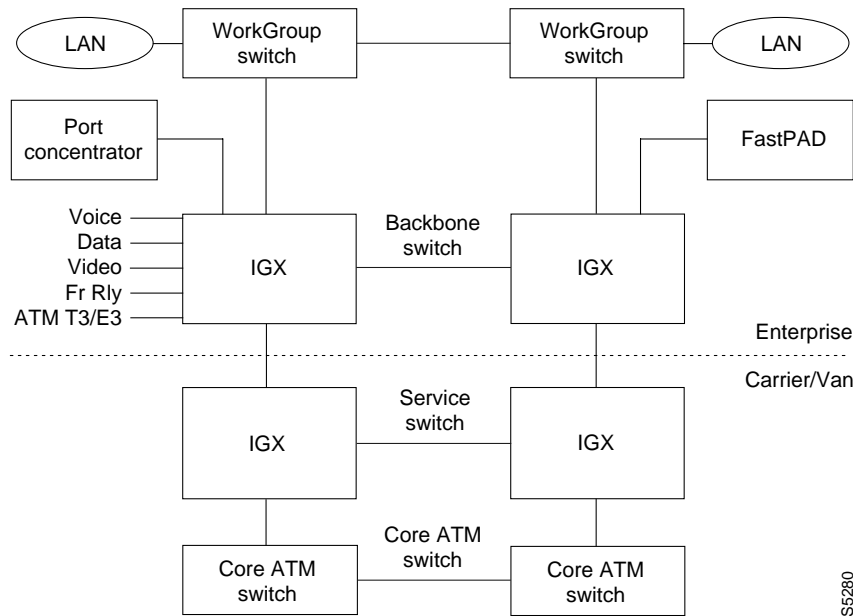
The Cisco IGX 8400 series multiband switch provides the capability to migrate wide area traffic from older time-division multiplexed (TCM) networks to more efficient, robust, and higher-speed ATM networks. The Cisco IGX 8400 series multiband switch is fully compatible with the Cisco IPX narrowband switch and the Cisco BPX 8600 series broadband switch.

The Cisco IGX 8400 series multiband switch uses standard ATM, together with efficient cell adaption to provide seamless conductivity across multiband networks using trunks ranging from 128 kbps to 155 Mbps. The Cisco IGX 8400 series multiband switch supports FastPAD access trunk connectivity from 9.6 Kbps up to 2 Mbps. The Cisco IGX 8400 series multiband switch also supports the Cisco IPX narrowband switch's narrowband cell relay protocol.

ELMI is used to inform a user (routers, bridges, etc.) about network parameters such as various quality of service (QoS) parameters that may be considered in user congestion control actions. The UFM-U and UFM-C frame relay cards support ELMI.

CAS switching enables the UVM/VNS to switch calls coming from CAS/DTMF-based PBX switches. The CAS signaling and DTMF tones are translated by the UVM into CCS call control messages which are processed by the VNS. The UVM card supports CAS switching.

Figur e1-10 A Cisco IGX 8400 Series Multiband Switch Configuration

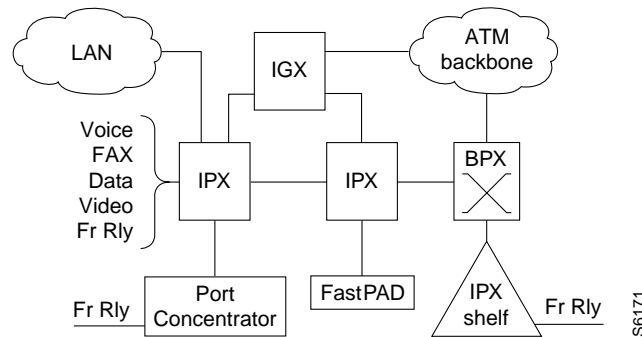


The Cisco IPX Narrowband Switc

The Cisco IPX narrowband switch is a narrowband cell switch (Figure 1-11) that accepts frame relay, digitized voice and FAX, encoded video, data streams, etc., and adapts these information streams into fixed length cells (Figure 1-11). These cells are then routed to appropriate network interfaces, either ATM or FastPacket. The Cisco IPX narrowband switch supports frame relay to ATM network interworking, which provides the advantages of transporting frame relay traffic across a high-speed ATM network.

The Cisco IPX narrowband switch is compatible with Cisco IGX 8400 series multiband switches and Cisco BPX 8600 series broadband switches, and supports FairShare, Opti-Class, frame relay to ATM Network Interworking-protocol translation, VAD, RPS, ABR with VS/VD and ForeSight congestion management, etc.

Figur e1-11 A Cisco IPX Narrowband Switch Configuration



Network Access Products

These products, located at the outer edges of a network, offer several substrate, narrowband, and broadband configurations such as multiplexers, frame relay access devices-FRADs, and routers with a wide range of interface options. They enable users to convert legacy and lower-speed traffic into fixed-length frames or cells for both narrowband and broadband switching.

FastPADs

The multi-media FastPADs are an OEM product that provides voice and data integration and Frame Relay switching over Frame Relay or leased line networks. A FastPAD connects to a Cisco IPX narrowband switch or a frame relay port on a Cisco IGX 8400 series multiband switch and is managed by Cisco WAN Manager. The multi-protocol FastPADs are OEM products that provide legacy protocol support.

Cisco 3810 Serie

The Cisco 3810 provides multi-service integration of voice, fax, video, legacy data, and LAN traffic over either a Frame Relay or ATM trunk. The 3810 typically connects to a frame relay port on a Cisco IGX 8400 series multiband switch configured as a routing node or to another 3810 over leased line or public frame relay. Cisco WAN Manager supports provisioning and management of the 3810.

INS-VNS and INS-DA

The Intelligent Network Server products, INS-VNS and INS-DAS use a robust high-powered processing platform to add several important capabilities to Cisco WAN switching networks. In addition, the INS-VNS and INS-DAS support some form of standards-based signaling between customer premise equipment (CPE) and a Cisco BPX 8600 series broadband switch, Cisco IGX 8400 series multiband switch, Cisco IPX narrowband switch network. Typically this signaling is a variation of common-channel, message-oriented Integrated Services Digital Network (ISDN) or Broadband ISDN (B-ISDN) signaling protocols. The INS-VNS and INS-DAS applications both interpret these

industry-standard signaling messages, translates the logical addresses to the appropriate physical endpoints of the network, and instructs the Cisco IGX 8400 series multiband switches and Cisco IPX narrowband switches to establish the connection required for the particular application. The switches then take over, dynamically establishing the optimum route through the network and maintaining the connection for its duration.

The two primary INS applications are:

- **DAS (Dial-Up Frame Relay) which is supported by INS-DAS**, provides semi-permanent virtual circuits (or soft PVCs) that can be used for dial-up and dial-backup services. A soft PVC is preconfigured in the network's data base but remains dormant until a call from the CPE signals the network to activate it. For information about dial-up frame relay, refer to the *INS Dial Access Server Installation* document and the *INS Dial-Up Frame Relay Operator* document.
- **VNS (Voice Network Switching)** which is supported by INS-VNS, provides voice switched virtual circuits (SVCs) across a Cisco WAN switching network for PBXes using Digital Private Network Signaling System (DPNSS) and QSIG Signaling. With Release 8.4, the name DNS was changed to Voice Network Switching (VNS). For additional information on VNS Release 2.0, including the break-out/break-in feature, which is associated with Switch Software Release 8.4, refer to the *Voice Network Switching Installation and Operation* document.

Each INS application uses one or more adjunct processors that are co-located with a node (that is a Cisco BPX 8600 series broadband switch/Cisco MGX 8220 edge concentrator, a Cisco IPX narrowband switch, or a Cisco IGX 8400 series multiband switch) and often installed in the same equipment rack. Available in either AC- or DC-powered models, the base INS processor is a scalable UNIX platform and contains:

- 140 MIPS CPU, with a 71 Mhz clock
- 64 Megabytes of memory
- 1 Gigabyte hard disk.

The base INS-VNS and INS-DAS processors are equipped with different interface modules, memory and disk configurations, and different application software.

For further information, refer to "Intelligent Network Server DAS and VNS" in Chapter 4, Network Services Overview, and to the *INS documents*.

System Software Description

The Cisco WAN switching cell relay system software shares most core system software, as well as a library of applications, between platforms. System software provides basic management and control capabilities to each node.

Cisco IPX narrowband switch, Cisco IGX 8400 series multiband switch, and Cisco BPX 8600 series broadband switch system software manages its own configuration, fault-isolation, failure recovery, and other resources. Since no remote resources are involved, this ensures rapid response to local problems. This distributed network control, rather than centralized control, provides increased reliability.

Software among multiple nodes cooperates to perform network-wide functions such as trunk and connection management. This multi-processor approach ensures rapid response with no single point of failure. System software applications provide advanced features that may be installed and configured as required by the user.

Some of the many software features are:

- Automatic routing of connections (AutoRoute feature).

- Various classes of service that may be assigned to each connection type (OptiClass feature).
- Bandwidth reservation on a time-of-day basis.
- Detection and control of network congestion with ABR with VSVD or ForeSight algorithms.
- Automatic self-testing of each component of the node.
- Automatic collecting and reporting of many network-wide statistics, such as trunk loading, connection usage, and trunk error rates, as specified by the user.

The system software, configuration database, and the firmware that controls the operation of each card type is resident in programmable memory and can be stored off-line in the StrataView Plus NMS for immediate backup if necessary. This software and firmware is easily updated remotely from a central site or from Cisco Customer Service, which reduces the likelihood of early obsolescence.

Connections and Connection Routing

The routing software supports the establishment, removal and rerouting of end-to-end channel connections. There are three modes:

- Automatic Routing—the system software computes the best route for a connection.
- Manual Routing—the user can specify the route for a connection.
- Alternate Routing—the system software automatically reroutes a failed connection.

The system software uses the following criteria when it establishes an automatic route for a connection:

- Selects the most direct route between two nodes.
- Selects unloaded lines that can handle the increased traffic of additional connections.
- Takes into consideration user-configured connection restrictions (for example whether or not the connection is restricted to terrestrial lines or can include satellite hops or routes configured for route diversity).

When a node reroutes a connection, it uses these criteria and also looks at the priority that has been assigned and any user-configured routing restrictions. The node analyzes trunk loading to determine the number of cells or packets the network can successfully deliver. Within these loading limits, the node can calculate the maximum combination allowed on a network trunk of each type of connection: synchronous data, ATM traffic, frame relay data, multi-media data, voice, and compressed voice.

Network-wide T3, E3, OC3, or OC12 connections are supported between Cisco BPX 8600 series broadband switches terminating ATM user devices on the Cisco BPX 8600 series broadband switch UNI ports. These connections are routed using the virtual path and/or virtual circuit addressing fields in the ATM cell header.

Narrowband connections, terminating on Cisco IPX narrowband switches, can be routed over high-speed ATM backbone networks built on Cisco BPX 8600 series broadband switches. FastPacket addresses are translated into ATM cell addresses that are then used to route the connections between Cisco BPX 8600 series broadband switches, and to ATM networks with mixed vendor ATM switches. Routing algorithms select broadband links only, avoiding narrowband nodes that could create a choke point.

Connection Routing Groups

The re-routing mechanism is enhanced so that connections are presorted in order of cell loading when they are added. Re-routing takes place by rerouting the group containing the connections with the largest cell loadings first on down to the last group which contains the connections with the smallest cell loadings. These groups are referred to as routing groups. Each routing group contains connections with loading in a particular range,

There are three configurable parameters for configuring the rerouting groups,

- total number of rerouting groups
- starting load size of first group
- load size range of each group

The three routing group parameters are configured with the **cnfcmparm** command.

For example, there might be 10 groups, with the starting load size of the first group at 50, and the incremental load size of each succeeding group being 10 cells. Then group 0 would contain all connections requiring 0-59 cell load units, group 1 would contain all connections requiring from 60-69 cell load units, on up through group 9 which would contain all connections requiring 140 or more cell load units.

Table 1-1 Routing Group Configuration Example

Routing gro	Connection cell loadi
0	0-59
1	60-69
2	70-79
3	80-89
4	90-99
5	101-109
6	110-119
7	120-129
8	130-139
9	140 and up

Network Synchronization

Cisco WAN switching cell relay networks use a fault-tolerant network synchronization method of the type recommended for Integrated Services Digital Network (ISDN). Any circuit line, trunk, or an external clock input can be selected to provide a primary network clock. Any line can be configured as a secondary clock source in the event that the primary clock source fails.

All nodes are equipped with a redundant, high-stability internal oscillator that meets Stratum 3 (Cisco BPX 8600 series broadband switch) or Stratum 4 requirements. Each node keeps a map of the network's clocking hierarchy. The network clock source is automatically switched in the event of failure of a clock source.

There is less likelihood of a loss of customer data resulting from re-frames that occur during a clock switchover or other momentary disruption of network clocking with cell-based networks than there is with traditional TDM networks. Data is held in buffers and packets are not sent until a trunk has regained frame synchronism to prevent loss of data.

Network Availability

Hardware and software components are designed to provide a node availability in excess of 99.99%. Network availability will be much more impacted by link failure, which has a higher probability of occurrence, than equipment failure.

Because of this, Cisco WAN switching switches are designed so that connections are automatically rerouted around network trunk failures often before users detect a problem. System faults are detected and corrective action taken often before they become service affecting. The following paragraphs describe some of the features that contribute to network availability.

System Diagnostics

Each node within the network runs continuous background diagnostics to verify the proper operation of all network trunks, active and standby cards, buses, switch paths, cabinet temperature, and power supplies. This background process is transparent to normal network operation.

Failures that affect system service are reported as major alarms. Failures that could affect service later (such as a failure of a standby card) are reported as minor alarms. For example, the following lists some of the failures that will generate an alarm:

- Failed network trunk.
- Failed system bus.
- Failed circuit line.
- Failed power supply.
- Failed or missing card.

When an alarm occurs in a network and a trunk or circuit line is in alarm, a loopback test for all cards associated with the failed trunk or line is automatically performed to verify proper operation of the node hardware.

Alarm Reporting

The Cisco BPX 8600 series broadband switch, the Cisco IGX 8400 series multiband switch, the Cisco IPX narrowband switch, and the Cisco MGX 8220 edge concentrator provide both software alarms displayed on operator terminal screens and hardware alarm indicators. The hardware alarms are LED indicators located on the various cards in the node. All cards have LEDs that indicate whether or not a fault is detected on the card. On interface cards, LEDs indicate whether or not there is a local or remote line failure.

Each power supply has front panel indications of proper output. Since the power supplies share the power load, redundant supplies are not idle. All power supplies are active. If one fails, then the others pick up its load. The power supply outputs are monitored as well as the cabinet internal temperature.

Statistical Alarms

The network manager can configure alarm thresholds on a per-trunk basis for any transmission problems that are statistical. Thresholds are configurable for a number of alarm types including frame slips, out of frames, bipolar errors, dropped packets, and packet errors. When an alarm threshold is exceeded, the screen displays an alarm message.

Failure Recovery

In a redundant system, if a hardware failure occurs, a redundant module is automatically switched into service, replacing the failed module. These systems provide redundant common control buses, redundant power supplies, and redundant cards. All cards have 1+1 redundancy option, which provides each card with a dedicated standby. If an active module fails, the system automatically switches to a standby.

Channel connections on a failed trunk are automatically rerouted to a different trunk if one is available. Rerouting time is a function of the complexity of the network, but normally the first will be completed within milliseconds and the last within several seconds.

Standards

The performance of the systems are compatible with most recent recommendations from various international standards committees and forums to assure seamless interworking with other network equipment. The following is a partial list of these standards committees:

- ATM Forum
- Frame Relay Forum
- ITU-T (CCITT)
- ANSI
- Bellcore

BPX, IGX, and IPX Architecture

The following paragraphs describe the Cisco WAN switches (nodes) used to build cell-based Wide Area Networks: the ATM broadband BPX Service Node, the multiband IGX, and the narrowband IPX. For further information, refer to the BPX, IGX, and IPX reference manuals. Also, for further information on associated products, refer to the *Release 4.1 Cisco MGX 8220 Installation and Configuration* and the *Release 4.1 Cisco MGX 8220 Command Reference* publications, and to the *FastPAD Reference* publication. For Network Management Information, refer to the *Cisco WAN Manager Operations* publication.

This chapter includes the following:

- BPX Service Node (broadband ATM)
- IGX (multi-band ATM)
- IPX (narrowband ATM and FastPacket)

BPX Service Node (Broadband ATM)

General

The BPX Service Node is a standards-based, broadband, scalable, ATM network switch utilizing very high-speed-switching and asynchronous transfer mode technologies. It supports large, high-capacity networks for existing users whose private networks have grown in size and for public service providers who demand the latest in high-speed technology for their large networks.

BXM card sets provide either service or trunk interfaces. The BXM card sets are compliant with ATM UNI 3.1 and Traffic Management 4.0 standards including ABR VS/VD and provide high scalability with up to 144 T3/E3 ports, up to 96 OC3c/STM-1 ports, or up to 24 OC-12c/STM-4 ATM ports.

The BPX provides the following features:

- Increased node switching capacity to facilitate the deployment of larger, higher-capacity networks.
- BXM card sets may be configured for either trunk or port (service access). The BXM-T3/E3 supports T3 and E3 trunk (network) and port (service access) interfaces. The BXM-155 supports OC3/STM-1 trunk (network) and port (service access) interfaces. The BXM-622 card set supports ATM OC-12/STM-4 trunk (network) and port (service access) interfaces.
- A platform for providing multiple services such as frame relay, ATM, SMDS, CES.
- Increased network capacity by using DS3, E3, OC3/STM-1, or OC-12c/STM-4 trunks to facilitate statistical bandwidth aggregation as frame relay Committed Information Rates increase.

- A clear migration path to networks utilizing the 53-byte cell relay network protocols defined by ATM standards.
- Ability to support higher-speed network access and CPE interfaces such as SONET, and SDH to access rates up to 622 Mbps.
- Advanced congestion management and routing features such as FairShare, ForeSight, AutoRoute, and Opticlass.
- Higher-speed network management interfaces and access for real-time, per-connection statistics collection.

With the BCC-4, the BPX employs a redundant 19.2 Gbps non-blocking crosspoint switch matrix for cell switching. The switch matrix can establish up to 20 million point-to-point connections per second between ports. A single BPX provides twelve card slots, with each card capable of operating at 800 Mbps for ASI and BNI cards. The BXM cards support egress at up to 1600 Mbps and ingress at up to 800 Mbps. Access to and from the crosspoint switch is through multi-port network and user access cards.

An arbiter is utilized to set up and release each cell path based on those ports with traffic data to switch. The arbiter can set up 20 million point-to-point connections per second between the BPX switch ports. For maximum reliability, the switch matrix, data paths, and all control and arbiter functions are provisioned in a fail-safe 1:1 redundant configuration.

Each BPX node is configured with twelve slots (switch ports) for network or user interfaces and three slots for common equipment cards.

The BXM cards may be configured to operate in either port (UNI) mode for connection to CPE or in trunk mode for connection to other BPX nodes or networks. The BXM-T3/E3, BXM-155, and BXM-622 provide T3, E3, OC3/STM-1, and OC12/STM-4 interfaces, respectively.

The Broadband Network Interface (BNI) card provides high-speed DS3, E3, or OC-3/STM-1 ATM network interfaces for trunking to other BPX nodes, to IPX nodes equipped with an ATM trunking cards or to other networks. The BNI high-speed trunk cards can be configured for 1:1 redundancy.

An ATM Service Interface (ASI) card provides two DS3, E3, or OC-3/STM-1 ATM ports for connection to ATM customer premise equipment.

The BPX is a virtual circuit switch. It can switch on VCI and/or VPI fields in the ATM cell header. Individual connections, as well as groups of connections within a path, may be switched. This greatly enhances the versatility of the BPX as a network switch.

Three techniques are used for transmitting user information over ATM cells.

- For pure cell-relay traffic, incoming ATM cells from ATM customer premise equipment are relayed and switched through the network in standards-based cell formats. Cells are routed based on address information in the ATM cell header.
- Frame Relay to ATM Interworking enables frame relay traffic to be connected across high-speed ATM trunks. 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. Two types of 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 Cisco MGX 8220 supports both Network Interworking and Service Interworking. On the Cisco MGX 8220, IPX, or IGX, Frame Relay data is received directly from the user equipment, converted from the Frame Relay protocol to the ATM protocol, and the data streamed into the ATM cell's 48-byte payload, sending ATM cell after cell, until the end of a frame is reached. This is repeated for each frame received from the user until the last frame is transmitted.

- If the traffic is voice or constant bit rate, low-speed data, which will terminate on an IPX or IGX, the 24-byte FastPackets are combined within the 48-byte ATM payload, normally two FastPackets to a cell. If the connection carries data or voice that is very delay sensitive, the Cisco WAN switch may, as an option, occasionally pack only one FastPacket into a cell payload to reduce delay to a minimum. Once data is segmented into cells, the cells can easily be converted from one size to another without requiring additional adaptation simply by changing the cell size and streaming the data into a different capacity payload.

The BPX may be used to provide a smooth migration path for network operators currently using IPX and IGX nodes in their networks who need additional capacity and would like to build on their existing equipment. Since the BPX can be integrated into an existing IPX/IGX network by adding only a single card to each of the IPX or IGX nodes that need to connect to the BPX, the hardware costs are kept to a minimum.

The ATM cell payload is ideally suited to transporting IPX FastPackets since both are fixed length and the size of the ATM cell is larger than the FastPacket cell length (no segmentation of the cell is necessarily required). Packets of data that are created by IPX and IGX nodes in a mixed environment are encapsulated in the ATM cells for transmission throughout the network.

For more details on the BPX hardware, refer to the *BPX Reference*. For information on the BPX commands, refer to the *Command Reference*.

Extended Services Processor

The Extended Services Processor (ESP) is an adjunct shelf that provides ATM and frame relay SVC signaling to the BPX Service Node. The ESP also provides PNNI routing management for the ATM and frame relay SVCs. For further information about the ESP, refer to the *BPX Service Node Extended Service Processor* document.

Cisco MGX 8220 Shelf

The Cisco MGX 8220 is a standards-based ATM interface shelf for the BPX that provides a service interface to multi-service networks and is usually co-located with a BPX. The Cisco MGX 8220 allows users to economically concentrate large numbers of PVC connections over high-speed ATM trunks. Release 4 of the Cisco MGX 8220 provides T1/E1 and subrate frame relay, T1/E1 ATM, and FUNI (frame based UNI over ATM), circuit emulation, inverse ATM multiplexing, and frame relay service interfaces for routing traffic over the ATM network via the BPX. For additional information on the Cisco MGX 8220, refer to the *Cisco MGX 8220 Installation and Configuration* document.

BPX Architecture

The BPX consists of the following elements:

- Common Core (Controller) Modules, which perform the node control and cell switching functions.
- Network Interface Modules, which provide the physical and electrical interface to the network trunks.
- Access Interface Modules, which provide standard interfaces for connecting the BPX directly to user devices.

ATM connections can be made between user devices equipped with DS3, E3, OC3, and OC12ATM interfaces. These user devices connect directly to a BPX switch port over an ATM User-Network Interface on a BXM card (service access mode) or on the ATM Service Interface (ASI) card and through Cisco MGX 8220 or IPX shelves. The BPX supports constant bit rate (CBR), variable bit rate (VBR), available bit rate (ABR), and unspecified bit rate (UBR) ATM connections.

Functions performed by the BXM (trunk mode) and BNI modules include queuing and servicing cells, implementing congestion control mechanisms, address translation and routing of cells to the switching matrix. Physical connection to the network is accomplished through simple line interface cards.

Segmentation and reassembly, cell switching, interfacing to the switch matrix, and system timing and status monitoring are performed by the BPX controller card. A non-blocking, crosspoint switching matrix in the BCC card is the heart of the BPX. With the BCC-3, which employs a 16 x 16 crosspoint switch, the BPX operates at 9.6 Gbps. With the BCC-4, which employs a 16 x 32 crosspoint switch, the BPX can operate at 19.6 Gbps when it is also equipped with BXM cards. A switch arbiter is used to determine which matrix input port needs to connect to which matrix output port to properly route the cells. This individual crosspoint switching technique provides superior performance at broadband speeds when compared with a system bus architecture.

Each BPX switch port has a capacity of operating at 800 Mbps and is capable of supporting OC-12 data rates in each of the twelve switch ports on the BPX

All data flow through the switch is monitored at every process to detect any corruption of the data or addressing. Flow-through parity recomputation and checking is performed within the individual cards and data paths on the backplane to detect data path and memory faults that would otherwise be difficult to detect with periodic checks.

The controller, switch matrix, backplane data paths, and all control and timing signals are redundant throughout the BPX switch to provide maximum reliability. Both Broadband Controller Cards (BCCs) are synchronized to each other and data paths are buffered so any switchover will be transparent.

IGX (Multi-band ATM Switch)

General

Cisco's IGX product family, consists of the IGX 8, IGX 16, and IGX 32 multiservice ATM switches. As multi-band members of the Cisco WAN Switching ATM portfolio, IGXs seamlessly integrate with Cisco's IPX, BPX, Cisco MGX 8220, INS, and FastPAD platforms under StrataSphere management, to provide multiband ATM solutions from the access to the core layer with integrated network management and call processing.

The IGX 16 and IGX 32 are network backbone node systems for large sites with multiple trunks and considerable local traffic requirements, where a large number of physical ports and gigabit-scale throughput are required. Both IGX systems are available as stand-alone units or can be rack mounted with other equipment.

All IGX services are supported by standard ATM narrowband and broadband trunk resources. IGX systems support trunk speeds from 128 Kbps to T3/E3 and FastPAD access trunk connectivity from 9.6 Kbps to 2 Mbps. IGX systems can be fully interconnected in a logical mesh through public ATM services, or provide edge switching in and out of such services, and can network to FastPAD systems over public frame relay services at up to 2 Mbps. IGX systems also support the IPX FastPacket trunk protocol for seamless connectivity with IPXs at smaller sites.

IGX Architecture

The IGX uses a 1.2 Gbps cell switching redundant bus to pass ATM cells between optionally redundant adaptation, trunking and gateway modules within the system. This architecture allows any amount of bandwidth to be assigned to any slot, and makes the IGX the only system in its class with more than 16 slots, for greater scalability.

Hardware, firmware and software are designed for maximum availability, non-stop networking, even during maintenance windows. Availability design features, common to all CiscoWAN switching systems, include:

- 100% Component Redundancy
- Extensive Background Diagnostics
- “Hot” Card Swapping
- Rapid power fail recovery
- Background software download
- Firmware and software upgrades (remote download)
- Class B EMI certified enclosures
- Hard and soft alarm interfaces including “call home”
- Minimum internal cabling

All switches use a mid-plane design with front cards performing processing functions and back cards providing interfacing and physical connectivity. This allows most system maintenance to be performed at the front cards, without disconnecting interface cables.

IPX (Narrowband ATM and FastPacket Switch)

General

The IPX narrowband switch is an intelligent network bandwidth manager. It accepts frame relay, digitized voice and FAX, encoded video, and data streams from user equipment and adapts these voice and data streams into fixed length cells. It then routes these cells to appropriate network trunk interfaces (ATM or FastPacket). The receiving node dis-assembles the received cells or packets and outputs the data streams to the proper voice and data ports. The IPX seamlessly integrates with the IGX, BPX, INS, FastPad, Port Concentrator, and StrataSphere NMS. The IPX routes traffic to appropriate network interfaces, either ATM or FastPacket.

The following are some of the features provided by the IPX narrowband switch:

- Adapts customer frame relay, voice/FAX, synchronous and asynchronous data, and digital video into high-speed, fixed length cells (either ATM cells or FastPackets).
- Transmits data utilizing high-speed digital transmission facilities and simple transmission protocol resulting in very low end-to-end delays.
- Maximizes network bandwidth utilization by compressing voice data, eliminating silent periods, eliminating repetitive patterns and idle codes in data streams, and transmitting packets only when there is actual circuit data to be sent.
- Connections are easily established and automatically routed.

- Users can assign a hierarchy of class of service and priorities for various types of traffic.
- Monitors all trunk activity, detects trunk failures, and automatically reroutes around failed trunks without loss of service (as long as sufficient spare bandwidth is available).
- Provides tools for detecting and dynamically avoiding and/or managing network congestion.
- Collects status and traffic statistics on all trunks, circuits, connections, frame relay ports.
- Display of network configuration and status from any node.

The IPX family consists of different configurations designed to meet the needs of small, medium, and large sites. There are three basic IPX packaging options: the IPX 8, IPX 16, and IPX 32. All three systems use virtually identical hardware and software; they differ only in the number of card slot provided. The IPX 8 extends the benefits of FastPacket technology to the edge of networks without losing bandwidth efficiency.

**Note**

The term IPX is used to refer to the generic system. Where necessary to differentiate between the various systems, the terms IPX 8, IPX 16, or IPX 32 are used.

The IPX also uses voice compression to provide increased capacity. Voice interfaces, such as PABX's and channel banks, transmit Pulse Code Modulation (PCM) voice streams to the IPX for processing. The IPX processes voice streams using Voice Activity Detection (VAD). With VAD, voice packets are transmitted only when speech is present (in typical phone conversations, speech is present only 40% of the time). This results in greater than 2-to-1 compression, with no degradation in voice quality. This compression ratio allows many more voice channels to share a single trunk.

In addition to Pulse Code Modulation (PCM), the IPX provides various levels of Adaptive Differential Pulse Code Modulation (ADPCM) to give another 2, 3, or 4-to-1 compression to voice. In this way, the IPX provides voice compression for as many as 256 (E1) or 192 (T1) voice channels on a single trunk.

The IPX uses Data Frame Multiplexing (DFM) to provide low-speed data compression. DFM is implemented using Repetitive Pattern Suppression (RPS). If data packets contain a repetitive data pattern, the IPX will suppress sending the packets across the network and merely regenerate the pattern at the far end. The net result of this suppression is a significant savings in bandwidth.

Non-disruptive diagnostics continuously monitor the performance of each component in the node. In the event of a failure, a backup component is automatically switched into service with no effect on the user. All cards and power supplies are hot-replaceable; they can be added or replaced on-line with no interruption in service.

IPX systems can be configured to operate in network applications at any world-wide network location. Options include interfaces to both North American and International standard trunks and ports. Power supply options include 208 VAC, 240 VAC, and 48 VDC modules. All cards are programmable, with downloadable firmware, to operate with various network parameters and are easily upgraded to take advantage of new features as they are developed.

For more details on the IPX hardware, refer to the *IPX Reference Manual*. For information on the IPX commands, refer to the *Command Reference Manual*.

IPX Architecture

Information enters an IPX from attached terminal equipment in the form of digital streams, through a circuit mode service interface, or in the form of LAP-D frames through a packet mode service interface. These interfaces segments the streams or frames into cells, addresses the cells, and transmits them onto a switching bus.

Each trunk interface monitors the system bus and receives any cells to be transmitted through it. Queuing of cells for transmission is done by the trunk interface. Cells that are being switched through an intermediate node are transmitted onto the system bus by the receiving trunk interface and picked up by another trunk for transmission out. When cells reach the destination node, they are put on the bus by the receiving trunk, picked up by the appropriate service interface, converted back to data streams or frames and transmitted to the attached user equipment.

All information entering an IPX, whether circuit or packet, streaming or framed, constant or variable bit rate, is converted to a single common multimedia format, a fixed-sized cell. These cells are switched and transmitted to their destination, where the original information is reconstructed and delivered. The benefits of IPX FastPacket networks arise from the power and flexibility of the universal switching architecture.

Cisco WAN Manager Network Management

StrataSphere combines embedded management intelligence distributed throughout the network elements (for fast implementation) along with centrally located NMS workstation advanced system applications and tools to provide integrated fault, performance, and configuration management functions unique to cell-based networks. StrataSphere includes the following applications:

- Cisco WAN Manager (formerly StrataView Plus)
A single unified management platform utilizing HP OpenView® to manage BPX, IPX, IGX, FastPAD, ESP, and INS devices.
- StrataSphere BILLder
Monitors traffic flow over a network and captures data per standard or customized billing periods and formats.
- StrataSphere Modeler
Network modeling tool used for preliminary design of new networks and for analysis and modification studies of existing networks.
- StrataSphere Adaptor
Exports network modeling information to external third party modeling systems.
- SNMP Service Agent
A service agent that provides an interface for automated provisioning and fault management to customers or Operations Support Systems (OSS).

This chapter discusses the role of Cisco WAN Manager in performing network management. For further information on network management, refer to the *Cisco WAN Manager Operations publication*.

Cisco WAN Manager

Cisco WAN Manager is a single unified management platform utilizing HP OpenView® to manage BPX, IPX, IGX, FastPAD, ESP, and INS devices. It provides a standards based multi-protocol management architecture. Regardless of the size or configuration of your network, CiscoWAN Manager collects extensive service statistics, tracks resource performance, and provides powerful remote diagnostic and control functions for WAN maintenance. Refer to the *Cisco WAN Manager Operations publication* for details.

On-line help screens, graphical displays, and easy command line mnemonics make Cisco WAN Manager a user-friendly system. A large amount of hard disk storage is provided to allow accumulating time of day statistics on many network parameters simultaneously. The data is accumulated by the node's controller card and transmitted to the Cisco WAN Manager workstation where it is stored, processed, and displayed on a large color monitor.

Cisco WAN Manager connects to the network over an Ethernet LAN connection. With Ethernet, Cisco WAN Manager connectivity to remote nodes can be established via frame relay over TCP/IP to the LAN connector on the local node, or via inband ILMI. Also, Cisco WAN Manager provides in-band management of network elements via SNMP agents and MIBs embedded in each node and Interface Shelf.

Each node contains an embedded SNMP agent interface. This agent allows a user to manage a StrataCom network or sub-network from any SNMP-based integrated network management system (INMS).

Network Topology

A map of the network is generated at system installation to graphically display all nodes, trunks, circuit lines, and access devices in the network. Various colors are used to indicate the status of each network item. An operator can zoom in to display specific network details while a small overview map continues to be displayed as a locator.

Network Performance

Cisco WAN Manager has powerful statistics gathering capability. From data gathered throughout the network, you can quickly view the operational integrity and deployment of installed network devices and communication media by activating and invoking statistics displays.

Statistics are collected and temporarily stored by each node in the network and released to Cisco WAN Manager when you enable polling, and in accordance with your configuration for specific information within reports. Cisco WAN Manager then stores statistics in a relational database; you retrieve and view these statistics by invoking a statistics display window from the Cisco WAN Manager GUI.

Equipment Management

The Cisco WAN Manager Equipment Manager provides the ability to perform equipment management functions such as adding lines and ports on a Cisco MGX 8220 edge concentrator shelf.

Connection Management

The Cisco WAN Manager Connection Manager provides the ability to perform connection provisioning such as adding, configuring, and deleting frame relay, ATM, and frame relay to ATM interworking connections

Alarm Reporting/Event Log

Cisco WAN Manager, in addition to providing network management capabilities, displays major and minor alarm status on its topology screen for all nodes in a network. It also provides an event log with configurable filtering of the log events by node name, start time, end time, alarm type, and user specified search string.

Software Updates

System software and software updates are supplied by StrataCom on magnetic tape or floppy disk. The system software files are then loaded onto the Cisco WAN Manager workstation where they can be downloaded to a buffer memory in each node in the network in a background mode without disturbing network operation. When the loading is complete for all nodes, a command can be issued to switch all nodes over to the new software. The previous software is preserved and can be recalled at any time.

Additionally, all configuration files for the network can be obtained from the network and stored on the Cisco WAN Manager workstation for backup purposes. In the event of a system update or a node failure, the configuration files can be downloaded to one or all nodes for immediate system restoration.

Network Services Overview

This chapter describes the standard and optional network services (features) supported by networks utilizing the Cisco BPX 8600 series broadband switches, Cisco IGX 8400 series multiband switches, and Cisco IPX narrowband switches. Standard features are bundled with the base nodes. Optional features can be purchased separately and must be enabled by Cisco Customer Service Support before they become operational.

This chapter contains the following:

- Broadband ATM Trunks
- Narrowband FastPacket Trunks
- ATM Connections
- ATM Switched Virtual Circuits and the ESP
- Intelligent Network Server (INS)
- Frame Relay Service
- Frame Relay Congestion Avoidance
- Point-to-Point Data Connections
- Voice Connections
- Bandwidth Control Features
- Graceful System Upgrades
- Network Synchronization
- International Bridging Trunks
- Reliability Features
- Network Management

Broadband ATM Trunks

Broadband trunks, trunks that operate at rates exceeding T1, provide the extra network bandwidth required by rapidly-expanding networks. For example, the 45 Mbps. T3 trunks can carry approximately 28 times as much traffic as a T1 trunk. The Cisco BPX 8600 series broadband switch supports even greater bandwidth with OC-3/STM-1 (155.52 Mbps) and OC-12/STM-4 (622.08 Mbps) rates.

Broadband trunks can be used as high-speed links between Cisco BPX 8600 series broadband switches to create ATM networks or they can connect directly to Cisco IGX 8400 series multiband switches, Cisco IPX narrowband switches, Cisco MGX 8220 edge concentrator shelves, and to CPE. Each broadband network trunk is monitored for signal quality with user-set thresholds for alarm reporting. Trunk load and utilization statistics are accumulated for display by the network administrator.

A broadband network trunk interface, operating at T3 or E3 rates, is provisioned on a Cisco IGX 8400 series multiband switch by equipping it with a BTM (Broadband Trunk Module) and associated backcard and on a Cisco IPX narrowband switch by equipping it with an AIT (ATM Interworking Trunk) card and associated back card. This feature allows users of existing CiscoWAN switching narrowband/multiband networks to economically migrate to ATM and higher bandwidths by adding one card to each IPX node.

The ATM trunk interface permits the Cisco IGX 8400 series multiband switches and Cisco IPX narrowband switches to transmit data using ATM cell relay at speeds up to 45 Mbps. Assigned trunk bandwidth can be as low as one T1 data rate and can be increased in T1 or E1 increments up to the full T3 or E3 rate.

The ATM trunk interface permits the IGX to transmit data using ATM cell relay at speeds up to 45 Mbps. Assigned trunk bandwidth can be as low as one T1 data rate and can be increased in T1 or E1 increments up to the full T3 or E3 rate.

The Cisco BPX 8600 series broadband switch provides BXM-622 and BXM-155 cards which may be configured for either trunk or port (service access) mode, as well as BNI cards.

The BXM-622 cards are either one or two port cards with each port operating at a 622.08 Mbps rate, corresponding to 1,412,830 cells per second at the full OC-12 rate.

By using a full complement of BNI cards, the Cisco BPX 8600 series broadband switch can provide up to thirty-two high-speed trunks per node making it an excellent candidate for an ATM tandem switch. Each T3 ATM port conforms to the Bellcore specifications for DS3 and has a bandwidth capacity of 96,000 cells per second. Each E3 port conforms to ITU-T G.804 interface specification and has a capacity of up to 80,000 cells per second. Each OC3/STM-1 port has a capacity of up to 353,208 cells per second and conforms to SONET/SDH specifications.

For users who may need T3 or E3 capacity in the future but currently can only justify, and have access to T2 trunks, Cisco can provide a T3 to T2 adapter to be used with the Cisco BPX 8600 series broadband switch or with the Cisco IGX 8400 series multiband or Cisco IPX narrowband switches. This adapter interfaces directly to a Cisco IGX 8400 series multiband or Cisco IPX narrowband switch or to a Cisco BPX 8600 series broadband switch T3 port on the node side and to a 6.312 Mbps T2 facility on the line side. The capacity of a T2 facility is 14,490 cells per second, approximately four times the bandwidth of a T1 line.

Narrowband FastPacket Trunks

The Cisco IGX 8400 series multiband and Cisco IPX narrowband switches support various narrowband network trunk interfaces:

- T1 (1.544 Mbps) trunks.
- Fractional T1 trunks.
- CEPT E1 (2.048 Mbps) trunks.
- Subrate (256 Kbps - 1.920 Mbps) trunks with X.21/V.11 trunk interfaces.
- Japanese 1.544 Mbps Y1 trunks.

These trunk types interface with the majority of standard digital trunking types available throughout the world. For linking sites with a small amount of traffic, the fractional T1 and subrate trunks provide an economical solution yet are fully compatible with nodes using higher bit rate trunks. T1 trunks provide a capacity of 8,000 packets/second while framed E1 can carry up to 10,333 packets/second. Fractional T1 trunks use only as many T1 64-Kbps channels as needed, instead of using a full T1 trunk. Subrate trunks use only a portion of the bandwidth of an E1 or T1 trunk and are used for special purposes such as satellite hub access.

ATM Connections

Broadband ATM network connections are switched by Cisco BPX 8600 series broadband switches and Cisco IGX 8400 series multiband switches. ATM connection services are established in CiscoWAN switching broadband networks by adding ATM connections between ATM service interface ports in the network. ATM permanent virtual circuits (PVCs) can originate and terminate on the ASI (ATM Service Interface) cards on the Cisco BPX 8600 series broadband switch, on BXM T3/E3, BXM-155 (OC-3) and BXM-622 (OC-12) cards configured for Port (UNI) operation, or on the Cisco MGX 8220 edge concentrator shelf (using the AUSM card for the Cisco MGX 8220 edge concentrator). Frame relay to ATM Network interworking connections are supported between either BXM or ASI cards to the Cisco IGX 8400 series multiband switch, to the Cisco IPX narrowband switch, or to the Cisco MGX 8220 edge concentrator. ATM T3 or E3 connections can originate and terminate on the IGX ALM-A cards. In addition, frame relay to ATM service interworking connections are supported between either BXM or ASI cards to FRSM cards on the Cisco MGX 8220 edge concentrator shelf.

The BXM cards conform to ATM Forum UNI v3.1 and support the full range of service types specified by ATM Forum Traffic Management 4.0. Traffic Management. The BXM cards are designed to support all the following service classes: Constant Bit Rate (CBR), Variable Bit Rate (VBR), Available Bit Rate (ABR with VS/VD, ABR without VS/VD, and ABR using Foresight), and Unspecified Bit Rate (UBR). ABR with VS/VD supports explicit rate marking and Congestion Indication (CI) control. The BXM cards support up to 16 classes of service. Additional queues are provided for ABR type connections. Egress queues include 32 virtual interface queues, each of which supports 32 VI Qbin queues.

For connections using the ASI-1 card, Cisco's dynamic closed-loop, rate-based ForeSight algorithm may be selected to provide a dynamic congestion control and to optimize bandwidth utilization. Ingress queue size for each card is available up to a maximum of 64,000 cells and can be allocated between the two ports in any manner. On the ASI-155 two separate egress queues are provided; one for CBR and one for VBR.

The ALM-A card on the Cisco IGX 8400 series multiband switch supports CBR and VBR connection types. Data buffers are allocated per connection, and provide buffering required for bursty traffic. The ALM-B card provides for trunk traffic across a network at full T3 or E3 rates.

Narrowband T1/E1 ATM connections terminate on the AUSM card on the Cisco MGX 8220 edge concentrator—and there can be connections between T1/E1, T3/E3, OC3, etc.

ATM Switched Virtual Circuits and the ESP

ATM switched virtual circuits (SVCs) are ATM connections set up and maintained by a standardized signaling mechanism between ATM CPE (that is ATM end systems) across a CiscoWAN switching network. ATM SVCs are created on user demand and removed when the call is over, thus freeing up resources.

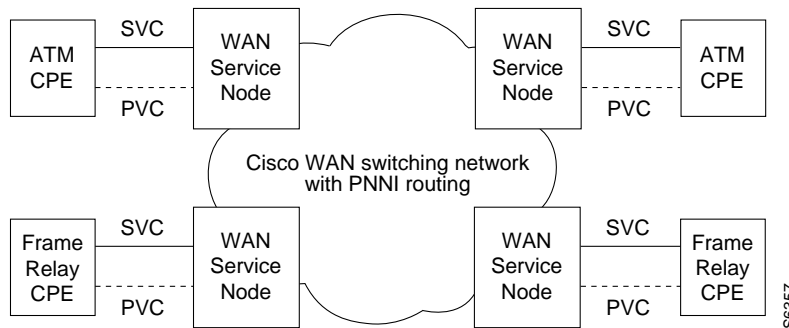
ATM SVCs are supported by the Cisco BPX 8600 series broadband switch when it is equipped with a collocated Extended Services Processor (ESP). When an ESP is installed in a Cisco BPX 8600 series broadband switch, the node resources, such as VPI range, queue allocations, and trunk bandwidth, are partitioned during initial node configuration between SVCs and PVCs.

This resource partitioning provides a firewall between the PVCs and SVCs so that problems occurring with the CPE or large call bursts do not affect the robustness and availability of PVC services.

Figure 4-1 depicts a network with Cisco BPX 8600 series broadband switches providing both ATM SVC and PVC service interfaces.

For further information about ATM SVCs and the ESP, refer to the *Cisco WAN Switching BPX Service Node Extended Services Processor Installation and Operation* publication.

Figur e4-1 ATM Switched Virtual Circuits



Intelligent Network Server DAS and VNS

The Intelligent Network Server (INS) DAS and VNS products use a robust high-powered processing platform to add several important capabilities to Cisco WAN switching networks. Although there are two distinct INS products, they are INS family members because they all use an adjunct processor that distributes selected processing and database-intensive functions across the Cisco WAN switching network. In addition, the INS products all support some form of standards-based signalling between customer premise equipment (CPE) and a Cisco WAN switching network. Typically this signalling is a variation of common-channel, message-oriented Integrated Services Digital Network (ISDN) or Broadband ISDN (B-ISDN) signalling protocols. The INS function interprets these industry-standard signalling messages, translates the logical addresses to the appropriate physical endpoints of the network, and instructs the Cisco WAN switches to establish the connection required for the particular application. The switches then take over, dynamically establishing the optimum route through the network and maintaining the connection for its duration.

The two INS applications are:

- **INS-DAS (Dial-Up Frame Relay)** providing semi-permanent virtual circuits (or soft PVCs) that can be used for dial-up and dial-backup services. A soft PVC is preconfigured in the network's data base but remains dormant until a call from the CPE signals the network to activate it.
- **INS-VNS (Voice Network Switching)** providing voice switched virtual circuits (SVCs) across a Cisco WAN switching network for PBXes using Digital Private Network Signalling System (DPNSS) and QSIG signalling. With Release 8.4, the name DNS has been changed to Voice

Network Switching (VNS). For additional information on VNS Release 2.0 including the breakout/break-in feature, which is associated with Switched Software Release 8.4, refer to the INS Voice Network Switching Installation and Operation Manual.

The INS-DAS and INS-VNS application use one or more adjunct processors that are collocated with a Cisco WAN switching node (that is a Cisco BPX 8600 series broadband switch/Cisco MGX 8220 edge concentrator, a Cisco IGX 8400 series multiband switch, or a Cisco IPX narrowband switch) and often installed in the same equipment rack. Available in either AC- or DC-powered models, the base INS processor is a scalable UNIX platform and contains:

- 140 MIPS CPU, with a 71 Mhz clock
- 64 Megabytes of memory
- 1 Gigabyte hard disk.

For both the INS-DAS and INS-VNS applications, the base INS processor is equipped with different interface modules, memory and disk configurations, and different application software.

INS-DAS Dial-Up Frame Relay

Released as part of Software Release 8.1, the INS Dial-Up Frame Relay application adds semi-permanent virtual circuits (or soft PVCs) to Cisco IPX narrowband, Cisco IGX 8400 series multiband, Cisco BPX 8600 series broadband switch/Cisco MGX 8220 edge concentrator networks. The INS Dial-Up Frame Relay application is implemented by the DAS Server Shelf (i.e., the adjunct processor) and INS Dial-Up Frame Relay software.

In this application, semi-permanent virtual circuits are frame relay PVCs that are configured in the network's data base but are not activated until requested by a dial-up call from end user's customer premise equipment (CPE), such as routers or frame-relay access devices (FRADs). These dial-up calls use ISDN signalling messages to activate the semi-permanent virtual circuits. ISDN signalling messages follow the Q.921 and Q.931 protocol standards.

The INS Dial-Up Frame Relay Application supports five ISDN signalling variations:

- **AT&T 4ESS** protocol variant for North American ISDN signalling
- **AT&T 5 ESS** protocol variant for North American ISDN signalling
- **NT DMS100** protocol variant for a Northern Telecom switch, used in North America
- **NT DMS250** protocol variant for a Northern Telecom switch, used in North America
- **ETSI EUROISDN** protocol variant for the ISDN signalling used primarily in Europe, the Middle-East, South America, and Africa
- **NTT** protocol variant used for ISDN signalling in Japan
- **AUSTEL** protocol variant used for ISDN signalling in Australia.

The INS Dial-Up Frame Relay application operates through both the DAS Server Shelf and SV+ Workstation. Pre-configured soft PVCs lie dormant in the database on the SV+ Workstation until they are activated by an ISDN call into the network. During configuration, the soft PVC has been associated to an ISDN Automatic Number Identification (ANI) field. Then when an ISDN call with that ANI is received and validated, Cisco WAN Manager activates the dormant soft PVC between two preconfigured endpoints. Once the dial-up connection is established, all of the routing, switching and (when needed) rerouting are handled by Cisco WAN switches in the same way that they are handled for PVCs that are connected with leased access lines. When the call is terminated, the Cisco WAN Manager returns the soft PVC to its dormant state.

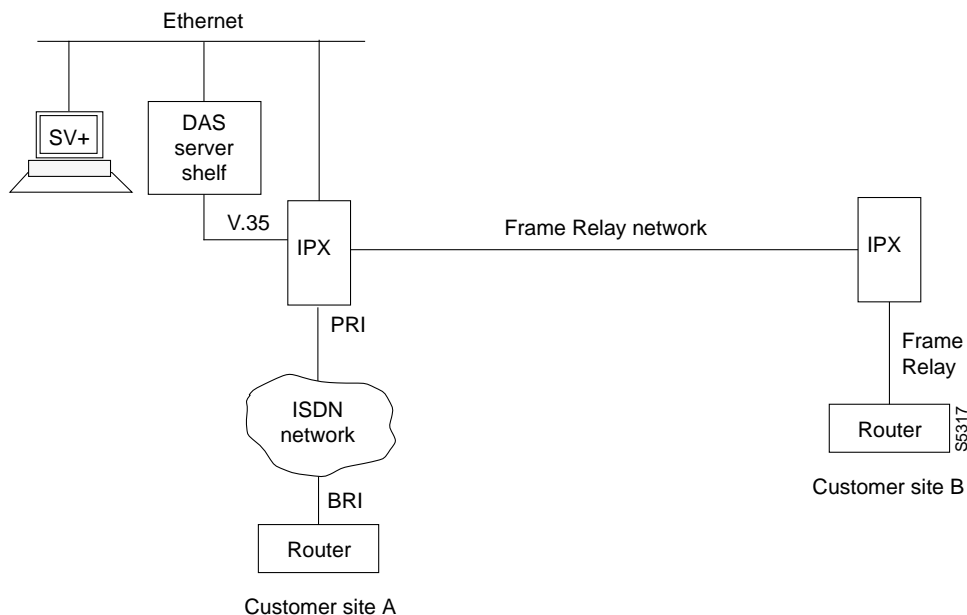
Currently, INS Dial-Up Frame Relay supports two types of connections (i.e., soft PVCs) in the SV+ database:

- Dial-Up connection
- Dial-Backup connection.

Dial-Up Connectio

For the Dial-Up connection, as shown in Figure 4-2, a soft PVC could be configured in the Cisco WAN Manager database between the routers at Customer Sites A and B. (In actuality, the PVC is configured between the ports on the Cisco WAN switching nodes which connect to these routers, i.e., the ports on the two Cisco IGX 8400 series multiband switches.)

Figur e4-2 INS Dial-Up Application



To begin a dialup process, customer premise equipment, such as the router at Customer Site A, issues call setup commands to the ISDN network. The ISDN network routes the incoming call to the Primary Rate Interface (PRI) access port of the Cisco WAN switching frame-relay network (i.e., a Cisco IGX 8400 series multiband switch). These call request and call setup messages are carried on the ISDN D channel, a dedicated signalling channel. The Cisco IGX 8400 series multiband switch passes the ISDN message, encapsulated in a V.35 frame-relay frame, to the DAS Server Shelf. The INS-DAS processes this message and initiates an INS server function on the CiscoWAN Manager workstation. The INS-DAS server searches the connection database for the appropriate connection information which has previously been configured (e.g., a PVC, virtual connection, to Customer Site B), then passes this information to the frame-relay network. In other words, the INS-DAS Dial-Up Frame Relay application determines the identity of the calling device (using the ANI) and directs the frame-relay network to activate the virtual circuits associated with that call source.

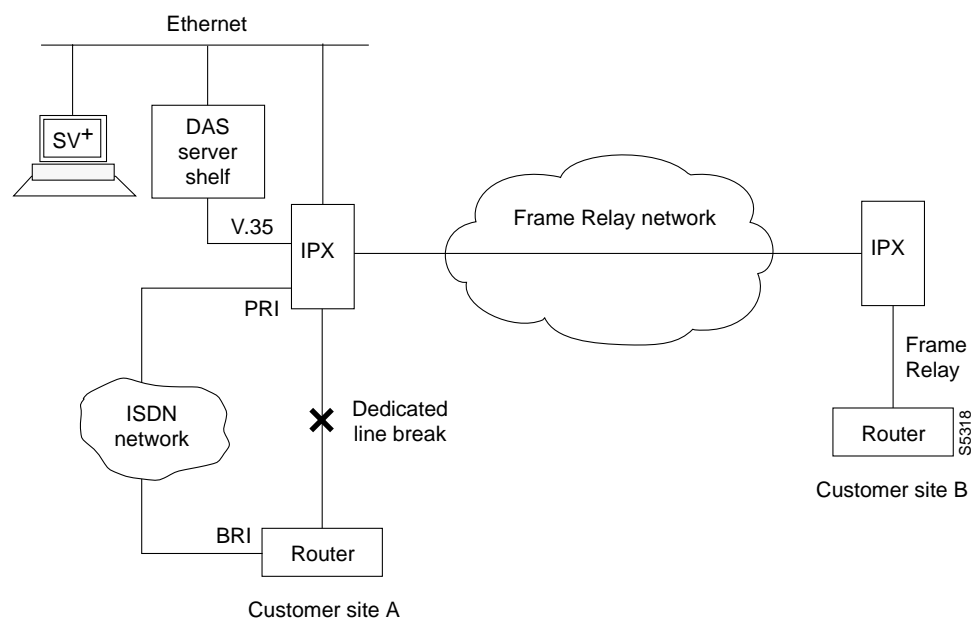
Once the dial up connection is established, all of the routing, switching and (when needed) rerouting are handled by the Cisco WAN switching switches (i.e., nodes) in the same way that they are handled for PVCs that are connected with leased access lines. When the call is terminated, the virtual circuit are torn-down in a similar manner.

Dial-Backup Connectio

Backup strategies for critical business applications are a concern in any network. Cisco's frame relay networks provide automatic re-routing around failed lines and equipment, but for many users the leased line connecting the remote site to the frame relay network, as well as the frame relay port to which the leased line is connected, represent single points of failure.

ISDN Dial-up Frame Relay provides a reliable, economical dial backup solution by using an ISDN switched connection as the backup link into the frame relay network. If the site's leased line connection to the network fails, a dial-up ISDN backup link, shown in Figure 4-3, can quickly and automatically re-establish communications. If the network offers the capability to provide a pool of ports to support dial-in users, the possibility of a single failed port disrupting communication is minimal.

Figur e4-3 INS Dial-Backup Application



The call processing is similar to that of the Dial-Up Call Process, except that the ISDN call is initiated automatically, and the original PVC for the dedicated line is torn down. Because Dial-Backup calls are associated with an active PVC, they must use the same DLCI as the original circuit. When the router at Customer Site A senses that the dedicated line to the frame-relay network is down, it will, after a configured delay, automatically dial into the ISDN network. This call will reach the INS Dial-Up Frame Relay application, which will have already been configured to process this backup call. First it will tear down the PVC that was supporting the dedicated line, then it will establish the dial-backup PVC. When the endpoints terminate the dial-backup call, the backup PVC will be torn down and the original PVC re-established.

Further information about the INS-DAS Dial-Up Frame Relay application can be found in the *INS DAS Server Shelf Installation Manual* and the *INS Dial-Up Frame Relay Operations Guide*.

INS-VNS

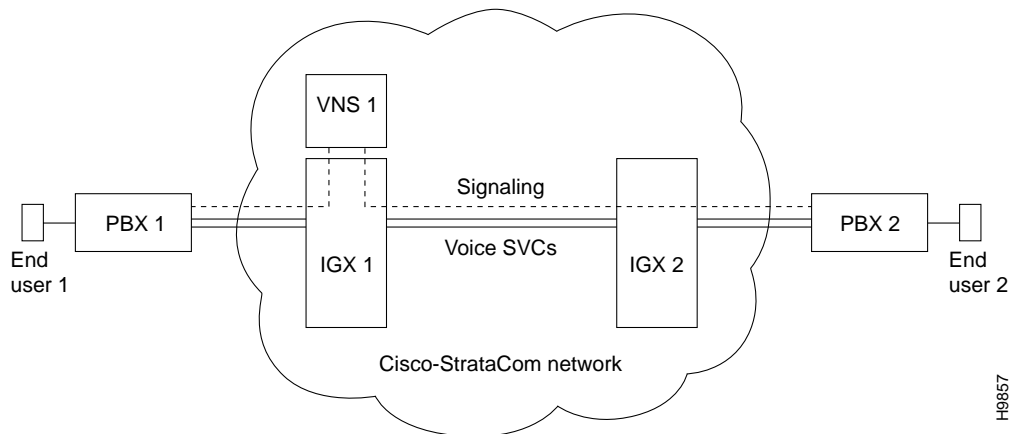
The INS Voice Network Switching (VNS) application, which is implemented with the Voice Network Switching Server Shelf (i.e., the adjunct processor, simply referred to as a VNS) and VNS software, provides switched virtual circuits (SVCs) for voice calls over a Cisco WAN switching network.

With the VNS application in a Cisco WAN switching network, private branch exchanges (PBXes) using Digital Private Network Signalling System (DPNSS) or QSIG signalling will be able to establish voice calls on demand, just as if they are dialing a public switched telephone network. Dynamic Network Switching provides for direct call-by-call routing for PBX voice, data, video, image, and fax connections, enabling this information to be transported across a CiscoWAN switching network efficiently and economically. Also by providing for the direct connection of DPNSS- or QSIG-based PBXes, Dynamic Network Switching eliminates the need for tandem PBX connections. This reduces the number of trunks (e.g., E1) required to interconnect PBXes.

A VNS network also saves network bandwidth by consolidating traffic over fewer physical interfaces, and through the use of Voice Activity Detection (VAD) and Adaptive Differential Pulse Code Modulation (ADPCM) voice compression provided by Cisco IGX 8400 series multiband/Cisco IPX narrowband switches. The VNS network allows the use of a Cisco WAN switching network's standard voice service features to be applied to switched voice circuits from DPNSS and QSIG PBXes. Cisco WAN switching standard voice services save network resources by providing a voice compression ratio of up to 10:1.

As shown in Figur e4-4, Voice Network Switching provides a signalling mechanism to establish and maintain SVCs between PBXes across a Cisco WAN switching frame relay network. In the figure, there is a separate signalling channel between the VNSes to manage the setup and disconnection of the voice SVC calls. The signalling channel actually stretches to the PBXes because the PBXes exchange signalling messages with the network and with each other. This signalling channel is indicated by the dashed line in the figure, and can be thought of as a virtual signalling network, or signalling plane overlaid on the traditional CiscoWAN switching network. The solid line indicates the end-user traffic, the actual voice SVCs, between the PBXes.

Figur e4-4 Basic VNS SVC Call



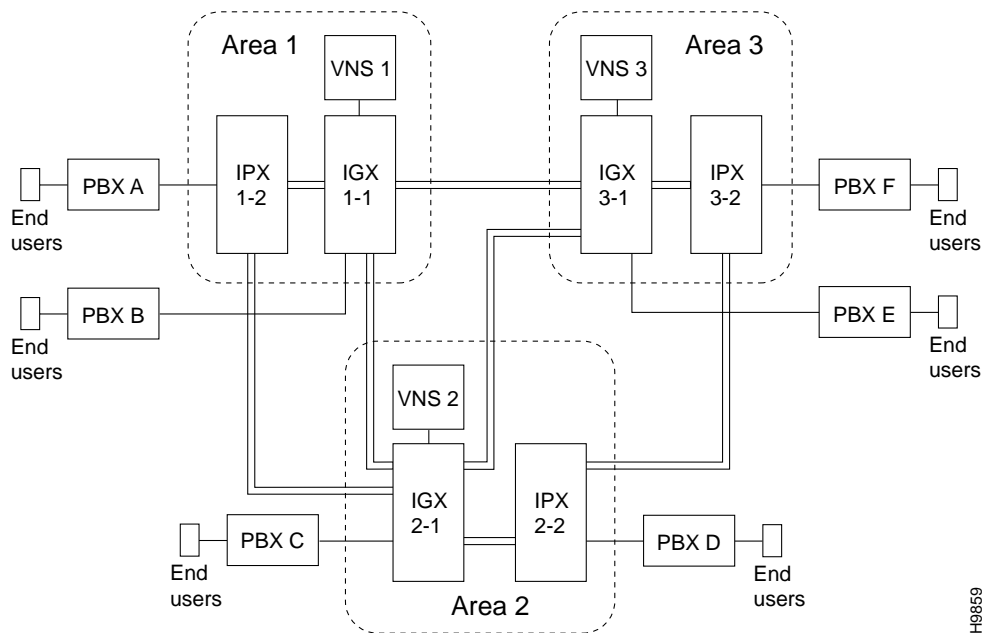
The signalling between the VNS and the PBXes is based on an ISDN variant. There are two signalling variants supported in DNS release 1.0:

- Digital Private Network Signalling System (DPNSS) in accordance with BTNR (British Telecommunications Network Requirement) No. 188, Issue 5, Volumes. 1 to 5, December 1989.

- QSIG based on ETSI QSIG standards.

In the VNS application, each VNS (or redundant pair of VNSes) is assigned to control a group of one or more nodes, i.e., a Cisco IGX 8400 series multiband switch or a Cisco IPX narrowband switch. These nodes are considered the VNS's area (or domain). The VNS will be directly attached to one of the nodes in its area as shown in Figure 4-5. The VNS processes call setup or release requests for calls originating in its area, or calls received from another area but destined for this one. VNS areas do not overlap.

Figur e4-5 VNS Areas



In Figure 4-5, each VNS has two Cisco IGX 8400 series multiband switches in its area. VNS 1 has IGX 1-1 and IGX 1-2, and VNS 2 has IGX 2-1 and IGX 2-2, and so on. Each Cisco IGX 8400 series multiband switch is connected to one PBX with its complement of end users. So in Area 1, VNS 1 would be responsible for processing call setups for all calls from PBX A and PBX B. It would also handle calls destined for PBX A and PBX B. Similarly, VNS 2 would be responsible call setups for PBX C and PBX D, and VNS 3 would be responsible for PBX E and PBX F.

Further information about Voice Network Switching can be found in the *INS VNS Installation and Operations* manual(s).

Frame Relay Service

This description is keyed to the IPX, but the general principles apply to the Cisco IGX 8400 series multiband switch and to the Cisco MGX 8220 edge concentrator. See those respective reference manuals for details. With Frame Relay to ATM Interworking, the Cisco MGX 8220 edge concentrator, the Cisco IGX 8400 series multiband switch, and the Cisco IPX narrowband switch all provide efficient transport of frame relay traffic via Cisco BPX 8600 series broadband switches across ATM networks.

**Note**

Frame relay service is an optional feature on the Cisco IGX 8400 series multiband switch and the Cisco IPX narrowband switch that must be purchased and enabled on each node where it is to be used.

Emerging data applications are characterized by large volumes of high-speed, bursty data with diverse connectivity. Typical applications require connectivity to multiple destinations and data transfers of millions of bits at megabit speeds. Today's systems require access to high bandwidth on demand, direct connectivity to all points in the network, and access to only the minimum bandwidth actually required. ITU-T has defined frame relay services as the standard data interface to meet these needs.

Cisco's implementation of frame relay is designed to provide high-speed data transmission, minimal network delay, and efficient use of network bandwidth. The frame relay ports on a Cisco IGX 8400 series multiband switch allow a connected device to access any other frame relay device by specifying its address within the transmitted data frame. Thus, Cisco WAN switching cell relay frame relay network transmits data frames to their destinations at high speeds.

Advantages of using a Cisco WAN switching cell relay network for bursty data include:

- Circuit bandwidth is variable and allocated on demand.
- Each cell has an address that allows it to route itself through the network.
- Cells are sent only on demand so a large number of connections can use a few trunks in a statistical multiplexing arrangement.
- Fast delivery due to minimal processing requirements and the use of high speed transport media (T1 or E1).
- Congestion avoidance credit manager and ForeSight features prevent one PVC from adversely affecting other PVCs.
- There is no call setup delay because all connections are permanently allocated.
- Fast rerouting of traffic in the event of a network failure.
- All cells for a frame relay destination are sent in sequence along the same route; no sequencing needs to be done on the packets at the receiving end.

A number of software features have been included in the Cisco BPX 8600 series broadband switch, the Cisco IGX 8400 series multiband switch, the Cisco IPX narrowband switch, and the Cisco MGX 8220 edge concentrator to detect, control, and avoid internal network congestion. These include ForeSight, Explicit Congestion Notification, Local Management Interface, and the Credit Manager. The ForeSight feature allocates extra bandwidth to all connections fairly while insuring that no congestion exists.

These features allow network operators more flexibility in defining the performance characteristics of frame relay virtual circuits bases upon initial, minimum, and maximum data rates as well as network delay caused by congestion. ForeSight, for example, is a closed-loop mechanism that provides continuous feedback on trunk utilization across the network to adjust the data rate of each virtual circuit at the network access points.

Cisco WAN switching frame relay complies with the following Standards Committees specifications:

- ANSI T1.606 and ITU-T I.233.1 Frame Relay Service Description.
- ANSI T1.618 and ITU-T Q.922 Data Transfer Protocol.
- ANSI T1.606 and ITU-T I.370 Congestion Management.
- ANSI T1.617 Annex D and ITU-T Q.933 Annex A Signaling.
- ITU-T I.372 NNI Interface Requirements.

- Frame Relay Forum LMI and NNI implementations.
- Bellcore TR-TSV-1370 Generic Requirements

Basic Frame Relay Service

The basic Cisco IGX 8400 series multiband and Cisco IPX narrowband frame relay service provides permanent virtual circuits (PVCs) as defined in the frame relay standards. These circuits are used for interconnecting widespread LANs transmitting bursts of data. User device control is transmitted over Local Management Interface (LMI/ELMI) port. Features of basic frame relay service are listed below:

- Up to 1024 frame relay virtual circuits are allowed on a node.
- 4 ports/card with port speeds to 2 Mbps.
- Each port may be configured either as a DCE or a DTE.
- Accommodates frame sizes of 5 to 4506 bytes.
- Up to 252 frame relay permanent virtual circuits per frame relay port/card.
- DLCI used for addressing.
- ITU-T V.35 port interface or ITU-T X.21 port interface (Rel. 6.3 and later).

Frame Relay SVCs

When enhanced with an Extended Services Processor (ESP), the Cisco BPX 8600 series broadband switch supports frame relay switched virtual circuits (SVCs).

Frame relay SVCs are set up and maintained by a standard signaling mechanism between frame relay CPE (that is end users) across a CiscoWAN switching network. Frame relay SVCs are created upon user demand and are removed when the all is over, thus freeing network resources.

When an ESP is installed on a Cisco BPX 8600 series broadband switch, the node's resources, such as frame relay port DLCI range and bandwidth and trunk bandwidth, are partitioned between frame relay SVCs and PVCs. This prevents problems with frame relay SVCs from affecting frame relay PVCs. Frame relay SVCs are supported only on Cisco MGX 8220 edge concentrator FRSM cards. For further information refer to the *BPX Service Node Extended Service Processor* documents.

Frame Relay T1/E1 Ports

This feature is an extension of the basic frame relay service to provide a higher speed (T1 or E1) user interface to the Cisco WAN switching network. Frame relay T1/E1 is used to extend CiscoWAN switching cell relay network frame relay services to remote LANs using standard T1 trunks by using a higher speed frame relay interface port.

Two port capacities are available in this application. One provides 6 logical ports, which is convenient for fractional T1. The other provides a full 24 logical ports (T1) or 31 logical ports (E1) applications. For example, on a T1 line connection, all 24 DS0/timeslots could be used if each of the 6 logical ports had 4 DS0/timeslots assigned.

Frame relay DS0 timeslots are grouped into "logical ports". These logical ports can be a single DS0/timeslot or a group of contiguous DS0 timeslots. Logical ports that consist of multiple DS0/timeslots transmit at the full rate of 64 Kbps per timeslot.

A single frame relay DS0 logical port may be configured for 56 Kbps or 64 Kbps. The 56 Kbps rate is typically used for groomed DDS circuits that appear on a T1/E1 line. If a logical port is configured for 56 Kbps, the IPX will strip off the least significant bit (signalling bit) in the incoming octet and stuff a “1” in the least significant bit of the outgoing octet.

With E1 frame relay, timeslot 16 cannot be used as part of a logical port if the line is configured for CAS type signalling. However, it may be used for CCS configurations.

User-to-Network Interface (UNI)

Frame relay uses the data link layer address to route data frames across a network between a source and destination. This service is provided via a standard interface between the user device and the network. This standard interface is called the User Network Interface (UNI).

There are two protocols used across the UNI, the data transfer protocol and the control protocol. The data transfer protocol operates end-to-end (through the network, between a pair of user devices). The control protocol operates between the user device and the network. The service allows 1024 connections per node.

ANSI Annex A and ITU-T Annex D Frame Relay standards groups specify a protocol and set of procedures called the User-to-Network Interface (UNI) to provide a control function for frame relay PVCs. Cisco WAN switching Local Management Interface (LMI) is a specialized UNI that offers additional features over and above those defined by either the ANSI or ITU-T standards. These control features are imbedded in the software and the FRP firmware and operate between the IPX and the user device. The messages communicated via the UNI/LMI protocol provide the following features:

- Network notification to the user device of the active and available PVCs.
- Network notification to the user device 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 to the user device of a change in PVC status.
- Network notification to the user device of the minimum bandwidth allocated per PVC.
- Permanent virtual circuit priority (not yet standardized).

The Cisco WAN switching LMI frame relay interface includes several internal software controls to regulate the flow of frame relay data and prevent congestion on network trunks. Refer to the Frame Relay Congestion Avoidance, page 4-13 section later for a more complete discussion.

Network-to-Network Interface (NNI)



Note

Frame relay NNI is an optional feature that must be purchased and enabled on each node where it is to be used.

Frame relay networks utilizing Cisco IGX 8400 series multiband switches, Cisco IPX narrowband switches, or Cisco MGX 8220 edge concentrators can be seamlessly connected together and to other frame relay networks adhering to standards set forth by the Frame Relay Forum. Inter-network messages flow between Network-to-Network Interface (NNI) frame relay ports to report internetwork connection status to user devices wherever they are located.

NNI is a frame relay port that forms a boundary between two independent wide area networks, e.g., a Cisco WAN switching network and another network that may or may not consist of Cisco WAN switching equipment. There is no user device connected, only another network port. Each network interface in a Cisco WAN switching network consists of a port on a FRP, FRM, FRSM, or PCS shelf card.

Only ports that need to support the reporting of the A-bit status need be equipped with Model F or H FRP cards. Frame Relay NNI can operate in networks with a mix of older model and new model FRP cards. Although, connections may be established between these and earlier FRPs they will not support the NNI feature.

Bundled Connections

Frame relay virtual circuits between the same two nodes over the same network route may be treated as one routing entity. These connections are bundled connections. Bundled connections require that all virtual circuits in the bundle be numbered consecutively and located on the same frame relay card.

All bundled connections in the same logical group are treated as one physical circuit thereby increasing the capacity of the Cisco IGX 8400 series multiband switch for frame relay circuits. Without bundling, a Cisco IGX 8400 series multiband switch can accommodate 252 virtual circuits. With bundling, a node can handle 1024 circuits.

Frame Forwarding

Frame forwarding is a software feature of frame relay allowing point-to-point connections for various data applications that do not conform to the Frame Relay Interface Specification. This includes bridges, switches, front-end processors, and other equipment that support SDLC, HDLC, or LAP-B interfaces.

With frame forwarding, all frames received on a local FRM port are transmitted via a single PVC to a remote FRM port, and all frames received on a remote FRM port are transmitted via a single PVC to a local FRM port.

Frames containing any data are interpreted as valid frames. These frames must use 7E (hex) flags, 5-4096 bytes, and a valid CRC as defined in ITU-T standards. If a frame relay connection exists, the frames are relayed via a PVC to their destination. If no connection exists, the frames are discarded and the invalid frame statistics are updated.

Frame Relay Congestion Avoidance

Cisco WAN switching cell relay networks have a number of features incorporated to monitor the network frame relay connections for congestion, to notify the user device of this congestion, and to take action to minimize the congestion where possible.

Credit Manager

The Cisco IGX 8400 series multiband switch assigns a control function to each frame relay port to control the transmission of user data from the port. In many frame relay applications, if the user is permitted to send a short “burst” of data, it will fulfill many user's requirements with a delay that is almost unnoticeable. Often, a short burst of data from any one port can be accommodated by the network without any undue stress.

However, there must be some means to throttle back the data rate so it doesn't persist at a high rate long enough to cause congestion. The long-term committed information rate must not be more than the user has contracted for. This feature is the Credit Manager, and, unlike ForeSight, which is optional, the Credit Manager is a standard operating feature of the IPX when equipped with frame relay.

The Credit Manager allows an initial high-bandwidth burst of data to be sent, but throttles back the data rate after the short initial period based on the number of system-assigned “credits” that have been accumulated by an individual circuit. This permits short transactions to be sent immediately for quick response but prevents long transactions from overloading the network.

“Credits” are allocated to a frame relay port at a constant rate, depending on the Committed Information Rate assigned to the port, up to a certain maximum. When a port is idle, it “saves” credits, when it is transmitting, it “spends” credits. If a port has been idle for a time, it has accumulated credits and can spend them as fast as it wants (up to the Peak Information Rate set for the port). After accumulated credits are spent, the port is throttled back to its Committed Information Rate for the duration of the session to minimize the possibility of congestion and for fairness to other users.

Congestion Notification (FECN/BECN)

Explicit Congestion Notification is a form of flow control for frame relay networks accepted by the Frame Relay standards committee as a way of preventing network congestion. This standard feature consists of two bits carried in the frame header that are set by the network in the event of congestion and it is up to the user device to react by reducing the data rate applied to the network. It reacts primarily to congestion at the source or destination of the frame relay connection.

When the IPX FRP card at the source end of the circuit detects its input buffer fill approaching a preset threshold, it sets the Forward ECN (FECN) bit. This bit is carried “forward” through the network, towards the frame destination. At the same time, the FRP sets the Backwards ECN (BECN) bit, which is sent back to the transmitting user device.

Similarly, when the IPX FRP card at the sink end of the circuit detects its output buffer fill approaching a preset threshold, it sets the Forward ECN (FECN) bit. This bit now is transmitted to the receiving user device. At the same time, the sink FRP sets the Backwards ECN (BECN) bit, which is sent through the network back to the transmitting user device. FECN and BECN are of no value in reducing congestion if the user device does not, or is incapable, of reacting to the change of state of the FECN and BECN bits.

Point-to-Point Data Connections

Circuit mode data ports are used for today’s conventional data communications applications that expect to communicate via a point-to-point bit-transparent connection, such as is provided by a data modem. Cisco WAN switching cell relay networks typically allow twice as many data connections to be carried compared to TDM-type circuit switching networks.

Synchronous Data Connections

For circuit mode data applications, the IPX Synchronous Data PAD (SDP) provides direct connection to the IPX from standard data communication interfaces. IPX data interfaces are full-duplex, protocol transparent circuits. Synchronous, binary synchronous, and bit synchronous protocols are supported.

The IPX supports high-speed synchronous data ports at rates up to 1.344 Mbps with the following port interfaces:

- EIA RS-422/449.
- EIA RS-232C/D (V.24).
- ITU-T V.35.
- ITU-T X.21/V.11.

The IPX supports low speed data applications on the LDP/LDI front card/back card set with 8 ports/card (up to 19.2 Kbps) or 4 ports/card (up to 56 Kbps) with the following port interfaces. Asynchronous data rates up to 19.2 Kbps are supported by oversampling.

- EIA RS 232C/D (ITU-T V.24).
- AT&T 56 Kbps Digital Data Service (DDS).

DS0A Data Connections

DS0A connections allow a channelized 64 Kbps interface (T1 or E1) to carry a single sub-rate data channel. This channel can operate at 2.4, 4.8, 9.6, or 56 Kbps. DS0A connections that originate on a CDP can terminate on either another CDP (output with parallel 8-bit bytes) or on a SDP or LDP card (serial data output) at the far end.

Imbedded signalling information is carried over the network in one of two modes, transparent or interpretive. For CDP to CDP connections, the transparent mode transmits signalling bits directly end-to-end and lets the connecting data device interpret the signalling. The interpretive mode is used with CDP to SDP or LDP data connections. The near-end CDP decodes the DS0A signalling and converts them to standard EIA control leads that can be mapped and observed using the software “breakout box” feature.

Nx56K, Nx64K Connections

This standard feature enables bundles of 56 or 64 Kbps time slots to be aggregated to make up a single data channel. This data channel can be terminated on another CDP or SDP. If the link is sensitive to long strings of zeros, the IPX will assure a minimum pulse density. If the link uses B8ZS or other means of suppressing long strings of zeros, the IPX can be set to forgo bit stuffing to save bandwidth.

Data Frame Multiplexing (DFM)

**Note**

Data Frame Multiplexing is an optional feature that must be purchased and enabled on each node where it is to be used.

Data Frame Multiplexing is a compression technique for data channels to reduce network bandwidth requirements. DFM monitors the input data and replaces repetitive patterns (such as idle codes, etc.) with an initial pattern and a repetition code number that indicates to the far end how many times to repeat the initial pattern in the output data stream. During long intervals of repetitious data, the initial data pattern is periodically transmitted for confirmation at the far end. DFM is only available for data rates up to 128 Kbps.

Data Clocking Options

In a network where the clocks are not synchronous, the card's character receive rate differs from the character playout rate. This results in continuously variable data delays and periodic bursts of bit errors. If a data card's remote partner is clocking at a slightly higher rate, delay increases as extra characters build overruns and all characters in the buffer are lost. Under these conditions, up to two seconds of delay may be introduced at the lowest baud rate before the overrun occurs.

The Pleisiochronous Clocking feature limits the amount of data (delay) allowed to accumulate in the Low-Speed Data PAD (LDP) card receive buffers by resynchronizing at an acceptable delay threshold that is less than the size of the complete receive buffer. This feature provides an upper limit to the receive buffer delay. It is always active with no user interface or configuration required. The delay variation imposed by the feature is between 12.5 and 50 milliseconds for all supported baud rates.

Synchronous Data Control Lead Options

There are three modes for transmitting control lead status for synchronous data channels. The normal mode periodically samples up to 12 control leads per data circuit and sends the lead status using the internode communications channel. This mode uses very little trunk bandwidth but trades delay for bandwidth.

Alternatively, the fully interleaved mode transmits a byte of data followed by a byte of control lead status (up to 7 control leads). This reduces the delay between the change of status and when it is recognized at the far end to one byte interval but doubles the bandwidth required for each data circuit.

A third mode, the partially interleaved mode, permits sending the status of one control lead in each direction (RTS or CTS) as the eighth data bit. All other lead status is transmitted as normal mode. This provides quick response to status changes on this lead without requiring much additional bandwidth. However, this mode is restricted to 19.2 Kbps and under.

Voice Connections

Voice interfaces, such as PABX's, channel banks, transmit Pulse Code Modulation (PCM) voice streams to the IPX for processing. The voice interface to T1 trunks is based on the D4 frame format. Via the T1-D4 format interface, T1 IPX systems are certified as fully compatible with major PBXs, channel banks, and other D4 equipment.

The IPX receives PCM voice from the T1 or E1 circuit lines as 64 Kbps channels, which is packetized and compressed by the CDP card. The IPX can provide up to 4:1 compression, which means it can support up to 96 active toll-quality channels on one T1 FastPacket trunk, and up to 120 voice channels on one E1 FastPacket trunk.

In addition to Pulse Code Modulation (PCM), the IPX provides Adaptive Differential Pulse Code Modulation (ADPCM) to give another 2-to-1 or 4-to-1 compression to voice. In this way, the IPX provides 4-to-1 or 8-to-1 voice compression for as many as 240 (E1) or 192 (T1) voice channels on a single trunk.

The IPX also compresses voice streams using Voice Activity Detection (VAD). With VAD, voice packets are transmitted only when speech is present (in typical phone conversations, speech is present only 40% of the time). This results in greater than 2-to-1 compression, with no degradation in voice quality. This compression ratio allows many more voice channels to share a single trunk.

The IPX supports Robbed Bit signalling for T1 networks, and Channel Associated Signalling (CAS) for CEPT (E1) international networks. Initially, the IGX CEPT (E1) supports a transparent pass-through (that is, no interpretation of the data) of Common Channel Signalling (CCS) for CEPT (E1) international and ISDN networks.

The CEPT E1 and Japanese J1 interfaces to voice equipment consists of 32 time slots, where one channel is used for framing (time slot 0) and one channel is used for signalling (time slot 16), leaving 30 available voice channels. The frames are transmitted on 2.048 Mbps trunks.

Voice SVCs

For information about voice SVCs, as supported by the INSVoice Network switching, refer to INS-VNS, page 4-8 of this chapter and to the *INS-VNS Installation and Operation* document.

PCM Voice Connections

Pulse Code Modulation is the current world-wide standard for converting analog signals to digital voice. This standard IPX feature carries voice in 64 Kbps DS0 channels from one circuit line to another with toll-grade voice quality. No voice compression techniques are applied to the encoded data bits.

ADPCM Compression for Voice Connections

Adaptive Differential Pulse Code Modulation (ADPCM) is a technique for encoding only changes in amplitude between PCM samples rather than the actual amplitude. Since a smaller range of amplitudes is generally encoded, it requires fewer bits to represent the change than it does the actual sample.

ADPCM compression reduces network bandwidth requirements for transmitting voice by encoding 64 Kbps voice channels as 32, 24, or 16 Kbps ADPCM. The 32 Kbps is available with all software releases. 16 and 24 Kbps became available with Release 6.0 and the CDP.

32 Kbps provides 2:1 compression, 24 Kbps provides 3:1 and 16 Kbps provides 4:1 compression. Generally, there is a trade-off between the amount of compression and voice quality. The CDP provides two new options, 32z, and 16z using a new compression algorithm with inherent zero code suppression (limits strings of zeros to a maximum of 6).

Voice Activity Detection Feature

The Voice Activity Detection (VAD) feature detects periods of silence when the speaker is quiet and the silence between words and suppresses the transmission of packets/cells during these periods. It provides a 2:1 savings in the amount of packets otherwise sent. VAD can be used in conjunction with any of the ADPCM compression rates. VAD is a standard feature of all IPX nodes equipped for voice.

VAD can be automatically enabled or disabled depending on connection bandwidth available. Refer to the associated AdaptiveVoice feature described next in the “Adaptive Voice” section.

Adaptive Voice

Adaptive Voice is a feature that provides dynamic enabling/disabling of VAD. As long as the network has available, unused bandwidth, VAD is automatically disabled on all voice connections using it to maximize the quality of voice connections. When there is no more network bandwidth available, VAD is automatically reapplied. Priority for using connections is by class of service (COS) or activity. See also the section, Voice Activity Detection Feature, page 4-17.

Instafax and Enhanced Instafax

FAX/modem upgrade (Instafax) is a standard feature that automatically disables ADPCM compression for voice connections whenever a V.25 protocol echo canceller disabling tone is detected from a high-speed modem or a facsimile (FAX) application.

Normally voice circuits employ some type of compression (ADPCM and/or VAD). When a high-speed modem or FAX is used on these circuits, the compression may interfere with error-free transmission of the modem tones. Instafax is a feature that disables VAD on the connection and forces the circuit to use the full 64 Kbps circuit bandwidth when the CDP card detects a V.25 echo canceller disabling tone from the modem or FAX.

Enhanced Instafax operates in a similar manner except that the user may choose to revert to either the full 64 Kbps PCM connection or to a CiscoWAN switching-proprietary 32 Kbps ADPCM that is optimized to pass modem/FAX transmission. This ADPCM provides a 2:1 compression that otherwise would be unavailable for these applications without affecting the modem/FAX performance.

The 32 Kbps ADPCM encoding used, however, is not compatible with the ITU-T-based 32 Kbps ADPCM used on other types of IPX voice circuits. Circuit defaults to 64 Kbps during Instafax transmission.

Bandwidth Control Features

Many users consider voice as a low priority service on their network that may be preempted in the case of trunk failures. This section describes features that allow this kind of routing. These features include the following:

- Courtesy Downing.
- Bandwidth Reservation.

Courtesy Downing

Courtesy Downing connections is the process of monitoring a set of voice connections to see when they go inactive. When the connections go inactive, they are downed or disabled. This frees up network bandwidth for other uses, generally for bandwidth reservation.

In the case of inter-node connections, downing connections deroutes these connections and frees up the network bandwidth. In the case of local (DAX) connections, these connections are only disabled. In both local and distant connections, downed connections are conditioned at both ends and do not generate network alarms.

Only voice-type connections (“v”, “c”, “a”, “p”, and “t”) can be monitored for activity, and then only when the on-hook status is configured by the user. All other connection types and voice types without their on-hook state defined are treated as active and cannot be courtesy downed. System software allows manual control over Courtesy Downning. This allows the user

- To immediately down and up specific connections.
- To immediately down and up all connections with a specified range of COS.

Bandwidth Reservation

Bandwidth Reservation is the process of configuring a currently inactive connection to become active at a specified time and accumulating the necessary bandwidth in advance. This is particularly useful when a periodic uploading of a large file, for example a database update, requires a large amount of bandwidth but only for a short period of time.

The user may request Courtesy Downning of specific connections or network-side connections matching the specified COS range at a particular time. At a later time the user may activate a downed connection, add a new connection, or change the COS of an existing connection to use this bandwidth. If sufficient bandwidth is available, then the connection is routed, completing the reservation.

An inactive connection is a connection that was previously downed manually. This state reserves the endpoint resources for the connection but does not route the connection. In addition, the connection is conditioned. Commands for reserving bandwidth can be used in conjunction with scheduled jobs to schedule connections on a daily or weekly basis.

Graceful System Upgrades

All the IPX, Cisco IGX 8400 series multiband, and Cisco BPX 8600 series broadband controller cards (NPC, NPM, and BCC, respectively) incorporate flash memory, which can be programmed yet retain their memory when power is removed. This permits upgrading both the system software as well as the firmware for individual cards. The following paragraphs describe these standard features. These upgrading operations require SuperUser access in conjunction with Customer Service support.

Software Downloading

Software downloading provides a quick and simple method of upgrading the node operating software from one revision to another and it does not require dedicated personnel at each site. If the node is equipped with redundant controllers, the upgrade can be accomplished with no disruption in the operation of the node.

Updated software is supplied on a CD supplied by Cisco for loading onto a CiscoWAN Manager workstation. From here, it is downloaded, upon command, to the local node, over the NMS connection to the network.

For nodes equipped with redundant controller cards, the new release is loaded into the standby controller while the active controller continues to run the existing release. Nodes are designed to run properly with different software releases in active and standby controllers.

In similar manner, each node in the network is updated with the new software release. Software for nodes that are not directly connected to a Cisco WAN Manager terminal is transmitted over network trunks. As each node completes its update, the software is automatically downloaded to the next adjacent node.

Upon command, each node can be instructed to run the newly downloaded release software in which case the standby controller with the new release becomes active. If, for any reason, a problem arises, the network can revert back to the prior release immediately by switching back to the other controller.

Firmware Download

Most of the hardware modules used in the Cisco IPX narrowband switch, the Cisco IGX 8400 series multiband switch, and the Cisco BPX 8600 series broadband switch operate under control of software imbedded in ROM (firmware). In many of these modules, a flash EPROM is used. When new firmware becomes available from Cisco (often with new features), this firmware can be downloaded from Cisco WAN Manager NMS in a manner similar to that just described for system software.

New firmware is first loaded onto the Cisco WAN Manager terminal from a CD supplied by Cisco. It is downloaded to each node controller, which temporarily stores it in a buffer. Upon command, the controller transfers the new firmware, card by card, to each of the appropriate module(s) in the node. A second command causes the new firmware to be “burned” into the module’s memory.

Firmware download permits easy upgrade of existing modules without the need to remove and replace them. The modules are taken out of service only during the download procedure, which lasts only a few minutes. A redundant module pair can be updated with no downtime. This feature extends to all nodes in the network.

System Software Restore

Software Restore provides a means of automatically restoring node operations after a node power failure, a momentary interruption in node controller processing, or following controller card replacement.

Each controller card has two memory areas for the system software, one is ROM and one is RAM. The controller is shipped from the factory with the current release of the system software loaded in ROM, which is transferred to RAM on power-up.

If a node power fail should completely erase the contents of the RAM or a momentary failure should corrupt the normal operation of the controller, a boot routine in ROM begins to immediately transfer its copy of the system software to RAM. Within a few minutes, when the transfer is complete, the node is up and running again.

The node configuration data (information about the network) is stored in BRAM (battery supported RAM) and is preserved during power interruptions.



Note

All nodes in the network carry the configuration for each node in the network to assist in determining proper rerouting during trunk failure.

Network Synchronization

Cisco WAN switching networks feature pleisiochronous network clocking, which provides a highly flexible and robust clocking system. An IPX network automatically configures its clock synchronization plan when a network is constructed, and reconfigures its clock routing as the network grows. Automatic clocking reconfiguration results in the selection of a new clock source in the event of line failures, minimizing the impact of loss of clock.

The IPX can automatically derive its clock source from an internal clock, a port, a trunk, or an external clock device. In pleisiochronous networks, different regional clocks may not be synchronized, so frame slips occur on trunks connecting the regions to compensate. The IPX is able to slip frames on empty packets to maintain network synchronization without information loss.

Network synchronization internal to IPX node sources corresponds to Stratum 4 (2.048 MHz, ± 10 ppm). Cisco BPX 8600 series broadband switch clocks are Stratum 3 sources. Network synchronization external to network sources can be provided by dynamic primary, secondary, and tertiary clocking sources, synchronized to the nearest highest-stratum clock available. Any network trunk or circuit line or 1.544/2.048 Mbps external clock can be used for the clock source.

International Bridging Trunks

European networks can be bridged to North American networks via T1, CEPT E1, or subrate trunks. The bridging trunks are digital trunks that carry voice, signalling, and data between circuits using different standards. Data connections can be made freely between any compatible end ports operating at the same speed (e.g. X.21 circuit in Europe and RS232 circuit in North America).

For voice data, the International IPX does an A-law/ μ -law conversion of the encoded voice samples in the CDP card. Signalling data requires two translations: A and B bit states and signalling timing. The International IPX system converts an E1 system SSDC5-A trunk A and B bit signalling states to a T1 system A and B signalling states in the CDP port card.

Reliability Features

The Cisco BPX 8600 series broadband, Cisco IGX 8400 series multiband, and Cisco IPX narrowband hardware and software architecture and features are designed to protect the node in the event of a major network outage, system disruption, or node power outage.

Distributed Network Intelligence

The intelligence for network routing with Cisco BPX 8600 series broadband, Cisco IGX 8400 series multiband, and Cisco IPX narrowband switches is distributed around the network to maximize robustness. Each switch contains an image of the network topology and a list of all network connections. This is used for automating routine network administration. The only centralized system functions are user interface and statistics collection.

Redundant Node Powering

Each node is initially supplied with at least one more power supply than is needed to provide full power to the node. All power supplies are on-line so that should a power supply fail, the remaining power supplies can easily carry the system load with no switching of supplies required. Inputs to each supply are separately protected with a circuit breaker.

All dc outputs of each power supply are monitored separately and automatically cause an alarm when a failure of any output voltage is detected. Status and output voltages for each power supply equipped are displayed on a NMS screen making it easy to isolate a failed unit.

Power supplies are mounted on slides and held in place with a captive mechanism. Any one supply may be replaced without powering down the node. Cisco BPX 8600 series broadband, Cisco IGX 8400 series multiband, and Cisco IPX narrowband nodes are available with dual power input configurations to enhance reliability.

Redundant Card Sets

Cisco BPX 8600 series broadband, Cisco IGX 8400 series multiband, and Cisco IPX narrowband nodes are generally equipped with dual controller cards for redundancy. The smaller IGX 8410 multiband nodes have provision for dual controllers but may be equipped with a single controller card if desired for cost savings. All other card sets may be provisioned either singly or with 1:1 redundancy (1:1 redundancy was provided for lines and trunks with the DTI card set found in earlier release systems).

All cards are monitored for normal operation in a background mode and, if redundancy is provided, the system automatically switches to the hot-standby card if the active card fails. Often this switchover is accomplished with no interruption in service.

Inter-Node Controller Communications

Each node maintains a database of the configuration the node elements as well as a configuration of the overall network. Each node is aware of all other nodes, trunks, trunk bandwidths used and available, as well as any special routing details entered by the user. As a result, each node is capable of deciding when and how to reroute connections around network failures.

The advantage of this is two-fold. It speeds rerouting decisions since a single node detecting a network failure can make the decision to reroute. But, more importantly, this architecture eliminates the possibility of a single node failure or a failure of the network management station from affecting the operation of the whole network.

With structured networks, the situation is similar in that each local node has knowledge of all other nodes within the structured network domain. The network database in junction nodes contains entries for routing for all interdomain connections in the network.

Configuration Save/Restore



Note

Configuration Save/Restore is an optional feature that must be purchased and enabled for each node where used.

The Configuration Save/Restore feature provides a means of storing the configuration of each Cisco BPX 8600 series broadband, Cisco IGX 8400 series multiband, or Cisco IPX narrowband node in a Cisco WAN Manager workstation and downloading it to one or more nodes upon command. Data that defines the configuration of the network is stored in the controller card for each node in the network. This configuration data is variable and changes with each modification in the network topology.

The node configuration is stored in non-volatile memory in the controller cards. This memory, battery backup RAM (BRAM), allows unlimited reading/writing of data while providing protection against power fail.

Configuration Save/Restore allows the network administrator to send a copy of the configuration database to the Cisco WAN Manager NMS where it is archived for ultimate protection against a catastrophic failure. When needed, the archived copy is downloaded to the selected node.

In normal operations, the configuration data should be saved on a periodic basis so that the backup database is current. The save routine runs as a background process and is non-invasive. Should the node configuration database require restoring, the first part of the restore is also non-invasive. However, the final writing to BRAM and database rebuild is disruptive and will take the node out of service for several minutes.

Alarm Summary Relay Outputs

The optional alarm summary feature provided by the Alarm Relay cards provides both a front panel visual indication of a Cisco BPX 8600 series broadband, Cisco IGX 8400 series multiband, or Cisco IPX narrowband node alarm as well as a set of relay outputs (dry-contact) for indicating node and network alarms. A visual alarm history indication is also provided. These outputs are intended to be connected to the site alarm system. This alarm reporting is separate and in addition to the alarm output at the node's control port, which provides a data output to a control terminal, e.g., the Cisco WAN Manager NMS.

To implement this feature on Cisco IPX narrowband nodes requires an Alarm Relay Front Card (ARC) and an Alarm Relay Interface (ARI) back card. Cisco BPX 8600 series broadband switches come standard with an Alarm Status Monitor (ASM) front card and a corresponding Line Monitor (LM-ASM) back card. The features provided by both card sets are quite similar.

One set of alarm relays is used to display a major alarm or minor alarm for the node. One pair of contacts on each relay is used for audible alarms. These contacts are in series with a front panel alarm cut-off (ACO) switch. The other set of relay contacts is used for visual alarms and is not affected by the ACO switch. Any time the ACO switch is operated, a front panel ACO indicator is lit as a reminder to the operator. If the ACO switch is operated to disable the node audible alarm output and a second alarm is received, the audible alarm is reactivated. Two front panel LEDs provide local indication of network major or minor alarms.

Network Management

Network management can be performed from a single location using a CiscoWAN Manager Network Management Station. The NMS is connected to a single node in a flat network (or a single node in each domain of a structured network)) and communications throughout the network is provided by the internode controller channel that connects to each node.

Cisco WAN Manager collects network status information and operating statistics accumulated by each node and stores this data in its own on-line database. It is graphically displayed in real-time on a topological map on the NMS console with status displayed in color. In addition, a maintenance log is kept on the NMS terminal of all network changes and maintenance activity. In addition, system software updates and firmware revisions can be remotely downloaded to each node from the CiscoWAN Manager NMS.

Users with other SNMP network administration systems can gain direct access to Cisco BPX 8600 series broadband switch, Cisco IGX 8400 series multiband switch, Cisco IPX narrowband switch, and Cisco MGX 8220 edge concentrator status and statistics via optional software features. Contact Customer Service for more information. Each node in the network (or domain for structured networks) maintains an image of the network configuration in memory. A backup copy of the system software and configuration can be manually stored on the CiscoWAN Manager terminal.

NMS Interfaces

All Cisco BPX 8600 series broadband, Cisco IGX 8400 series multiband, and Cisco IPX narrowband switches provide a serial RS-232, communications port for either local or remote connection to a terminal for troubleshooting or local use. In addition, all Cisco BPX 8600 series broadband, Cisco IGX 8400 series multiband, and Cisco IPX narrowband switches and the Cisco MGX 8220 edge concentrator shelf may be connected via a high-speed Ethernet LAN port to NMS workstations. The StrataView Plus workstation requires a LAN connection to the Cisco WAN switching network. Network Management is implemented via in-band ILMI through the use of SNMP agents and proprietary and ATM MIBs in each node. Cisco's SNMP proxy agent is also available to connect Cisco WAN Manager NMS to SNMP-based management integrators.

Node Hardware Administration

The standard node hardware administration feature allows the Cisco BPX 8600 series broadband, Cisco IGX 8400 series multiband, and Cisco IPX narrowband software to recognize each card installed in the node. Hardware Administration also allows the system to test a card and switch to a redundant card or take the card off-line if the test fails. It provides the following alarm functions:

- Marking of cards as “failed” if a self test or background test fails.
- Generation of a major alarm if a card fails with no backup (non-redundant, or other redundant backup card has already failed).
- Generation of a minor alarm if a redundant standby card fails.
- Generation of a major alarm if a card with no backup is removed.

Network Trunk Administration

Network trunk administration is a standard node feature of the Cisco BPX 8600 series broadband switch, the Cisco IGX 8400 series multiband switch, and the Cisco IPX narrowband switch, allowing network trunks to be configured using software commands from a remote user interface. The trunk administration routines monitor all network FastPacket and ATM trunks for alarm conditions and collect performance statistics to assist in evaluating network performance. This data is routed to the Cisco WAN Manager terminal where it is stored in the NMS database.

Embedded SNMP Agent

**Note**

Embedded SNMP Agent is an optional feature that must be purchased and enabled on each node where it is to be used.

All Cisco BPX 8600 series broadband, Cisco IGX 8400 series multiband, and Cisco IPX narrowband switches include a Simple Network Management Protocol (SNMP) Agent as part of the system software. The SNMP Agent permits a SNMP Manager other than Cisco WAN Manager (e.g. HP OpenView) to view and set certain network objects in Management Information Bases (MIBs) maintained in each node allowing integrated network management systems (INMS) the capability to manage a CiscoWAN switching network with SNMP.

The SNMP Manager connects to one of the Cisco BPX 8600 series broadband, Cisco IGX 8400 series multiband, or Cisco IPX narrowband switches through the Ethernet LAN port. This node then acts as the gateway for the SNMP Manager to communicate with all other nodes. In flat networks, this communication extends to the whole network. With structured networks, the SNMP Agent maintains a MIB that extends throughout a single domain.

The embedded SNMP Agent supports the standard Internet MIB II and a proprietary MIB. Currently, the proprietary MIB contains status and configuration data for each frame relay connection in the managed Cisco BPX 8600 series broadband, Cisco IGX 8400 series multiband, and Cisco IPX narrowband network.

Multiple User Access to NMS



Note

Multiple User Access is an optional feature that must be purchased and enabled on each node where it is to be used.

A system software default allows only one user access to a node at a time on one of three terminal access ports (auxiliary, control, or LAN). The network manager may allow up to 40 simultaneous user command interfaces on a node. A second option of this feature is to allow user access on any or all of the terminal access ports simultaneous.

This option can be configured for multiple simultaneous users and multiple active terminal ports. Access can be either local or remote via the `vt` command. Default is one direct user and one `vt` user. In addition, a job may be running while a user session.

Integration of FastPAD into NMS

The FastPAD feeder multiplexer provides analog voice and data channels that are presented to the user as extensions of Cisco IGX 8400 series multiband and Cisco IPX narrowband facilities. Local user interface to the FastPAD for configuration, connection management, and status is via front panel controls and display. This can be inconvenient in a centrally managed network. This release provides a remote user interface to the FastPAD from the Cisco WAN Manager NMS workstation.

The FastPAD interfaces to Cisco WAN switching networks via a V.35 frame relay PVC. The Cisco IGX 8400 series multiband switch or Cisco IPX narrowband switch sends control information and gathers status via a special frame relay port on an FTM/FTI card set. It forwards this information to Cisco WAN Manager to integrate the FastPAD into the NMS topology map, maintenance log, administration window, and statistics databases. The associated Cisco IGX 8400 series multiband switch or Cisco IPX narrowband switch also collects statistics on the PVCs associated with the FastPAD.

One of the four ports on a FTM card is available to interface to a FastPAD, the other three can be used as normal frame relay ports like those found on the FRM card. A maximum of 64 FastPADs can be interfaced to a node.

Included in the FastPAD NMS are commands for card management, port management, connection management, and FastPAD management.

- Card management includes software to detect and manage the FTM/FPM cards.
- Port management includes upping and downing the FTM port, LMI alarms, EIA signalling, and port statistics.

- Connection management includes setting up PVCs between FastPADs, gathering statistics, connection maintenance such as testcon and loopback.
- FastPAD management commands are used to map a FastPAD onto an FTM port and establishing control sessions.

PART 2
NETWORKS

Networking Architecture

This chapter describes various types of networks or network applications that may be implemented with the Cisco WAN switching cell relay network switches, the Cisco BPX 8600 series broadband switch, the Cisco IGX 8400 series multiband switch, the Cisco IPX narrowband switch, the Cisco MGX 8220 edge concentrator, and associated equipment such as the Cisco 3800 and the FastPAD.

The Cisco BPX 8600 series broadband switch, Cisco IGX 8400 series multiband switch, and the Cisco IPX narrowband switch, which are sometimes referred to as nodes, are the basic network building blocks. The nodes are interconnected by transmission facilities called trunks. User data inputs to the network at nodes through various user devices.

This chapter contains the following:

- Broadband ATM Networks
- Cisco WAN Switching ATM Network Architectures
- Narrowband FastPacket Networks
- Frame Relay Networks
- Networks with the FastPAD
- International Networks
- Network Synchronization

Broadband ATM Networks

Cisco WAN switching 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.

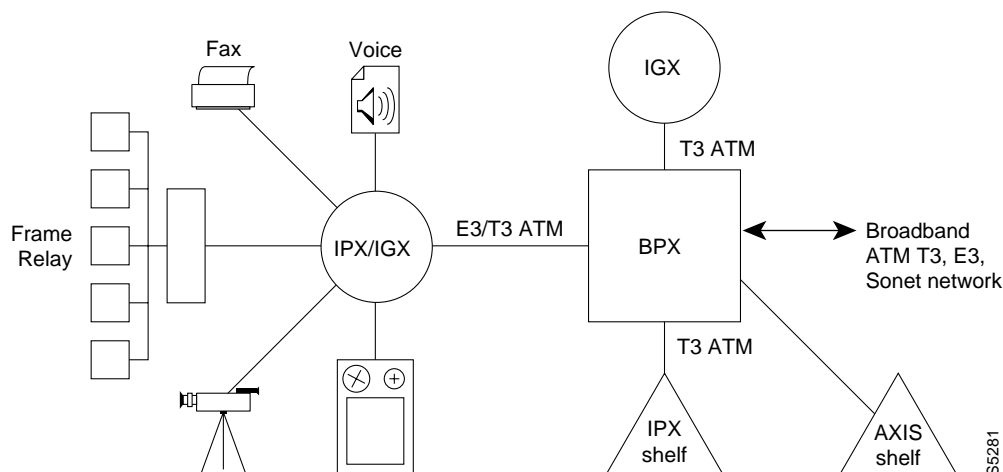
Multi-Media Networking

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 Cisco IPX narrowband switch, the Cisco 3800, the Cisco MGX 8220 edge concentrator, Cisco IGX 8400 series multiband switch, and the ATM connections to the Cisco BPX 8600 series broadband switch (). Since the Cisco IPX narrowband switch has always been a cell-based node, it handles multimedia inputs with a maximum of efficiency. The Cisco IGX 8400 series multiband switch provides all the capabilities of the Cisco IPX narrowband switch but with additional multiband and multiservice capabilities.

Figur e5-1 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 Cisco BPX 8600 series broadband switch currently supports T3, E3, OC3, and OC12 rates.

Figure 5-2 illustrates a possible backbone configuration using T3, E3, OC3, or OC12 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/OC12 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

Growth Path for Existing Cisco IPX Narrowband/Cisco IGX 8400 Serie Multiband Networks

Another application for the Cisco BPX 8600 series broadband switch is to provide a growth path for existing Cisco IPX narrowband or Cisco IGX 8400 series multiband networks. As network operators experience rapid growth in data communications requirements, existing Cisco IPX narrowband or Cisco IGX 8400 series multiband networks may become heavily meshed with tandem T1 connections approaching the capacity of the network.

Cisco BPX 8600 series broadband switches may be added to existing Cisco IPX narrowband or Cisco IGX 8400 series multiband networks to greatly expand the network capacity and relieve congestion. The Cisco IPX narrowband or Cisco IGX 8400 series multiband switches 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 Cisco IGX 8400 series multiband or Cisco IPX narrowband network ports. See for a typical example.

Figur e5-2 High-Capacity Backbone Configuration

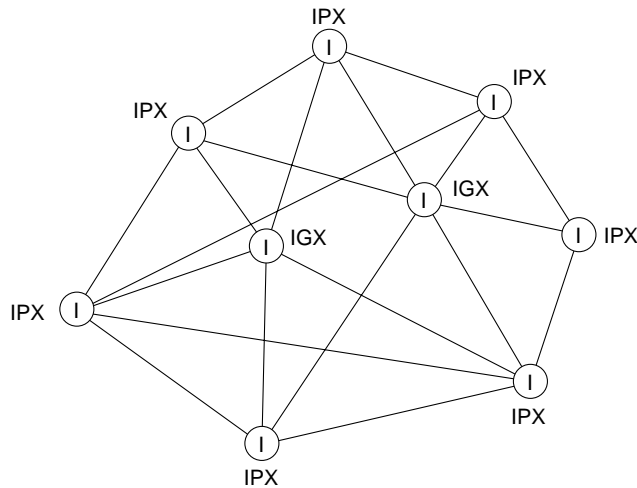


S5282

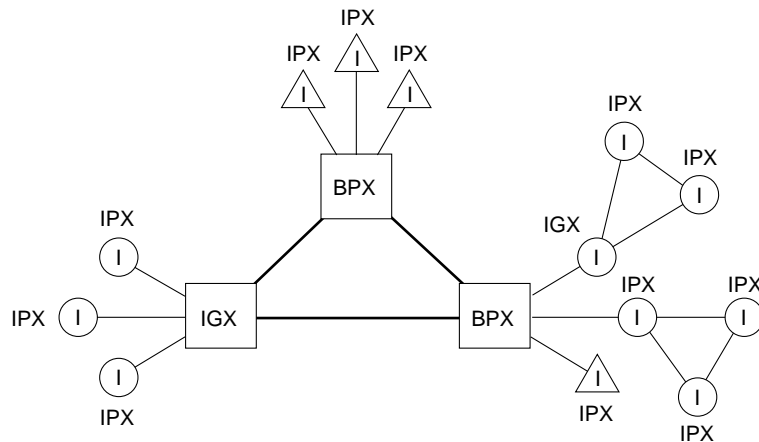
A Cisco BPX 8600 series broadband switch operates as a peer with other Cisco IGX 8400 series multiband switches and Cisco IPX narrowband switches in a mixed Cisco WAN switching network. From the network operator's perspective, the Cisco BPX 8600 series broadband switch appears merely as a higher capacity Cisco IGX 8400 series multiband switch or Cisco IPX narrowband switch.

Cisco IPX narrowband nodes are connected to a Cisco BPX 8600 series broadband switch with the addition of an AIT card set. Similarly, Cisco IGX 8400 series multiband switches are connected to a Cisco BPX 8600 series broadband switch with the addition of a BTM card set. The AIT card set operates at T3 or E3 rates and converts the narrowband 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 Cisco BPX 8600 series broadband switch will be installed at sites where there are already several Cisco IGX 8400 series multiband switches or Cisco IPX narrowband switches located, for example, at a network hub or large switching center and connect to these nodes with a short wire connection.

Figur e5-3 Upgrading Existing Cisco IGX 8400 Series Multiband/IPX Networks



a. Heavily-meshed IPX/IGX network



b. BPX Add to IPX/IGX network

S5172

Integrating a Cisco BPX 8600 series broadband backbone with its higher speed trunking and switching capacity expands the network capacity with minimal disruption of the existing network. The Cisco BPX 8600 series broadband switch provides expansion capabilities with full compatibility with the Cisco IGX 8400 series multiband switch and the Cisco IPX narrowband switch. The Cisco BPX 8600 series broadband switch also utilizes the same network management platform as the Cisco IGX 8400 series multiband switch and the Cisco IPX narrowband switch.

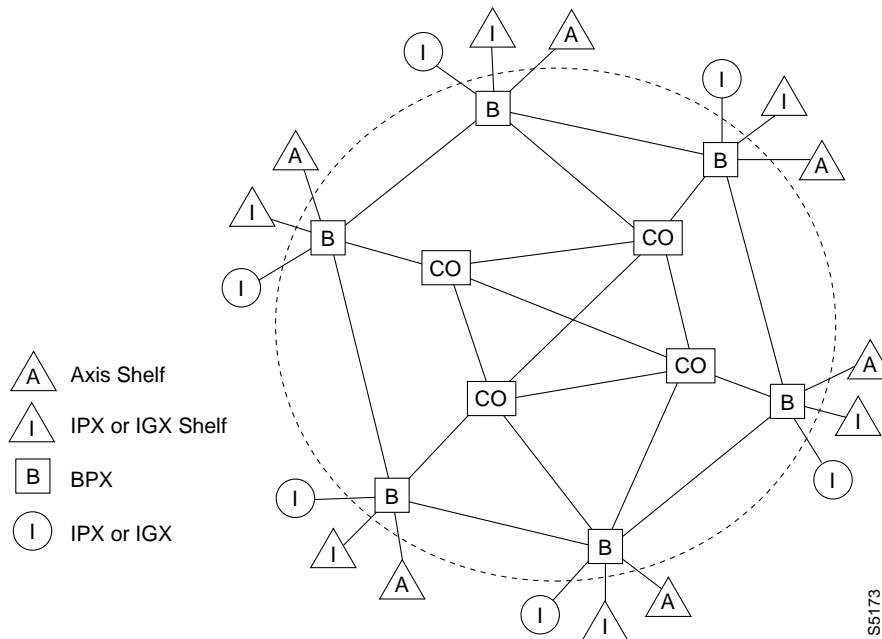
The actual bandwidth used on the trunks between a Cisco BPX 8600 series broadband switch and an Cisco IGX 8400 series multiband switch or Cisco IPX narrowband switch may be incrementally increased in T1 or E1 size steps. Therefore, existing Cisco IGX 8400 series multiband T1 and E1 trunks can be decommissioned one at a time and their cell traffic routed across the new Cisco BPX 8600 series broadband network backbone. The ATM interface to the Cisco IGX 8410 multiband switch economically serves small feeder locations and provides all of the same interfaces as the larger Cisco IGX 8420 or 8430 multiband switches.

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. Cisco BPX 8600 series broadband switches 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 ATM service nodes, Cisco BPX 8600 series broadband switches with Cisco MGX 8220 edge concentrator shelves, Cisco IGX 8400 series multiband switches and Cisco IPX narrowband switches configured as shelves, as well as Cisco IGX 8400 series multiband switches and Cisco IPX narrowband switches, provide ATM access and termination closer to the user and transport ATM cells to central office ATM platforms. Each Cisco BPX 8600 series broadband switch acts as a concentrator for the narrowband inputs from a number of local user devices. The Cisco IGX 8400 series multiband switch, the Cisco IPX narrowband switch, and the Cisco MGX 8220 edge concentrator convert these inputs to cells that are forwarded to the Cisco BPX 8600 series broadband switches over T3 or E3 trunks. Local connections can be switched back to the outer layer of Cisco IGX 8400 series multiband switches and Cisco IPX narrowband switches, and longer-distance connections can be switched to the public network at a suitably equipped central office.

Figur e5-4 Service Node for Large C.O. Switches

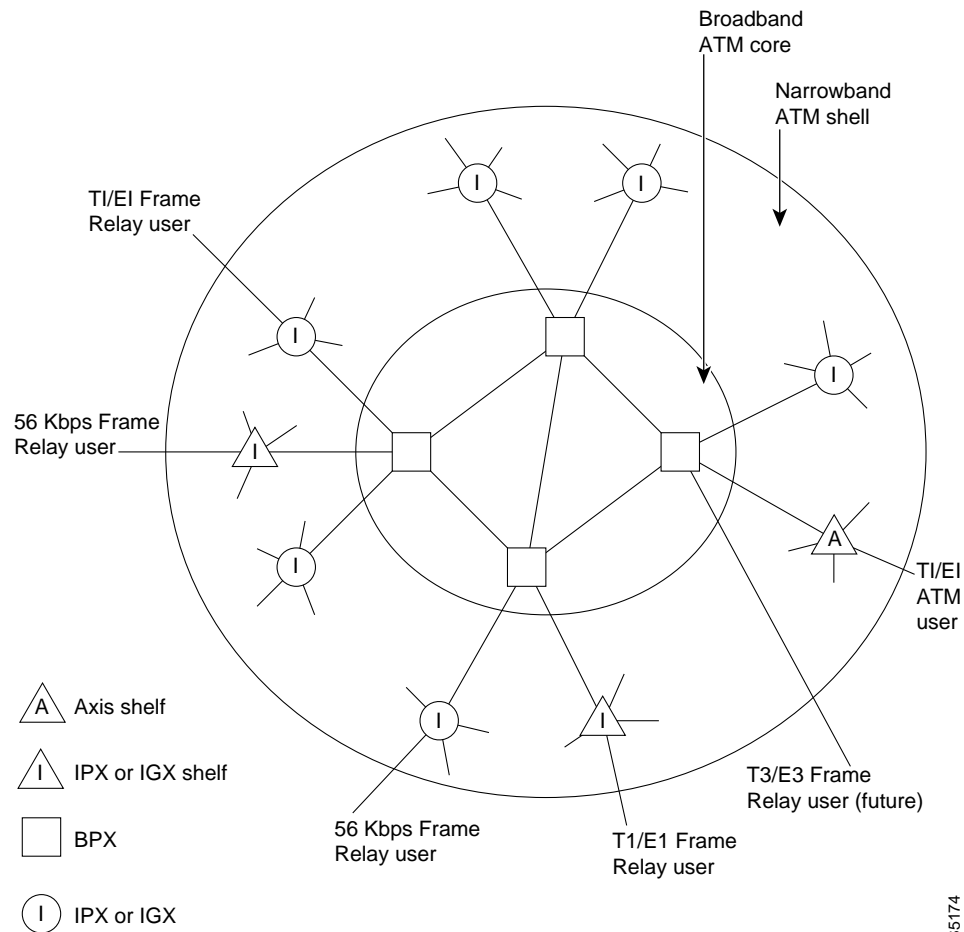


Expansion of Wide-Area Data Networks

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 . The narrowband shell, consisting primarily of Cisco IPX narrowband tributaries, Cisco IGX 8400 series multiband switches and Cisco IPX narrowband switches configured as shelves, and Cisco IGX 8400 series multiband switches configured as routing 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 Cisco IGX 8400 series multiband switches and Cisco IPX narrowband switches.

Figur e5-5 ATM Network of the Future



For frame relay connections terminating at more remote destinations, the Cisco IGX 8400 series multiband switches and Cisco IPX narrowband switches would feed this traffic to the higher-speed broadband core network utilizing the very high-speed tandem switching characteristic of the Cisco BPX 8600 series broadband switches. 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 Cisco IGX 8400 series multiband switches can communicate with adjacent Cisco IGX 8400 series multiband switches 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, Cisco 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 DS1, E1, DS3, and E3 rates. The UFM cards on the Cisco IGX 8400 series multiband switch support frame relay T1/E1, HSSI, X.21, and V.35 interfaces.

T2 ATM Network Connections

A 6.312 Mbps T2 digital trunk interface is available from Cisco. This is for markets where 6 Mbps trunks are more prevalent than T3 trunks (e.g., the Japanese markets). T2 trunks provide four times the bandwidth as T1 trunks. Since the Cisco IGX 8400 series multiband, Cisco IPX narrowband, or Cisco BPX 8600 series broadband network interface is either T3 or E3, a rate adapter is required.

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

Figur e5-6 T2 ATM Trunking



Cisco WAN Switching ATM Network Architectures

Using the standardized UNI header allows the Cisco IGX 8400 series multiband switch, the Cisco IPX narrowband switch, or the Cisco BPX 8600 series broadband switch to be used in a generic ATM network (cloud) with a mix of various equipment. With the use of BXM cards in trunk mode, Cisco BPX 8600 series broadband switches can be connected using the standard NNI header. When using BNI cards, to take advantage of several features unique to CiscoWAN switching 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 Figur e5-7.


These are the five types of ATM network connections supported:

- Direct connection of two or more Cisco BPX 8600 series broadband switches via T3/E3/OC3/OC12 lines. Provides broadband networking for connections originating on Cisco IGX 8400 series multiband, Cisco IPX narrowband, and Cisco BPX 8600 series broadband service interfaces.
- The connection of Cisco BPX 8600 series broadband switches to a generic, standards-based, ATM network. This allows the Cisco BPX 8600 series broadband switch to be a part of a non-Cisco WAN switching network, typically using virtual trunking.
- Direct connection of two Cisco IGX 8400 series multiband (or Cisco IPX narrowband) switches via a T3/E3 line. This provides a point-to-point broadband trunk between Cisco IGX 8400 series multiband switches. Uses StrataCom Interface (STI) cell header to provide ForeSight.
- Indirect connection of Cisco IGX 8400 series multiband (or Cisco IPX narrowband) switches via an ATM network. This provides a broadband interface into a non-Cisco WAN switching ATM network. Uses User Network Interface (UNI) cell header.

- Indirect connection of two Cisco IGX 8400 series multiband (or Cisco IPX narrowband) switches via a Cisco BPX 8600 series broadband network. This provides a switched broadband trunk between Cisco IGX 8400 series multiband switches. Uses STI cell header to provide ForeSight.

Figur e5-7 Cisco WAN Switching ATM Network Architectures


BPX Direct connection




BPX Indirect connection



AIT Direct connection




 The ATM Trunk Card is an AIT on the IPX and a BTM on the IGX.

AIT Indirect connection



IPX to BPX ATM



S6154

ATM Service Interfaces, BXM, ASI, and Cisco MGX 8220 Edge Concentrator Shelves

The Cisco BPX 8600 series broadband/Cisco IGX 8400 series multiband ATM networks have a flexible architecture and configuration interface that allows the user to define service parameters to tailor their connection types, including CBR, VBR, ABR, and UBR. Up to 32 different classes of service can be predefined for various connection types.

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.

ATM connections can be provisioned to support 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.

BXM and ASI Interface

ATM service (connections) can be provisioned directly on a Cisco WAN switching ATM network through the use of a BXM T3/E3, 155, or 622 card configured for port mode (service access, UNI), or of an ASI-1 or ASI-155 card. The ATM Service Interfaces provides direct connection of high-speed user devices (such as routers or hubs) with ATM standard interfaces to Cisco WAN switching ATM wide-area networks. The BXM cards support ATM Traffic Management 4.0 VSVD congestion flow control for ABR connections.

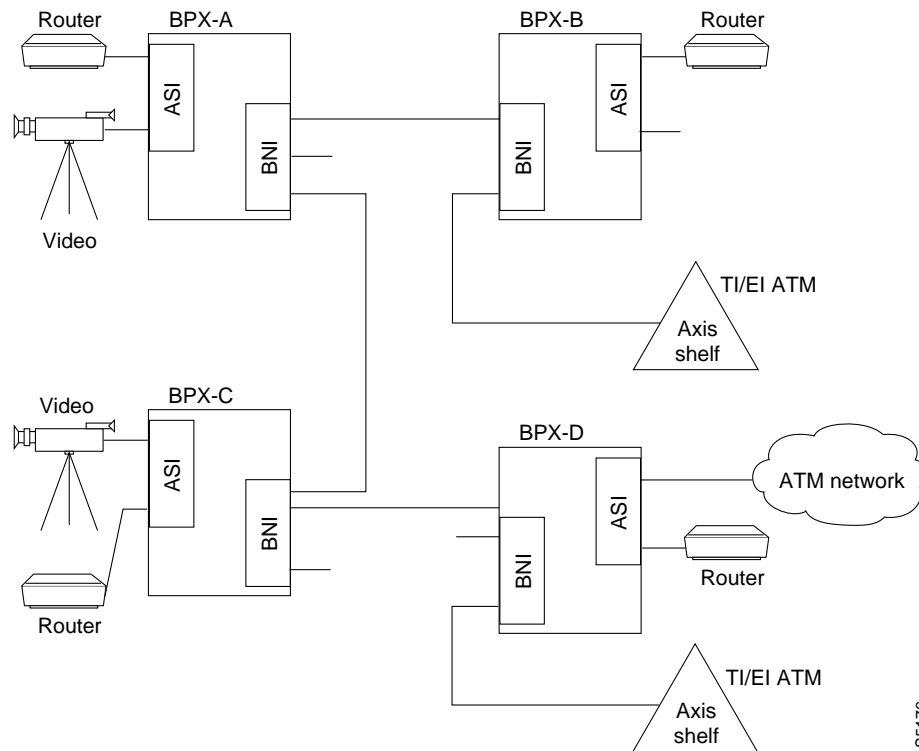
Each ASI-1 card set, installed in the Cisco BPX 8600 series broadband switch, provides two T3 or DS3 ATM service interfaces (ports), and each ASI-155 card set provides two OC3/STM-1 ATM service interfaces (ports). BXM T3/E3 card sets are available with either 8 or 12 service interfaces, BXM-155 card sets are available with either 4 or 8 ports with MMF, SMF, or SMFLR interfaces. BXM-622 card sets are available with either 1 or 2 ports with SMF or SMLR interfaces.

The ports on the BXM or ASI cards can be configured as a User-to-Network (UNI) interface for directly interfacing user ATM devices to the ATM network as shown in . Or the ports may be configured as a Network-to-Network Interface for interface between two dissimilar ATM networks (Figure 5-8).

Cisco MGX 8220 Edge Concentrator Shelf

The Cisco MGX 8220 edge concentrator shelf provides an interface to the Cisco BPX 8600 series broadband switch for T1/E1 frame relay, T1/E1 ATM, CES, and FUNI connections. With interworking, the frame relay connections may be terminated at another Cisco MGX 8220 edge concentrator, an ASI on a Cisco BPX 8600 series broadband switch, a Cisco IGX 8400 series multiband switch, or a Cisco IPX narrowband switch. The Cisco MGX 8220 edge concentrator shelf T1/E1 ATM connections may be routed to another Cisco MGX 8220 edge concentrator or to a BXM or an ASI on a Cisco BPX 8600 series broadband switch.

Figur e5-8 Direct ATM User-to-Network Connections



Narrowband FastPacket Networks

The Cisco IGX 8400 series multiband switch and the Cisco IPX narrowband switch 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

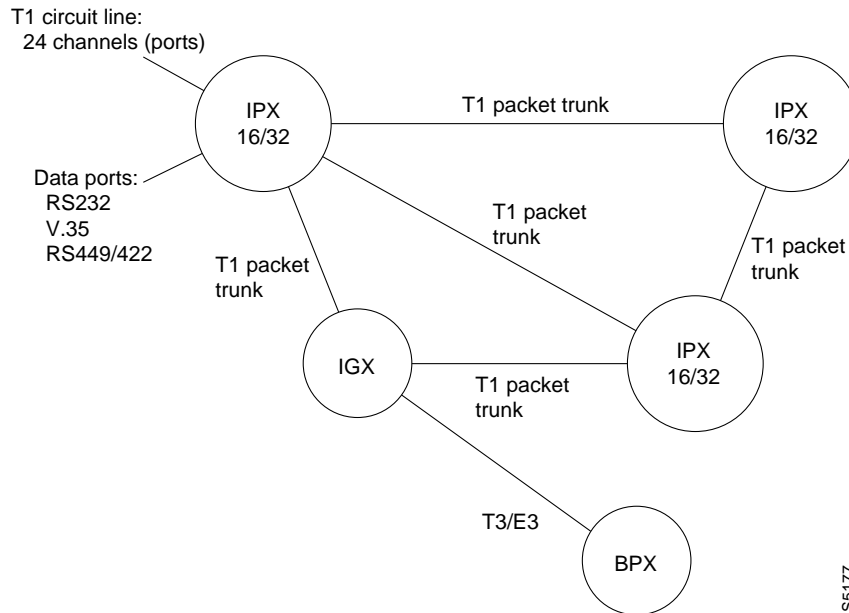
A simplified example of a Cisco IPX 16/32 narrowband T1 network is shown in Figure 5-9. The various node, consisting of Cisco IPX 16/32 narrowband T1 network switching bandwidth managers, are interconnected with T1 packet trunks that carry packets of data at 1.544 Mbps. The data input to the

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 Cisco IGX 8400 series multiband switch and the Cisco IPX narrowband switch:

Comment: In the following, the FRM, HDM, and CVM cards perform the same functions in the Cisco IGX 8400 series multiband switch as the FRP, SDI, and CDP do in the Cisco IPX narrowband switch.

- 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 on the Cisco IPX narrowband switch using a FRP with either a FRI-V.35, FRI-X.21, or FRI-T1/E1 back card.
- The T1/E1 frame relay port interfaces are implemented on the Cisco IGX 8400 series multiband switch using a FRM-C with either a UFI-8T1-DB15, UFI-8E1-DB-15, or a UFI-8E1-BNC back card.
- The HSSI, V.35, and X.21, frame relay port interfaces are implemented on the Cisco IGX 8400 series multiband switch using a UFM-U with a variety of HSSI, V.35, and X.21 back cards.
- 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.

Figur e5-9 Typical T1 IPX Network

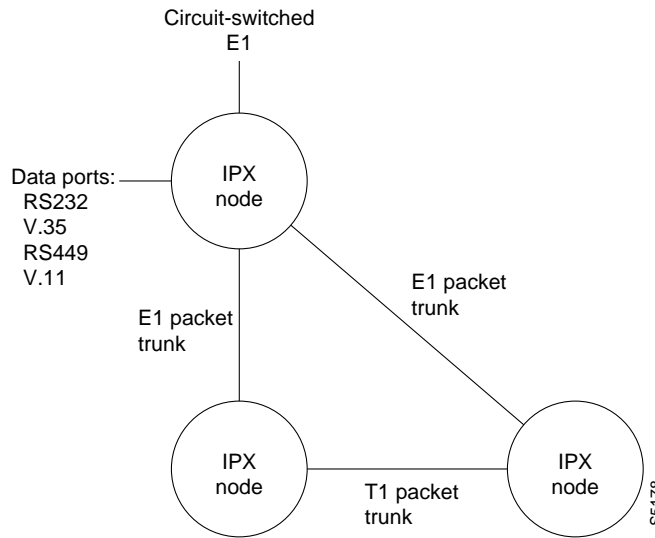
E1 Networks

A simplified example of a Cisco IGX 8400 series multiband E1 network is shown in Figure 5-10. Each Cisco IGX 8400 series multiband switch 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 Cisco IGX 8400 series multiband and Cisco IPX narrowband E1:

Comment: In the following, the FRM, HDM, and CVM cards perform the same functions in the Cisco IGX 8400 series multiband switch as the FRP, SDI, and CDP do in the Cisco IPX narrowband switch.

- The 2.048 Mbps standard interface as specified in G.703. This interface is available on the circuit line and trunk side of the Cisco IGX 8400 series multiband switch.
- 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.

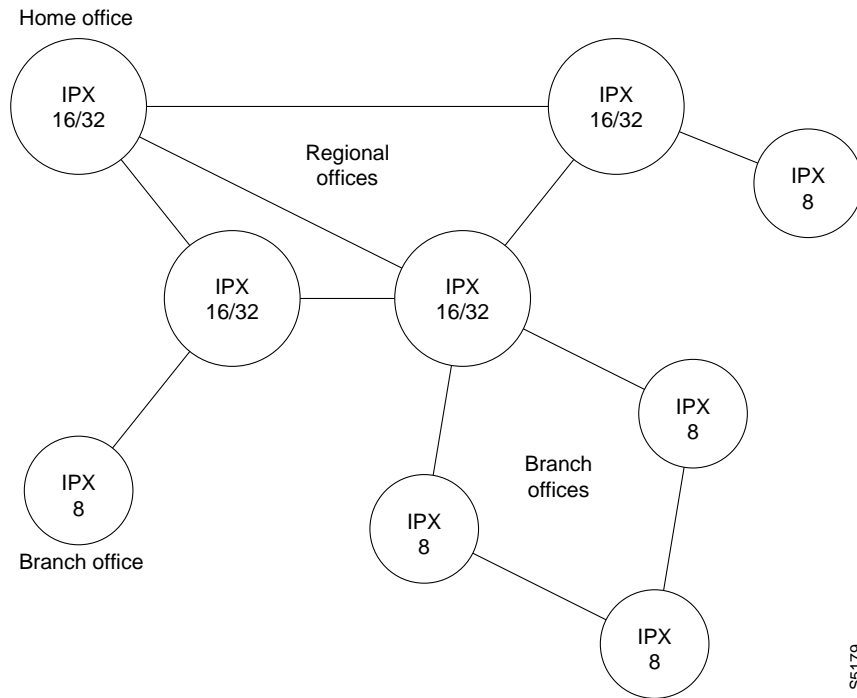
Figur e5-10 E1 IPX Network




Networks with an Eight Slot Cisco IPX Narrowband Switch

The eight slot Cisco IPX narrowband switch operates seamlessly in existing and new networks with Cisco'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 eight slot Cisco IPX narrowband switch can be used in large networks where there is a mixture of eight slot Cisco IPX narrowband switches, sixteen slot Cisco IPX narrowband switches, sixteen slot Cisco IPX narrowband switches and Cisco IGX 8400 series multiband switches (). It can be used as a feeder () to a larger node, or it can be configured as a Cisco IPX narrowband switch shelf and connected to a Cisco BPX 8600 series broadband switch, as can a Cisco IGX 8400 series multiband switch. 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 a Cisco IPX narrowband backbone network.
- Connection of a redundant loop of branch offices to a regional office on a Cisco IPX narrowband backbone network and one or two T1 ports. Redundant systems provide automatic switchover on equipment or network failure.

Figur e5-11 Typical IPX Network with IPX 8 Remote Nodes

S5179

Figur e5-12 IPX 8 Used as a Feeder in Large Networks

S5180

Japanese J1/Y1 Networks

The BC-J1 and BC-Y1 back cards provide Cisco IGX 8400 series multiband switches and Cisco IPX narrowband switches 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 NTM in an Cisco IGX 8400 series multiband switch (NTC in a Cisco IPX narrowband switch) front card to provide an interface for a trunk in a Cisco IGX 8400 series multiband or the Cisco IPX narrowband 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 CVM on a Cisco IGX 8400 series multiband switch (CDP in the Cisco IPX narrowband switch) front

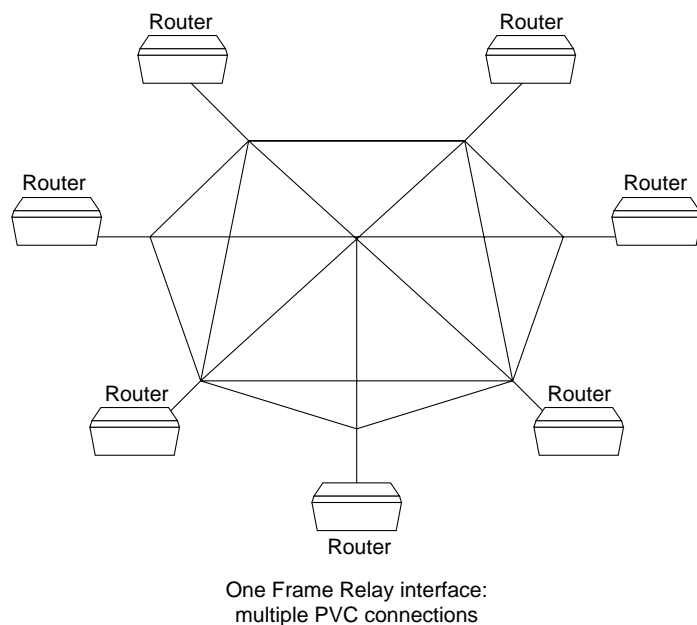
card to provide an interface for a circuit line to a PBX using a Japanese standard TTC-JJ-20, 2.048 Mbps circuit (J1 circuit). In all other aspects, these networks are similar to the T1 and E1 networks just described.

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, 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.

Figur e5-13 Frame Relay PVCs Provide Full Mesh Connectivity



Cisco WAN switching frame 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 Interface

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 CiscoWAN switching node can operate at speeds from 56 Kbps up to 2 Mbps.

Frame relay ports are added to the Cisco IPX narrowband switch using a Frame Relay PAD (FRP) front card and associated Frame Relay Interface (FRI) back card or to the Cisco IGX 8400 series multiband switch using the Frame Relay Module (FRM) and FRI back card.

Frame relay ports are added to Cisco IGX 8400 series multiband switches using FRM, UFM-C, or UFM-U front cards and a variety of T1/E1, HSSI, X.21, and V.35 back cards, as applicable.

On the Cisco MGX 8220 edge concentrator, the frame relay interface is provided by the FRSM card. The Cisco IGX 8400 series multiband and the Cisco IPX narrowband card sets provide four or more UNI ports for interface to frame relay routers, hubs, or bridges as indicated in the following:

- 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 Cisco IGX 8400 series multiband switch or the Cisco IPX narrowband switch, 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 (FRM or FRP/FRI) to achieve maximum capacity.

Port Concentrator

The Port Concentrator provides up to 44 low-cost frame relay ports and connects to the Cisco IPX narrowband switch via an FRP-2 card or to the Cisco IGX 8400 series multiband switch 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.

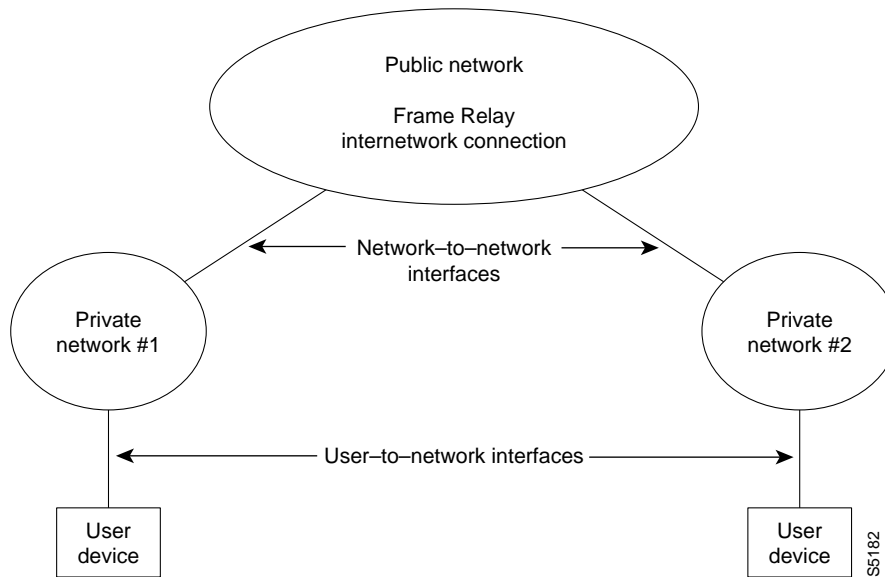
Frame Relay Network Interfaces, Cisco IPX Narrowband Switch, Cisco IGX 8400 Series Multiband Switch, FastPAD, Cisco MGX 8220 Edge Concentrator

There are two types of network interfaces protocols that are supported by CiscoWAN switching 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 Cisco WAN switching wide area network carrying the frame relay traffic. Currently, the Cisco IGX 8400 series multiband switch and the Cisco IPX narrowband switch 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 *Cisco MGX 8220 Installation and Configuration* and *Cisco MGX 8220 Command Reference publications* and the *FastPAD Reference publication* for further information on the protocols they support.

Figur e5-14 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 Cisco WAN switching network and another network and may or may not consist of Cisco WAN switching 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, both the Cisco IGX 8400 series multiband switch and the Cisco IPX narrowband switch support UNI via the following protocols: StrataCom LMI/ELMI, ITU-T Q.933 Annex A, and ANSI T1.617 Annex D.



Note

ELMI is an enhancement to LMI. ELMI adds capabilities that are not currently supported in LMI so that network switches, e.g., Cisco BPX 8600 series broadband switches, Cisco IGX 8400 series multiband switches, etc., can inform a user (routers, bridges, etc.) about network parameters such as various quality of service (QoS) parameters. Depending on the implementation, these might be such parameters as Committed Information Rate (CIR), Committed Burst Size (Bc), Excess Burst Size (Be), maximum Frame Size, etc.

LMI transmits on a logical connection between the Cisco IGX 8400 series multiband switch and the Cisco IPX narrowband switch and the user device (router) separate from the data path using DLCI 1023. This connection is a special PVC, carrying messages between the Cisco IGX 8400 series multiband switch or the Cisco IPX narrowband switch 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 Cisco IGX 8400 series multiband/Cisco IPX narrowband network.

Network to Network Interface (NNI)

Currently, Cisco IGX 8400 series multiband and the Cisco IPX narrowband 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 Cisco WAN switching frame relay network and another independent, Cisco WAN switching or a non-Cisco WAN switching 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 Cisco WAN switching 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 the first paragraph of this section. 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 Cisco WAN switching 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) port 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 . 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.

Figur e5-15 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 Cisco IGX 8400 series multiband switch or a Cisco IPX narrowband switch. 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 Cisco WAN switching 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.

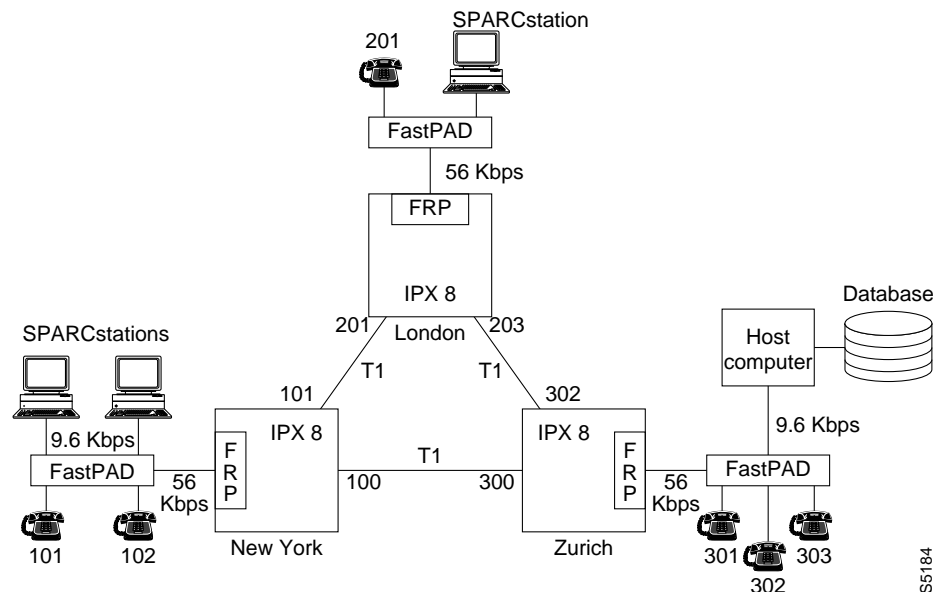
Similarly, a FastPAD network provides a cost-effective way of linking a number of small LANs, some of which may consist of only one or two terminals, without having to purchase a separate router at each location. The FastPAD acts very much like a router itself.

FastPAD Access to IGX Multiband or IPX Narrowband 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 the Cisco IPX narrowband switch, or the FRM/FRI interface of the Cisco IGX 8400 series multiband switch, Permanent Virtual Circuits (PVCs) can be established between Cisco IGX 8400 series multiband switches or the Cisco IPX narrowband switches to connect many FastPADs using a frame relay network.

An example of a simple network using FastPADs across a Cisco IGX 8400 series multiband or a Cisco IPX narrowband frame relay network is illustrated in . 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.

Figur e5-16 Cisco IPX Narrowband Network with FastPAD Access



In the example of , 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 Cisco IGX 8400 series multiband switch via an FRM, the FastPAD is considered to be a frame relay user device attached to a frame relay network device. When linking a FastPAD to a Cisco IGX 8400 series multiband switch, 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 Cisco IGX 8400 series multiband frame relay port include 56, 64, 112 or 128 Kbps. Refer to the *FastPAD User's Guide* for additional configuration information.

International Networks

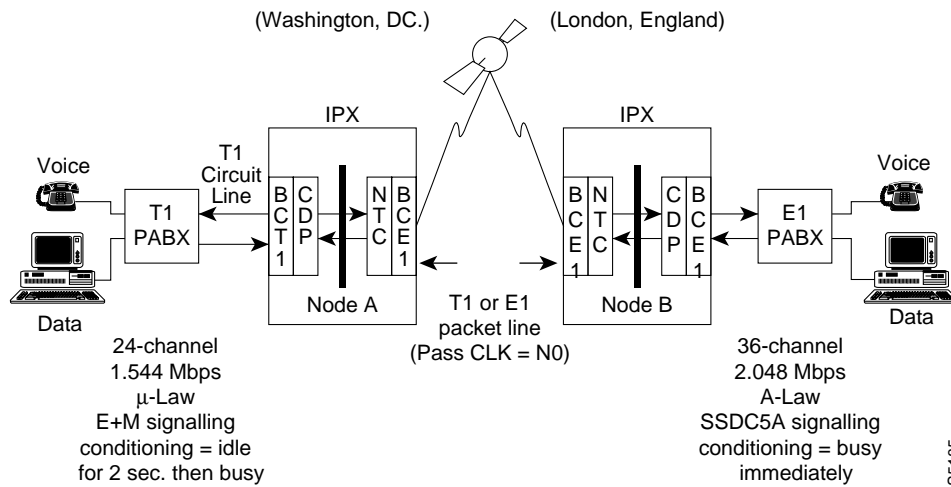
Cisco WAN switching cell-relay networks can easily be configured to provide direct communication links between domestic and international locations each using the differing transmission standards, e.g., T1 and CEPT E1.

For example, in , the Cisco IPX narrowband 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 Cisco IGX 8400 series multiband switch as the FRP, SDI, and CDP do in the Cisco IPX narrowband switch.

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 can be programmed by the network operator within each CDP.

Figur e5-17 IPX International Network



Example of an 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 E1 - 24 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 Example of a Pleisiochronous Network). the Cisco IPX narrowband and Cisco IGX 8400 series multiband switches 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 Cisco IGX 8400 series multiband switch as the FRP, SDI, and CDP do in the Cisco IPX narrowband switch.

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. Both Cisco IPX narrowband and Cisco IGX 8400 series multiband networks handle the code conversions, varying data rates, and different signalling formats through software conversion and inherent transparency of the packet network.

The Cisco 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 Cisco IPX narrowband 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 Cisco IPX narrowband node, and a CDP must also be installed in the CEPT Cisco IPX narrowband 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 (cnfxmtsiz) 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

Cisco WAN switching 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 Cisco IGX 8400 series multiband switch, the Cisco IPX narrowband switch, or the Cisco BPX 8600 series broadband switch reverts to its internal oscillator.
6. If no clocking hierarchy is defined in the network, the Cisco IGX 8400 series multiband switch 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 (**dsprtk**) 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.

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.

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.

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.

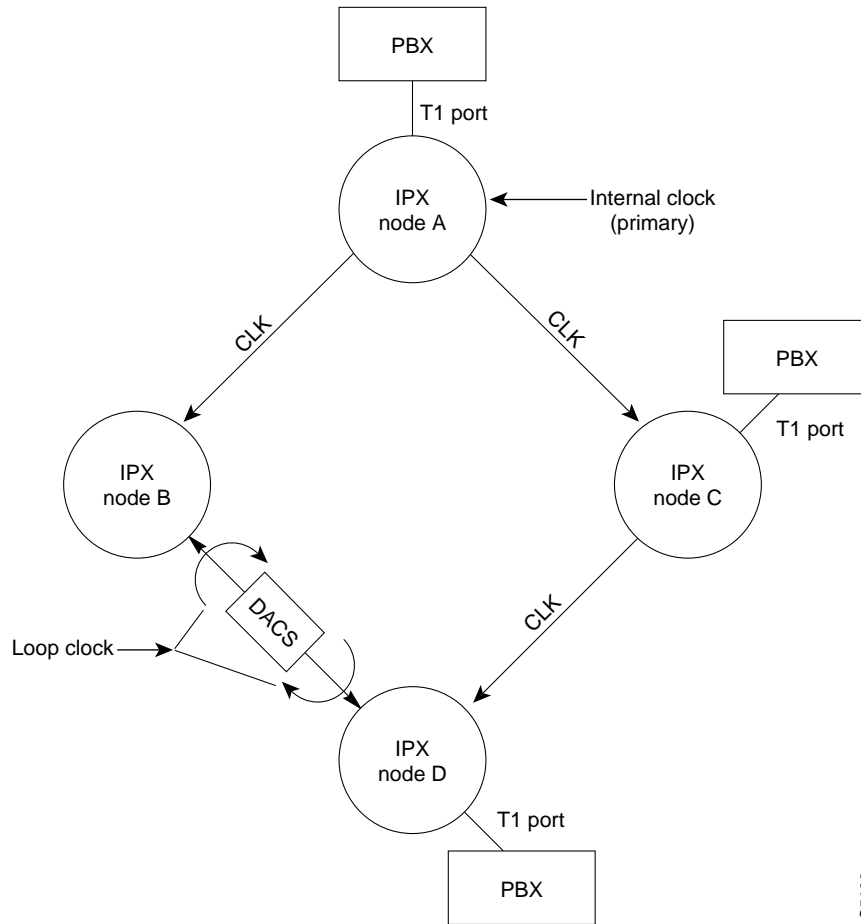
Pleisiochronous Network

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.

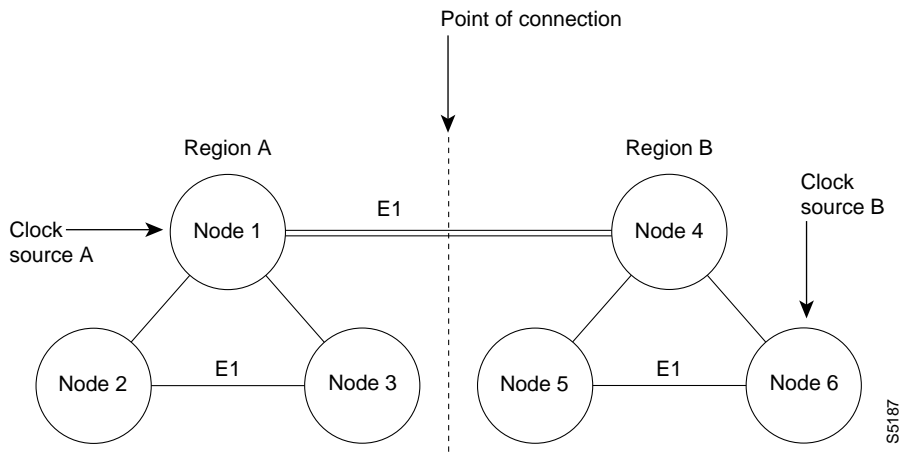
The Common Carriers each have their own master clocks and, to date, they are not synchronized to each other. For this reason, it is best to avoid, if possible, mixing packet trunks from different carriers. If this is unavoidable, consider using the CDP or NTC cards. These cards can turn clock around 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 A 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.

Figur e5-18 Example of IPX Network Clock Architecture



Figur e5-19 Example of a Pleisiochronous Network



Tiered Networks

This chapter describes the tiered network architecture that supports interface shelves (non-routing nodes) connected to Cisco WAN switching routing network.

The chapter contains the following:

- Introduction
- Cisco IGX 8400 Series Multiband Switches as Routing Hubs in a Tiered Network
- User Interface Commands
- Cisco BPX 8600 Series Broadband Switches as Routing Hubs in a Tiered Network
- Cisco WAN Manager NMS

With Release 8.5 and beyond, tiered networks now support voice and data connections as well as frame relay connections. With this addition, a tiered network can now provide a multi-service capability (frame relay, circuit data, voice, and ATM). By allowing CPE connections to connect to a non-routing node (interface shelf), a tiered network is able to grow in size beyond that which would be possible with only routing nodes comprising the network.

Routing Hubs and Interface Shelves

In a tiered network, interface shelves at the access layer (edge) of the network are connected to routing nodes via feeder trunks (Figure 6-1). Those routing nodes with attached interface shelves are referred to as routing hubs. The interface shelves, sometimes referred to as feeders, are non-routing nodes. The routing hubs route the interface shelf connections across the core layer of the network.

The interface shelves do not need to maintain network topology nor connection routing information. This task is left to their routing hubs. This architecture provides an expanded network consisting of a number of non-routing nodes (interface shelves) at the edge of the network that are connected to the network by their routing hubs.

Cisco BPX 8600 Series Broadband Switches and Cisco IGX 8400 Series Multiband Switches Configured as Routing Hubs

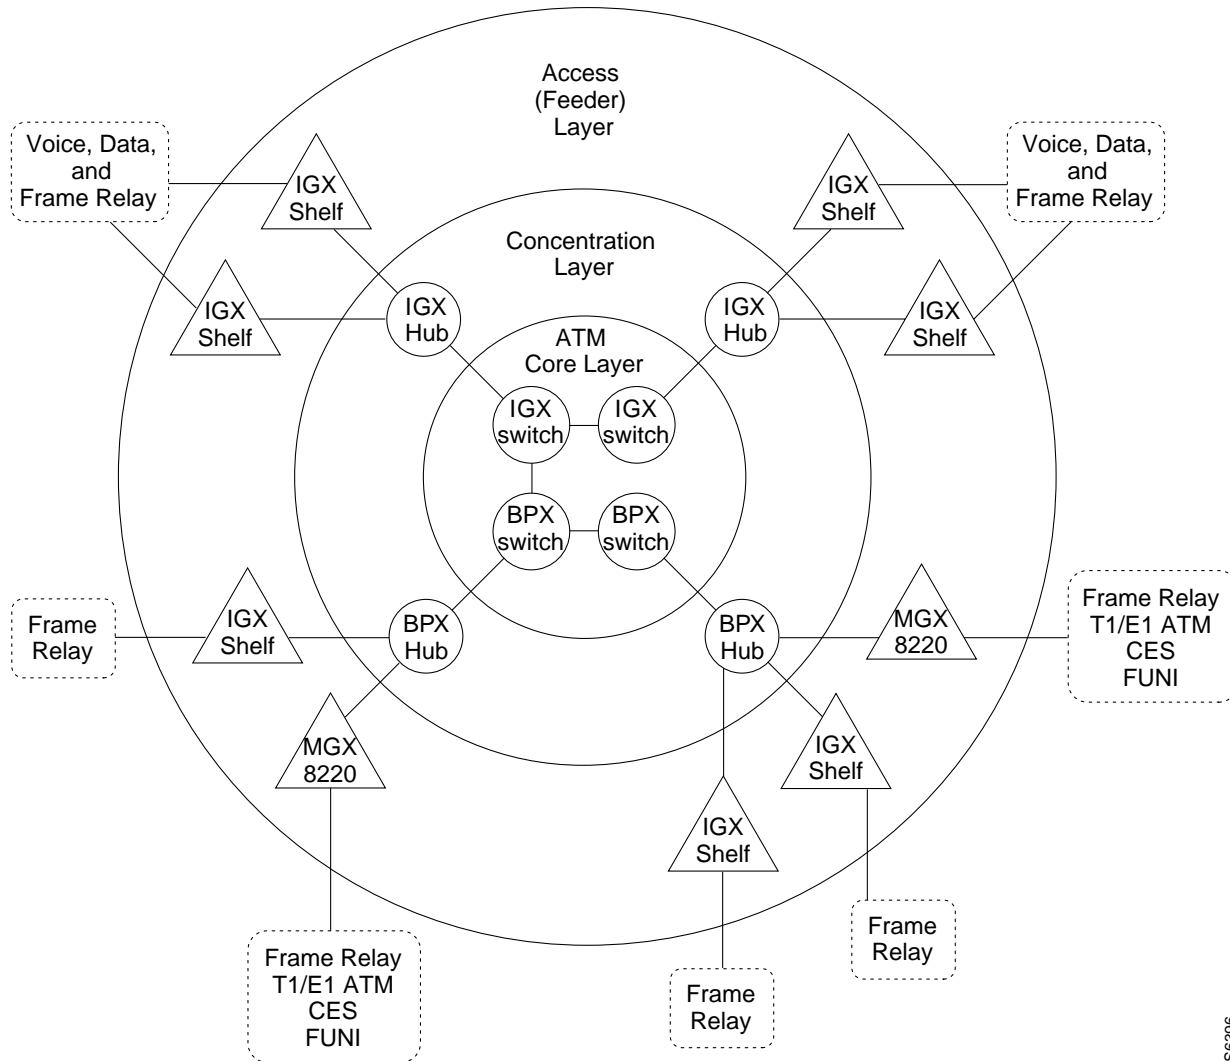
Voice and data connections originating and terminating on Cisco IGX 8400 series multiband switches configured as interface shelves (feeders) are routed across the routing network via their associated Cisco IGX 8400 series multiband switches configured as routing hubs. Intermediate routing nodes must be Cisco IGX 8400 series multiband switches.

Frame relay connections originating at IPX interface shelves and frame relay, ATM, CESM, and FUNI connections originating at Cisco MGX 8220 edge concentrator interface shelves are routed across the routing network via their associated Cisco BPX 8600 series broadband routing hubs.

**Note**

The Cisco IGX 8400 series multiband switch may also be configured as an interface shelf feeding frame relay connections to a Cisco BPX 8600 series broadband routing hub.

Figur e6-1 Tiered Network with Cisco BPX 8600 Series Broadband and Cisco IGX 8400 Series Multiband Routing Hubs

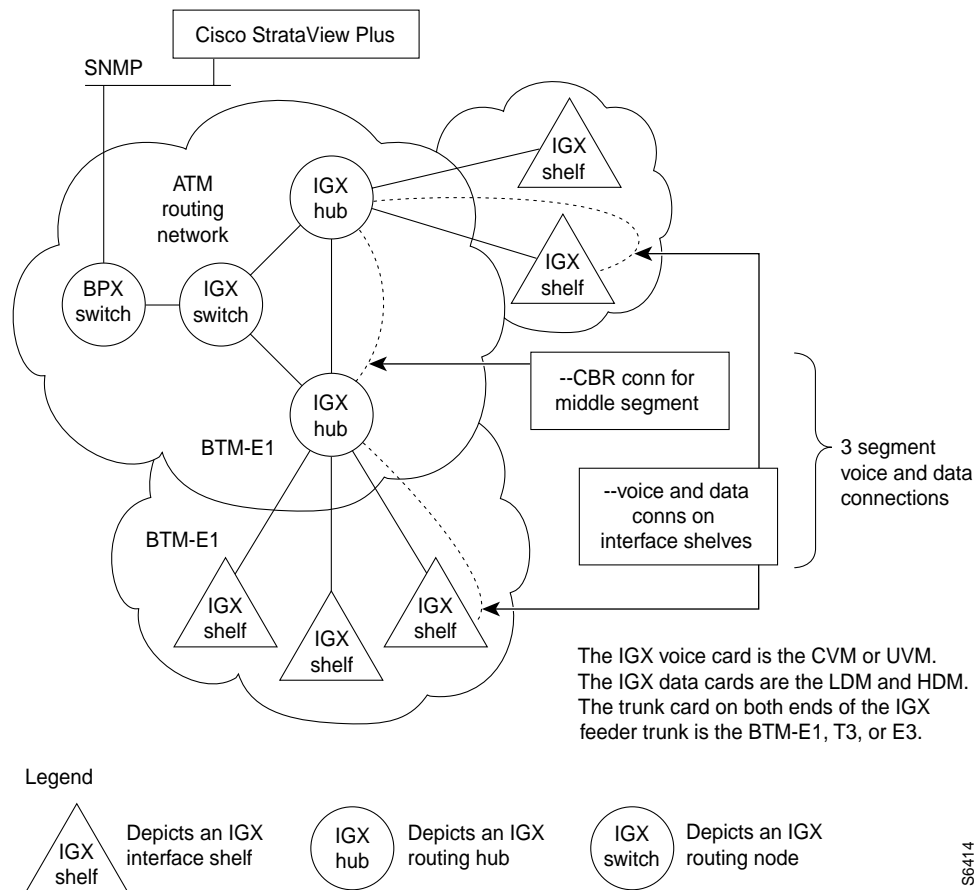


S6396

Cisco IGX 8400 Series Multiband Routing Hubs in a Tiered Network

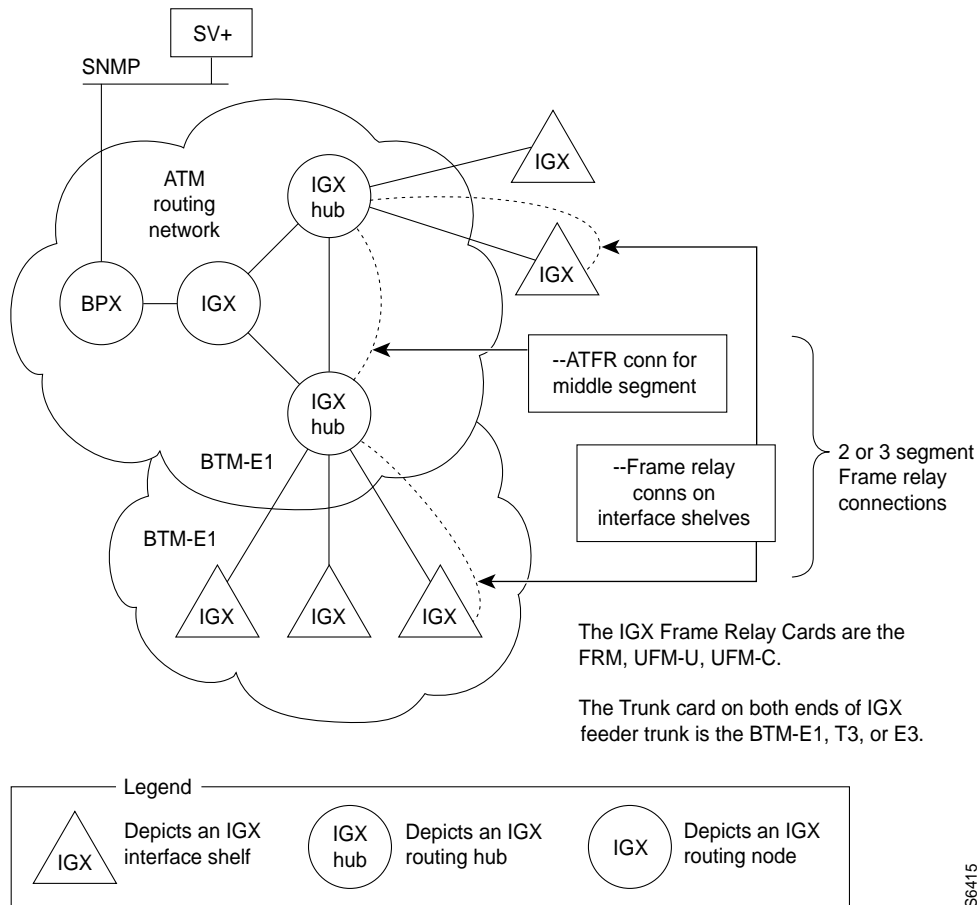
With tiered networks, Cisco IGX 8400 series multiband switches on the edge of the network are configured as interface shelves and are connected to Cisco IGX 8400 series multiband switches configured as router hubs. The interface shelves allow the network to support additional voice, data and frame relay connections without adding additional routing nodes. An example of 3-segment voice and data connections via a Cisco IGX 8400 series multiband switch configured as an interface shelf and Cisco IGX 8400 series multiband routing hubs is shown in Figure 6-2. An example of a frame relay connection via Cisco IGX 8400 series multiband switches configured as interface shelves and routing hubs is shown in Figure 6-3.

Figur e6-2 Cisco IGX 8400 Series Multiband Shelves and Routing Hubs, Voice and Data Connections



S6414

Figur e6-3 Cisco IGX 8400 Series Multiband Shelves and Routing Hubs, Frame Relay Connections



S6415

Tiered Network Implementation

The following applies to Cisco IGX 8400 series multiband routing hubs and interface shelves:

- A Cisco IGX 8400 series multiband switch configured as a routing hub supports up to 4 Cisco IGX 8400 series multiband interface shelves.
- A Cisco IGX 8400 series multiband interface shelf can have only one feeder trunk to the routing network.
- A Cisco IGX 8400 series multiband interface shelf is the only type of interface shelf that can connect to Cisco IGX 8400 series multiband routing hubs.
- No direct trunking between interface shelves is supported.
- No routing trunk is supported between the routing network and interface shelves.
- The feeder trunks between Cisco IGX 8400 series multiband hubs and Cisco IGX 8400 series multiband interface shelves are connected to a BTM-E1 backcard on each end of the trunk.
- Voice and data connection management to a Cisco IGX 8400 series multiband switch interface shelf is provided by Cisco WAN Manager.
- Telnet is supported to an interface shelf; the vt command is not.

- Remote printing by the interface shelf via a print command from the routing network is not supported.

The following applies to voice and data connections over Cisco IGX 8400 series multiband interface shelves:

- 3-segment connections are supported, that is: originating Cisco IGX 8400 series multiband interface shelf data or voice card to Cisco IGX 8400 series multiband routing hub, across Cisco IGX 8400 series multiband intermediate switches, as applicable, to a Cisco IGX 8400 series multiband routing hub, to a terminating Cisco IGX 8400 series multiband interface shelf data or voice card.
- 2-segment connections are not supported, (Cisco IGX 8400 series multiband interface shelf voice or data card to routing hub).
- Routing through the middle segment of the three segment connection is done via Cisco IGX 8400 series multiband routing nodes using CBR mode and simple gateway over the Cisco IGX 8400 series multiband trunks.
- Connection statistics are supported at user endpoints only.
- Adaptive voice is not supported.

The following applies to frame relay connections over Cisco IGX 8400 series multiband interface shelves via a Cisco IGX 8400 series multiband hub.

- 3-segment connections are supported, that is: originating Cisco IGX 8400 series multiband interface shelf frame relay card to Cisco IGX 8400 series multiband routing hub, across Cisco IGX 8400 series multiband intermediate switches, as applicable, to a Cisco IGX 8400 series multiband routing hub, to terminating Cisco IGX 8400 series multiband interface shelf data or voice card.
- 2-segment connections are supported, (Cisco IGX 8400 series multiband interface shelf frame relay card to routing hub).
- Routing through the middle segment of the three segment connection is done via Cisco IGX 8400 series multiband routing nodes using ATFR mode and simple gateway over the Cisco IGX 8400 series multiband trunks.
- Connection statistics are supported at user endpoints only.

General

Annex G, a bi-directional protocol, defined in Recommendation Q.2931, is used for monitoring the status of connections across a UNI interface. Tiered Networks use the Annex G protocol to pass connection status information between a hub node and attached shelf.

Definitions

Cisco IGX 8400 Series Multiband Routing Hub	A Cisco IGX 8400 series multiband switch in the routing network which has attached Cisco IGX 8400 series multiband switches configured as interface shelves. Also referred to as a hub node or IGX hub.
IGX Interface Shelf	A special configuration of a Cisco IGX 8400 series multiband switch that is connected as a shelf to an Cisco IGX 8400 series multiband switch configured as a routing hub. A Cisco IGX 8400 series multiband interface shelf is sometimes referred to as IGX A/F or feeder. The Cisco IGX 8400 series multiband interface shelf does not perform routing functions nor keep track of network topology.
Feeder Trunk	Refers to a trunk which interconnects a Cisco IGX 8400 series multiband interface shelf with the routing network via a Cisco IGX 8400 series multiband routing hub. A feeder trunk is sometimes referred to as an interface shelf trunk.
IGX/AF	Another name for the IGX interface shelf
Routing Network	The portion of the tiered network which performs automatic routing between connection endpoints.
VPI	Virtual Path Identifier
VCI	Virtual Connection Identifier

Upgrades

Converting an Cisco IGX 8400 series multiband switch to an interface shelf requires re-configuring connections on the node, as no upgrade path is provided in changing a routing node to an interface shelf.

Only Cisco IGX 8400 series multiband switches are able to act hub nodes for IGX interface shelves for voice and data transport over the IGX tiered network. A Cisco IGX 8400 series multiband switch, acting as a hub node, is not restricted from providing any other feature which is normally available on IGX nodes. An IGX hub supports up to 4 IGX interface shelves.

Connections within tiered networks consist of three distinct segments within each tier. A routing segment traverses the routing network, with an interface shelf segment at each end providing connectivity to the interface shelf end-point. Each of these segments are added, configured and deleted independently of the other segments. The CiscoWAN Manager Connection Manager provides management of these individual segments as a single end-to-end connection.

Interface shelves are attached to the routing network via an IGX node using a BTM E1 trunk. The connection segments within the routing network are terminated on IGX feeder trunks.

Co-locating Routing Hubs and Shelves

The feeder trunk between an interface shelf and the routing network is a single point of failure, therefore, the interface shelves should be co-located with their associated hub node. Card level redundancy is supported by the Y-Cable redundancy for the CVM, LDM, and HDM.

Network Management

Communication between CPE devices and the routing network is provided in accordance with Annex G of Recommendation Q.2931. This is a bidirectional protocol for monitoring the status of connections across a UNI interface.

Communication includes the real time notification of the addition or deletion of a connection segment and the ability to pass the availability (active state) or unavailability (inactive state) of the connections crossing this interface.

A proprietary extension to the Annex G protocol is implemented which supports the exchange of node information between an interface shelf and the routing network. This information is used to support the IP Relay feature and the Robust Update feature used by network management.

Network Management access to the interface shelves is through the IP Relay mechanism supported by the SNMP and TFTP or by direct attachment to the interface shelf. The IP Relay mechanism relays traffic from the routing network to the attached interface shelves. No IP Relay support is provided from the interface shelves into the routing network.

IGX routing hubs are the source of the network clock for its associated feeder nodes. Feeders synchronize their time and date to match their routing hub.

Robust Object and Alarm Updates are sent to a network manager which has subscribed to the Robust Updates feature. Object Updates are generated whenever an interface shelf is added or removed from the hub node and when the interface shelf name or IP Address is modified on the interface shelf. Alarm Updates are generated whenever the alarm state of the interface shelf changes between Unreachable, Major, Minor and OK alarm states.

An interface shelf is displayed as a unique icon in the SV+ Network Management topology displays. The colors of the icon and connecting trunks indicate the alarm state of each. Channel statistics are supported by CVM, HDM, and LDM endpoints. Trunk Statistics are supported for the feeder trunk and are identical to the existing IGX trunk statistics.

Preferred Routing

Preferred routing within the routing network can be used on all connections. Priority bumping is supported within the routing network, but not in the interface shelves. All other connection features such as conditioning, **rrtcon**, **upcon**, **dncon**, etc. are also supported.

Local and Remote Loopback

Connection local and remote loopbacks are managed at the user interface of the voice or data endpoint Routing Node or interface shelf. The existing IGX voice and data port loopback features are supported on the IGX interface shelf.

Testcon and Testdly

Tstcon is supported at the voice and data endpoints in a non-integrated fashion and is limited to a pass/fail loopback test. Fault isolation is not performed. Intermediate endpoints at the BTM cards do not support the tstcon feature. Tstdelay is also supported for the in a non-integrated fashion similar to that of the tstcon command.

IGX Interface Shelf Description

The IGX interface shelf supports the termination of voice and data connection segments to a BTM. The IGX interface shelf connects to the routing network via a BTM and associated BMT-E1 back card on both the interface shelf and the IGX routing hub.

IGX interface shelves support the following network management features:

- Interval Statistics enable/disable/collection
- IP Relay
- Robust Object Updates
- Robust Alarm Updates
- Real-time Counters
- Event Logging
- Software/Firmware Downloads
- Configuration Save/Restore
- SNMP

Configuration and Management

The interface shelves attached to each hub must have unique names. Each interface shelf must also be assigned a unique IP address.

An interface shelf communicates with a routing hub over a new type of NNI. It is similar to the existing Frame Relay NNI in purpose and function, and is based on the ATM LMI message set described in Recommendation 2931, Annex G. A routing hub and interface shelf use this NNI to maintain a control session with each other. Any change to the status of the feeder trunk affects this control session.

Feeder trunks are the communication path between the routing hub and the Feeder. These feeder trunks are supported by the BTM trunk card on both the IGX interface shelf and the IGX routing hub. Feeder trunks are upped using the “**uptrk**” command. Feeder trunks must be upped on both the routing hub and the interface shelf before it can be joined to the routing network.

An IGX node must be converted to an interface shelf by entering the appropriate command at the node. Once an IGX has been converted to an interface shelf, it can be joined to the IGX routing hub, by executing the **addshelf** command at the IGX routing hub. The **addshelf** command has the following syntax:

Shelf Management

```
addshelf <trunk> <shelf_type>
trunkslot.port
```

shelf_typeI (IGX/AF)

```
delshelf <trunk> | <shelf_name> deletes interface shelf
```

dsptime: Displays feeder trunk status. IGX routing hubs display the status of all attached IGX interface shelves. IGX interface shelves display a single status item, that of the attached IGX hub node.

Alarm Management of Interface Shelf on the IGX Hub Node

dsptime: The field, interface shelf alarms, shows a count of the number of interface shelves which are Unreachable, in Minor Alarm, or in Major Alarm. The nnn-A bit status failures for interface shelf connections are also shown.

Alarm Management on the IGX Interface Shelf

dsptime: The field, routing network Alarms, shows a count of major and minor alarms in the routing network. Feeder A-bit connection status reported by feeder NNI is shown in the “Connection A-Bit Alarms” field.

dsptime: Shows if the routing network is reachable and the attached IGX hub node.

Port Management

Uses existing commands.

Connection Management

Parameters entered at the Cisco WAN Manager workstation when adding connections.

Bandwidth Management

Parameters entered at the Cisco WAN Manager workstation when adding connections. Bandwidth performance is monitored by viewing selected statistics at the Cisco WAN Manager NMS.

Bandwidth Efficiency

Since voice traffic is time sensitive, and low-speed voice connections can result in SGW cells being sent with only a single packet placed in the cell in order to avoid excessive delay between cells. It may be necessary to use the **cnfmb** command on the interface shelves in order to configure the packet combining timeout rate for a particular application.

Statistics

Enabled and monitored via CiscoWAN Manager.

User Interface Commands

Refer to the Command Reference manual for additional information on commands associated with tiered networks. The following is a list of most often used commands with IGX routing hubs and IGX interface shelves supporting voice and data connections.

Shelf

addshelf
delshelf
dspnode
dspalms
dsptrks

Data Connection Commands

addcon
dspcon
dspcons

Data Channel Command

cnfchdfm
cnfcheia
cnfkdir
cnfdchtp
cnfdclk
cnfict

Voice Connection Commands

addcon

Voice Channel Commands

cnfchadv

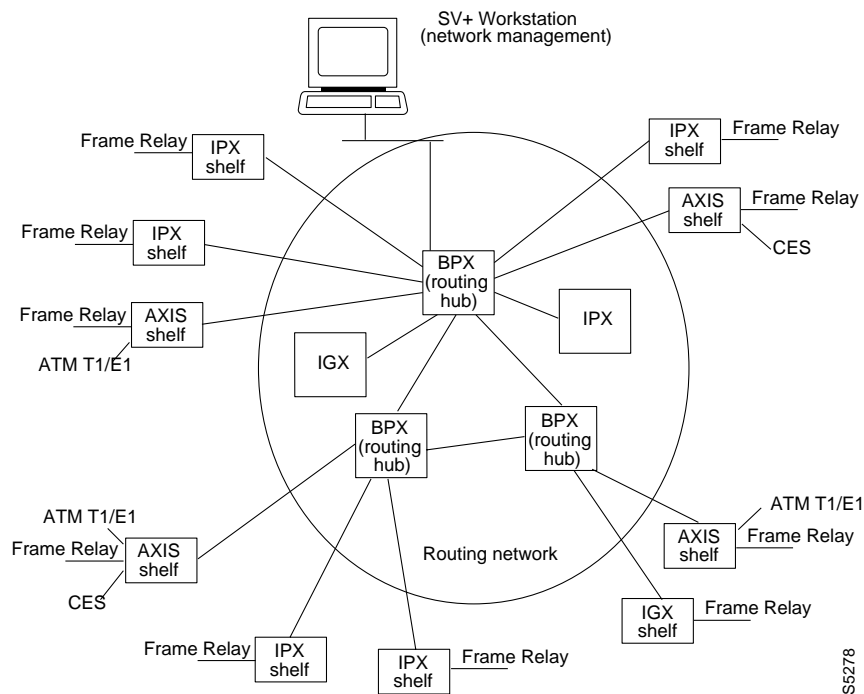
cnfchutil
cnfchkdl
cnfcos
cnfchech
cnfchgn
cnfcond
cnfrevsig
cnfvchtp
cnfxmtisig
cnfcmb

Cisco BPX 8600 Series Broadband Routing Hubs in a Tiered Network

Tiered networks with Cisco BPX 8600 series broadband routing hubs have the capability of adding interface shelves/feeders (non-routing nodes) to an IPX/IGX/Cisco BPX 8600 series broadband routing network (Figure 6-4). The Cisco MGX 8220 edge concentrator interface shelf, and IPX or IGX nodes configured as interface shelves are connected to Cisco BPX 8600 series broadband routing hubs. Interface shelves allow the network to support additional connections without adding additional routing nodes.

The Cisco MGX 8220 edge concentrator supports frame T1/E1, X.21 and HSSI frame relay, ATM T1/E1, and CES, and is designed to support additional interfaces in the future. The IPX interface shelf supports frame relay ports, as does the IGX (option is available to configure as an interface shelf).

Figur e6-4 Tiered Network with Cisco BPX 8600 Series Broadband Routing Hubs



Tiered Network Implementation

The following requirements apply to Cisco BPX 8600 series broadband routing hubs and their associated interface shelves:

- Cisco MGX 8220 edge concentrator Release 4 level is required on all Cisco MGX 8220 edge concentrator interface shelves.
- Only one feeder trunk is supported between a routing hub and each interface shelf.
- No direct trunking between interface shelves is supported.
- No routing trunk is supported between the routing network and interface shelves.
- The feeder trunks between Cisco BPX 8600 series broadband hubs and IPX or IGX interface shelves are either T3 or E3.
- The feeder trunks between Cisco BPX 8600 series broadband hubs and Cisco MGX 8220 edge concentrator interface shelves are T3, E3, or OC3-c/STM-1.
- Frame Relay Connection management to an IPX interface shelf is provided by Cisco WAN Manager.
- Frame Relay and ATM connection management to an Cisco MGX 8220 edge concentrator interface shelf is provide by Cisco WAN Manager.
- Telnet is supported to an interface shelf; the vt command is not.
- Remote printing by the interface shelf via a print command from the routing network is not supported.

General

Annex G, a bi-directional protocol, defined in Recommendation Q.2931, is used for monitoring the status of connections across a UNI interface. Tiered Networks use the Annex G protocol to pass connection status information between a hub node and attached interface shelf.

Definitions

Cisco BPX 8600 series broadband Routing Hub	A Cisco BPX 8600 series broadband switch in the routing network which has attached interface shelves. Also referred to as a hub node or Cisco BPX 8600 series broadband hub.
Cisco MGX 8220 edge concentrator Interface Shelf	A standards based service interface shelf that connects to a Cisco BPX 8600 series broadband routing hub, aggregates and concentrates traffic, and performs ATM adaptation for transport over broadband ATM networks
IPX Interface Shelf	A special configuration of the Cisco IPX narrow band node designated as a interface shelf that supports frame relay connections.
IGX Interface Shelf	A special configuration of the IGX multiband node designated as a interface shelf that supports frame relay connections.
Feeder Trunk	Refers to a trunk which interconnects an interface shelf with the routing network via a Cisco BPX 8600 series broadband switch configured as a routing hub. A feeder trunk is sometimes referred to as an interface shelf trunk.
IPX/AF	Another name for the IPX interface shelf
IGX/AF	Another name for the IGX interface shelf
Routing Network	The portion of the tiered network which performs automatic routing between connection endpoints.
VPI	Virtual Path Identifier
VCI	Virtual Connection Identifier

Upgrades

Converting an IPX or IGX node to an interface shelf requires re-configuring connections on the node, as no upgrade path is provided in changing a routing node to an interface shelf.

A Cisco BPX 8600 series broadband switch, acting as a hub node, is not restricted from providing any other feature which is normally available on Cisco BPX 8600 series broadband switches. A Cisco BPX 8600 series broadband switch configured as a routing hub supports up to 16 interface shelves.

Connections within tiered networks consist of distinct segments within each tier. A routing segment traverses the routing network, and an interface shelf segment provides connectivity to the interface shelf end-point. Each of these segments are added, configured and deleted independently of the other segments. The SV+ Connection manager provides management of these individual segments as a single end-to-end connection.

Interface shelves are attached to the routing network via a Cisco BPX 8600 series broadband switch configured as a routing hub using a BXM trunk (T3/E3 or OC3) or BNI trunk (T3/E3). The connection segments within the routing network are terminated on the BNI feeder trunks.

All frame relay connection types which can terminate on the Cisco BPX 8600 series broadband ASI card are supported on the BNI feeder trunk (currently VBR, CBR, ABR, and ATF types). No check is made by the routing network to validate whether the connection segment type being added to a BNI feeder trunk is actually supported by the attached interface shelf.

Co-locating Routing Hubs and Interface Shelves

The trunk between an interface shelf and the routing network is a single point of failure, therefore, the interface shelves should be co-located with their associated hub node. Card level redundancy is supported by the Y-Cable redundancy for the BXM, BNI, AIT, and BTM.

Network Management

Communication between CPE devices and the routing network is provided in accordance with Annex G of Recommendation Q.2931. This is a bidirectional protocol for monitoring the status of connections across a UNI interface. (Note: the feeder trunk uses the STI cell format to provide the ForeSight rate controlled congestion management feature.)

Communication includes the real time notification of the addition or deletion of a connection segment and the ability to pass the availability (active state) or unavailability (inactive state) of the connections crossing this interface.

A proprietary extension to the Annex G protocol is implemented which supports the exchange of node information between an interface shelf and the routing network. This information is used to support the IP Relay feature and the Robust Update feature used by network management.

Network Management access to the interface shelves is through the IP Relay mechanism supported by the SNMP and TFTP projects or by direct attachment to the interface shelf. The IP Relay mechanism relays traffic from the routing network to the attached interface shelves. No IP Relay support is provided from the interface shelves into the routing network.

The Cisco BPX 8600 series broadband switch configured as a routing hub is the source of the network clock for its associated feeder nodes. Feeders synchronize their time and date to match their routing hub.

Robust Object and Alarm Updates are sent to a network manager which has subscribed to the Robust Updates feature. Object Updates are generated whenever an interface shelf is added or removed from the hub node and when the interface shelf name or IP Address is modified on the interface shelf. Alarm Updates are generated whenever the alarm state of the interface shelf changes between Unreachable, Major, Minor and OK alarm states.

An interface shelf is displayed as a unique icon in the SV+ Network Management topology displays. The colors of the icon and connecting trunks indicate the alarm state of each. Channel statistics are supported by FRP, FRM, ASI, and Cisco MGX 8220 edge concentrator endpoints. BNIs, AITs, and BTMs do not support channel statistics. Trunk Statistics are supported for the feeder trunk and are identical to the existing BNI trunk statistics.

ForeSight

Foresight for an IPX interface shelf terminated Frame Relay connections is provided end-to-end between Frame Relay ports, regardless as to whether these ports reside on an IPX interface shelf or within the routing network.

Preferred Routin

Preferred routing within the routing network can be used on all connections. Priority bumping is supported within the routing network, but not in the interface shelves. All other connection features such as conditioning, **rttcon**, **upcon**, **dncon**, etc. are also supported.

Local and Remote Loopback

Connection local and remote loopbacks are managed at the user interface of the FRP endpoint routing node or interface shelf. The existing IPX Frame Relay port loopback feature is supported on the IPX interface shelf. Remote loopbacks are not supported for DAX connections. A new command **addlocrmtlp** is added to support remote loopbacks at FRP DAX endpoints.

Tstcon and Tstdela

The **tstcon** command is supported at the FRP endpoints in a non-integrated fashion and is limited to a pass/fail loopback test. Fault isolation is not performed. This is the same limitation currently imposed on inter-domain connections. Intermediate endpoints at the AIT and BNI cards do not support the **tstcon** command. The **tstdelay** command is also supported for the FRP and ASI in a non-integrated fashion similar to that of the **tstcon** command.

IPX Interface Shelf Description

The IPX interface shelf supports the termination of Frame Relay connection segments to an AIT. DAX voice and low speed data connections are also supported, but they can't terminate on an AIT. The IPX interface shelf connects to the routing network via an AIT card on the IPX and a BNI card on the Cisco BPX 8600 series broadband switch configured as a routing hub.

Admission control and ForeSight rate control for IPX interface shelf terminated Frame Relay connections is performed at the FRP port on the IPX interface shelf. Only a single trunk line is supported between the IPX interface shelf and the routing network. Trunks on the IPX interface shelf linking other nodes are not supported.

Frame Relay type connections, remotely or locally terminated are supported on IPX interface shelves. Interface shelf connections for which both endpoints reside on the same interface shelf are not known to the routing network and will not route through the routing network.

IPX interface shelves support the following network management features:

- Interval Statistics enable/disable/collection
- IP Relay
- Robust Object Updates
- Robust Alarm Updates
- Real-time Counters

- Event Logging
- Software/Firmware Downloads
- Configuration Save/Restore
- SNMP

Configuration and Management

The interface shelves attached to each hub must have unique names. Each interface shelf must also be assigned a unique IP address.

An interface shelf communicates with a routing hub over a new type of NNI. It is similar to the existing Frame Relay NNI in purpose and function, and is based on the ATM LMI message set described in Recommendation 2931, Annex G. A routing hub and interface shelf use this NNI to maintain a control session with each other. Any change to the status of the feeder trunk affects this control session.

Feeder trunks are the communication path between the routing hub and the Feeder. These feeder trunks are supported by the AIT trunk card on the IPX interface shelf and the BNI trunk card on the Cisco BPX 8600 series broadband switch configured as a routing hub. Feeder trunks are upped using the **uptrk** command. Feeder trunks must be upped on both the routing hub and the interface shelf before it can be joined to the routing network.

Once an IPX has been converted to an interface shelf, it can be joined to the Cisco BPX 8600 series broadband routing hub, by executing the **addshelf** command at the Cisco BPX 8600 series broadband routing hub. The **addshelf** command has the following syntax:

Interface Shelf Management

```
addshelf <trunk> <shelf_type>
trunkslot.port
shelf_typeI (IPX/AF) or A (Cisco MGX 8220 edge concentrator)
```

```
delshelf <trunk> | <shelf_name> deletes interface shelf
```

dsnode: Displays feeder trunk status. Cisco BPX 8600 series broadband switches configured as hub nodes display the status of all attached interface shelves. IPX interface shelves display a single status item, that of the attached Cisco BPX 8600 series broadband switch.

Alarm Management of Interface Shelf on the Cisco BPX 8600 Series Broadband Switch Configured as a Hub Node

dspalms A new field, interface shelf alarms, shows a count of the number of interface shelves which are Unreachable, in Minor Alarm, or in Major Alarm. The nnn-A bit status failures for shelf connections are also shown.

Alarm Management on the IPX Interface Shelf

dspalms	A new field, routing network Alarms, shows a count of major and minor alarms in the routing network. Feeder A-bit connection status reported by Feeder NNI is shown in the “Connection A-Bit Alarms” field.
dspnode:	Shows if the routing network is reachable and the attached Cisco BPX 8600 series broadband hub node.

Port Management

Uses existing commands

Connection Management

Parameters entered at the Cisco WAN Manager workstation when adding connections.

Bandwidth Management

Parameters entered at the Cisco WAN Manager workstation when adding connection. Bandwidth performance monitored by viewing selected statistics at the CiscoWAN Manager workstation.

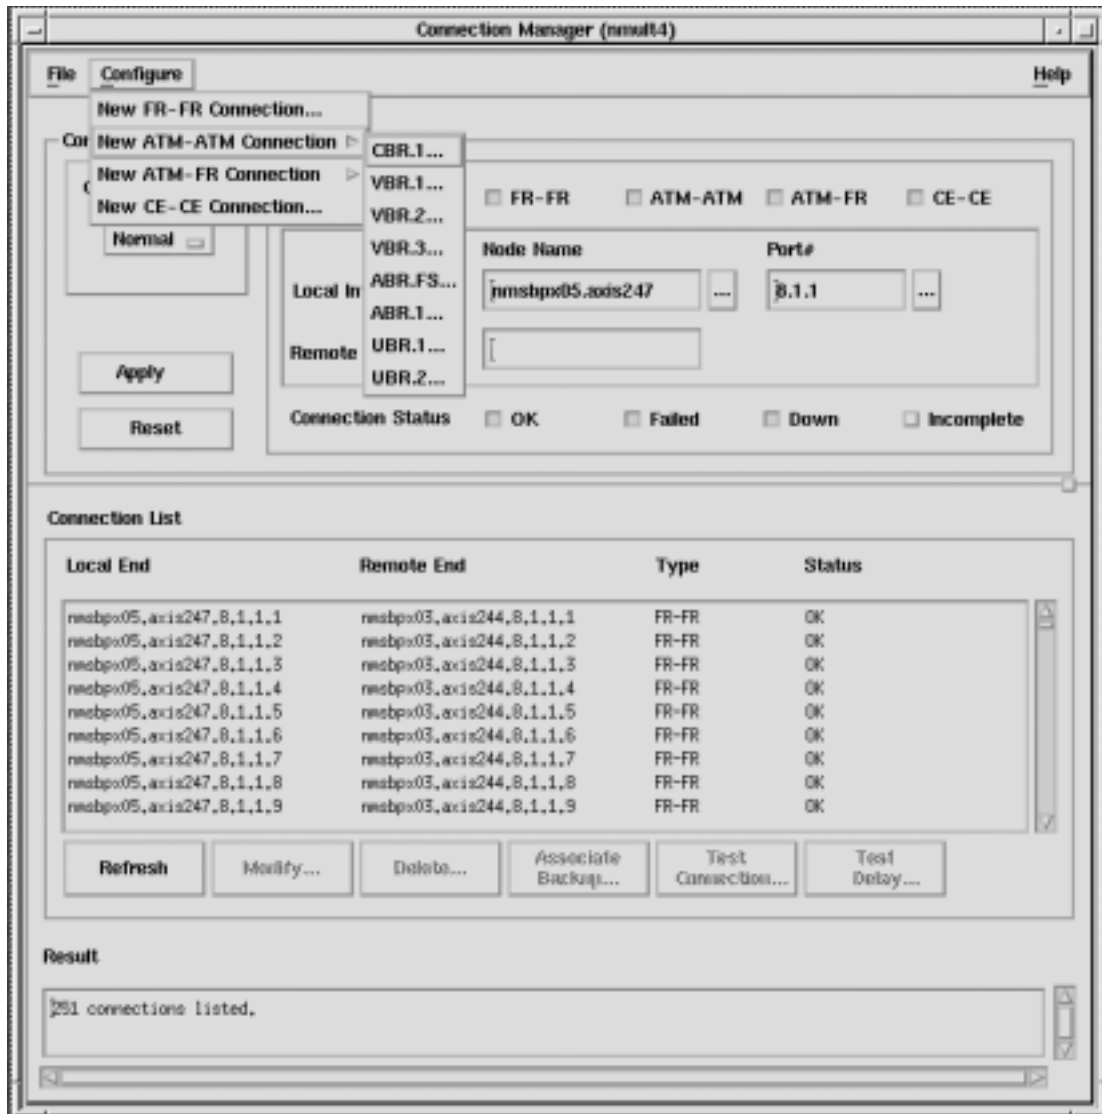
Statistics

Enabled and monitored via CiscoWAN Manager.

Cisco WAN Manager NMS

Interface shelf and feeder trunk information is reported to Cisco WAN Manager by the routing hub and interface shelf. Cisco WAN Manager can virtually connect to any node in the network via a TCP/IP connection. The Cisco WAN Manager Connection Manager is used to add, delete, and monitor voice and data connections for tiered networks with IGX 8400 series multiband switches configured as hubs. It is also used to add, delete and monitor frame relay connections for tiered networks with Cisco BPX 8600 series broadband switches. A sample of the Connection Manager GUI is shown in Figur e6-5.

Figur e6-5 Cisco WAN Manager Connection Manager



Network Maintenance

This chapter describes some of the tools provided for detecting and identifying network and/or equipment problems that are available to the network operator.

There are considerably more advanced tools built into the system software that are available only to Cisco Customer Service personnel. These advanced tools require in-depth knowledge of the hardware and software and are used generally to locate the less common types of system problems.

This chapter contains the following:

- Automatic Alarm Reporting
- Network Troubleshooting
- System Troubleshooting Tools
- Network Statistics

Automatic Alarm Reporting



Caution

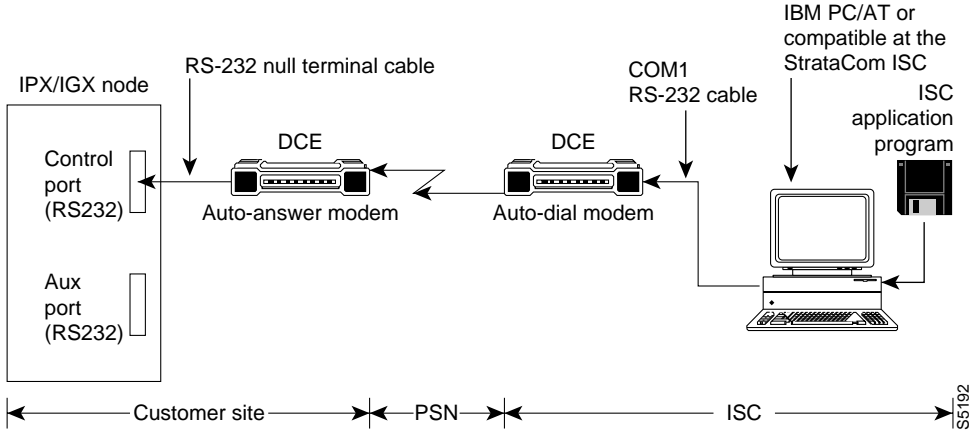
Do not perform any disruptive tests or repairs to the network on your own. Before commencing with troubleshooting, call Cisco Customer Service so that they can provide you with assistance in locating a fault.

In a network with Cisco BPX 8600 series broadband switches, Cisco IGX 8400 series multiband switches, and Cisco IPX narrowband switches, it is recommended that at least one node be configured to transmit alarms automatically to Cisco Customer Service. [Figure 7-1](#) illustrates the hardware configuration required for implementation. This can be a Cisco IGX 8400 series multiband switch, or a Cisco IPX narrowband switch.

When an alarm occurs on the network, the autodial modem automatically dials the specified telephone number. An auto-answer modem at Cisco Customer Service answers the call and directs it to a dedicated personal computer. The alarm is logged under the network ID (an ASCII character string) specified by the network administrator and approved by Cisco Customer Service personnel.

If the auto-answer modem at Customer Service is busy when an alarm arrives, then the autodial modem will keep dialing until the call is completed. A suggested modem is the Codex V.34 RSA 28.8 Kbaud modem. Connections to the node are detailed in the appropriate *Cisco BPX 8600 Series Reference Manual* or *Cisco IGX 8400 Series Reference Manual*.

Figur e7-1 Automatic Alarm Reporting



Network Troubleshooting

The Cisco BPX 8600 series broadband, Cisco IGX 8400 series multiband, and Cisco IPX narrowband system software provides a variety of tools for monitoring the network and detecting and locating network and system problems. Almost all network troubleshooting can be performed from the console of the Cisco WAN Manager Network Management Station (NMS). These tools are described in general in the following paragraphs.

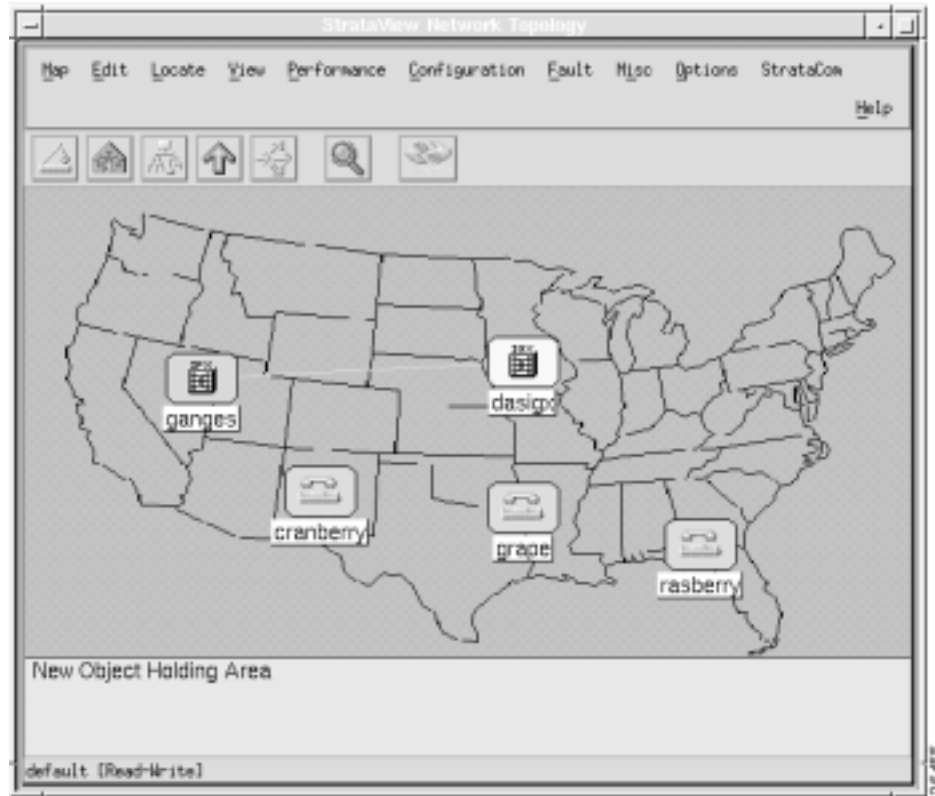


Note For more specific information on using the various commands referenced in this section, refer to the *Cisco WAN Switching Command Reference* publication. For details on the Cisco WAN Manager NMS, refer to the *Cisco WAN Manager Operations* publication.

The Cisco WAN Manager NMS displays a map of the network topology indicating where the nodes are located and the packet trunks that provide the connections between the nodes (refer to). This topolog map displays, in color, any network faults detected by the system.

All nodes in the network are accessible from the Cisco WAN Manager NMS and any network or system fault detected will be reported to the NMS. In addition the network operator can access any routing node in the network by using a Virtual Terminal (vt) command to create a virtual connection from the NMS console directly to the desired node. Non-routing nodes are accessed via “telnet” command on a LAN interface, and are managed via in-band SNMP. This allows the operator to enter commands and vie system status as if he were locally connected to the node.

Figur e7-2 Example of a Cisco WAN Manager Topology Map



Displaying Network Alarms

Whenever an alarm condition is detected in the network, the NMS console will sound an alarm tone, a Major Alarm or Minor Alarm notice will flash, and one or more of the icons on the topology map will change color to alert the network operator of the fault. Additionally, each node can be equipped to output a set of relay contact closures to operate customer alarm systems to notify network operators of an alarm condition in the network. See *the Cisco IPX Reference*, *the Cisco IGX 8400 Series Reference*, or *the Cisco BPX 8600 Series Reference publications* for further information.

A Major or Minor alarm indication after one or more of the node names will be displayed if there is an alarm anywhere in the network. This screen will identify the node(s) that detected the alarm condition. Another method of identifying the location of a network alarm is to issue a Display Networks (**dspnw**) command as shown in the following example:

Example: Display Network Screen

```

bpxb          VT    YourID:1          BPX 15    9.1    Mar. 9 1999  14:34 PST

NodeName      Alarm  Trunk          Trunk          Trunk
bpxa          Alarm  1.1-1.1/bpxe   6.2-3.2/bpxc   5.1-6.1/bpzb
              6.1-11.2/bpzd   1.2-5.1/bpzb   1.3-1.3/bpzc
              5.2-12.2/bpxg
bpxb          Alarm  6.1-5.1/bpza   5.1-1.2/bpza   6.2-3.1/bpzc
              4.1-4.1/bpzc   4.2-4.2/bpzc   4.3-4.3/bpzc
              5.3-1.2/bpxe
bpxc          Alarm  3.2-6.2/bpza   3.1-6.2/bpzb   2.1-14.1/bpzd
              2.2-16/ipxb    2.3-14.3/bpzd   1.3-1.3/bpza
              4.1-4.1/bpzb   4.2-4.2/bpzb   4.3-4.3/bpzb
bpxd          Alarm  14.1-2.1/bpzc  14.2-3/ipxa    14.3-2.3/bpzc
              11.2-6.1/bpza
bpxe          Alarm  1.1-1.1/bpza   1.2-5.3/bpzb   1.3-2.1/bpzf
bpxf          Minor  2.1-1.3/bpxe

```

This Command: dspnw

Continue?

Next Command:

Once the location (node) where the alarm has been detected is identified, the network operator can use the Virtual Terminal feature to observe the status at that location. The Display Alarms (**dspalms**) command and associated screen will be used here. The following example illustrates a Minor Alarm at a node named D1.beta:

Example: Display Alarms Screen

```

bpxf          VT    YourID:1          BPX 15    9.1    Mar. 9 1997  14:36 PST

Alarm summary  (Configured alarm slots: None)
Connections Failed:      None
Groups Failed:           None
TRK Alarms:              None
Line Alarms:             None
Cards Failed:            None
Slots Alarmed:           None
Missing Cards:           None
Remote Node Alarms:     None
Remote Domain Alarms:   None

Interface Shelf Alarms:  1 Minor
ASM Alarms:              None

```

Last Command: dspalms

Displaying Trunk and Line Status

If the Display Nodes (**dspnds**) or Display Network (**dspnw**) command screen indicates a service-affecting problem with either a network trunk or CPE service access line, there are two similar screens that will display their status. The following example illustrates the Display Trunk (**dsptrks**) command screen with a Major Alarm on trunk 5.2. It also indicates whether the alarm is a Red (Local) or Yellow (Remote) alarm.

Example: Display Trunk Status Screen

```
bpxb          VT    YourID:1          BPX 15    9.1    Mar. 9 1999  14:44 PST

TRK  Type      Current Line Alarm Status          Other End
4.1  E3        Clear - OK                          bpxc/4.1
4.2  E3        Clear - OK                          bpxc/4.2
4.3  E3        Clear - OK                          bpxc/4.3
5.1  T3        Clear - OK                          bpxa/1.2
5.2  T3        Major - Loss of Sig (RED)           axis31(Axis)
5.3  T3        Clear - OK                          bpxe/1.2
6.1  OC3       Clear - OK                          bpxa/5.1
6.2  OC3       Clear - OK                          bpxc/3.1
```

Last Command: dsptrks

The following example illustrates the Display Lines (**dsplns**) command screen with a Major Alarm on line 12.2. It also indicates whether the line alarm is a Red (Local) or Yellow (Remote) alarm.

Example: Display Line Status Screen

```
bpxb          VT    YourID:1          BPX 15    9.1    Mar. 9 1999  14:44 PST

Line Type      Current Line Alarm Status
10.1 E3        Clear - OK
10.2 E3        Clear - OK
12.1 T3        Clear - OK
12.2 T3        Major -
13.1 T3        Clear - OK
13.2 T3        Clear - OK
14.1 OC3       Clear - OK
14.2 OC3       Clear - OK
```

Last Command: dsplns

Network trunk and line inputs to each node are constantly monitored for error conditions such as loss of signal, bipolar errors, frame errors, etc. These trunk and line errors may not be severe enough to cause a Red or Yellow Alarm but can be observed using the Display Trunk Errors (**dsptrkerrs**) and Display

Line Errors (**dsplnerrs**) command screens. The following example illustrates a typical summary trunk errors screen. Errored tens of seconds statistics are also kept and are displayed by a second trunk or circuit line screen.

Example: Display Trunk Errors Screen

bpxb VT YourID:1 BPX 15 9.1 Mar. 9 1999 14:44 PST

Total Errors

TRK	Code Errors	Rx Cell Dropped	Out of Frames	Loss of Signal	Frame BitErrs	HCS Errors	Tx Cell Dropped	Cell Errors	Cell Oofs
1.1	0	0	0	0	0	-	0	0	-
1.3	0	0	0	0	0	-	0	0	-
2.1	0	0	0	0	0	-	0	0	-
2.2	0	0	0	0	0	-	0	0	-
2.3	0	0	0	0	0	-	0	0	-
3.1	-	0	0	0	0	-	0	0	-
3.2	-	0	0	0	0	-	7	0	-
4.1	0	0	0	0	0	-	0	0	-
4.2	0	0	0	0	0	-	0	0	-
4.3	0	0	0	0	0	-	0	0	-

Last Command: dsptrkerrs

Next Command:

Displaying Node Status

The system software automatically determines the card configuration of the node and displays the card type in each shelf slot. Each card is shipped with a model and revision number burned into Read-Only Memory to assist in tracking each individual card for inventory purposes.

Each card in a node is periodically tested in a background mode so as not to disrupt normal traffic. A short, automatic self-test is performed to identify card failures and, if the node is equipped with a redundant card set, transfer operation to the alternate card before any downtime is detected. The Display Cards (**dspcds**) screen, in the following example, shows what cards are equipped and their status: active, failed, or standby (for redundant card sets).

Example: Display Cards Screen

```
bpxb          VT   YourID:1          BPX 15    9.1   Mar. 9 1999  14:44 PST
```

FrontCard						BackCard					
Type	Rev	Type	Rev	Status		Type	Rev	Type	Rev	Status	
1	Empty					9	Empty				
2	Empty					10	ASI-E3	JA11	E3-2	BE	Active
3	Empty					11	Empty				
4	BNI-E3	CD03	E3-3	BE	Active	12	ASI-T3	JA11	T3-2	BE	Active
5	BNI-T3	CB04	T3-3	AA	Active	13	ASI-T3	JX11	T3-2	BE	Active
6	BNI-OC3BHA		MMF-2	AH	Active	14	ASI-OC3BWA		MMF-2	DJ	Active
7	BCC	AVC	LMBCC	P01	Standby	15	ASM	ACA	LMASM	P01	Active
8	BCC	A0405	LMBCC	AC	Active						

Last Command: dspcds

Next Command:

If any card is removed for any reason, including replacement, the node configuration database will remember what type of card was initially installed in each slot to minimize the possibility of installing the wrong card type in a slot.

Additionally, all node power supply voltage outputs are monitored and an alarm is generated if any output fails or if any voltage falls out of tolerance. The status of each power supply in the node is displayed on the Display Power Supply Status (**dsppwr**) command screen.

Each node is cooled by various fan assemblies. In the Cisco IGX 8400 series multiband switch and in the Cisco IPX narrowband switch cabinets and in the *Cisco BPX 8600 Series wideband switch* enclosure, there is a temperature sensor that monitors the ambient temperature inside the enclosure. If there should be a failure of one or more of the fans, causing the temperature to rise, an alarm will be generated so the fan or assembly can be replaced before the temperature exceeds the maximum allowed. The cabinet temperature can be observed in the Cisco IPX narrowband switch, the Cisco IGX 8400

series multiband switch, and the Cisco BPX 8600 series broadband switch using the **dsppwr** command screen as shown in the following example. In addition, the Cisco BPX 8600 series broadband switch provides a Display Assembly Modules (**dspsasm**) screen.

Example: Display Power Supply Status Screen

```

bpxb          VT    YourID:1          BPX 15    9.1    Mar. 9 1999  14:44 PST

                Power Status                                Cabinet Temperature

ASM Status: Active                                31            87
Power voltage A/B:          0 / 48 V              C 60 | | 140 F
                                     e 50 | | 122 h
PSU  Ins Type Rev SerNum Failure                t 40 | | 104 e
A    N  N/A N/A  N/A   N/A                       g 30 | | 86  h
B    Y  240V 0C 26229 None                          a 20 | | 68  e
                                     e 20 | | 68  i
                                     e 20 | | 68  t

                Fan Status

FAN   1    2    3
      3300 3240 3240 RPM

Last Command: dsppwr
Next Command:
    
```

Displaying Connection Status

The endpoints and status of each connection that terminates on a node can be displayed using the Display Connections (**dspscons**) command and associated screen. This lists the connections by number and remote node. It also displays the status, OK or failed, and the Class of Service assigned. If desired, a filter can be applied so that this screen displays only certain types of connections, for example (Options: -v,-d,-a,-atfr,-g,+d,-abit,-fabit,-fail, nodename, -down,start_channel). The following is an example of a typical Display Connections (**dspscons**) command screen.

Example: Display Connections Screen

```

bpxb          VT    YourID:1          BPX 15    9.1    Mar. 9 1999  14:44 PST

Local          Remote          Remote          Only          Route
Channel        NodeName      Channel        State atfr      Avoid COS O
5.2.6.100      bpxb         13.2.10.1000   Failed cbr
5.2.6.101      bpxb         13.2.10.1001   Failed abr
10.1.1.1       bpxa         13.1.1.1       Ok     abr           0 L
10.1.1.5       bpxb         10.1.1.5       Ok     vbr
10.1.1.7.*     bpxd         4.1.1.7.*     Ok     abr           0 L
10.1.9.*       bpxa         13.1.9.*       Ok     vbr           0 L
10.1.10.10     bpxa         13.1.10.10    Ok     cbr           0 L
10.1.11.*      bpxa         13.1.11.*     Ok     abr           0 L
10.1.17.200    bpxa         13.1.17.201   Ok     abr           0 L
10.1.23.102    bpxd         4.1.23.102    Ok     abr-Grp      0 R
10.1.23.103    bpxd         4.1.23.103    Ok     abr-Grp      0 R
10.1.23.104    bpxd         4.1.23.104    Ok     abr-Grp      0 R
10.1.23.105    bpxd         4.1.23.105    Ok     abr-Grp      0 R

This Command: dspscons -atfr

Continue?
    
```


Maintenance (Events) Log

A record of all events is kept in each node as well as in the Cisco WAN Manager NMS. This Event Log is useful in maintaining the network as it lists in chronological order, all network events, and displays the alarm/event category, major, minor, information, etc.

You can use the HP OpenView Event Browsers to display a network-wide list of events and alarms or filter the list as desired. Alternatively, the **dspllog** command can be used at any selected node to observe an event log for that particular node as indicated in the following example. The node's log database can generally hold events for approximately 30 days after which the most recent log entry replaces the oldest entry. An example of the **dspllog** command is shown in the following example:

Example: Event Log Screen for a Node

```
bpxb          VT    YourID:1          BPX 15    9.1    Mar. 9 1999  14:44 PST

Most recent log entries (most recent at top)
Class  Description                                     Date      Time
Info   User YourID logged in (Virtual Terminal)       12/09/95  15:33:29
Info   User richard logged out (Virtual Terminal)      12/09/95  15:05:55
Info   User richard logged in (Virtual Terminal)       12/09/95  15:05:32
Info   User rachel logged out (Virtual Terminal)       12/09/95  14:57:42
Info   User rachel logged in (Virtual Terminal)        12/09/95  14:57:17
Clear  Communication Break with bpxb Cleared           12/08/95  16:31:22
Minor  Communication Break with bpxb                  12/08/95  16:31:12
Info   User StrataCom charlie in (Local)               12/08/95  16:25:14
Info   User StrataCom nancy out (Local)                12/08/95  16:09:06
Info   User StrataCom evans out (Virtual Terminal)    12/08/95  15:46:15
Info   User StrataCom nancy in (Virtual Terminal)     12/08/95  15:46:04
Info   Clock switch to Line 5.1 of bpxb via TRK 5.1   12/08/95  15:34:59
Clear  TRK 5.1 OK                                     12/08/95  15:34:28
```

This Command: dspllog

Continue?

Display of Network Trunk Loading

Cisco WAN switching system software automatically assigns connections to the network trunks as they are added. If there is a choice between several trunks to the same destination, the software will attempt to balance the loading on the trunks. If an attempt to add a connection fails, it often is the result of inadequate bandwidth available for the connection.

Trunk loading on each trunk can be displayed by using the Display Trunk Loading (**dspload**) command, which displays the loading in both directions of transmission for each of the packet types. This screen is useful for determining the capacity of each trunk to accommodate additional frame relay traffic. The following example illustrates a typical trunk loading screen:

Example: Display Trunk Loading

```
bpxb          VT    YourID:1          BPX 15      9.1    Mar. 9 1999  14:44 PST
```

Trunk loads for node 'bpx1'

TRK	Units		Used		Available		Reserved	
	Xmt	Rcv	Xmt	Rcv	Xmt	Rcv	Xmt	Rcv
1.1	Cell	Cell	88320	88304	6688	6704	992	992
1.2	Cell	Cell	63808	63808	31200	31200	992	992
1.3	Cell	Cell	0	0	95008	95008	992	992
5.1	Cell	Cell	273008	273008	79200	79200	992	992
5.2	Cell	Cell	0	0	352208	352208	992	992
6.1	Cell	Cell	0	0	352208	352208	992	992
6.2	Cell	Cell	326592	326592	25616	25616	992	992

Last Command: dspload

Next Command:

System Troubleshooting Tools

There are a number of manually-initiated tests that can be performed from the Cisco WAN Manager NMS console to assist in system troubleshooting. These tests may be included in a job so they can be scheduled to run remotely at a specified time if desired.

User-initiated Tests

There are several user-initiated tests that can be used to diagnose system problems. These tests are self-contained in that they do not require the use of external test equipment. They also do not require the operator to place a loopback at the far end to test both directions of transmission. These tests are listed in Table 7-1.

There are also several display commands that can be used to obtain information that may be helpful in troubleshooting system problems. These are also listed in Table 7-1.

Table 7-1 System Troubleshooting Commands Available

Command	Descriptio
Test Connection (tstcon)—frame relay	Performs a bi-directional test of the specified frame relay connection or range of connections by inserting an IPX-generated test pattern and comparing the returned pattern with the pattern transmitted. A pass or fail indication appears next to the tested connection in the Display Connections screen.
Test Connection (tstcon)—data	Same as above except for synchronous data connections.
Test Connection (tstcon)—voice	Same as above except for voice connections.
Test Delay (tstdelay)—frame relay	Measures the round-trip delay over the selected frame relay connection.
Test Port (tstport)—frame relay	Tests the operation of the selected frame relay port on the node.
Test Port (tstport)- data	Same as above except for synchronous data ports
Display Connection States (dspconst)	Displays in real-time the status of all voice connections terminating at a specified node.
Display Breakout Box (dspbob)—frame relay	Displays in real-time the status of data and control leads on selected frame relay connection.
Display Breakout Box (dspbob)—data	Same as above for synchronous data connections
Display Breakout Box (dspbob)—trunk	Same as above for network subrate trunks
Display Buses (dspsbuses)	Displays the status of the IPX system buses.
Display Slot Errors (dspsloterrs)	Displays any data errors associated with the slots in a BPX node.
Display Slot Alarms (dspslotalms)	Displays any alarms associated with the slots in a BPX node.
Display Trunk Errors (dsptrkerrs)	Displays any data errors associated with the network trunks connected t a node.

Loopback Tests

There are also various loopback paths that can be set up to help diagnose transmission problems. They rely on using external test equipment to provide the source of a test signal. The available loopback commands are listed in Table 7-2. illustrates the various loopback paths using, in this example, a frame relay card set (FRP/FRI) in an IPX node.

A local loopback path (LL) is set up in the local node at the PAD card (FRP) associated with the port or connection to be tested. A test signal applied to the input passes through the associated Interface Card (FRI), is sent to the Frame Relay PAD card (FRP) over the system bus where it is looped back toward the input. This tests the cabling and the local node processing of the signal.

Table 7-2 System LoopbackTests

Command	Descriptio
Add Local Loopback (addloclp)—frame relay port	Adds a loopback path at the frame relay port from the transmit side back to the receive side at the local node.
Add Local Loopback (addloclp)—frame relay connection	Does the same as above only for an individual frame relay connection.

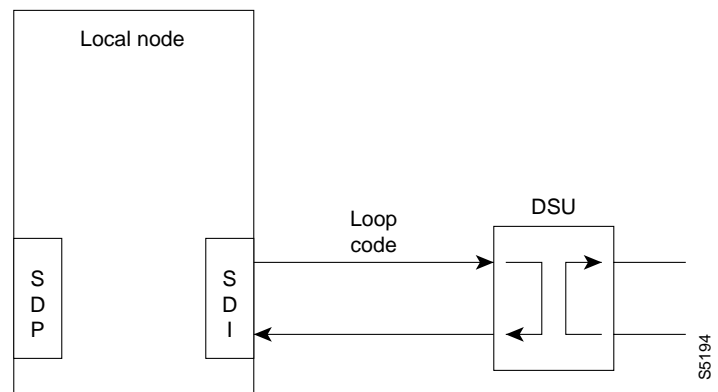
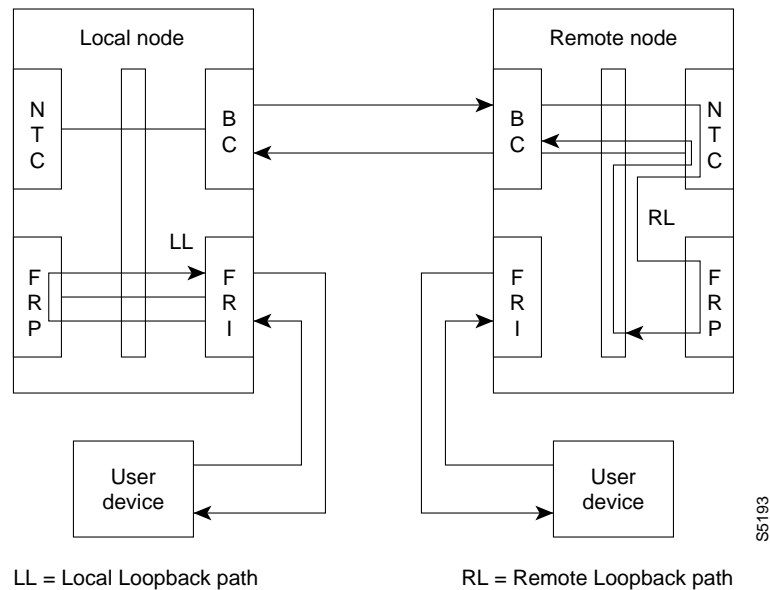
Table 7-2 System LoopbackTests (continued)

Command	Descriptio
Add Local Loopback (addloclp)—data	Adds a loopback path at the synchronous data port from the transmit side back to the receive side at the local node.
Add Local Loopback (addloclp)—voice	Adds a loopback path for an individual voice channel on a circuit line at the local node.
Add Remote Loopback (addrmtlp)—frame relay port	Adds a loopback path at the frame relay port from the transmit side back to the receive side at the remote node.
Add Remote Loopback (addrmtlp)—frame relay connection	Does the same as above only for an individual frame relay connection.
Add Remote Loopback (addrmtlp)—data	Adds a loopback path at the synchronous data port from the transmit side back to the receive side at the remote node.
Add Remote Loopback (addrmtlp)—voice	Adds a loopback path for an individual voice channel on a circuit line at the remote node.
Add External Loopback (addextlp)—data	Activates a near end or far end loopback on an external device, such as a DSU, connected to a synchronous data port.

A remote loopback path (RL) is set up in the remote node also at the PAD card (FRP). But, in this case, the signal travels over the network and through the remote node processing equipment but does not include the remote node Interface Card (FRI) or associated cabling. These components would be tested using another local loopback at the remote node.

The external loopback command finds limited use in data applications where an external data interface unit (DSU or CSU) is attached to the local node data interface card, illustrated by the SDI card in . The local node transmits the appropriate loopback codes out the circuit line towards the external device and then sets up the appropriate loopback path.

Figur e7-3 Network Loopback Paths



Connection Testin

System software includes a Test Connection (**tstcon**) command for testing network connections. This test is initiated by the network operator from the NMS console and can be performed at any time but it momentarily interrupts traffic on the connection during the test. Testing a connection should be performed only when an alarm has been reported from the connection or during off-hours.

Test Connection tests both directions of transmission from end-to-end and displays a pass or fail indication for each connection tested. The test may be specified for a single connection or for a group of connections. The operator may specify a single connection, all connections, all connections of a particular type (voice, data, or frame relay), or a starting and ending connection number.

In addition to testing the connection, the Test Connection routine will attempt to isolate and repair any failure it detects. The controller card at the node where the Test Connection (**tstcon**) command is issued instructs the service card to build packets containing special test frames. These packets are sent across the network to the terminating node, which depacketizes them, repacketizes the frame, and sends them back to the originating node where the returned frame is analyzed.

If the returned test pattern is incorrect, the system goes into an automatic fault isolation mode. Controllers in the various nodes along the connection route communicate with each other over an overhead message channel separate from the normal circuits.

The test pattern continues to be transmitted and analyzed at each node along the path as it is transmitted and as it is received until the failed network element is identified. Redundant cards may be switched into operation and routing tables in associated network trunk cards may be reprogrammed in an attempt to correct the problem. If all else fails, the suspected path and/or network component is then reported to the network manager (NMS).

External Device Window

External devices connected to network nodes, such as bridges, routers, or sub-rate multiplexers may be accessed through the NMS Window command. This feature provides a direct command line interface to external devices from the NMS console. Depending on the capability of the external device, it is often possible to report status and alarms and to control or configure the device through an RS232 port connection.

The following example illustrates a Window display of a router connected to the local node. In this example, the window is used to initiate a ping of the router connection.

```
Example: NMS Window to a Local Router
Protocol [ip]:

Target IP address: 192.9.202.1

Repeat count [5]:
Datagram size [100]:

Timeout in seconds [2]:

Extended commands [n]:

Type escape sequence to abort. ^^

Sending 5, 100-byte ICMP Echos to 192.9.202.1, timeout is 2 seconds:
. . . . .

Success rate is 100 percent
```

Network Statistics

Cisco WAN Manager collects network statistical data on the operation of the network and stores them in its database. They are available for display on the Cisco WAN Manager console in either tabular form or as bar charts. Statistics can be a useful source of information for troubleshooting problems that do not necessarily cause a major or minor alarm indication or for locating intermittent failures that may occur at random.

There are four classes of statistics:

- Trunk statistics.
- Line statistics.
- Connection statistics.
- Frame Relay port statistics.

Table 7-3 lists the statistics categories and the general nature of the statistics collected in each category. Note this is not a complete list of statistics but merely indicates some of the various conditions monitored. Refer to the *Cisco WAN Manager Operations document* for a complete listing.

Most statistics are collected on-demand and must be enabled by the system operator. The operator can set the collection interval, the sampling times, and the number of collection buckets to tailor the statistics for either long-term network performance evaluation or short term for network troubleshooting.

Table 7-3 Typical Statistics Collected

Statistics Category	Types of Statistic
Trunk statistics	Various trunk errors, bipolar violations, frame bit errors, loss of signal, etc.
	Packet errors and out of frame
	FastPackets and ATM cells of various types transmitted/dropped
	Transmitted ATM cell counts
	Received ATM cell counts
	Cells with CLP and EFCN set
	ATM header error counts
	DS3 PLCP error counts
	Bdata queue dropped cells.
Line statistics	Various circuit line errors, bipolar violations, frame bit errors, loss of signal, etc.
Connection statistics	Packets transmitted and received
	Transmitted and received data bytes
	Frame relay frames transmitted/discarded
	Frames transmitted with FECN or BECN or DE set
	Packets with CLP bit set dropped
	Seconds in service
Frame Relay Port	Frames transmitted and received
	Bytes transmitted and received
	Frames received with CRC or other errors
	Frames discarded at the connection ingress
	Frames discarded at the connection egress
	Frames discarded at the port egress

Table 7-3 Typical Statistics Collected (continued)

Statistics Category	Types of Statistic
	LMI messages sent or dropped for various errors
	DE frames dropped

PART 3

**NARROW AND BROADBAND
TRUNKS**

ATM and Broadband Trunks

This chapter is provided for users who wish to have an in-depth knowledge of the ATM and broadband trunks functions. It discusses ATM concepts and the various high-speed digital trunks that are used to carry ATM connections.

This chapter contains the following:

- Asynchronous Transfer Mode
- Broadband (ATM) Trunk Formats
- Virtual Trunks

Asynchronous Transfer Mode

Asynchronous Transfer Mode (ATM) uses a very flexible method of carrying broadband information between devices on a local or wide area network. It transmits this information using small, fixed length (53-byte) data packets, called cells, over high-speed digital transmission facilities. A key advantage to ATM is that it has been designed to carry voice, video, and data equally well on a single network.

ATM was developed to provide large amounts of bandwidth for network connections economically and on demand. When a user does not need access to a network connection, the bandwidth is available for use by another connection that does require it. ATM allows the bandwidth to be easily scaled to the requirements of the individual user.

ATM provides a distinct separation between the preparation of the customer data and the transportation of the data over the network. This allows a network operator to migrate to ATM using only the amount of bandwidth initially required. As bandwidth requirements increase, the user can scale up by using higher speed data links.

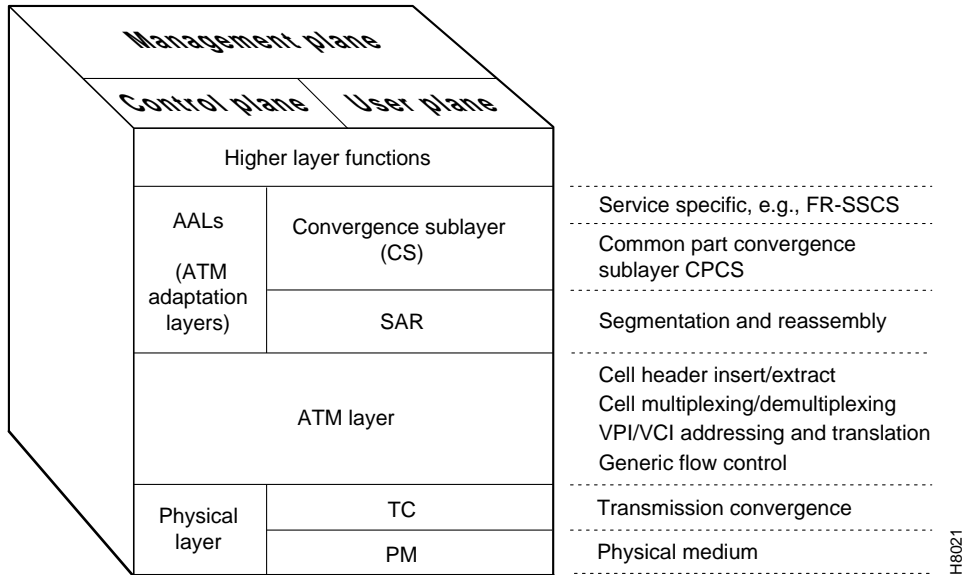
ATM is an outgrowth of Broadband Integrated Services Digital Network (B-ISDN), which in itself is an extension of ISDN that provides a definition for services and interfaces for public telecommunications networks.

ATM Mode

B-ISDN utilizes a layered architecture similar to the Open Systems Interconnection (OSI) 7-layer model. ATM redefines the lower three layers as indicated in Figure 8-1. These layers, the Physical Layer, ATM Layer, and ATM Adaptation Layer will be described in detail in the following paragraphs.

By bypassing the OSI network layer, ATM is able to process cells much more quickly and efficiently than current packet-based routing. The higher-order four layers are associated with specific user applications served by the ATM layers.

Figur e8-1 B-ISDN Model



Physical Layer

The Physical Layer defines the interface with the transmission media. It concerns itself with the physical interface, transmission rates, and how the ATM cells are converted to the line signal. Unlike many LAN technologies, such as Ethernet, which specify a certain transmission media, ATM cells can be carried on many different physical layers. The speed and bandwidth of the physical media will be the primary determining factor in selecting the transmission media used with ATM.

Initially, ATM will utilize existing physical transport media like the North American DS3 and CEPT E3 facilities and may be carried on T1 and E1 for users with low initial bandwidth requirements. For higher bandwidth requirements, the Synchronous Optical Network (SONET) provides a well defined and well accepted set of data rates (see Table 8-1). For example, the Cisco BPX 8600 series broadband switch provides 45 Mbps T3, 34 Mbps E3, and 155 Mbps interfaces but is designed to be easily expanded to include up to OC-12 port interfaces.

There are two sub-layers to the Physical Layer that separate the physical transmission medium and the extraction of data:

- Physical Medium Dependent (PMD) sub-layer.
- Transmission Convergence (TC) sublayer.

The Physical Medium Dependent concerns itself with the details specific to a particular physical layer, transmission rate, physical connector type, clock extraction, etc. For example, the SONET data rate utilized is part of the PMD. Transmission Convergence is involved with extracting the information content from the physical layer data transmission. This includes HEC generation and checking, extracting cells from the incoming bit stream, and processing of idle cells.


Table 8-1 SONET Data Rates

Data Rate	OC Level	SONET Designati	ITU-T Designati
52 Mbps	OC-1	STS-1	
155 Mbps	OC-3	STS-3	STM-1
466 Mbps	OC-9	STS-9	STM-3
622 Mbps	OC-12	STS-12	STM-4

ATM Layer

The ATM layer processes ATM cells. The format of the ATM cell consists of a 5-byte header and a 48-byte payload (Figure 8-2). The header contains the ATM cell address and other important information. The payload contains the user data being transported over the network. Cells are transmitted serially and propagate in strict numeric sequence throughout the network.

The payload length was chosen as a compromise between a long cell length, which is more efficient for transmitting long frames of data and a short cell length, which minimizes the end-to-end processing delay, and is good for voice, video, and delay sensitive data protocols. Although not specifically designed as such, the cell payload length conveniently accommodates two 24-byte IPX FastPackets.

Figur e8-2 ATM Cell Format

ATM Cell Headers

There are two basic header types defined by the standards committees, a UNI header and a NNI header; both are quite similar. Cisco has expanded on these header types to provide additional features above those proposed for basic ATM service. Usage of each of the various cell header types is described as follows:

- The UNI header (Figure 8-3) must be specified for each *User-to-Network Interface*. A UNI is an interface between a user device, such as an ATM router, and an ATM network.
- The NNI header (Figur e8-4) must be specified for each *Network-to-Network Interface*. This is used, for example, at the interface between a user's private ATM network and a service provider's public ATM network.
- The STI header (Figur e8-5) is an extension of these two header types and is a Cisco WAN switching interface. This header type is used between Cisco WAN switching nodes to provide advanced network features, including ForeSight, that improve performance and efficiency.

The most important fields in all three ATM cell header types are the **Virtual Path Identifier (VPI)** and a **Virtual Circuit Identifier (VCI)**. The VPI identifies the route (path) to be taken by the ATM cell while the VCI identifies the circuit or connection number on that path. The VPI and VCI are translated at each ATM switch, they are unique only for a given physical link.

A 4-bit **Generic Flow Control (GFC)** field in the UNI header is intended to be used for controlling user access and flow control. At present, it is not defined by the standards committees and is generally set to all zeros.

A 3-bit **Payload Type Indicator (PTI)** field indicates the type of data being carried in the payload. The first bit is a “0” if the payload contains user information and is a “1” if it carries connection management information. The second bit indicates if the cell experienced congestion over a path. If the payload is user information, the third bit indicates if the information is from Customer Premises Equipment.

In the STI header (Figure 8-5), these bits are used to indicate data from certain Cisco BPX 8600 series broadband queues corresponding to various classes of service e.g. OptiClass, the enhanced class of service feature of the Cisco BPX 8600 series broadband switch.

The **Cell Loss Priority (CLP)** bit follows the PTI bits in all header types. When set, it indicates this cell is subject to discard if congestion is encountered in the network. For frame relay connections, the frame Discard Eligibility is carried by the CLP bit. The CLP bit is also set at the ingress to the network for all cells carrying user data transmitted above the minimum rate guaranteed to the user.

Figur e8-3 UNI Header

Bit —	8	7	6	5	4	3	2	1
Byte 1	Flow control				Virtual path identifier			
Byte 2	Virtual path identifier				Virtual circuit identifier			
Byte 3	Virtual circuit identifier							
Byte 4	Virtual circuit identifier				Payload type		Cell loss priority	
Byte 5	Header Error Control (HEC)							


H8147

Figur e8-4 NNI Header

Bit —	8	7	6	5	4	3	2	1
Byte 1	Virtual path identifier							
Byte 2	Virtual path identifier				Virtual circuit identifier			
Byte 3	Virtual circuit identifier							
Byte 4	Virtual circuit identifier				Payload type		Cell loss priority	
Byte 5	Header Error Control (HEC)							

H8148

Figur e8-5 STI Header



HCF: Header Control Field, a 01 indicates an STI Cell
 VPI/VCI: Virtual Path/Virtual Channel Identifiers, same as UNI and NNI.

Payload Class:

0001 Non-Timestamped Data/Constant Blt Rate
 0010 High Priority/Variable Bit Rate
 0011 Voice/Constant Blt Rate
 0100 Bursty Data A/Variable Blt Rate
 0101 Time-Stamped Data/Constant Blt Rate
 0110 Bursty Data B/Variable Blt Rate

CC: Congestion Control

00: No report 10: Congestion
 01: Uncongested 11: Severe Congestion

F: ForeSight Forward Congestion Indication (FFCI).

Set to 1 if FECN in Frame is a 1, or if incoming cell FFCI is a 1, or egress queue experiences congestion.

R: Reserved

PTI: Payload Type Indicator

CLP: Cell Loss Priority. Same as for UNI or NNI. The CLP bit is set to 1 if the DE is set for a frame, or if the first FastPacket in a frame has its CLP set.

PTI, bits 4,3, and 2: bit 4 = 0, user data cell; bit 4 = 1, connection management cell bit 3 = 0, No congestion experienced bit 3 = 1, Congestion experienced bit 2 = 0, for user data cell, indicates CPE information bit 2 = 1, not used			
PTI Bits	Description		
432			
000	User Data Cell	no congestion experienced	SDU Type 0 (CPE information)
001	User Data Cell	no congestion experienced	SDU Type 1
010	User Data Cell	congestion experienced,	SDU Type 0 (CPE information)
011	User Data Cell	congestion experienced,	SDU Type 1
100	Connection Management Cell, OAM F5 Segment Flow Related cell		
101	Congestion Management Cell, OAM F5 End-to-End Flow related cell		
110	Connection Management Cell, reserved for future use.		
111	Connection Management Cell, reserved for future use.		

H8149

Some of the VCI bits have been reserved in the STI header type to implement ForeSight congestion management control, unique to Cisco WAN switching networks, refer to Figure 8-5 and Table 8-2.

Table 8-2 STI Congestion Control Bits

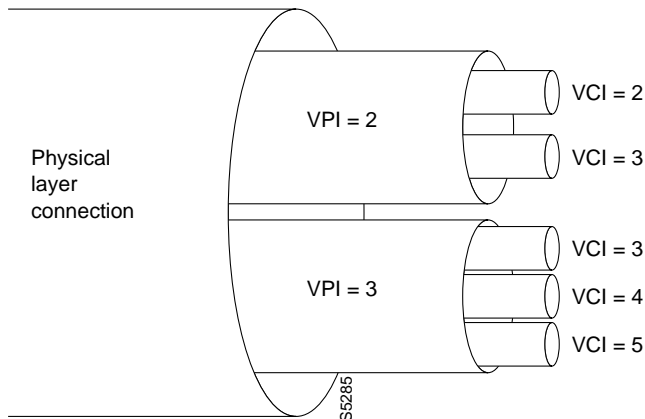
Function	Possible Status
Congestion Control	No report
	Uncongested
	Congestion
	Severe Congestion
Forward Congestion Indicator	No ForeSight congestion indication.
	Congestion indication in incoming cell or congestion detected locally.
Reserved	Not used.

To assure reliable delivery of each ATM cell, a **Header Error Correction (HEC)** checksum field is included as the last field of the header. This is an 8-bit result of a CRC check on the header bits (the payload bits are not checked). The HEC is calculated and inserted after all other fields in the header have been inserted. When ATM cells are carried on unframed digital facilities, the HEC is used for cell delineation.

ATM Cell Addressing


Each ATM cell contains a two-part address, VPI/VCI, in the cell header. This address uniquely identifies an individual ATM virtual connection on a physical interface. Virtual Channel Indicator (VCI) bits are used to identify the individual circuit or connection. Multiple virtual circuits that traverse the same physical layer connection between nodes are grouped together in a virtual path (Figure 8-6). The virtual path address is given by the Virtual Path Indicator (VPI) bits. The Virtual Path can be viewed as a trunk that carries multiple circuits all routed the same between switches.

Figur e8-6 Virtual Paths and Virtual Channels

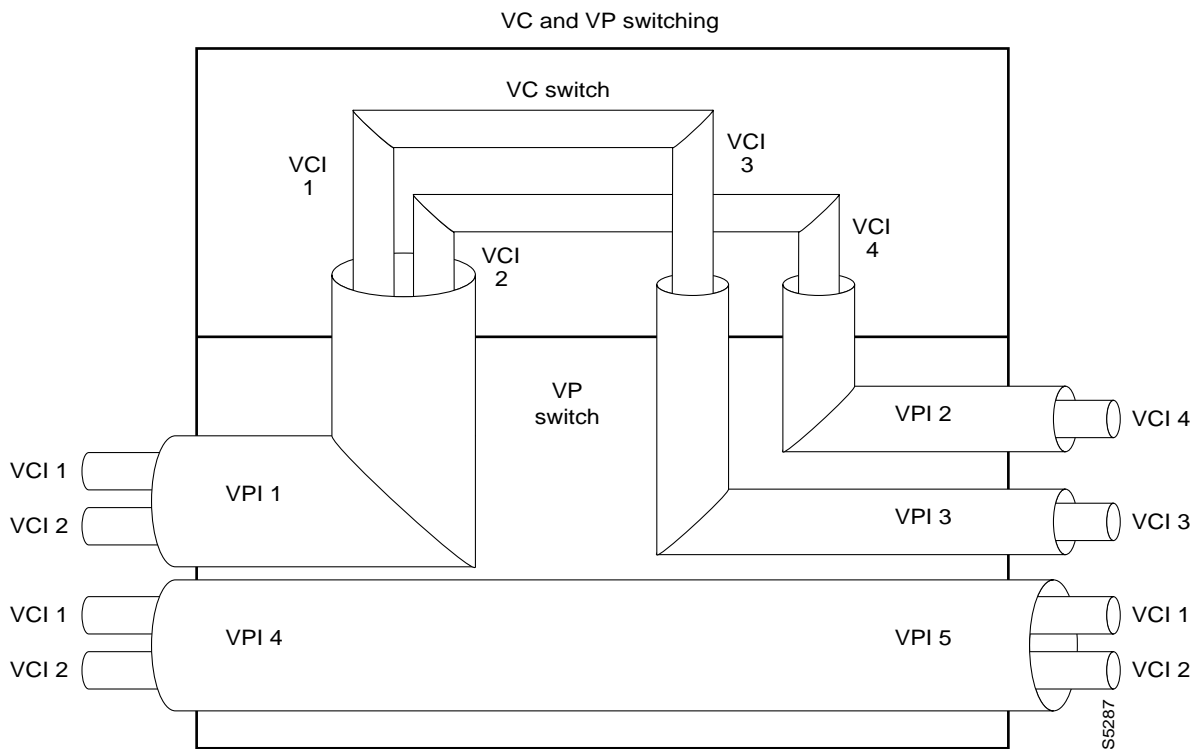


The VPI and VCI addresses may be translated at each ATM switch in the network connection route. They are unique only for a given physical link. Therefore, they may be reused in other parts of the network as long as care is taken to avoid conflicts. Figure 8-7 illustrates switching using VP only, which may be done at tandem switches while Figure 8-8 illustrates switching on VC as well as VP.

Figur e8-7 VP-only Switching



Figur e8-8 VP and VC Switching



The VCI field is 16 bits wide with UNI and NNI header types described earlier. This allows for a total possible 65, 535 unique circuit numbers. The UNI header reserves 8 bits for VPI (256 unique paths) while the NNI reserves 12 bits (4,096 unique paths) as it is likely that more virtual paths will be routed between networks than between a user and the network. The STI header reserves 8 bits for VCI and 10 bits for VPI addresses.

ATM Adaptation Layer

The purpose of the ATM Adaptation Layer (AAL) is to receive the data from the various sources or applications and convert, or adapt, it to 48-byte segments that will fit into the payload of an ATM cell. Since ATM benefits from its ability to accommodate data from various sources with differing characteristics, the Adaptation Layer must be flexible.

Traffic from the various sources have been categorized by the standards committees into four general classifications, Class A through Class D, as indicated in Table 8-3. This categorization is somewhat preliminary and initial developments have indicated that it may be desirable to have more than these initial four classes of service.

Table 8-3 *Classes of Traffic and Associated AAL Layers*

Traffic Class	Class A	Class B	Class C	Class D
Adaption Layer (AAL)	AAL-1	AAL-2	AAL-3/4 AAL-5	AAL-3/4
Connection Mode	Connection-oriented	Connection-oriented	Connection-oriented	Connectionless
End-to-End Timing Relationship	Yes	Yes	No	No
Bit Rate	Constant	Variable	Variable	Variable
Adaptation Layer (AAL)	1	2	3/4, 5	3/4
Examples	PCM voice, constant bit-rate video	Variable bit-rate voice and video	Frame relay, SNA, TCP-IP, E-mail	SMDS

Initially four different adaptation layers (AAL1 through AAL4) were envisioned for the four classes of traffic. However, since AAL3 and AAL4 both could carry Class C as well as Class D traffic and since the differences between AAL3 and AAL4 were so slight, the two have been combined into one AAL3/4.

AAL3/4 is quite complex and carries a considerable overhead. Therefore, a fifth adaptation layer, AAL5, has been adopted for carrying Class C traffic, which is simpler and eliminates much of the overhead of the proposed AAL3/4. AAL5 is referred to as the Simple and Efficient Adaptation Layer, or SEAL, and is used for frame relay data.

Since ATM is inherently a connection-oriented transport mechanism and since the early applications of ATM will be heavily oriented towards LAN traffic, many of the initial ATM products will be implemented supporting the Class C Adaptation Layer with AAL5 Adaptation Layer processing for carrying frame relay traffic.

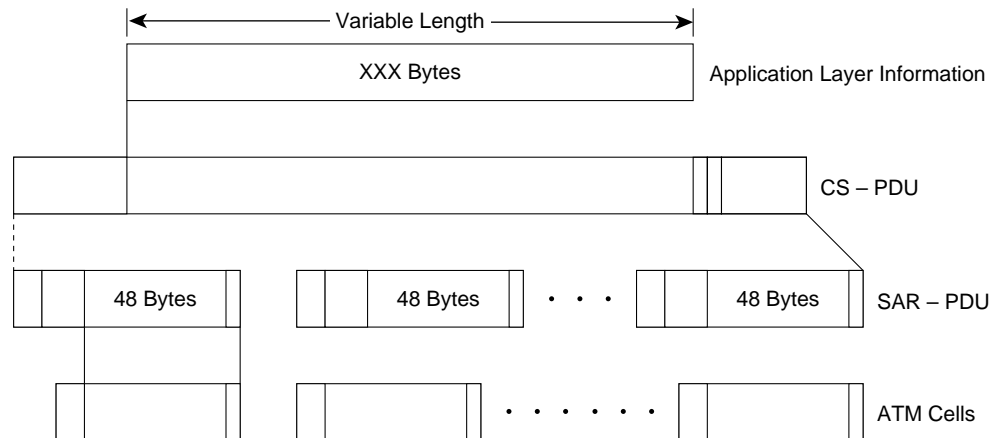
Referring back to Figure 8-1, the ATM Adaptation Layer consists of two sub-layers:

- Convergence Sub-Layer (CS)
- Segmentation and Reassembly Sub-Layer (SAR)

Data is received from the various applications layers by the Convergence Sub-Layer and mapped into the Segmentation and Reassembly Sub-Layer. User information, typically of variable length, is packetized into data packets called Convergence Sublayer Protocol Data Units (CS-PDUs). Depending on the Adaptation Layer, these variable length CS-PDUs will have a short header, trailer, a small amount of padding, and may have a checksum.

The Segmentation and Reassembly Sub-Layer (SAR) receives these CS-PDUs from the Convergence Sub-Layer and segments them into one or more 48-byte SAR-PDUs, which can be carried in the 48-byte ATM information payload bucket. The SAR-PDU maps directly into the 48-byte payload of the ATM cell transmitted by the Physical Layer. Figur e8-9 illustrates an example of the Adaptation Process.

Figur e8-9 Example of Adaptation Process



H8022

PVCs vs. SVCs

Current implementation of ATM is based on setting up permanent virtual circuits (PVCs). PVCs are defined when the user adds the connection to the network. The routing is programmed by the network operator into routing tables and the circuit operating parameters are assigned. However, the circuit does not utilize any network bandwidth until there is traffic to be carried. Since they resemble nailed up voice connections, there is no signaling requirements for setting up PVCs.

Switched virtual circuits, on the other hand, are established on request by the user and are removed from the network database(s) upon completion of the transmission. In this respect, they resemble dial-up voice connections and are under control of the user. ATM standards for support SVCs are under development.

Cisco IGX 8400 Series Multiband and Cisco IPX Narrowband Trunk Interfaces to ATM

The Cisco IGX 8400 series multiband switch or the Cisco IPX narrowband switch can connect to a Cisco IGX 8600 series broadband switch or other ATM switch via an AIT/BTM T3 or E3 trunk. The AIT (Cisco IPX narrowband switch) or BTM (Cisco IGX 8400 series multiband switch) can operate in several different addressing modes selected by the user (see Table8-4 and Figur e8-10). The BPX Addressing Mode (BAM) is used for all Cisco WAN switching ATM Networks. To allow the Cisco IGX 8400 series multiband switch or the Cisco IPX narrowband switch to be used in mixed networks with other ATM switches, there are two other addressing modes available, Cloud Addressing Mode (CAM) and Simple Addressing Mode (SAM).

BAM

In the BPX Addressing Mode (BAM), used for all Cisco WAN switching networks, the system software determines VPI and VCI values for each connection that is added to the network. The user enters the beginning and end points of the connection and the software automatically programs routing tables in each node that will carry the connection to translate the VPI/VCI address. The user does not need to enter anything more. This mode uses the STI header format and can support all of the optional Cisco WAN switching features.

SAM

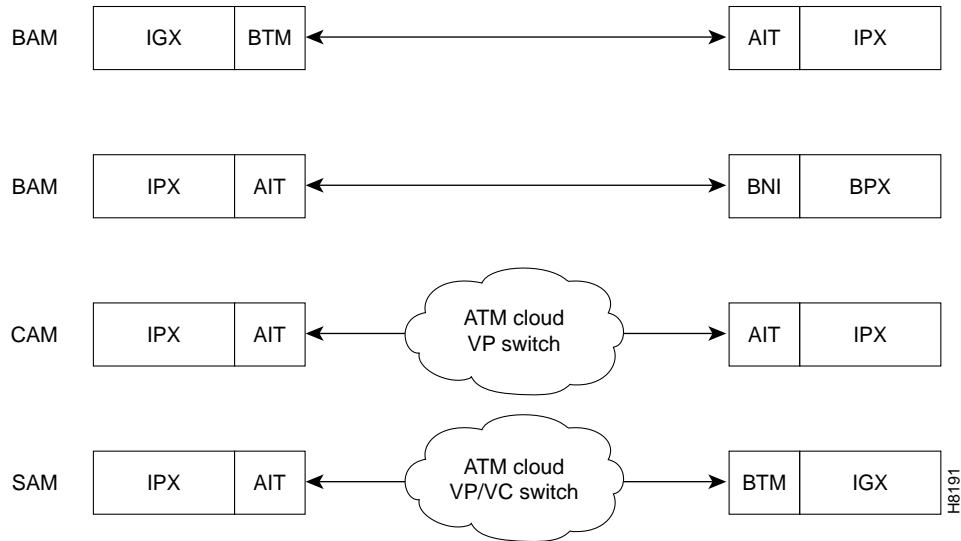
In the Simple Addressing Mode, the user must manually program the path whole address, both VPI and VCI values.

CAM

The Cloud Addressing Mode is used in mixed networks where the virtual path addresses are programmed by the user and the switch decodes the VCI address. Both CAM and SAM utilize the UNI header type.

Table 8-4 ATM Cell Addressing Modes

Addressing Mode	Hdr. Type	Derivation of VPI/VC	Where Used
BAM—BPX Addressing Mode	STI	VPI/VCI = Node Derived Address	Between Cisco IPX narrowband (or Cisco IGX 8400 series multiband) and Cisco BPX 8600 series broadband nodes, or between Cisco IPX narrowband (or Cisco IGX 8400 series multiband) nodes.
CAM—Cloud Addressing Mode	UNI	VPI = User Programmed VCI = Node Derived Address	Cisco IPX narrowband to Cisco IPX narrowband (or Cisco IGX 8400 series multiband) connections over networks using ATM switches that switch on VPI only. VPI is manually programmed by user. Terminating Cisco IPX narrowband switch converts VCI address to FastPacket address.
SAM—Simple Addressing Mode	UNI	VPI/VCI = User Programmed	Cisco IPX narrowband to Cisco IPX narrowband (or Cisco IGX 8400 series multiband) connections over networks using ATM switches that switch where all routing is manually programmed by user, both VPI and VCI.

Figur e8-10 BAM, CAM, and SAM Configurations

Note: IPX with AIT card are interchangeable with IGX with BTM card in this diagram.

FastPacket Adaptation to ATM

A specialized adaptation that is of particular interest to users of Cisco WAN switching equipment is the adaptation of Cisco IPX narrowband FastPackets to ATM cells. There are a large number of narrowband IPX networks currently in existence that are efficiently carrying voice, video, data, and frame relay. A means must be provided to allow these networks to grow by providing a migration path to broadband.

Since FastPackets are already a form of cell relay, the adaptation of FastPackets to ATM cells is relatively simple.

Simple Gateway

With the Simple Gateway protocol, the AIT card in the Cisco IPX narrowband switch (or the BTM card in the Cisco IGX 8400 series multiband switch) loads 24-byte FastPacket cells into ATM cells in ways that are consistent with each application. (Each of the two FastPacket cells loaded into the ATM cell is loaded in its entirety, including the FastPacket header.) For example, two FastPackets can be loaded into one ATM cell provided they both have the same destination. This adaptation is performed by the Cisco IPX narrowband AIT card or the Cisco IGX 8400 series multiband BTM card.

The AIT (or BTM) is configured to wait a given interval for a second FastPacket to combine in one ATM cell for each FastPacket type. The cell is transmitted half full if the wait interval expires. High priority and non-time stamped packets are given a short wait interval. High priority FastPackets will not wait for a second FastPacket. The ATM trunk interface will always wait for frame relay data (bursty data) to send two packets. NPC traffic will always have two FastPackets in an ATM cell.


Complex Gateway, Frame Relay to ATM Network Interworkin

Starting with Release 8.1, with the Complex Gateway capability, the FRSM card in the Cisco MGX 8220 edge concentrator, the AIT card in the Cisco IPX narrowband switch (or the BTM card in the Cisco IGX 8400 series multiband switch) streams the frame relay data into ATM cells, cell after cell, until the frame has been completely transmitted. Since only the data from the FastPacket is loaded, the Complex

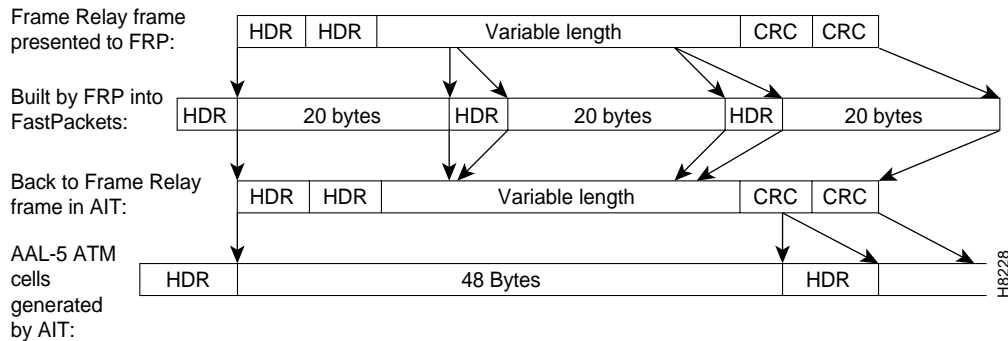
Gateway is an efficient mechanism. Also, discard eligibility information carried by the frame relay bit is mapped to the ATM cell CLP bit, and vice versa. See Chapter 14 for further information on frame relay to ATM interworking. A comparison of the simple gateway and complex gateway formats is shown in Figure 8-11.

Figur e8-11 Simple and Complex Gateway Formats

Simple gateway (AIT card) :



Complex gateway (AIT Card) :



ATM Cell Switching

An ATM switch is conceptually a simple device taking ATM cells from an input port and transferring them to an output. Routing of cells through an ATM switch is directed by tables that interpret the VPI/VCI addresses in the ATM cell header. The simplicity of ATM switching is one of the advantages that promotes its use in broadband networking.

There are two fundamental types of ATM switches, those utilizing a bus architecture and those using a matrix switching fabric. Bus-based ATM switches are primarily utilized in LAN equipment since they are limited in backplane speeds but easily support multicasting (a requirement of LAN equipment). Bus type switches are referred to as shared media switches as the bandwidth is shared among all users.


ATM switches for WAN networks, on the other hand, require the higher switching speeds available with a crosspoint matrix switch. These types of switches are often square, having the same number of output ports as input ports with sufficient crosspoints to be able to switch a cell from any input port to an output port. Routing of cells through an ATM crosspoint switch is directed by tables that interpret the VPI/VCI addresses in the ATM cell header. The total bandwidth of the crosspoint switch element is available to relay a cell from input to output. Buffering is used in matrix switches to avoid contention and blocking.

The BCC controller card in the Cisco BPX 8600 series broadband switch utilizes either a 16 X 16 or 16 X 32 crosspoint switch implemented with a very high speed VLSI switching device. With the BCC-3, which employs a 16 x 16 crosspoint switch, the Cisco BPX 8600 series broadband switch operates at

9.6 Gbps. With the BCC-4, which employs a 16 x 32 crosspoint switch, the Cisco BPX 8600 series broadband switch can operate at 19.6 Gbps when it is also equipped with BXM cards. The crosspoint switch is under control of an arbiter that takes requests from each port with traffic waiting to be switched and sets up the appropriate crosspoint. Multiple crosspoints are operated simultaneously at each switch cycle to connect various input and output ports as long as there is no contention for a particular crosspoint. Figure 8-12 illustrates several typical cycles of a 16 X 16 crosspoint matrix.

The arbiter polling is programmed by system software to insure each port is given equal access to the switch matrix. If there are cells from several input ports destined for the same output port, the arbiter selects one of the inputs on one switch cycle and grants access to another input at a later cycle to prevent switch contention (and resulting blocking). A small buffer at each input temporarily holds the ATM cells to prevent loss of cells in this situation.

Figur e8-12 Operation ofa Typical Crosspoint Switch Matrix




Cycle "N"

In this frame, arbiter has configured four parallel connections:

- Slot 1→ 2
- Slot 3→ 6
- Slot 6→ 9
- Slot 10→ 4

(Dots indicate connections made)



Cycle "N+1"

Arbiter reconfigures crosspoint to five new parallel connections:

- Slot 2→ 1
- Slot 3→ 12
- Slot 6→ 5
- Slot 8→ 14
- Slot 11→ 2

S5288

Broadband (ATM) Trunk Formats

The following paragraphs describe the digital line format for various types of digital transmission lines used to transmit ATM cells throughout Cisco WAN switching networks. These lines operate at data rates of typically 45 Mbps and higher and are referred to as broadband.

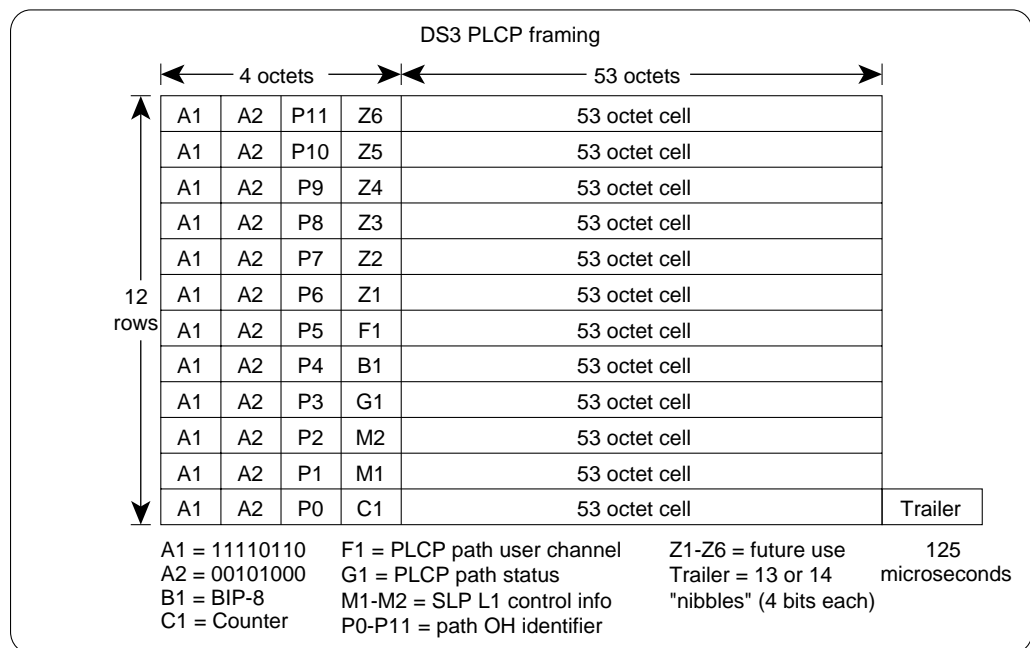
DS3 PLCP Frame Structure

T3 trunks can be used for transmission of packets on links requiring higher capacity than is available with T1 lines. They operate at the DS3 bit rate of 44.736 Mbps. Because of the higher bit rate, T3 trunks are generally carried over fiberoptic or digital microwave.

Transport of ATM cells at the DS3 rate is accomplished using the Switched Multimegabit Data Services (SMDS) Physical Layer Convergence Protocol (PLCP) framing structure as defined in IEEE 802.6 and Bellcore TR-TSV-000773 specifications. The DS3 M-frame pattern is observed, but there is no direct correlation between M-frames and PLCP framing. Figure 8-13 illustrates the DS3 PLCP Frame Sequence.

The DS3 PLCP frames occur at a rate of 8000 per second, or one every 125 μ sec. Since one DS3 PLCP frame can carry twelve 53-octet cells, the bandwidth capacity in cells per second rate for DS3 trunks is (8000 frames/sec x 12 cells/frame = 96,000 cells per sec.). Since the T3 signal is bipolar, it carries the clocking along with the data just as is done with T1. The node recovers receive clock and uses it to clock in the receive data. The signal is B3ZS encoded (similar to the B8ZS used on T1) to scramble the bit stream to eliminate long strings of "0"s or "1"s so that ones density requirement of T1 is not a consideration.

Figur e8-13 DS3 PLCP Frame Format




G.804 E3 Frame Structure

Transport of ATM cells at the E3 rate is accomplished using a framing structure as defined in ITU-T (CCITT) recommendations G.832 and G.804. The frame consists of 537 octets (8-bit bytes) with 7 bytes of overhead occurring every 125 μsec. Figure 8-14 illustrates this frame format. Cisco WAN switching nodes monitor the two frame alignment octets and set the MA payload bits to ATM and does nothing with the remaining overhead bits.

The G.804 frame can transmit 10 ATM cells and so the bandwidth of the E3 trunks is (8000 frames/sec. x 10 cells/frame = 80,000 cells/sec.) The ATM cells are arranged in 9 rows of 59 octets each and cell #1 is not constrained to begin immediately following the frame alignment octets. This frame structure is not compatible with earlier ITU-T recommendations for E3 lines.

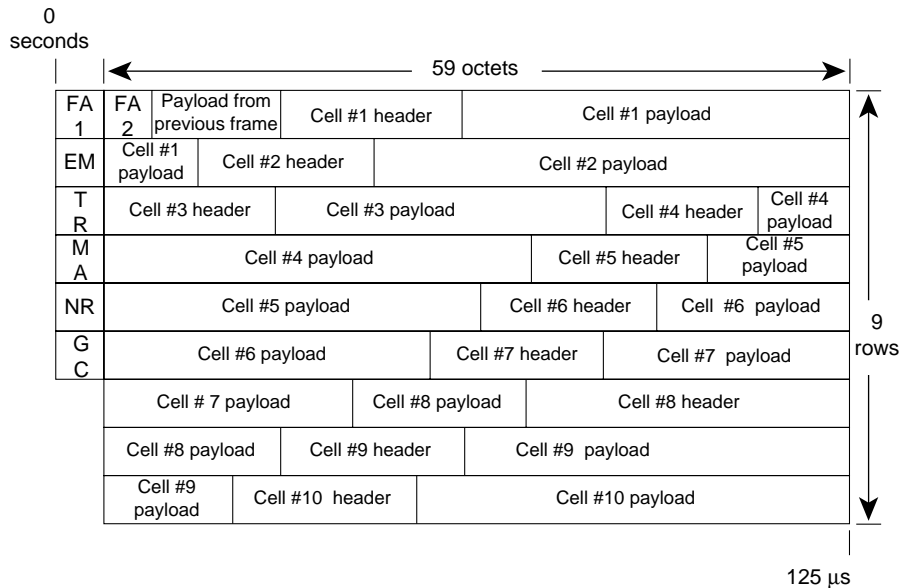
E3 trunks encode the data in a form called HDB3, which also eliminates long strings of zeros. Error monitoring is provided in the Error Monitoring octet which is an 8-bit number representing the Bit Interleaved Parity (BIP-8) for the bits in the previous frame.

Figur e8-14 G.804 E3 Frame Format



FA = Frame Alignment
 EM = Error Monitoring, BIP-8
 TR = Trail Trace
 MA = Maintenance & Adaptation
 NR = Network Operator Byte overhead channel
 GC = General purpose Communications channel

MA = Bit 1 = Far End Receive Failure (FERF)
 Bit 2 = Far End Block Error (FEBE)
 Bits 3 to 5 = payload type:
 Bits 6 & 7 = SDH TU payload dependent indications
 Bit 8 = Timing marker:



06290

Virtual Trunks

Virtual trunking provides the ability to define multiple trunks within a single physical trunk port interface. Virtual trunking benefits include the following:

- Reduced cost by configuring the virtual trunks supplied by the Public Carrier for only as much bandwidth as needed instead of the full T3, E3, or OC3 bandwidths.
- Utilization of the full mesh capability of the Public Carrier to reduce the number of leased lines needed between nodes in the CiscoWAN switching subnetworks.
- Or, choice of keeping existing leased lines between nodes, but using virtual trunks for backup.
- Ability to connect BNI trunk interfaces to a Public Network using the standard ATM UNI cell format.
- Virtual trunking can be provisioned via either a Public ATM Cloud or a Cisco WAN switching ATM cloud.

A virtual trunk may be defined as a “trunk over a public ATM service”. The trunk really doesn’t exist as a physical line in the network. Rather, an additional level of reference, called a **virtual trunk number**, is used to differentiate the virtual trunks found within a physical trunk port.

With only a single trunk port attached to a single ATM port in the cloud, a node uses the virtual trunks to connect to multiple destination nodes on the other side of the cloud.

Since a virtual trunk is defined within a trunk port, its physical characteristics are derived from the port. All the virtual trunks within a port have the same port attributes.

(Note: All port and trunk attributes of a trunk are configured with **cnftrk** or **cnftrkparm**)

Virtual Trunk Capacities

In Release 8.2, a BNI T3/E3 or BNI-155 (OC3) trunk from the WAN switching network connects to an ATM UNI interface at the Public Network ATM Cloud. If the cloud uses exclusively Cisco WAN switching equipment, this UNI interface is provided by an ASI-T3/E3 or ASI-155 (OC3). Trunk and channel capacities are as follows:

- Maximum number of virtual trunks per port
 - 32 virtual trunks per BNI-T3/E3 port
 - 12 virtual trunks per BNI-OC3 port
- Maximum number of logical (physical and virtual) trunks per node
 - 32 logical trunks per Cisco BPX 8600 series broadband node
- Total connection channels per port which are shared by all virtual trunks on the port. The number of channels used by all virtual trunks on a port cannot exceed the total number of channels on the port
 - 1771 per BNI-T3/E3
 - 15867 per BNI-OC3
- Total bandwidth per port is shared by all virtual trunks on the port. The sum of the bandwidth of all the virtual trunks on a port cannot exceed the bandwidth of the port.
 - 96,000 cells per second per T3 port
 - 80,000 cells per second per E3 port
 - 353,208 cells per second per OC3 port

VPC Configuration within the ATM Cloud

In order for a virtual trunk to successfully transmit data through the ATM cloud, the ATM equipment in the cloud must support Virtual Path switching and transmit incoming cells based on the VPI in the cell header.

A virtual path connection (VPC) is configured in the cloud to join two endpoints. The VPC can support either CBR, VBR, or ABR traffic. A unique VPI per each VPC is used to move data from one endpoint to the other. The Cisco WAN switching equipment at the edge of the cloud transmits cells which match the VPCs VPI value. As a result the cells are switched to the other end of the cloud.

Within the ATM cloud one virtual trunk is equivalent to one VPC. Since the VPC is switched with just the VPI value, the 16 VCI bits (from the ATM-UNI format) of the ATM cell header are passed transparently through to the other end.

If the public ATM cloud consists of Cisco BPX 8600 series broadband nodes, the access points to the cloud are ASI ATM-UNI ports. Since the cells transmitted to the ASI trunk interface are coming from a Cisco WAN switching device, e.g., a BNI card, the 16 VCI bits have already been left shifted by 4 bits and contain 12 bits of VCI information and 4 bits of ForeSight information. Therefore, the ASI cards at either end of the cloud and of the VPC are configured **not** to shift the VCI when formatting the cells with an STI header for transport through the cloud. (Note: The command **cnfport** is modified to allow the user to configure “no shifting” on the ASI port. Cisco BPX 8600 series broadband software Release 8.2 or higher is required to support this new configuration.)

If the ATM cloud consists of non-Cisco WAN switching nodes, then the 12 VCI + 4 Foresight bits in the cells coming from the BNI card in the Cisco BPX 8600 series broadband are passed through untouched as 16 VCI bits. Since it is a non-Cisco WAN switching network, the ForeSight bits are ignored.

Virtual Trunk Traffic Classe

All types of CiscoWAN switching traffic are supported over virtual trunks through an ATM cloud. Every trunk is defaulted to carry every type of traffic. The CBR, VBR, ABR, and UBR virtual trunks within the cloud should be configured to carry the correct type of traffic. The recommended traffi configurations are as follows:

- CBR Trunk: ATM CBR traffic, voice/data traffic
- VBR Trunk: ATM VBR traffic, frame relay traffic
- ABR Trunk: ATM ABR traffic, frame relay ForeSight traffic
- UBR Trunk: ATM UBR traffic

The CBR trunk is best suited to carry delay sensitive traffic such as Cisco IGX 8400 series multiband voice/data and Cisco BPX 8600 series broadband CBR traffic. The VBR trunk is best suited to carry Cisco IGX 8400 series multiband frame relay and Cisco BPX 8600 series broadband VBR traffic. The ABR trunk is best suited to carry Cisco IGX 8400 series multiband ForeSight and BPX ABR traffic. The user can change the types of traffic each trunk carries. However, to avoid unpredictable results, it is best to conform to the recommended traffic types for a given type of VPC.

A user can configure any number of virtual trunks between two ports up to the maximum number of virtual trunks per port and the maximum number of logical trunks per node. These trunks can be any of the three trunk types, CBR, VBR, or ABR.

Virtual Trunk Addressin

Cells transmitted to a virtual trunk use the standard ATM-UNI cell format. Because of the UNI format, two types of information found in the STI header are no longer available in cells received from a virtual trunk. The Header Control Field (HCF) is unavailable by definition of the UNI format. The payload information is removed to increase the number of VCI bits from 8 to 12 per VPI.

- No HCF in the cell means that only one “type” of cell is expected from the virtual trunk. The GFCI bits in the cell are not modified on entering the cloud and are non-deterministic exiting the cloud. Cell filtering at the receiving end treats every cell the same way, i.e. the trunk does not distinguish between VCC or VPC cells.
- With no payload field in the cell, software must manually program a fixed payload type for every connection. The payload type is already preprogrammed at all nodes for path connections and at terminating nodes only for non-path connections. With virtual trunks, the payload type must be preprogrammed at the edge of the cloud since the cell header does not contain payload information.

The BNI trunk at the edge of the ATM cloud is programmed to ignore the payload information in the cell header. The correct payload (queue) assigned to a cell is determined by the contents of the Bframe configured for the cells of that connection.

Several cases exist in which the payload field changes for a connection. But since the payload field no longer exists in the cell header, the payload type is fixed for the life of the connection. Therefore, the following have to be handled differently for virtual trunks.


- **Modem activity over voice traffic**—The payload type transitions from voice to Non-Timestamped (NTS) for modem activity in order to not drop the modem data if congestion is encountered. In order to allow this priority to occur on a CBR virtual trunk, the CLP bit is set for all normal priority voice traffic. When modem activity is detected, the CLP bit is not set, thereby giving the modem traffic the servicing advantage.
- **Initial voice spurt**—The initial voice spurt is sent as High-Priority traffic to ensure the other end receives the cells in a timely manner. Similar to the modem activity, the voice spurt cells do not have the CLP bit set. This reduces the chance of dropping initial voice spurt cells.
- **ForeSight Round Trip Delay (RTD)**—The RTD measurement currently uses the High-Priority queue to determine the transmission delay on a ForeSight connection route. For routing over a virtual trunk the payload type for a ForeSight connection is hard-configured, so the RTD measurement uses this queue. By calculating the ForeSight delay after a connection is routed, a good approximation of the uncongested delay can be obtained. The software periodically checks the delay and stores the smallest value. If the true uncongested RTD is less than the smallest measured delay, the ForeSight adjustments may not be as precise as before, but the measurements should be adequate for the feature to work correctly.

The trunk card at the edge of the cloud ensures that cells destined for a cloud VPC have the correct VPI/VCI. The VPI is an 8-bit value ranging from 1–255. The VCI is a 12-bit value ranging from 1–4095. The standard UNI VCI is 16 bits, but the 4 least significant bits are used as ForeSight bits by the Cisco WAN switching trunks.

A variety of cell types may exist at a port (CBR, VBR, ABR). Each of the virtual trunks sends cells with the same VPI/VCI format for all the cell types.

Connection (non-path) traffic—The connection identifier is stored in the VCI as shown in Figure 8-14

Figur e8-15 Connection Identifier



Virtual Trunk Example

The following example describes a typical scenario of adding one virtual trunk across an ATM network (. On one side of the cloud is BPX_A with a BNI trunk card in slot 4. On the other side of the cloud is BPX_B with a BNI trunk card in slot 10. A virtual trunk is added between port 4.3 on the BNI in BPX_A and port 10.1 on the BNI in BPX_B.

A VPC within the cloud must be configured first.

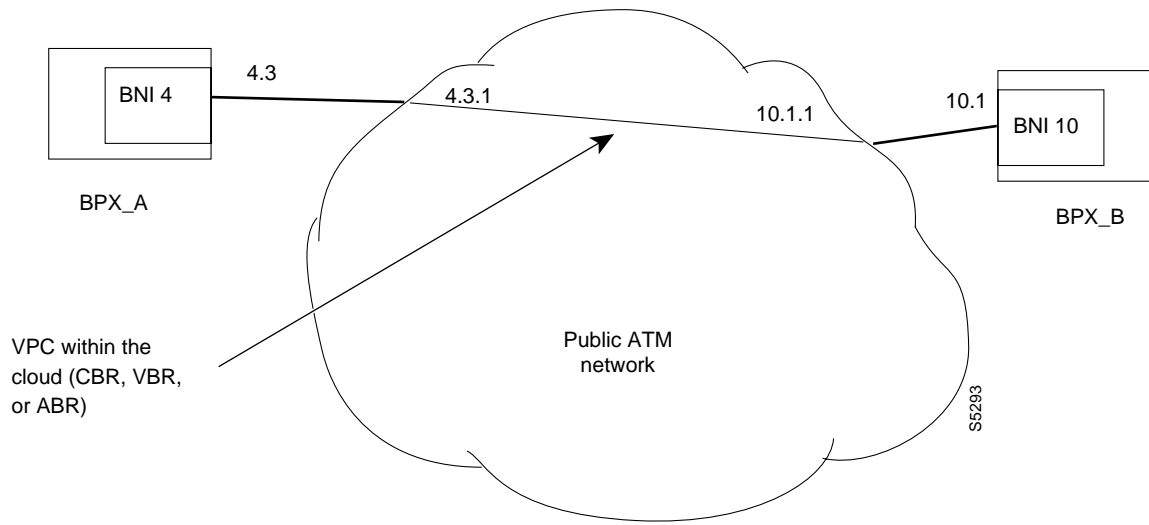
- BPX_A **uptrk 4.3.1** up virtual trunk #1 on BNI trunk port 4.3
- BPX_A **cnftrk 4.3.1** configure VPI, VPC type, traffic classes, # of connection channels
- BPX_B **uptrk 10.1.1** up virtual trunk #1 on BNI trunk port 10.1
- BPX_B **cnftrk 10.1.1** configure VPI, VPC type, traffic classes, # of connection channels
- BPX_A **addtrk 4.3.1** add the virtual trunk between the two nodes
(addtrk 10.1.1 at BPX_B would do the same)

The VPI values chosen during **cnftrk** must match those used by the cloud VPC. In addition both ends of the virtual trunk must match on VPC type, traffic classes supported, and number of connection channels supported. The **addtrk** command checks for matching values before allowing the trunk to be added to the network topology.

The network topology from BPX_A's perspective after the trunk addition would be

BPX_A 4.3.1 to 10.1.1/BPX_B

Figur e8-16 Single Virtual Trunk Addition



Full Mesh with Virtual Trunks

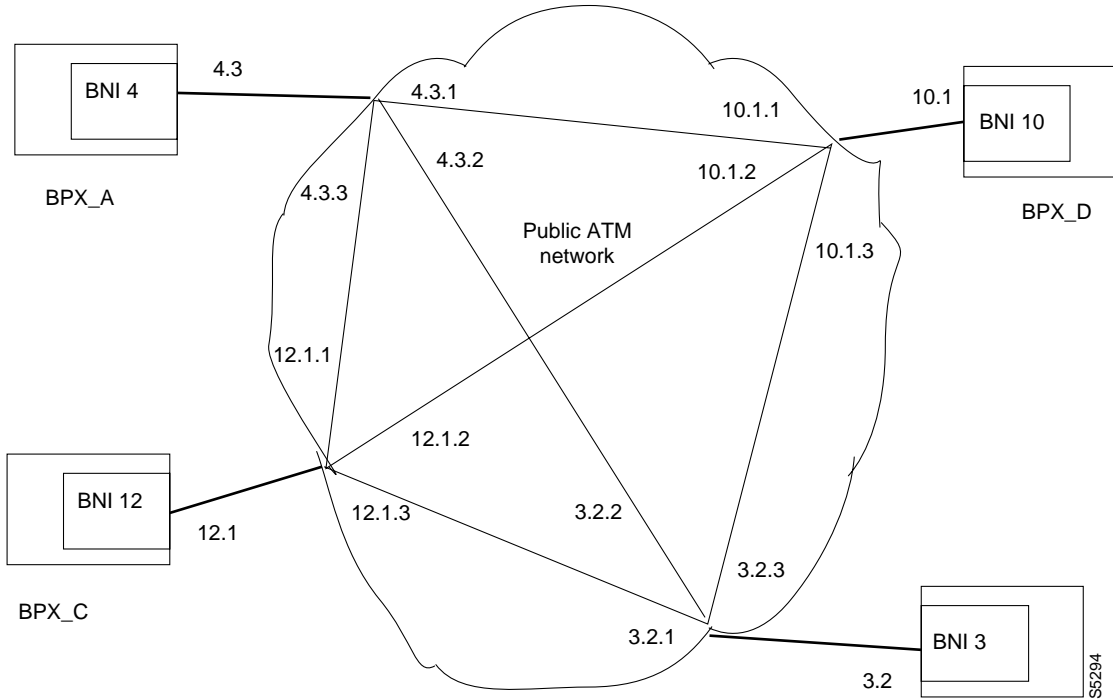
One of the purposes of virtual trunking is to increase the efficiency of connectivity through an ATM cloud. The following example describes how virtual trunks may be used to fully mesh multiple nodes by attaching them to a cloud.

In this 4-node example, Figure 8-17, only four trunk ports are used to link into the cloud, yet all four nodes are directly connected through the cloud with six virtual trunks. The fanout of three virtual trunk endpoints per port produces a savings of two ports per node.

Adding an additional node to this network would require adding one physical link to the cloud. By increasing the fanout of virtual trunks at each port by one, all the nodes would still be fully connected.

This savings of trunk ports provides a much lower resource cost in using an ATM service to connect a network.

Figur e8-17 Four Node Example of Virtual Trunking



BNI One Stage Queueing

Each virtual trunk consists of one queue (Q_BIN) on the BNI. This queue corresponds to the configured VPC type for the virtual trunk; CBR, VBR or ABR. No distinction is made between the different types of CBR, VBR, and ABR traffic. For example, voice and NTS data traffic are placed into the same CBR queue.

The BNI-T3/E3 contains 32 cell queues per port. The BNI-OC3 contains 12 cell queues per port.

The user commands **cnftrk** and **cnftrkparm** are used to configure the one queue within the virtual trunk.

Virtual Trunk Statistics

BNI Virtual Trun

The following statistics are collected for a BNI virtual trunk.

- Cells Sent
- CLP Cells Dropped
- Overflow Cells Dropped
- Max Queue Depth
- Total cells transmitted
- Total CLP dropped
- Total overflow cells dropped
- Maximum queue depth

Card Redundancy (Y-Redundancy)

Y-Cable redundancy is supported for BNI-T3/E3 trunk cards at the edge of the ATM cloud. For BNI-OC3, Y-redundancy is not supported.

Virtual Trunk Alarms

Trunk Specific Alarm

A virtual trunk has alarms which may be generated solely from the trunk itself. These are statistical alarms only.

BNI Virtual Trunk Alarm

The following queue statistical alarms are available.

- Tx HP Cell Drops
- Tx Voice Cell Drops
- Tx TS Cell Drops
- Tx NTS Cell Drops
- Tx Bursty Data A Cell Drops
- Tx Bursty Data B Cell Drops
- Tx CBR Cell Drops
- Tx VBR Cell Drops
- Tx ABR Cell Drops

Trunk Port Alarm

A virtual trunk also has trunk port alarms which are shared with all the other virtual trunks on the port. These alarms are cleared and set together for all the virtual trunks sharing the same port.

Feeder Trunk Support

A virtual trunk cannot be used as a feeder trunk. Connections cannot be terminated on a feeder trunk. Both of these are restricted at the user interface.

Connection Management

Routing VPCs over Virtual Trunk

The routing algorithm excludes VPCs from being routed over a virtual trunk. The reason for this restriction is due to how the virtual trunk is defined within the ATM cloud.

The cloud uses a VPC to represent the virtual trunk. Routing an external VPC across a virtual trunk would consist of routing one VPC over another VPC. This use of VPCs is contrary to its standard definition. A VPC should contain multiple VCCs, not another VPC. In order to avoid any non-standard configuration or use of the ATM cloud, VPCs cannot be routed over a virtual trunk through the cloud.

Structured Networks Support

Structured networks and virtual trunking are not allowed to coexist in the same network.

Error Messages

New error messages for virtual trunks.

- “Port does not support virtual trunking”
- “Port configured for virtual trunking”
- “Invalid virtual trunk number”
- “Maximum trunks per node has been reached”
- “Invalid virtual trunk VPI”
- “Invalid virtual trunk traffic class”
- “Invalid virtual trunk VPC type”
- “Invalid virtual trunk conid capacity”
- “Mismatched virtual trunk configuration”
- Port is not configured for virtual trunks
- Port is not configured for a physical trunk
- Virtual trunk number is invalid
- Trunk limit per node has been reached
- Virtual trunk VPI is invalid
- Virtual trunk traffic class is invalid
- Virtual trunk VPC type is invalid
- Virtual trunk conid capacity is invalid
- Ends of virtual trunk have different configurations

Commands

Syntax

BNI virtual trunk:

```
addrtrk <slot>.<port>.<vtrunk>
```

where <slot> is the BNI slot number

<port> is the BNI port number

<vtrunk> is the virtual trunk number.

FastPackets and Narrowband Trunks

This chapter is provided for users who wish to have an in-depth knowledge of FastPackets and the narrowband digital trunks used to carry them. It describes the various FastPacket types that may be used with the Cisco IGX 8400 series multiband switches and Cisco IPX narrowband switches. It also discusses the various trunks and circuit lines and differences between T1, fractional T1, subrate, E1, and the various narrowband digital line formats.

This chapter contains the following sections:

- FastPacket Formats
- Narrowband Trunk Formats

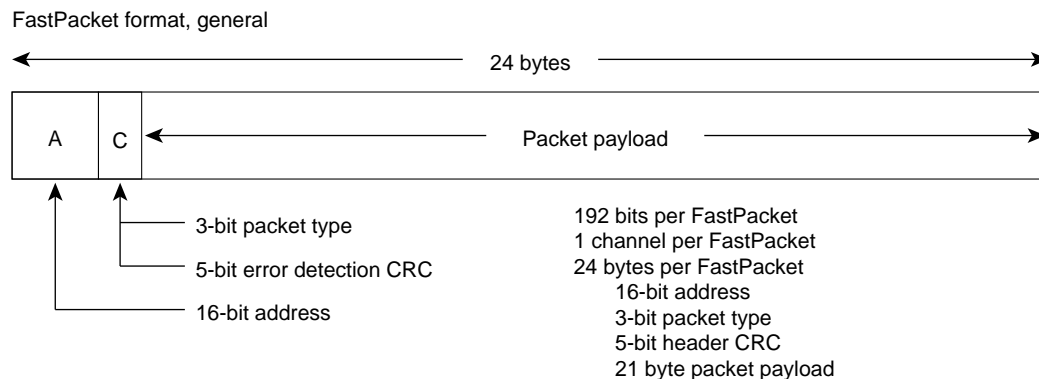
FastPacket Formats

All FastPackets, regardless of the type of information they contain, have precisely the same length and general structure as shown in Figure 9-1. A particular packet contains information from only one source and for only one destination.

The basic FastPacket frame consists of:

- 2-byte (16-bit) Address field.
- 1-byte (8-bit) Control field consisting of a 3-bit packet type and a 5-bit CRC error detection code to check the address field.
- 21-byte (168-bit) Packet Payload for user data.

Figur e9-1 General FastPacket Format



S5197

All FastPackets have a similar 3-byte header that is monitored by almost every card in every Cisco IGX 8400 series multiband or Cisco IPX narrowband node. The address is the destination address of the packet that defines the route to the far end node as well as the terminating card in the far node.

The packet type defines the content of the packet and the CRC is a cyclical redundancy check of the three header bytes to detect any corruption of these bits. The remaining bits in the packet are used to carry payload data. The various packet types may have additional overhead bytes and will be describe later.

There are six main packet types. They are identified by a 3-bit packet type code following the packet address. Packet types are summarized as indicated in Table 9-1.

Table 9-1 Packet Types

Queue	Traffic type
High Priority	NPC/PCC, First 2 packets of voice
Non-Time-Stamped Data	Data > 56 Kbps, "a", "t" voice, modem
Time-Stamped Data	Data ≤ 56 Kbps
Voice, PCM and ADPCM	"c", "v" voice
Bursty Data without ForeSight	Frame relay
Bursty Data with ForeSight	Frame relay

High Priority Packets

Highest priority is given to system control packets and initial bursts of voice packets. These high priority type packets are queued with priority over any other packet type and pass through the network with minimum delay. Cards that generate high priority packets are the NPC and CDP.

High priority packets (Figure 9-2) are generated by the NPC to communicate with other nodes in the network for system-level message interchange. To make this overhead message channel as robust as possible, control fields are added to the header. The receiving NPC checks that all packets received in a message are in sequence. If a packet is missing, the receiving NPC requests a retransmission of that packet. Also, a second CRC byte is added to the end of the packet to perform error-checking on the 20-byte NPC packet message.

Initial voice packets, or talk-spurt packets as they are sometimes referred to, are also transmitted with minimum delay. This is to establish a minimum delay datum for filling the null-timing buffers in the far end packet receivers.

Figure 9-2 High Priority Packet Format

PCC Packet format

Address (16)	Type (3) & CRC (5)	Sequence (6 bytes)	Message (14 bytes)	CRC (8)	S5198

PCM and ADPCM Voice Packets


With the exception of the first two packets of voice, which are sent via high priority packets, digitized voice channel information from circuit lines is carried in the voice packet type (Figure9-3). Voice packets have a 168-bit payload message field. PCM-encoded packets carry twenty-one 64 Kbps (8-bit) PCM samples in this message field.

Figur e9-3 Voice Packet Format

PCM VoicePacket format



ADPCM VoicePacket format



In most cases, the voice channels will be compressed by a factor of either 2:1, 3:1, 4:1, or 8:1 by the CDP card to increase the transmission efficiency. This is illustrated in the ADPCM voice packet format, which can carry 42 compressed voice samples with 32 Kbps compression, 56 samples with 24 Kbps, or 84 samples with 16 Kbps compression. Consider also that many of the voice channels will contain “silence” that will be discarded by the Voice Activity Detector algorithm resulting in no packets being sent.


The “one’s density” requirement for T1 lines (packet line) is inherent in 32 Kbps, 24 Kbps, and 3-level 16 Kbps ADPCM (c16 type) so any T1 facility can be used to transmit this code. However the 4-level 16 Kbps ADPCM (c16z type) may require a clear channel T1 or E1 facility. The CDP generates this packet type.

Non-Timestamped Data Packets

Most data at rates over 56 Kbps is sent using non-timestamped data packets (Figure9-4). The packet payload is filled up with 21 data bytes data from a queue associated with a particular data channel. The data is loaded in a first-in, first-out (FIFO) basis and when the packet payload is full, the FastPacket is transmitted to the destination node.

Figur e9-4 Non-Timestamped Data Packet Format

Non-timestamped data packet format



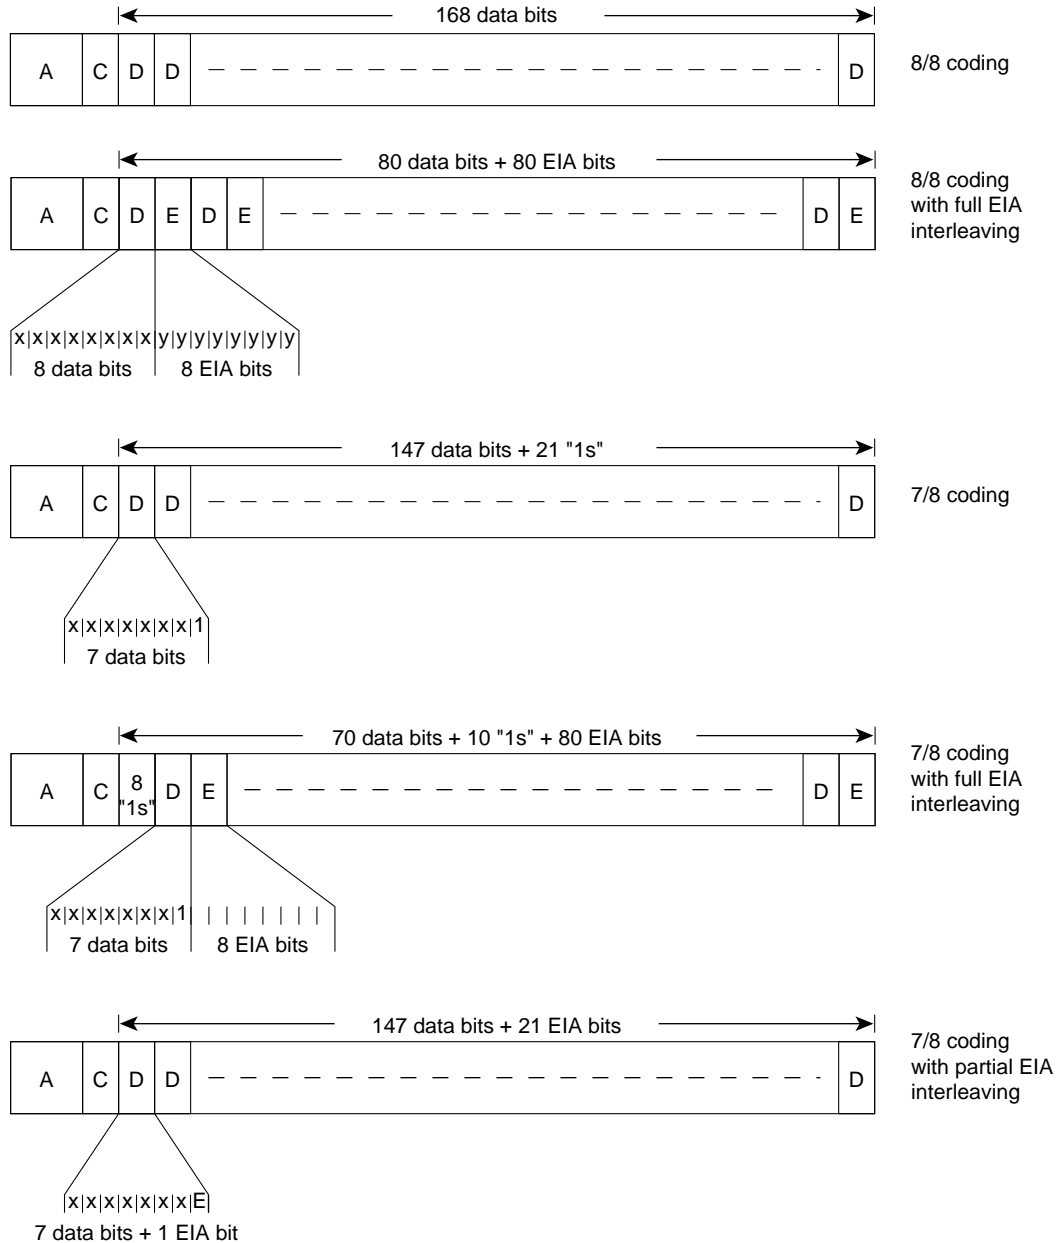
The structure of the data words in the packet message may be one of several types:

- 8/8 encoding, non-interleaved control leads (EIA).
- 8/8I encoding, non-interleaved control leads (EIA).
- 8/8 or 8/8I, fully interleaved control leads (EIA).

- 7/8 coding, non-interleaved control leads (EIA).
- 7/8 coding, fully interleaved control leads (EIA).

The data frame formats for these various types of non-timestamped packets are illustrated in Figure 9-5 and described in the following paragraphs.

Figur e9-5 Various Non-Timestamped Data Message Formats



S5201

The 8/8 coding can be specified in the software for data channels that carry data whose protocol ensures that there are no long strings of zeros in the data. For this data, the packet message carries 168 data bits. EIA control leads are sent in separate packets only when a change of state on one or more control lead

is detected (referred to as non-interleaved data). The data bits are inverted by the Cisco IGX 8400 series multiband switch or Cisco IPX narrowband switch when 8/8I encoding is specified, and is used when the data protocol (such as SDLC) never uses the “FF” character.

When Fast EIA option is specified, the bytes in the message alternate between 8 data bits followed by 8 control lead bits (fully interleaved data). This cuts the throughput in half but ensures the control lead information leads the data within 8 bits. The 8/8 option is used where the attached customer data device strobes the receive data with one of the control leads and there may not be sufficient buffering to avoid loss of data when the control lead change of state is late. Since only 160 bits are used, each packet message starts with a byte loaded with “1s” that is discarded at the far end.

Another format that can be specified is the 7/8 coding. This coding is used when there is no checking for zeros by the connecting data device and there is possibility of long strings of zeros. With 7/8 coding, only the first 7 bits of the message byte carries data, the eighth bit is always forced to a one and gets discarded at the far end. If E1 or subrate packet trunks only are used or if the NTC using B8ZS is used to carry these non-timestamped data packets, the 8/8 coding may be used.

Non-timestamped data is also used by voice circuits that do not employ VAD. These channels are likely to carry data from voice/data multiplexers or from high-speed modems or FAX machines.


Non-timestamped data packets are generally given priority behind high-priority packets and are characterized by low but variable delay. The variability in delay is removed in buffers at the far end Cisco IGX 8400 series multiband switch or Cisco IPX narrowband switch so each channel experiences a consistent amount of end-to-end delay. Cisco IGX 8400 series multiband and Cisco IPX narrowband cards that use non-timestamped data include the HDM for the Cisco IGX 8400 series multiband switch and the SDP and CDP for the Cisco IPX narrowband switch.

Timestamped Data Packets

Timestamped data packets (Figur e9-6) have a 1-byte timestamp field leaving 20 bytes for user data. The timestamp is set by software to a maximum allowable delay figure (typically 40 ms.). This timestamp is checked at every node by the NTC to set the priority of the packet in the outgoing queue. If the packet has accumulated more delay than the maximum indicated in the timestamp, the packet is discarded. If this happens, the user device will have to retransmit this data.

Figur e9-6 Timestamped Data Packet Format

Timestamped data packet format



Timestamped data packets generally have low priority in queueing. However, packets with older timestamps are given higher priority with the various packet line queues. Over a network path with multiple hops, timestamped data generally experiences longer delays than non-timestamped data.

The timestamp is used to reduce the worst case maximum data delay by giving older packets priority. This is especially useful for data channels operating at less than 56 Kbps that might otherwise encounter excessive end-to-end delay when sent over links of five or more hops. There is additional overhead to timestamped data packets. But when the channel data rate is reasonably low, the additional overhead is not so much a penalty.

All of the packet structures described for non-timestamped data, 8/8, 7/8I etc., are applicable to timestamped data. In addition, there is provision for partially interleaved data. This is a modification on the 7/8 coding with Fast EIA described in “Interleaved (Fast EIA) Control Leads” section in Chapter 14,

Synchronous Data Connections. With partially interleaved data, each message byte consists of seven data bits followed by one EIA control lead in the eighth bit. The control lead to be sent with the data is specified by the user. All other EIA control leads are sent in a separate packet as is done for the non-interleaved data.


All Data Frame Multiplexed (DFM) data connections (64 Kbps and below) are timestamped. DFM discards data packets that contain a repetitive data pattern (up to 15 consecutive packets can be discarded) at the transmitting node. The receiving node needs to know how many packets were discarded so it can recreate the packets and maintain the integrity of the transmitted data. For these packets, the timestamp byte is used to specify a sequence number, indicating the number of suppressed packets. Data channels using DFM are prevented from using fast EIA option. Timestamped data packets are generated by HDM, SDP, LDP, and CDP cards.

Frame Relay Data Packets

The frame relay data packet format (sometimes called bursty data) is similar to the timestamped data format. The fourth byte is now used by the frame relay data format for a second control byte (C2) as shown in Figure 9-7. The first four bits of this control byte contains a code indicating whether this packet is the start of the frame relay frame, the end, or one of the frames in the middle. Two bits are reserved for congestion control. The last four bits of this control byte is the hop count. The remaining 20 bytes is payload frame relay data as shown in Figure 9-7.

Figur e9-7 Frame Relay Data Packet Format

Frame Relay data packet format



Idle Packets

Idle packets from an NTC are only four bytes long since the framing is based on the CRC in the header rather than the T1 frame bit. NTC idle packets start with the first byte filled with all ones followed by three bytes such that a valid CRC is contained in the third byte as is found in a normal packet. The following message byte contains a hex DD. The repeating four-byte idle packet pattern contains only a moderate ones density and puts a minimum stress on the repeated packet line.

Packet framing for NTC packet trunks is based on the CRC code rather than either the T1 or E1 framing. Since idle packets are short, recovery from an out-of-frame condition is faster on an NTC packet line. In this condition, almost all packets will be idle and, since they are short packets, the CRC will repeat much more often speeding up the reframe process.

Remote Alarm Packets

Remote alarm packets are used to transmit the yellow alarm code to the node at the far end of the packet line. They are similar to idle packets and have the same all ones initial byte and a valid CRC. The message byte pattern, a hex DB, is unique to the remote alarm code. The remote alarm packets are used only with the NTC.

Narrowband Trunk Formats

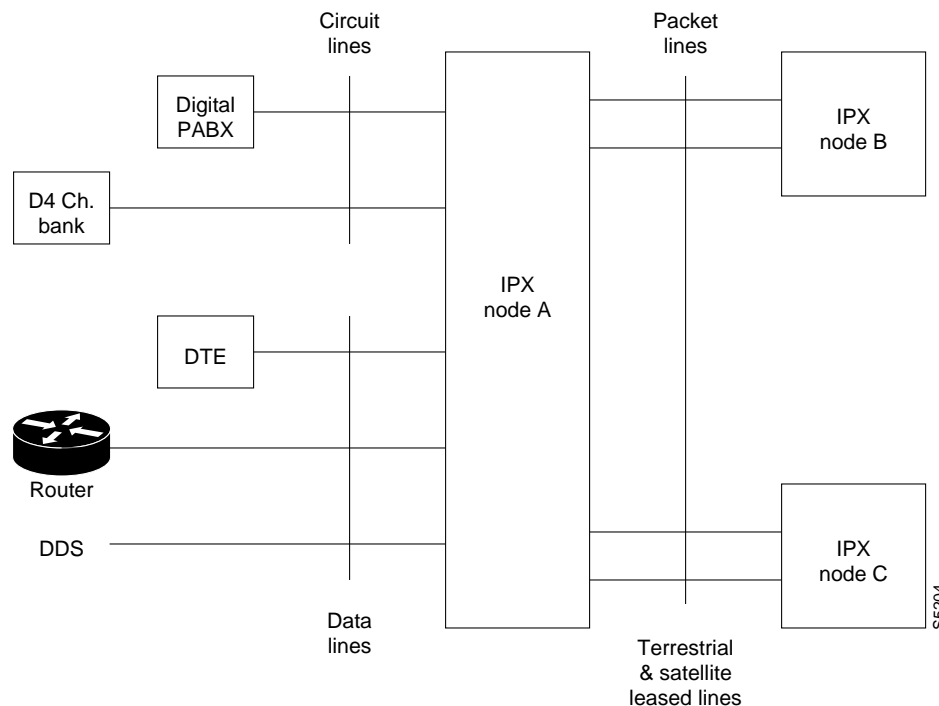
The following paragraphs describe the digital line format for various types of digital transmission lines used to transmit FastPackets throughout Cisco WAN switching networks. These lines operate at data rates of typically 2 Mbps and less and are referred to as narrowband in contrast to trunk rates described previously.

Lines and Trunks

Throughout this manual and other CiscoWAN switching user documentation, the terms lines (sometimes called circuit lines) and trunks (sometimes called packet lines) are used when describing the narrowband interfaces to Cisco IGX 8400 series multiband or Cisco IPX narrowband switches. The following paragraphs describe the subtle differences between trunks and lines.

Refer to Figure 9-8, which illustrates a very simple three-node network. As indicated on the figure and by definition, lines connect local sources of digital data and terminate on the Cisco IGX 8400 series multiband or Cisco IPX narrowband switch. These sources of digital signals may include digital FastPADs, PABX's, and D4 channel banks for voice transmission and Data Terminal Equipment (DTEs), bridges, routers, video compression devices, or AT&T Digital Data Service[®] lines for data transmission.

Figur e9-8 Lines and Trunks (sometimes called Circuit Lines and Packet Lines)



The Cisco IGX 8400 series multiband or Cisco IPX narrowband node connects these local sources of data to the digital transmission lines of the long-haul, high-speed, digital networks of various telecommunication common carriers for transmission to other nodes at locations around the country or around the world. This may be accomplished using terrestrial lines of copper, fiber optics, or digital microwave radio. Often, for international circuits, satellite links are used.

One of the main purposes of the Cisco IGX 8400 series multiband or Cisco IPX narrowband switch is to condense the data and transmit it more efficiently over these lines. This is accomplished by packetizing the incoming data before transmitting it. Packet lines, as the term is used in this application, are digital transmission lines, carrying packetized data in the form of FastPackets, connecting the various nodes.

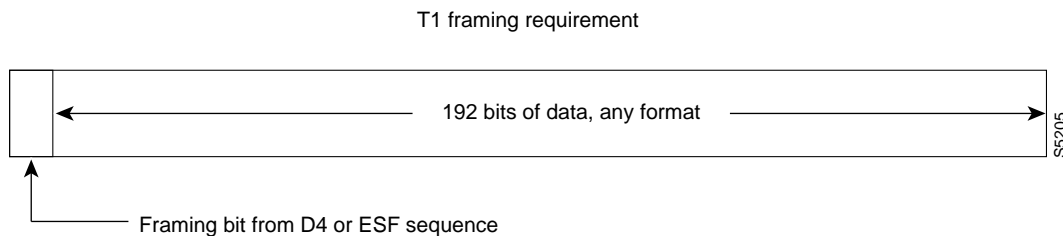
In digital transmission systems using the AT&T defined digital hierarchy, both circuit lines and packet trunks are T1 and fractional T1, and T3 transmission lines. In locations employing the CEPT standards, E1 digital transmission lines are used. Packet trunks connect to the network side of the Cisco IGX 8400 series multiband or Cisco IPX narrowband switch whereas circuit lines connect to the customer side of the Cisco IGX 8400 series multiband or Cisco IPX narrowband switch. In addition to the connection point, the primary difference between circuit lines and packet trunks is in the format of the transmitted data on these lines.

In summary, packet trunks are used to link Cisco IGX 8400 series multiband or Cisco IPX narrowband nodes together in a network configuration and use FastPacket frame formats for the data. Circuit lines are used to interface the Cisco IGX 8400 series multiband or Cisco IPX narrowband switch to customer equipment or circuits using one of several industry-standard digital transmission data formats.

T1 Packet Trunks

T1 lines transmit data at 1.544 million bits per second (Mbps). This is called the Digital Signal Level 1 (DS1) rate. To keep the transmitting and receiving devices at each end of the line synchronized, a framing bit is periodically inserted in the data stream (see “T1 Packet Trunks” section). This framing bit is given a unique pattern that is easily recognized by the terminating equipment. The data bits between each framing bit is referred to as a frame. On a T1 line, there are 8000 frames/second with each frame consisting of 193 bits (192 bits of data and one framing bit).

Figur e9-9 T1 Framing Requirement

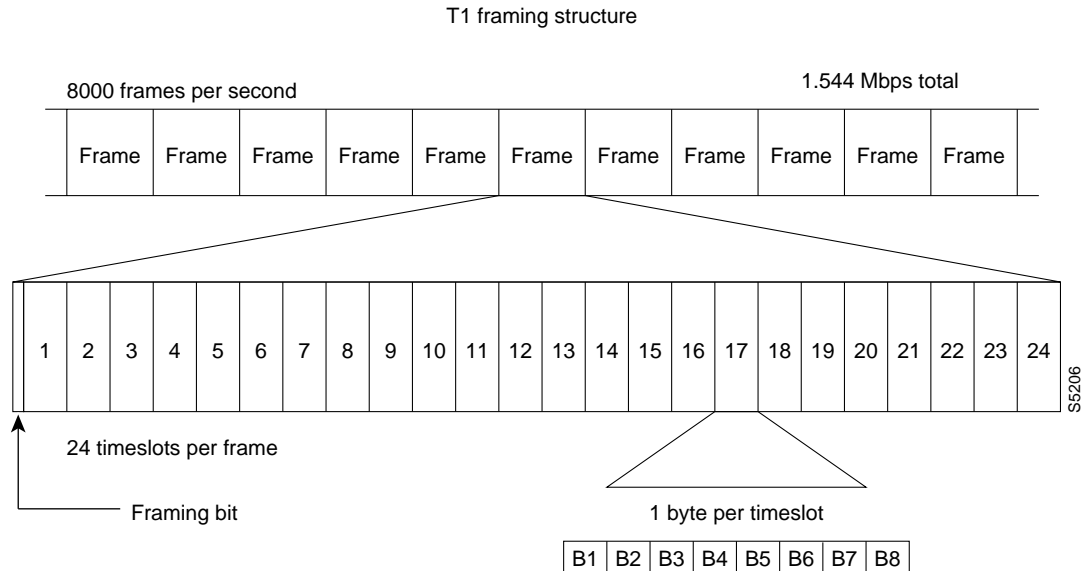


T1 transmission networks are not concerned with the actual channelization of the T1 signal, only with the frame bit. Packet lines use these 192-bit T1 frames in a somewhat different manner by organizing them into a 192-bit FastPacket format. A FastPacket is a group of 192 bits and on a T1 line the FastPacket is aligned with the 192-bit D4 frame. Of the total 1.544 Mbps available with a T1 line, 1.536 Mbps is available for transmitting user data as there are 8000 framing bits per second overhead.

On T1 circuit lines, the data in each frame is organized in a standard format (Figur e9-10). There are twenty-four bytes of data, each byte consisting of 8-bits. Frames of data are transmitted 8000 times a second. Each of the 24 bytes represent a single customer voice channel or data channel. These channels

operate at a rate of 64 Kbps each are referred to as DS0 channels. The channels in this format, called D4 frame format, are numbered in a numeric sequence 1 through 24. The start of each frame is a single framing bit that defines the start of channel 1.

Figur e9-10 T1 Frame Format



Fractional T1 Packet Trunks

A fractional T1 line can also be used as a narrowband trunk. A fractional T1 line, like a normal T1 line, operates at 1.544 Mbps and consists of 8000 frames per second. The fractional trunks use only as many 64-Kbps channels (DS0s) as needed, instead of using a full E1 or T1 trunk. These bytes may be arranged either adjacent or alternating and are usually groups of multiples of four channels.

Fractional trunks have the following characteristics:

- Allows for the definition and management of Cisco IGX 8400 series multiband or Cisco IPX narrowband trunks with reduced aggregate bandwidths.
- Views the bandwidth on the trunk side at the software level as a set of 64-Kbps channels (or DS0s) called a bundle. The Cisco IGX 8400 series multiband switch or Cisco IPX narrowband switch supports any trunk speed from 256 Kbps to an upper limit of 1.544 Mbps (T1) or 2.048 Mbps (E1) in 64 Kbps increments.
- Provides bundle management (alarms and events must be detected/asserted at the bundle level).
- Affects the Bandwidth Management and Connection Management software with the introduction of an extended set of Fractional Trunk aggregate bandwidths.
- Adjusts the sizes of the Network Interface (NTC) queues when the trunk bandwidth varies. The new sizes guarantee an acceptable network delay with a reliable congestion control.

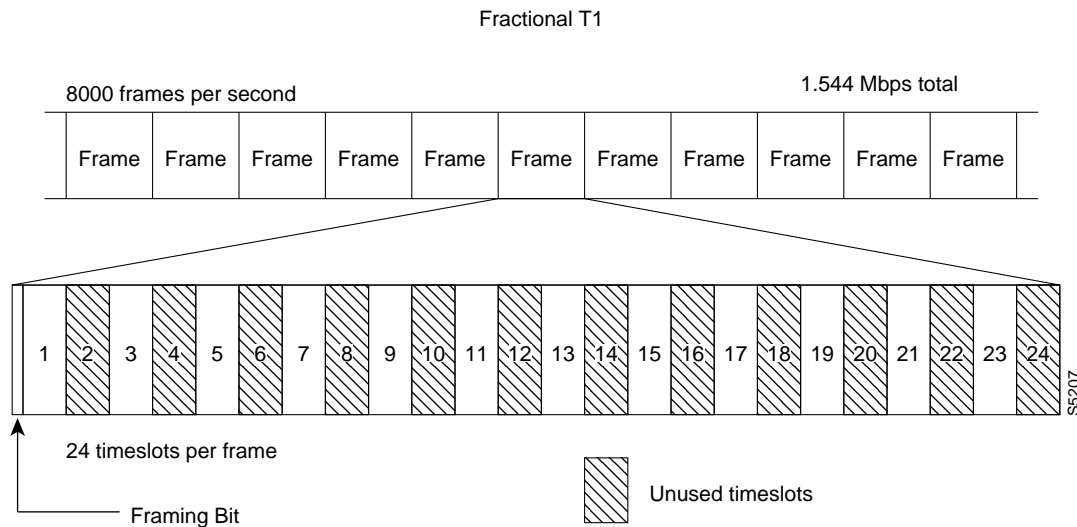


Note

The relative percentage of bandwidth used by the high priority traffic between Cisco IGX 8400 series multiband or Cisco IPX narrowband nodes is significant in the case of small trunk aggregate bandwidths.

When using a fractional T1 trunk, a FastPacket is sent in 24 consecutive channels, regardless of how many frames are spanned in the process. In Figure 9-11, the fractional T1 line will support user data in 12 of the 24 bytes in each frame. In this example, it will take two frames to transmit a full FastPacket. Fractional T1 is used when only a limited capacity is required between some nodes such as a tail circuit node.

Figur e9-11 Fractional T1 Frame Format



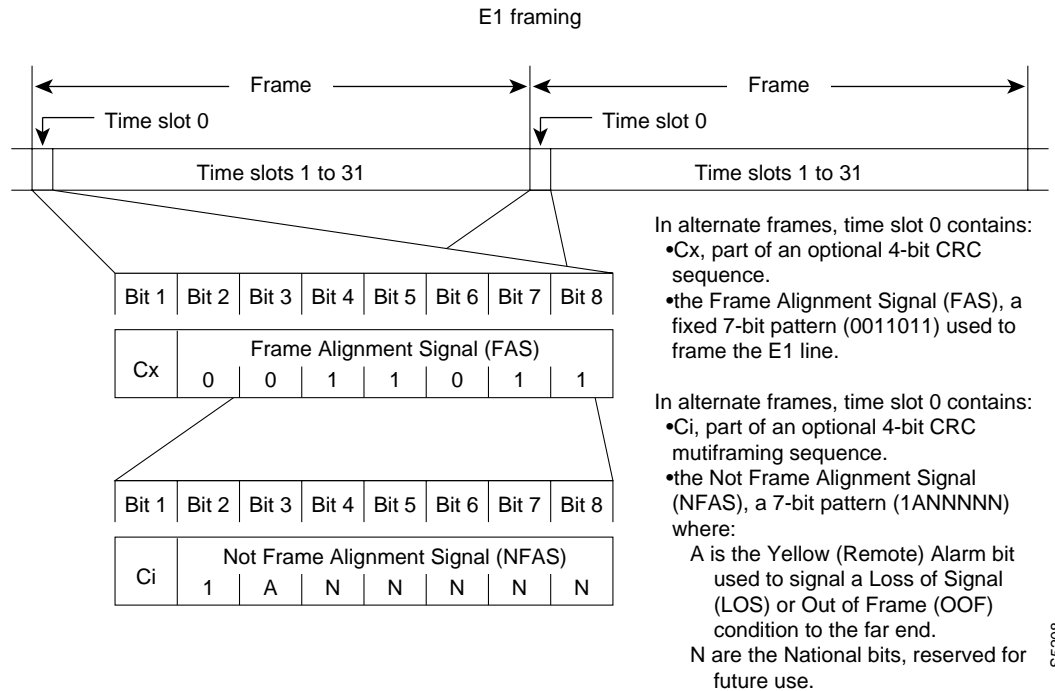
E1 Packet Trunks

In many international locations, the CEPT E1 is commonly used for digital transmission lines. E1 lines can also be used as narrowband trunks. While the 8-bit byte is still used, as in T1 lines, the bit rate, the number of channels and the frame format differs between the two lines. E1 lines carry 30, 31, or 32 channels at a transmission rate of 2.048 Mbps.

There are 32 channels in the E1 frame as shown in Figure 9-12. Channels are generally numbered from 0 to 31 with channel 0 reserved for frame alignment and CRC error checking. Channel 16 is used to carry circuit signalling information in many (but not all) applications and often cannot be used for customer data. However, all other channels 1–15 and 17–31 may be used for user data. The E1 frame size is 8 bits/byte X 32 bytes = 256 bits.

Since the E1 frame size is not the same as it is with T1 or with a FastPacket, the FastPacket frame boundaries are not aligned on an E1 line as they are on a T1 line. For earlier Cisco IGX 8400 series multiband or Cisco IPX narrowband systems, it was convenient to use the T1 framing bit to synchronize not only the T1 line but the FastPacket. This cannot be used in the E1 world so a different method of framing is used for both E1 and T1 in current Cisco IGX 8400 series multiband or Cisco IPX narrowband systems. Refer to the “Packet Framing (NTM and NTC)” section for further details.

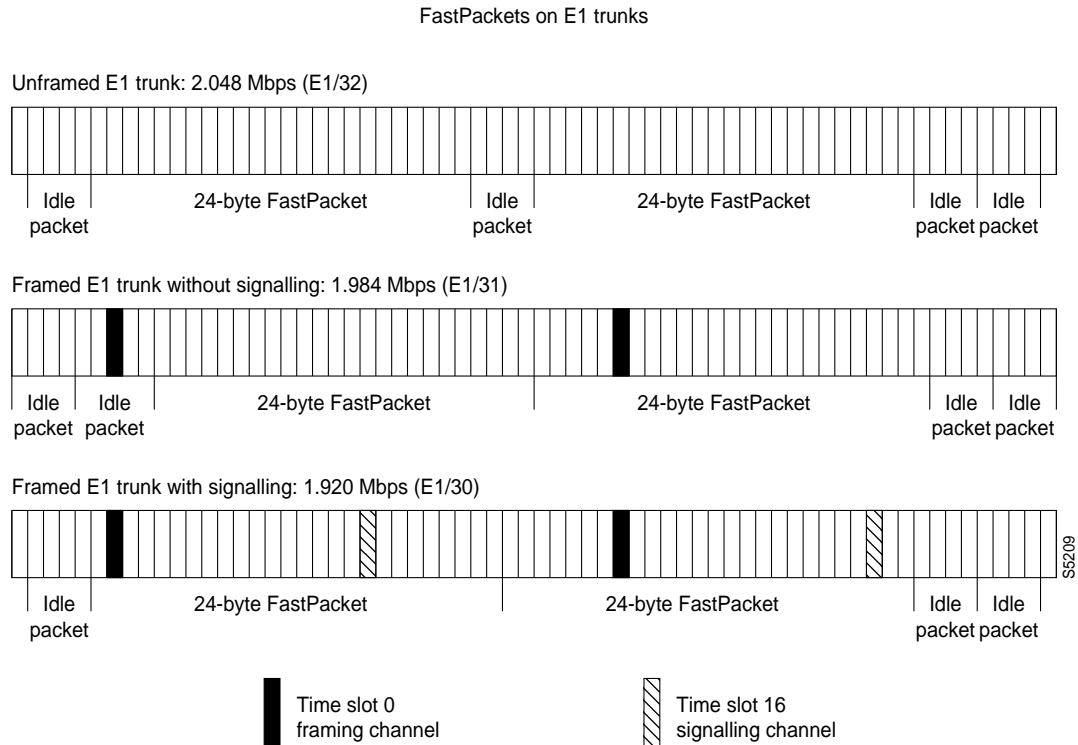
Figur e9-12 E1 Frame Format



The 24-byte FastPacket format cannot be changed but the 192 bits are transported in all available user transport bits in the E1 line. To determine where the FastPacket begins and ends, a synchronizing technique is used that looks at all user data bits for a valid FastPacket CRC code. When this is located among all the user data bits, the Cisco IGX 8400 series multiband switch or Cisco IPX narrowband switch now knows where the start of the FastPacket frame is and ignores the E1 frame synchronizing information.

On most E1 lines 30 channels are available for FastPackets. These lines are referred to as E1/30. There are 1.920 Mbps available for FastPacket data versus 1.536 Mbps available with T1. If channel 16 is not used for signalling (E1/31), it is available for data bringing the available data rate up to 1.984 Mbps. On a totally unframed E1 line (E1/32), the full 2.048 Mbps. data rate can be used for FastPackets (Figure 9-13).

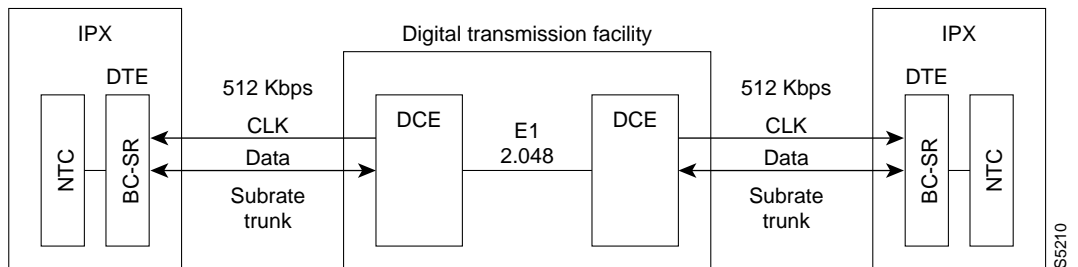
Figur e9-13 FastPackets with Various E1 Formats



Subrate Packet Trunks

Subrate lines can also be used as a packet trunks. Subrate lines usually utilize a portion of an E1 or T1 transmission facility (see Figure 9-14). They are available mainly in Europe and other international locations, but in North America they are used for special purposes such as satellite hub access. Subrate trunks, like fractional T1 trunks, are used when only a limited capacity is required between some nodes such as a tail circuit.

Figur e9-14 Typical SubrateTrunk Setup



Subrate packet trunks function like a synchronous data circuit and operate at specified data rates from 64 Kbps to 1.920 Mbps. A penalty in using the lower rates is that the network overhead becomes a significant percentage of the overall transmission rate.

Unlike T1 and fractional T1 trunks, subrate trunks carry their own clock with the data. As such, the BC-SR card synchronizes to the subrate trunk with looped clock. Allowed clock rates are: 64 Kbps, 128 Kbps, 256 Kbps, 384 Kbps, 512 Kbps, 768 Kbps, 1.024 Mbps, 1.536 Mbps, and 1.920 Mbps. In

addition, there is no frame format in the subrate trunk and no framing bit(s) defined. The Cisco IGX 8400 series multiband switch or Cisco IPX narrowband switch uses its own proprietary synchronizing technique to synchronize end-to-end transmission on subrate trunks.

Cisco IGX 8400 series multiband or Cisco IPX narrowband cards that are used to interface to subrate trunks are the NTM and the BC-SR back card on the Cisco IGX 8400 series multiband switch or the NTC and the BC-SR back card on the Cisco IPX narrowband switch. The subrate trunk is always a **DCE** so the BC-SR interface is configured as a **DTE** as indicated in Figure 9-14. Three interfaces, V.11/X.21, V.35, and MIL188/RS-449, are available with the BC-SR card. Most of the standard EIA RS-232 and ITU-T V.35 type control leads are supported and can be monitored and/or conditioned.

The receive clock is monitored for out of range or clock loss. Because data timing is carried on separate clock leads along with the data, there is no requirement for sufficient ones density as in the T1 packet trunks.

Packet Framing (NTM and NTC)

Both the NTM and the NTC cards use a unique Cisco WAN switching frame acquisition algorithm for providing end-to-end data transmission synchronization independent of any of the standard digital line frame formats previously discussed in this chapter. This is the reason the NTM and the NTC cards will operate with various types of line interface (back) cards such as the BC-T1, BC-E1, and BC-SR.

With the NTM or NTC mode of packet line framing, the concept of line framing and packet framing have been disassociated. A FastPacket may start at any byte boundary of the standard E1 or T1 line format. Even a line that is completely unframed (for example unframed E1 lines or subrate lines) can be used for a packet line. Unframed lines provide no byte boundaries; they act just like a large data pipe. In these cases, a FastPacket can begin at any arbitrary bit position.

Refer to back to Figure 9-10, illustrating the general FastPacket data format. All FastPacket types have a 2-byte address code followed by a 1-byte packet type and CRC for these bytes only. NTM or NTC framing is done by searching for a four-byte sequence in which the third byte contains a valid CRC for the sequence. The four-byte sequence instead of a three-byte sequence was chosen to accommodate FastPacket types that contain certain information in the fourth byte.

With this framing mode, it is no longer necessary to use a full 24-byte idle packet on the packet line any time there is no FastPacket ready to send. Consequently, the NTC uses a 4-byte idle packet that reduces the queuing delay of a FastPacket in the NTC.

Occasionally a packet trunk will have a large percentage of the available bandwidth filled with idle packets. In order to reduce the negative effects of a repeating 4-byte pattern on the packet line, there are four different idle packet codes that are used in sequence. Each packet starts with a hex "FF" followed by three bytes such that a valid CRC is contained in the third byte.

If the receiving node loses track of the CRC (causing it also to lose frame sequence), it issues a Loss of Packet alarm to indicate the packet line is no longer serviceable. If this occurs, the far node is notified of this failure by transmitting a remote alarm code. If the failure happens to be a one-way fault, the remote alarm code makes certain the far node will avoid using this packet line and reroute traffic to another. If the problem is a two-way fault, the far end will already have detected a Loss of Packet Frame on its own and will have already initiated a reroute.

There are four additional packet codes used to send a remote (yellow) alarm to the far end node when a failure on the incoming direction of the packet line is detected. These packets also start with the hex "FF" in a manner similar to the idle packets. Both idle packets and remote alarm packets are discarded at the far node since they contain no customer data.

Because the NTM or NTC framing does not depend on any transmission media protocol, it is very easy for FastPackets to travel along a number of different interconnecting trunk types on hops that may span many different physical level trunks. There are no protocol converters required to pass from a subrate tail circuit, for example, to a full trans-continental T1 or even T3 toll trunk.

PART 4

**ATM AND FRAME RELAY
CONNECTIONS**

ATM Connections

This chapter describes how ATM connection services are established by adding ATM connections between ATM service interface ports in the network using ATM standard UNI 3.1 and Traffic Management 4.0. It describes BXM and ASI card operation and summarizes ATM connection parameter configuration

The chapter contains the following:

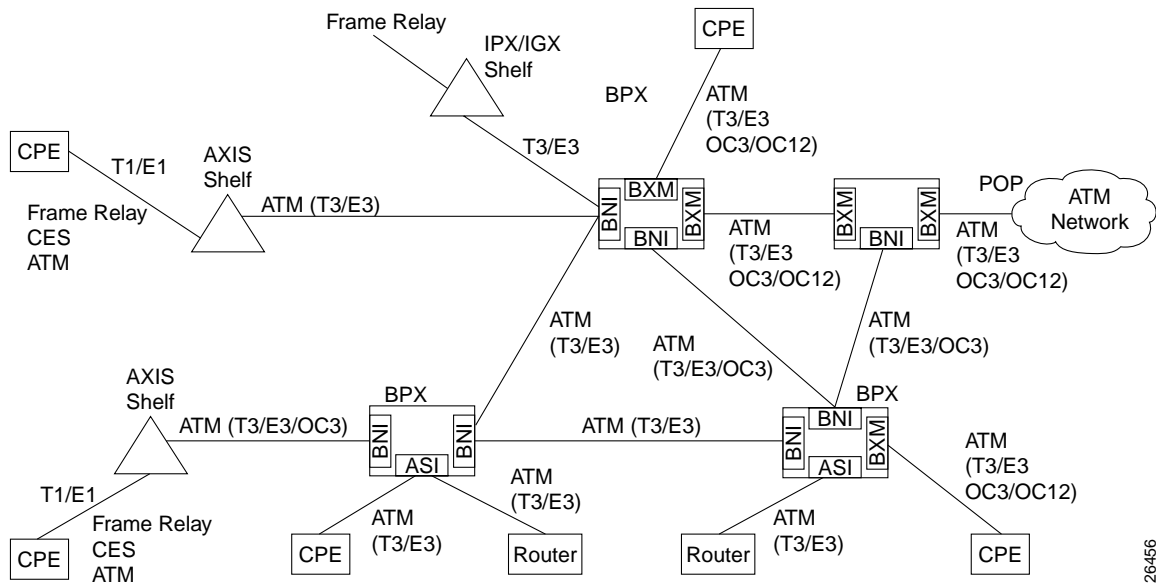
- ATM Connection Services
- SVCs
- Traffic Management Overview
- ATM Connection Requirements
- ATM Connection Configuration
- Traffic Policing Examples
- LMI and ILMI Parameters

ATM Connection Services

ATM connection services are established by adding ATM connections between ATM service interface ports in the network. ATM connections can originate and terminate on the ASI (ATM Service Interface) cards, on BXM-T3/E3, BXM-155 (OC-3), and BXM-622 (OC-12) cards configured for port (service access) operation on the Cisco BPX 8600 series broadband switch, or on the Cisco MGX 8220 edge concentrator shelf (using the AUSM card for the Cisco MGX 8220 edge concentrator). Frame relay to ATM network interworking connections are supported between either BXM or ASI cards to the Cisco IGX 8400 series multiband switch, Cisco IPX narrowband switch, or Cisco MGX 8220 edge concentrator. Frame relay to ATM service interworking connections are supported between either BXM or ASI cards to FRSM cards on the Cisco MGX 8220 edge concentrator shelf.

Figure 10-1 is a depiction of ATM connections over a Cisco BPX 8600 series broadband network. It shows ATM connections via BXM-T3/E3, BXM-155, BXM-622, ASI-1, and ASI-155 cards, as well as over Cisco MGX 8220 edge concentrator shelves. It also shows frame relay to ATM interworking connections over Cisco MGX 8220 edge concentrator, Cisco IGX 8400 series multiband, and Cisco IGX narrowband shelves. For further information on the Cisco MGX 8220 edge concentrator shelf, refer to the *Cisco MGX 8220 Installation and Configuration* document.

Figur e10-1 ATM Connections over a Cisco BPX 8600 Series Broadband Network



26456

SVCs

When an Extended Services Processor (ESP) is collocated with a Cisco BPX 8600 series broadband switch, you have what is called a Cisco BPX 8600 series extended services node. ATM and Frame Relay Switched Virtual Circuits (SVCs) are supported in addition to Permanent Virtual Circuits (PVCs). For further information on ATM SVCs, refer to **Chapter 10**, SVCs, ATM and Frame Relay and to the *BPX Service Node Extended Processor Installation and Operation document*.

Traffic Management Overview

The ATM Forum Traffic Management 4.0 Specification defines five basic traffic classes:

- CBR (Constant Bit Rate)
- rt-VBR (Real-Time Variable Bit Rate)
- nrt-VBR (Non-Real Time Variable Bit Rate)
- UBR (Unspecified Bit Rate)
- ABR (Available Bit Rate)

Table 10-1 summarizes the major attributes of each of the traffic management classes:

Table 10-1 Standard ATM Traffic Classes

Attribut	CBR	rt-VBR	nrt-VBR	UBR	ABR
Traffic Parameters					
PCR & CDVT	x	x	x	x	x

Table 10-1 Standard ATM Traffic Classes (continued)

Attribut	CBR	rt-VBR	nrt-VBR	UBR	ABR
SCR & MBS		x	x		
MCR					x
QoS Parameters					
Pk-to-Pk CDV	x	x			
Max CTD	x	x			
CLR	x	x	x		nw specific
Other Attributes					
Congestion Control Feedback					x

Traffic parameters are defined as:

- PCR (Peak Cell Rate in cells/sec): the maximum rate at which a connection can transmi
- CDVT (Cell Delay Variation Tolerance in usec): establishes the time scale over which the PCR is policed. This is set to allow for jitter (CDV) that is introduced for example, by upstream nodes.
- MBS (Maximum Burst Size in cells): is the maximum number of cells that may burst at the PCR but still be compliant. This is used to determine the BT (Burst Tolerance) which controls the time scale over which the SCR (Sustained Cell Rate) is policed.
- MCR (Minimum Cell Rate in cells per second): is the minimum cell rated contracted for delivery by the network.

QoS (Quality of Service) parameters are defined as:

- CDV (Cell Delay Variation): a measure of the cell jitter introduced by network elements
- Max CTD (Cell Transfer Delay): is the maximum delay incurred by a cell (including propagation and buffering delays.
- CLR (Cell Loss Ratio): is the percentage of transmitted cells that are lost.

Congestion Control Feedback:

- With ABR, provides a means to control flow based on congestion measurement.

Standard ABR notes:

Standard ABR uses RM (Resource Management) cells to carry feedback information back to the connection's source from the connection's destination.

ABR sources periodically interleave RM cells into the data they are transmitting. These RM cells are called forward RM cells because they travel in the same direction as the data. At the destination these cells are turned around and sent back to the source as Backward RM cells.

The RM cells contain fields to increase or decrease the rate (the CI and NI fields) or set it at a particular value (the explicit rate ER field). The intervening switches may adjust these fields according to network conditions. When the source receives an RM cell it must adjust its rate in response to the setting of these fields.

VSVD Descriptio

ABR sources and destinations are linked via bidirectional connections, and each connection termination point is both a source and a destination; a source for data that it is transmitting, and a destination for data that it is receiving. The forward direction is defined as from source to destination, and the backward direction is defined as from destination to source. shows the data cell flow in the forward direction from a source to its destination along with its associated control loop. The control loop consists of two RM cell flows, one in the forward direction (from source to destination) and the other in the backward direction (from destination to source).

The data cell flow in the backward direction from destination to source is not shown, nor are the associated RM cell flows. However, these flows are just the opposite of that shown in the diagram for forward data cell flows.

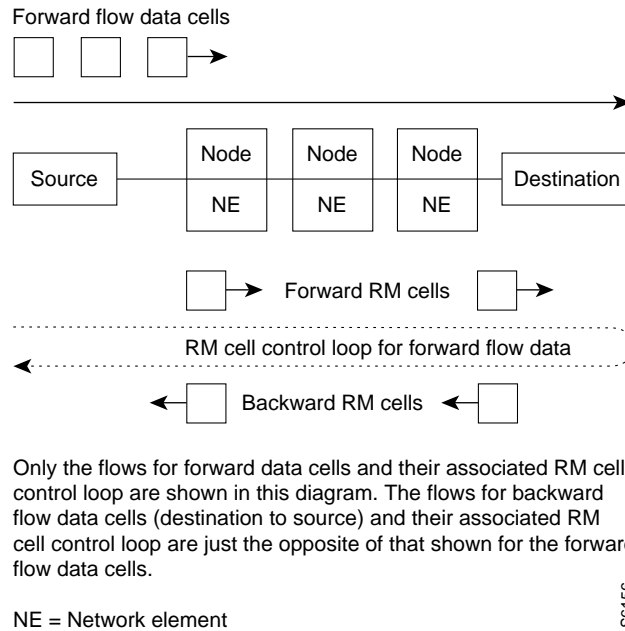
A source generates forward RM cells which are turned around by the destination and returned to the source as backward RM-cells. These backward RM-cells may carry feedback information from the network elements and/or the destination back to the source.

The parameter *Nrm* is defined as the maximum number of cells a source may send for each forward RM cell, i.e., one RM cell must be sent for every *Nrm*-1 data cells. Also, in the absence of *Nrm*-1 data cells, as an upper bound on the time between forward RM cells for an active source, an RM cell must be sent at least once every *Trm* msec.

BXM Connections

The BXM-T3/E3, BXM-155, and BXM-622 cards support ATM Traffic Management 4.0. The BXM cards are designed to support all the following service classes: Constant Bit Rate (CBR), Variable Bit Rate (VBR), Available Bit Rate (ABR with VS/VD, ABR without VS/VD, and ABR using Foresight), and Unspecified Bit Rate (UBR). ABR with VS/VD supports explicit rate marking and Congestion Indication (CI) control.

Figur e10-2 ABR VSVD Flow Control Diagram



Only the flows for forward data cells and their associated RM cell control loop are shown in this diagram. The flows for backward flow data cells (destination to source) and their associated RM cell control loop are just the opposite of that shown for the forward flow data cells.

ForeSight Congestion Control

ForeSight may be used for congestion control across Cisco BPX 8600 series broadband/Cisco IGX 8400 series multiband/Cisco IPX narrowband switches for connections that have one or both end points terminating on ASI-T3/E3 or BXM cards. The ForeSight feature is a proprietary dynamic closed-loop, rate-based, congestion management feature that yields bandwidth savings compared to non-ForeSight equipped trunks when transmitting bursty data across cell-based networks. The BXM cards also support the VSVD congestion control mechanism as specified in the ATM Traffic Management 4.0 standards.

ATM Connection Requirements

There are two connection addressing modes supported. The user may enter a unique VPI/VCI address in which case the Cisco BPX 8600 series broadband switch functions as a virtual circuit switch. Or the user may enter only a VPI address in which case all circuits are switched to the same destination port and the Cisco BPX 8600 series broadband switch functions as a virtual path switch. The full ATM address range for VPI and VCI is supported. Virtual Path Connections are identified by an * in the VCI field. Virtual Circuit Connections specify both the VPI and VCI fields.

The VPI and VCI fields have significance only to the local Cisco BPX 8600 series broadband switch, and are translated by tables in the Cisco BPX 8600 series broadband switch to route the connection. Connections are automatically routed by the AutoRoute feature once the connection endpoints are specified.

ATM connections can be added using either the Cisco WAN Manager Connection Manager or a node's command line interface (CLI). Typically, the Cisco WAN Manager Connection Manager is the preferred method as it has an easy to use GUI interface. The CLI may be the method of choice in some special cases or during initial node setup for local nodes.

When adding ATM connections, first the access port and access service lines connecting to the customer CPE need to be configured. Also, the trunks across the network need to be configured appropriately for the type of connection. Following that the **addcon** command may be used to add a connection, first specifying the service type and then the appropriate parameters for the connection.

For example, when configuring a BXM for CPE connections, the BXM is configured for port mode, a line is upped with the **upln** command and configured with the **cnfln** command. Then the associated port is configured with the **cnfport** command and upped with the **upport** command. Following this, the ATM connections are added via the **addcon** command with the syntax.

Connection Routing

ATM connections for a BXM or ASI card are identified as follows:

- slot number (in the Cisco BPX 8600 series broadband shelf where the BXM or ASI is located)
- port number (one of the ATM ports on the BXM or ASI)
- Virtual Path Identifier (VPI)
- Virtual Circuit Identifier (VCI) – (* for virtual path connections).

The slot and port are related to the Cisco BPX 8600 series broadband switch hardware. Virtual path connections (VPCs) are identified by a “*” for the VCI field. Virtual circuit connections (VCCs) are identified by both a VPI and VCI field.

Connections added to the network are automatically routed once the end points are specified. This AutoRoute feature is standard with all Cisco BPX 8600 series broadband, Cisco IGX 8400 series multiband, and Cisco IPX narrowband nodes. The network automatically detects trunk failures and routes connections around the failures.

Addcon Command Syntax

The following parameters are entered for BXM **addcon** command. Depending upon the connection type, the user is prompted for the appropriate parameters as shown in the following:

```
addcon local_addr node remote_addr traffic_type...extended parameters
```

Field	Value	Descripti
local/remote_addr	slot.port.vpi.vci	desired VCC or VPI connection identifier
node		slave end of connection
traffic_type		Type of traffic, chosen from CBR, VBR, ABR, and UBR
extended parameters		The traffic management and performance parameters associated with an ATM connection.

**Note**

The range of VPIs and VCIs reserved for PVC traffic and SVC traffic is configurable using the **cnfport** command. While adding connections, the system checks the entered VPI/VPC against the range reserved for SVC traffic. If there is a conflict, the **addcon** command fails with the message “VPI/VCI on selected port is reserved at local/remote end”.

ATM Connection Configuration

The following figures and tables describe the parameters used to configure ATM connections:

- Table 10-2, Traffic Policing Definitions.
 - This table describes the policing options that may be selected for ATM connection types: CBR, UBR, and VBR. The policing options for ABR are the same as for VBR.
- Table , Connection Parameter Default Settings and Ranges
 - This table specifies the ATM connection parameter ranges and defaults. Not all the parameters are used for every connection type. When adding connections, you are prompted for the applicable parameters, as specified in the prompt sequence diagrams included in Figure e10-3 through Figure 10-10.
- Table 10-4, Connection Parameter Descriptions
 - This table defines the connection parameters listed in Table e10-2.

The following figures list the connection parameters in the same sequence as they are entered when a connection is added:

- , CBR Connection Prompt Sequence
- Figure 10-4, VBR Connection Prompt Sequence
- , ATFR Connection Prompt Sequence
- , ABR Standard Connection Prompt Sequence

The following figure shows the VSVD network segment and external segment options available when ABR Standard or ABR ForeSight is selected. ForeSight congestion control is useful when both ends of a connection do not terminate on BXM cards. At present, FCES (Flow Control External Segment) as shown in Figure 10-7 is not available for ABR with ForeSight.

- , Meaning of VSVD and Flow Connection External Segments

The following figures list the connection parameters in the same sequence as they are entered when a connection is added:

- , ABR ForeSight Connection Prompt Sequence
- , ATFST Connection Prompt Sequence
- , UBR Connection Prompt Sequence

Table 10-2 Traffic Policing Definitions

ConnectionTy	ATM Forum TM spec. 4.0 conformance definiti	PCR Flo (1st leaky bucket)	CLP tagging (for PCR flow)	SCR Flow (2nd leaky bucket)	CLP tagging (for SCR flow)
CBR	CBR.1 when policing set to 4 (PCR Policing only)	CLP(0+1)	no	off	n/a
CBR	When policing set to 5 (off)	off	n/a	off	n/a
UBR	UBR.1 when CLP setting = no	CLP(0+1)	no	off	n/a
UBR	UBR.2 when CLP setting = yes	CLP(0+1)	no	CLP(0)	yes
VBR, ABR, ATFR, ATFST	VBR.1 when policing set to 1	CLP(0+1)	no	CLP(0+1)	no
VBR, ABR, ATFR, ATFST	VBR.2 when policing set to 2	CLP(0+1)	no	CLP(0)	no
VBR, ABR, ATFR, ATFST	VBR.3 when policing set to 3	CLP(0+1)	no	CLP(0)	yes
VBR, ABR, ATFR, ATFST	when Policing set to 4	CLP(0+1)	no	off	n/a
VBR, ABR, ATFR, ATFST	when Policing set to 5 for off)	off	n/a	off	n/a

Note 1: - For UBR.2, SCR = 0

Note 2:

- CLP = Cell Lost Priority
- CLP(0) means cells that have CLP = 0
- CLP(1) means cells that have CLP = 1
- CLP(0+1) means both types of cells: CLP = 0 & CLP = 1
- CLP(0) has higher priority than CLP(1)
- CLP tagging means to change CLP = 0 to CLP = 1, where CLP= 1 cells have lower priority

Table 10-3 Connection Parameters with Default Settings and Ranges

PARAMETER WITH [DEFAULT SETTING]	BXM T3/E3, OC3 & OC12 RANGE	ASI T3/E3 RANGE	ASI-155 RANGE
PCR(0+1)[50/50]	50- T3/E3 cells/sec 50 - OC3 50 - OC12	T3: MCR – 96000 E3: MCR – 80000 Limited to MCR – 5333 cells/sec for ATFR connections.	OC3 (STM1): 0 – 353200
% Util[100/100] for UBR [1/1]	0 - 100%	1 - 100%	1 - 100%
MCR[50/50]	cells/sec 6 - T3/E3OC3/OC12	T3: 0 – 96000 cells/sec E3: 0 – 80000 cells/sec	N/A
FBTC (AAL5 Frame Base Traffic Control): for VBR [disable] for ABR/UBR [enable] for Path connection [disable]	enable/disable	enable/disable	enable/disable
CDVT(0+1): for CBR [10000/10000], others [250000/250000]	0 - 5,000,000 usec	T3/E3 1 – 250,000 usecs.	OC3/STM1: 0 – 10000 usecs.
VSVD[disable]	enable/disable	enable/disable	Select disable, as only ABR w/o VSVD is supported.
FCES (Flow Control External Segment) [disable]	enable/disable	enable/disable	N/A
Default Extended Parameters[enable]	enable/disable	enable/disable	N/A
CLP Setting[enable]	enable/disable	enable/disable	enable/disable
SCR [50/50]	cells/sec 50 - T3/E3OC3/OC12	T3: MCR – 96000:T3 E3: MCR – 80000: E3 Limited to MCR – 5333 cells/sec for ATFR connections.	OC3/STM1: 0 – 353200
MBS [1000/1000]	1 - 5,000,000cells	T3/E3: 10 – 24000 cells	OC3 (STM1): 10 – 1000 cells

Table 10-3 Connection Parameters with Default Settings and Ranges (continued)

PARAMETER WITH [DEFAULT SETTING]	BXM T3/E3, OC3 & OC12 RANGE	ASI T3/E3 RANGE	ASI-155 RANGE
Policing[3] For CBR: [4]	1 - VBR.1 2 - VBR.2 3 - VBR.3 4 - PCR policing onl 5 - off	1 - VBR.1 2 - VBR.2 3 - VBR.3 4 - PCR policing onl 5 - off	1 - VBR.1 2 - VBR.2 3 - VBR.3 4 - PCR policing onl 5 - off
ICR: max[MCR, PCR/10]	MCR - PCR cells/sec	MCR - PCR cells/sec	N/A
ADTF[1000]	62 - 8000 msec	1000 - 255000 msec.	N/A
Trm[100]	ABRSTD: 1 - 100 msec ABRFST: 3 - 255 msec	20 - 250 msec.	N/A
VC QDepth [16000/16000] For ATFR/ATFST [1366/1366]	0 - 61440 cells	Applies to T3/E3 only ABR: 1 - 64000 cells ATFR: 1 - 1366 cells	ATFR: 1 - 1366 cell
CLP Hi [80/80]	1 - 100%	1 - 100%	N/A
CLP Lo/EPD [35/35]	1 - 100%	1 - 100%	N/A
EFCI [30/30] For ATFR/ATFST [100/100]	1 - 100%	1 - 100%	0 - 100%
RIF: For ForeSight: max[PCR/128, 10]	If ForeSight, then in absolute (0 - PCR)	If ForeSight, then in absolute (0 - PCR)	N/A
For ABR STD[128]	If ABR then 2 ⁿ (1 - 32768)	If ABR, then 2 ⁿ (1 - 32768)	
RDF: For ForeSight [93]	If ForeSight, then % (0% - 100%)	If ForeSight, then % (0% - 100%)	N/A
For ABR STD [16]	If ABR then 2 ⁿ (1 - 32768)	If ABR, then 2 ⁿ (1 - 32768)	
Nrm[32], BXM only	2 - 256 cells	N/A	N/A
FRTT[0], BXM only	0 - 16700 msec	N/A	N/A

Table 10-3 Connection Parameters with Default Settings and Ranges (continued)

PARAMETER WITH [DEFAULT SETTING]	BXM T3/E3, OC3 & OC12 RANGE	ASI T3/E3 RANGE	ASI-155 RANGE
TBE[1,048,320], BXM only	0 - 1,048,320 cells (different max range from TM spec. but limited by firmware for CRM(4095 only) where CRM=TBE/Nrm	N/A	N/A
IBS[1/1]	0 - 24000 cells	T3/E3 ABR: 0 - 24000 cells ATFR: 1 - 107 cells	0 - 999 cells

Table 10-4 Connection Parameter Descriptions

Parameter	Descriptio
PCR	Peak cell rate: The cell rate which the source may never exceed
%Util	% Utilization; bandwidth allocation for: VBR, CBR, UBR it's PCR*%Util, for ABR it's MCR*%Util
MCR	Minimum Cell Rate: A minimum cell rate committed for delivery by network
CDVT	Cell Delay Variation Tolerance: Controls time scale over which the PCR is policed
FBTC (AAL5 Frame Basic Traffic Control)	To enable the possibility of discarding the whole frame, not just one non-compliant cell. This is used to set the Early Packet Discard bit at every node along a connection.
VSVD	Virtual Source Virtual Destination: (see Meaning of VSVD and Flow Control External Segments, Figur e10-7)
FCES (Flow Control External Segments)	(see Meaning of VSVD and Flow Control External Segments, Figur e10-7)
SCR	Sustainable Cell Rate: Long term limit on the rate a connection can sustain
MBS	Maximum Burst Size: Maximum number of cells which may burst at the PCR but still be compliant. Used to determine the Burst Tolerance (BT) which controls the time scale over which the SCR is police
Policing	(see definitions of Traffic Policing, Table 10-2)
VC QDepth	VC Queue Depth
CLP Hi	Cell Loss Priority Hi threshold (% of VC QMax)


Table 10-4 Connection Parameter Descriptions (continued)

Parameter	Descriptio
CLP Lo/EPD	Cell Loss Priority Low threshold (% of VC QMax)/Early Packet Discard. If AAL5 FBTC = yes, then for the BXM card this is the EPD threshold setting. For ASI cards, regardless of the FBTC setting, this is the CLP Lo setting.
EFCI	Explicit Forward Congestion Indication threshold (% of VC QMax)
ICR	Initial Cell Rate: The rate at which a source should send initially and after an idle period
ADTF (ATM Forum TM 4.0 term)	The Allowed-Cell-Rate Decrease Factor: Time permitted between sending RM-cells before the rate is decreased to ICR
Trm (ATM Forum TM 4.0 term)	An upper bound on the time between forward RM-cells for an active source, i.e., RM cell must be sent at least every Trm msec
RIF (ATM Forum TM 4.0 term)	Rate Increase Factor: Controls the amount by which the cell transmission rate may increase upon receipt of an RM cell
RDF (ATM Forum TM 4.0 term)	Rate Decrease Factor: Controls the amount by which the cell transmission rate may decrease upon receipt of an RM cell
Nrm (ATM Forum TM 4.0 term), BXM only.	Nrm Maximum number of cells a source may send for each forward RM cell, i.e. an RM cell must be sent for every Nrm-1 data cells
FRTT (ATM Forum TM 4.0 term), BXM only.	Fixed Round Trip Time: the sum of the fixed and propagation delays from the source to a destination and back
TBE (ATM Forum TM 4.0 term), BXM only.	Transient Buffer Exposure: The negotiated number of cells that the network would like to limit the source to sending during start-up periods, before the first RM-cell returns.
IBS	Initial Burst Size

CBR Connections

The **CBR** (constant bit rate) category is a fixed bandwidth class. CBR traffic is more time dependent, less tolerant of delay, and generally more deterministic in bandwidth requirements. CBR is used by connections that require a specific amount of bandwidth to be available continuously throughout the duration of a connection. Voice, circuit emulation, and high-resolution video are typical examples of traffic utilizing this type of connection. A CBR connection is allowed to transmit cells at the peak rate, below the peak rate, or not at all. CBR is characterized by peak cell rate (PCR).

The parameters for a CBR connection are shown in in the sequence in which they occur during the execution of the **addcon** command. The CBR policing definitions are summarized in Table 10-5.

Figur e10-3 CBR Connection Prompt Sequence

- ① For policing prompt:
 4 = PCR policing only
 5 = policing off

Note: BW allocation = (PCR)x(%Util) S6159

Table 10-5 CBR Policing Definitions

ConnectionTy	ATM Forum TM spec. 4.0 conformance definiti	PCR Flo (1st leaky bucket)	CLP tagging (for PC flow)	SCR Flo (2nd leak bucket)	CLP tagging (for SCR flow)
CBR	CBR.1 when policing set to 4 (PCR Policing only)	CLP(0+1)	no	off	n/a
CBR	When policing set to 5 (off)	off	n/a	off	n/a

VBR and ATFR Connections

VBR Connection

VBR (variable bit rate) connections may be classified as rt-VBR or nrt-VBR connections.


The rt-VBR (real-time variable bit rate) category is used for connections that transmit at a rate varying with time and that can be described as bursty, often requiring large amounts of bandwidth when active. The rt-VBR class is intended for applications that require tightly constrained delay and delay variation. An example of rt-VBR is video conferencing which requires real-time data transfer with bandwidth requirements that can vary in proportion to the dynamics of the video image at any given time. The rt-VBR category is characterized in terms of PCR, SCR (sustained cell rate), and MBS (maximum burst size).

The nrt-VBR (non-real time variable bit rate) category is used for connections that are bursty but are not constrained by delay and delay variation boundaries. For those cells in compliance with the traffic contract, a low cell loss is expected. Non-time critical data file transfers are an example of an nrt-VBR connection. A nrt-VBR connection is characterized by PCR, SCR, and MBS.

Configuring VBR connections. The characteristics of rt-VBR or nrt-VBR are supported by appropriately configuring the parameters of the VBR connection.

The parameters for a VBR connection are shown in Figure 10-4 in the sequence in which they occur during the execution of the **addcon** command. The VBR policing definitions are summarized in Table 10-6.

Figur e10-4 VBR Connection Prompt Sequence



- ① For policing prompt:
- 1 = VBR.1
 - 2 = VBR.2
 - 3 = VBR.3
 - 4 = PCR policing only
 - 5 = policing off

Note: BW allocation = (PCR)x(%Util)

S6160

Table 10-6 VBR Policing Definitions


ConnectionTy	ATM Forum TM spec. 4.0 conformance definiti	PCR Flo (1st leaky bucket)	CLP tagging (for PCR flow)	SCR Flow (2nd leaky bucket)	CLP tagging (for SCR flow)
VBR, ABR, ATFR, ATFST	VBR.1 when policing set to 1	CLP(0+1)	no	CLP(0+1)	no
VBR, ABR, ATFR, ATFST	VBR.2 when policing set to 2	CLP(0+1)	no	CLP(0)	no
VBR, ABR, ATFR, ATFST	VBR.3 when policing set to 3	CLP(0+1)	no	CLP(0)	yes
VBR, ABR, ATFR, ATFST	when Policing set to 4	CLP(0+1)	no	off	n/a
VBR, ABR, ATFR, ATFST	when Policing set to 5 for off)	off	n/a	off	n/a

ATFR Connection

An **ATFR** (ATM to Frame Relay) connection is a frame relay to ATM connection and is configured as a VBR connection, with a number of the ATM and frame relay connection parameters being mapped between each side of the connection.

The parameters for an ATFR connection are shown in in the sequence in which they occur during the execution of the **addcon** command.

Figur e10-5 ATFR Connection Prompt Sequence



① For policing prompt:

- 1 = VBR.1
- 2 = VBR.2
- 3 = VBR.3
- 4 = PCR policing only
- 5 = policing off

② VC QDepth maps to VC Queue Max for frame relay

- EFCI maps to ECN for frame relay
- IBS maps to Cmax for frame relay

Note: FBTC (Frame based traffic control - AAL5, same as FGCRA) is automatically set to yes.

36161

ABR Notes

The term ABR is used to specify one of the following:

- ABR standard without VSVD (This is ABR standard without congestion flow control.)
 - Supported by BXM, ASI-T3 (& ASI-E3), and ASI OC3 cards.
- ABR standard with VSVD. (This is ABR standard with congestion flow control as specified by the ATM Traffic Management, Version 4.0)
 - Also, referred to as ABR.1
 - Supported only by BXM cards
 - Feature must be ordered.
- ABR with ForeSight congestion control
 - Also, referred to as ABR.FST

- Supported by BXM and ASI-T3 (& ASI-E3) cards
- Feature must be ordered.

ABR and ATFST Connections

ABR Connection

The **ABR** (available bit rate) category utilizes a congestion flow control mechanism to control congestion during busy periods and to take advantage of available bandwidth during less busy periods. The congestion flow control mechanism provides feedback to control the connections flow rate through the network in response to network bandwidth availability. The ABR service is not restricted by bounding delay or delay variation and is not intended to support real-time connections. ABR is characterized by: PCR and MCR.

Policing for ABR connections is the same as for VBR connections which are summarized in Table 10-6.

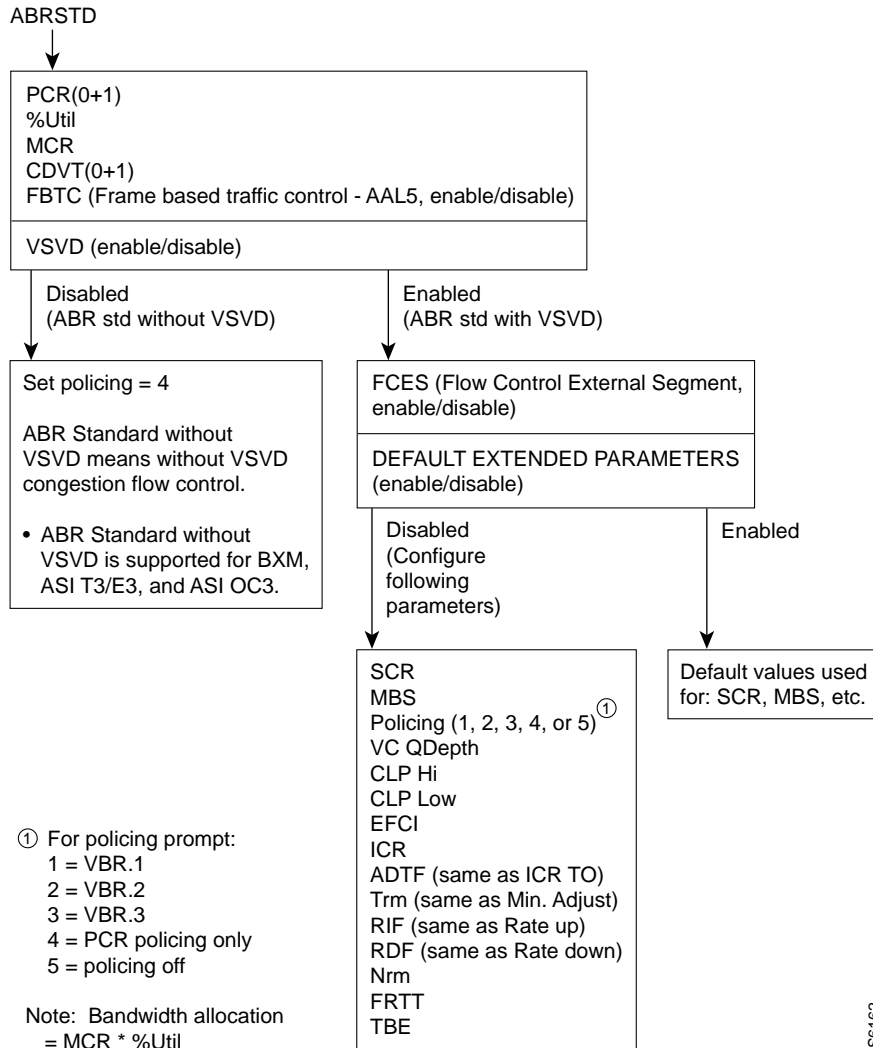
The ABR connections are configured as either ABR Standard (**ABRSTD**) connections or as ABR ForeSight (**ABRFST**) connections.

The parameters for an ABRSTD connection are shown in in the sequence in which they occur during the execution of the **addcon** command.

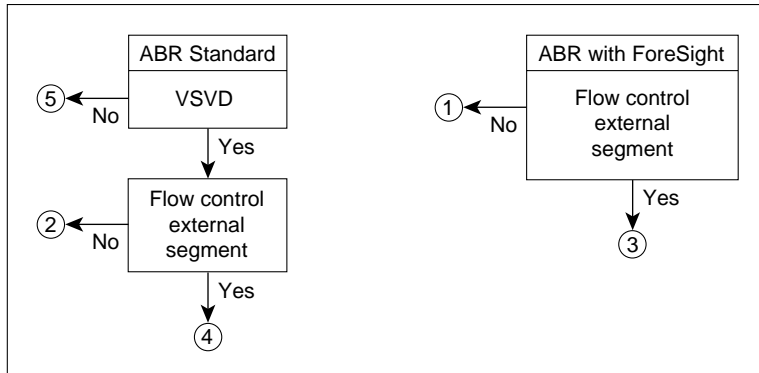
The ABRSTD connection supports all the features of ATM Standards Traffic Management 4.0 including VSVD congestion flow control.

VSVD and flow control with external segments are shown in .

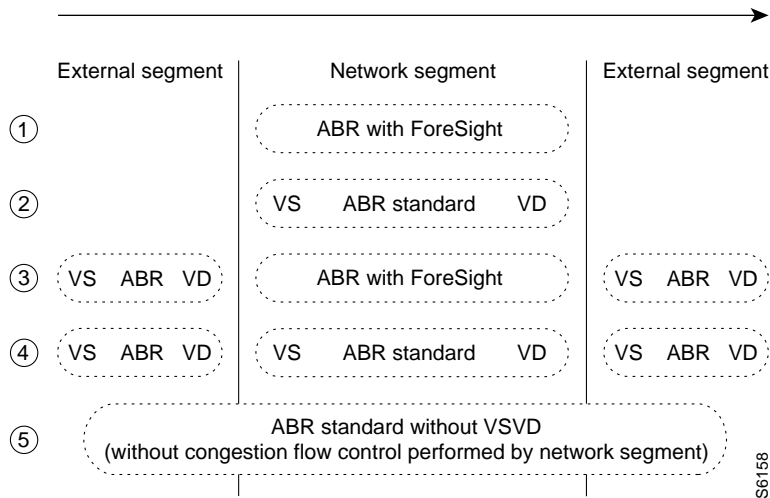
Figur e10-6 ABR Standard Connection Prompt Sequence



Figur e10-7 Meaning of VSVD and Flow Control External Segments



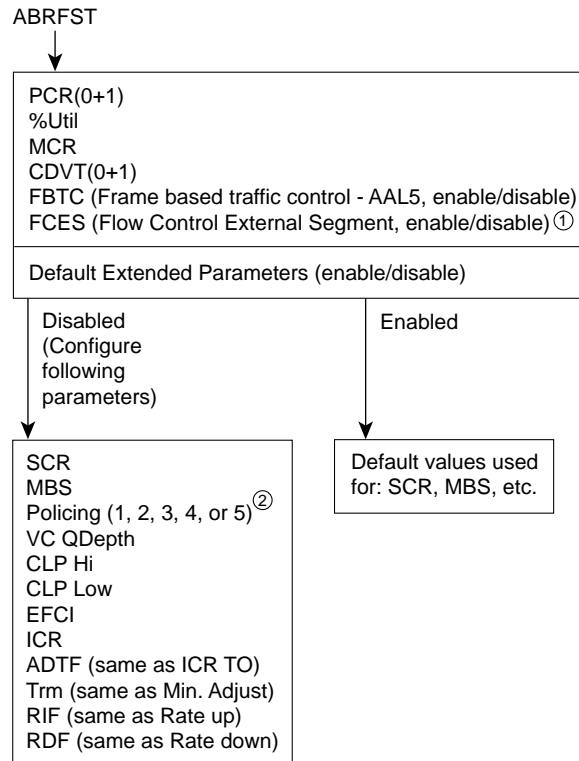
VS and VD shown below are for traffic flowing in direction of arrow. For the other direction of traffic, VS and VD are in the opposite direction.



ATFST Connection

The **ABRFST** connection uses the propriety ForeSight congestion control and is useful when configuring connections on which both ends do not terminate on BXM cards.

The parameters for an ABRFST connection are shown in in the sequence in which they occur during the execution of the **addcon** command.

Figur e10-8 ABR ForeSight Connection Prompt Sequence

① At present, FCES is not available for ABR with ForeSight

② For policing prompt:

- 1 = VBR.1
- 2 = VBR.2
- 3 = VBR.3
- 4 = PCR policing only
- 5 = policing off

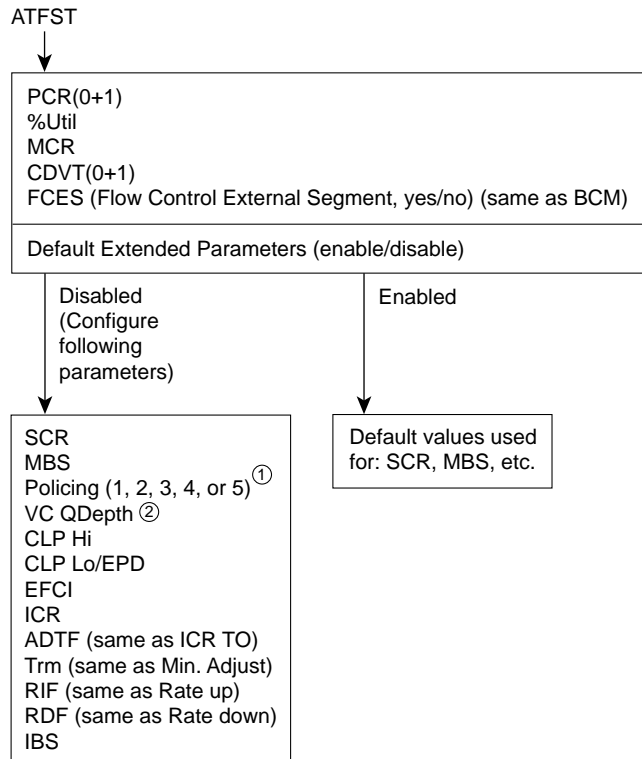
Note: Bandwidth allocation
= (MCR)x(%Util)

56163

An **ATFST** connection is a frame relay to ATM connection that is configured as an ABR connection with ForeSight. ForeSight congestion control is automatically enabled when connection type ATFST is selected. A number of the ATM and frame relay connection parameters are mapped between each side of the connection.

The parameters for an ATFST connection are shown in in the sequence in which they occur during the execution of the **addcon** command.

Figur e10-9 ATFST Connection Prompt Sequence



① For policing prompt:

- 1 = VBR.1
- 2 = VBR.2
- 3 = VBR.3
- 4 = PCR policing only
- 5 = policing off

② VC QDepth maps to VC Queue max for frame relay.

- EFCI maps to ECN for frame relay.
- IBS maps to C max for frame relay.

Note: FBTC (Frame based traffic control - AAL5, same as FGCR) is automatically set to yes.

86164

UBR Connections

The unspecified bit rate (**UBR**) connection service is similar to the ABR connection service for bursty data. However, UBR traffic is delivered only when there is spare bandwidth in the network. This is enforced by setting the CLP bit on UBR traffic when it enters a port.

Therefore, traffic is served out to the network only when no other traffic is waiting to be served first. The UBR traffic does not affect the trunk loading calculations performed by the switch software.

The parameters for a UBR connection are shown in the sequence in which they occur during the execution of the **addcon** command.

The UBR policing definitions are summarized in Table 10-7.

Figur e10-10 UBR Connection Prompt Sequence

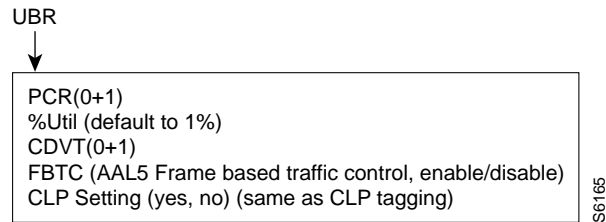


Table 10-7 jUBR Policing Definitions

ConnectionTy	ATM Forum TM spec. 4.0 conformance definiti	PCR Flo (1st leaky bucket)	CLP tagging (for PCR flow)	SCR Flow (2nd leaky bucket)	CLP tagging (for SCR flow)
UBR	UBR.1 when CLP setting = no	CLP(0+1)	no	off	n/a
UBR	UBR.2 when CLP setting = yes	CLP(0+1)	no	CLP(0)	yes

Traffic Policing Examples

Traffic Policing, also known as Usage Parameter Control (UPC), is implemented using either an ATM Forum single or dual-leaky bucket algorithm. The buckets represent a GCRA (Generic Cell Rate Algorithm) defined by two parameters:

- Rate (where I, expected arrival interval is defined as 1/Rate)
- Deviation (L)

If the cells are clumped too closely together, they are non-compliant and are tagged or discarded as applicable. If other cells arrive on time or after their expected arrival time, they are compliant, but there is no accrued credit.

Dual-Leaky Bucket (An Analogy)

A GCRA viewpoint is as follows:

- For a stream of cells in an ATM connection, the cell compliance is based on the theoretical arrival time (TAT).
- The next TAT should be the time of arrival of the last compliant cell plus the expected arrival interval (I) where $I = 1/\text{rate}$.
- If the next cell arrives before the new TAT, it must arrive no earlier than new TAT - CDVT to be compliant.
- If the next cell arrives after the new TAT, it is compliant, but there is no accrued credit.

CBR Traffic Policing Examples

CBR traffic is expected to be at a constant bit rate, have low jitter, and is configured for a constant rate equal to Peak Cell Rate (PCR). The connection is expected to be always at peak rate.

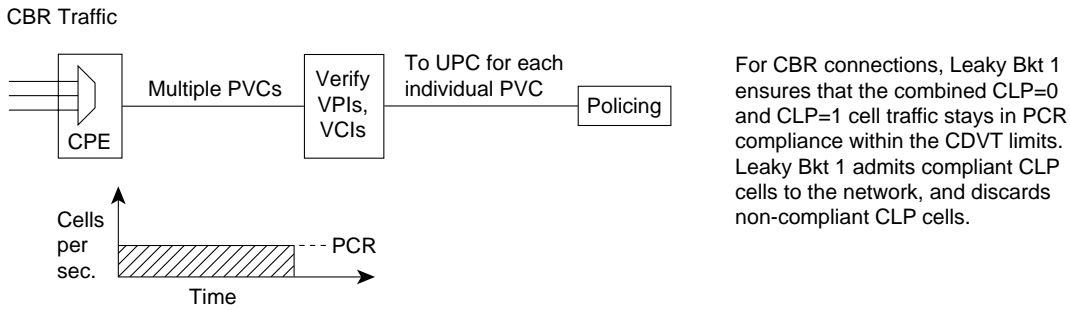
When a connection is added, a VPI.VCI address is assigned, and the UPC parameters are configured for the connection. For each cell in an ATM stream seeking admission to the network, the VPI.VCI addresses are verified and each cell is checked for compliance with the UPC parameters. The CBR cells are not enqueued, but are processed by the policing function and then sent to the network unless discarded.

For CBR, traffic policing is based on:

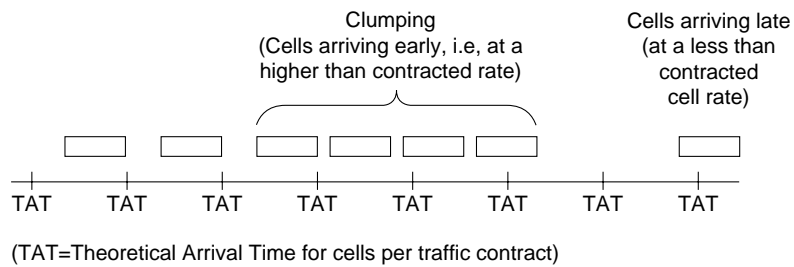
- Bucket 1
 - PCR(0+1), Peak Cell Rate
 - CDVT(0+1), Cell Delay Variation

The CBR connection may be configured with policing selected as either 4 or 5. With policing set to 5, there is no policing. With policing set to 4, there is single leaky bucket PCR policing as shown in Figure 10-11. The single leaky bucket polices the PCR compliance of all cells seeking admission to the network, both those with CLP = 0 and those with CLP = 1. Cells seeking admission to the network with CLP set equal to 1 may have either encountered congestion along the user's network or may have lower importance to the user and have been designated as eligible for discard in the case congestion is encountered. If the bucket depth CDVT (0+1) limit is exceeded, it discards all cells seeking admission. It does not tag cells. If leaky bucket 1 is not full, all cells (CLP = 0 and CLP = 1) are admitted to the network.

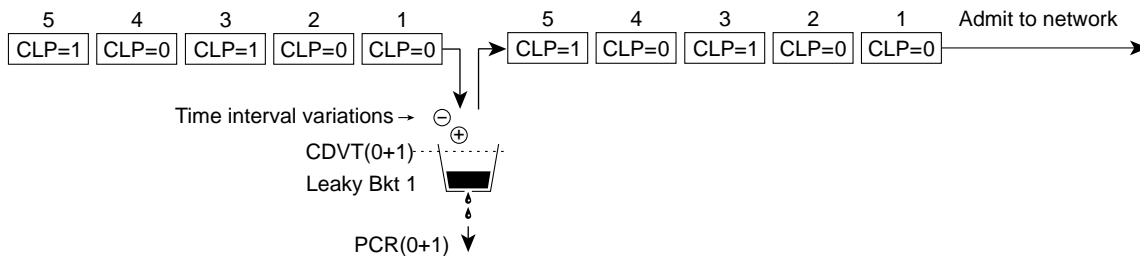
Figur e10-11 CBR Connection, UPC Overview



Policing: 4 = PCR Policing only
5 = off



Example: Policing = 4



Discards incoming CLP(0+1) cells if Bkt 1 depth > CDVT(0+1). Does not tag cells. If Bkt 1 depth < CDVT(0+1), passes CLP=0 and CLP=1 cells on to network.

Note: The notation 0, 1, and 0+1 refers to the types of cell being specified: cells with CLP set to 0, CLP set to 1, or both types of cells, respectively. For example, CLP(0), CLP(1), and CLP(0+1).


S6341

Figure 10-12 shows a CBR.1 connection policing example, with policing set to 4, where the CDVT depth of the single leaky bucket is not exceeded, and all cells, CLP(0) and CLP(1) are admitted to the network.

Figur e10-12 CBR.1 Connection with Bucket Compliant

Connection setup
and compliance status:

CBR.1
policing=4
Bkt 1 depth < CDVT (0+1)



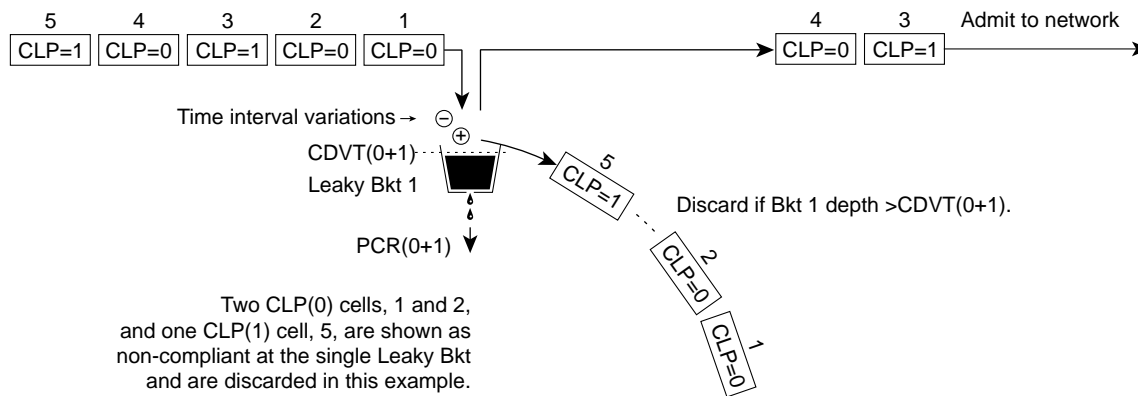
S6342

Figure 10-13 shows a CBR connection policing example, with policing = 4, where the CDVT(0+1) of the single leaky bucket is exceeded and non-compliant cells are discarded. The leaky bucket only discards cells; it does not tag them.

Figur e10-13 CBR.1 Connection, with Bucket Discarding non-Compliant Cells

Connection setup
and compliance status:

CBR.1
policing=4
Bkt 1 depth > CDVT (0+1)



S6343

VBR Dual-Leaky Bucket Policing Examples

The contract for a variable bit rate connection is set up based on an agreed upon sustained cell rate (SCR) with allowance for occasional data bursts at a Peak Cell Rate (PCR) as specified by maximum burst size MBS.

When a connection is added, a VPI.VCI address is assigned, and UPC parameters are configured for the connection. For each cell in an ATM stream, the VPI.VCI addresses are verified and each cell is checked for compliance with the UPC parameters as shown in Figur e10-14.

The VBR cells are not enqueued, but are processed by the policing function and then sent to the network unless discarded.


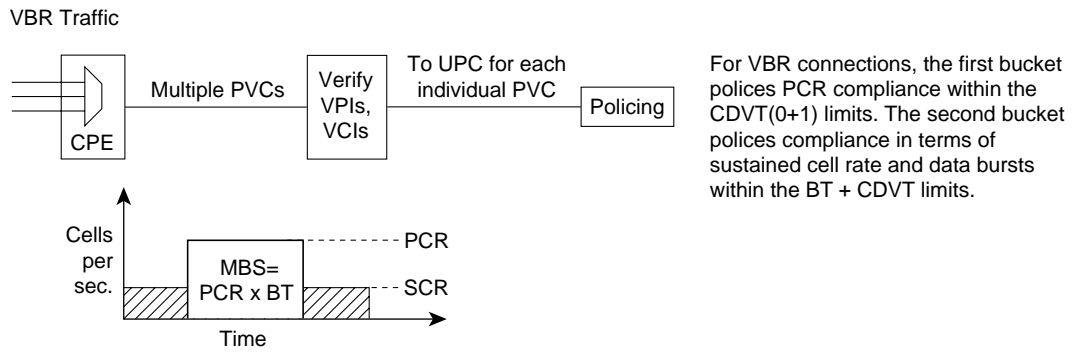
For VBR, traffic policing, depending on selected policing option, is based on:

- Leaky bucket 1, PCR and CDVT
- Leaky bucket 2, SCR, CDVT, and MBS

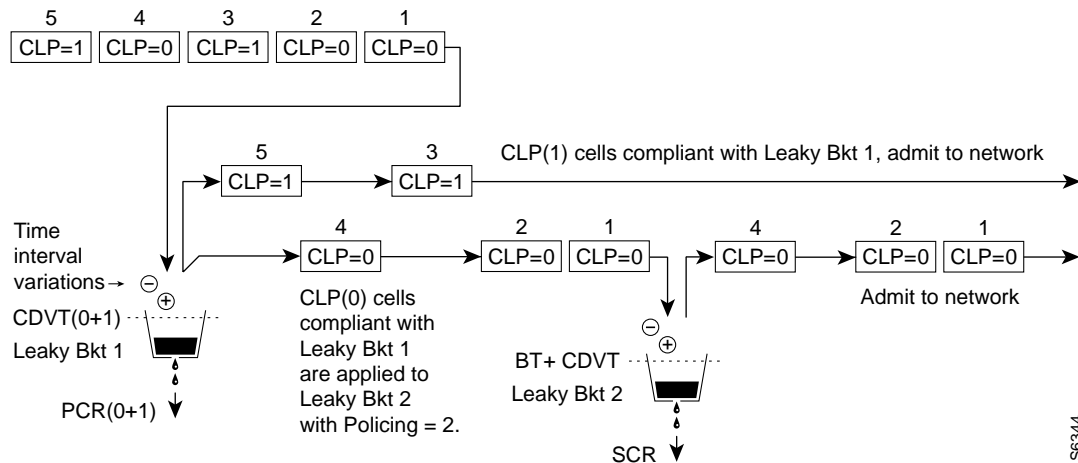
The policing options, selected by entering 1-5 in response to the policing choice prompt, are as follows for VBR connections:

VBR.1 VBR with policing set to 1.	CLP(0+1) cells compliant with leaky bucket 1 are passed to leaky bucket 2; non-compliant cells are discarded. CLP(0+1) cells compliant with leaky bucket 2 are admitted to the network; non-compliant cells are discarded.
VBR.2 VBR with policing set to 2.	CLP(1) cells compliant with leaky bucket 1 are admitted to the network; non-compliant CLP(0+1) cells are dropped. CLP(0) cells compliant with leaky bucket 1 are applied to leaky bucket 2; non-compliant cells are dropped. CLP(0) cells compliant with leaky bucket 2 are admitted to the network; non-compliant cells are dropped.
VBR.3 VBR with policing set to 3.	CLP(1) cells compliant with leaky bucket 1 are admitted to the network; non-compliant CLP(0+1) cells are dropped. CLP(0) cells compliant with leaky bucket 1 are applied to leaky bucket 2; non-compliant cells are dropped. CLP(0) cells compliant with leaky bucket 2 are admitted to the network; non-compliant cells are tagged and admitted to the network.
VBR with policing set to 4.	CLP(0+1) cells compliant with leaky bucket 1 are admitted to the network; non-compliant cells are dropped. Leaky bucket 2 is not active.
VBR with policing set to 5.	Policing is off, so there is no policing of cells on ingress.

Figur e10-14 VBR Connection, UPC Overview



Example: VBR.2
Policing = 2



S6344

Leaky Bucket 1

Leaky bucket 1 polices for the PCR compliance of all cells seeking admission to the network, both those with CLP = 0 and those with CLP = 1. For example, cells seeking admission to the network with CLP set equal to 1 may have either encountered congestion along the user's network or may have lower

importance to the user and have been designated as eligible for discard in the case congestion is encountered. If the bucket depth in the first bucket exceeds CDVT (0+1), it discards all cells seeking admission. It does not tag cells.

With policing set to 1 (VBR.1), all cells (CLP=0 and CLP=1) that are compliant with leaky bucket 1, are sent to leaky bucket 2. With policing set to 2 (VBR.2) or to 3 (VBR.3), all CLP=1 cells compliant with leaky bucket 1 are admitted directly to the network, and all CLP=0 cells compliant with leak bucket 1 are sent to leaky bucket 2.

Leaky Bucket 2

For VBR connections, the purpose of leaky bucket 2 is to police the cells passed from leaky bucket 1 for conformance with maximum burst size MBS as specified by BT and for compliance with the SCR sustained cell rate. The types of cells passed to leaky bucket 2 depend on how policing is set:

- For policing set to 5, cells bypass both buckets.
- For policing set to 4, leaky bucket 2 sees no traffic.
- For policing set to 2 or 3, the CLP(0) cells are admitted to the network if compliant with BT + CDVT of leaky bucket 2. If not compliant, cells may either be tagged (policing set to 3) or discarded (policing set to 2).
- For policing set to 1, the CLP(0) and CLP(1) cells are admitted to the network if compliant with BT + CDVT of leaky bucket 2. If not compliant, the cells are discarded. There is no tagging option.

Examples

Figure 10-15 shows a VBR connection policing example, with policing set to 4, leaky bucket compliant, and all cells being admitted to the network

Figur e10-15 VBR Connection, Policing = 4, Leaky Bucket 1 Compliant

Connection setup
and compliance status:

VBR
Policing = 4
Bkt 1 depth < CDVT(0+1)

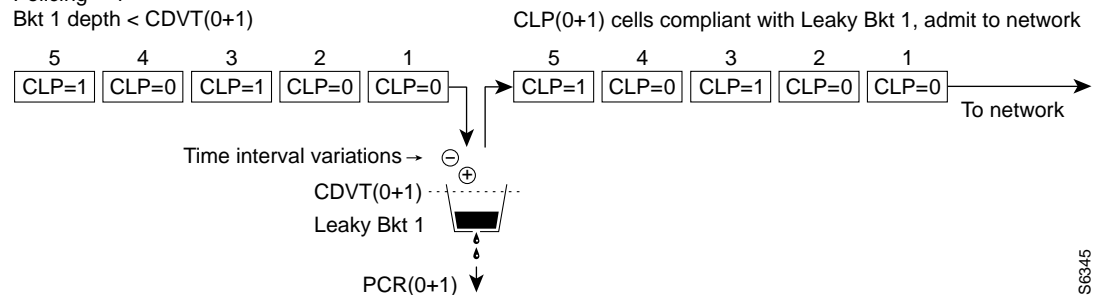


Figure 10-16 shows a VBR connection policing example, with the policing set to 4, and leaky bucket 1 non-compliant which indicates that the connection has exceeded the PCR for a long enough interval to exceed the CDVT (0+1) limit. Non-compliant cells with respect to leaky bucket 1 are discarded

Figur e10-16 VBR Connection, Policing = 4, Leaky Bucket 1 Non-Compliant

Connection setup and compliance status:

VBR
Policing = 4
Bkt 1 depth > CDVT(0+1)

CLP(0+1) cells compliant with Leaky Bkt 1, admit to network




Figure 10-17 shows a VBR.2 connection policing example, with policing = 2, and both buckets compliant. Leaky bucket two is policing the CLP(0) cell stream for conformance with maximum burst size MBS (as specified by BT), and for compliance with the SCR sustained cell rate.

Figur e10-17 VBR.2 Connection, Policing = 2, with Buckets 1 and 2 Compliant

Connection setup and compliance status:

VBR.2
Policing = 2
Bkt 1 depth < CDVT(0+1)
Bkt 2 depth < BT + CDVT

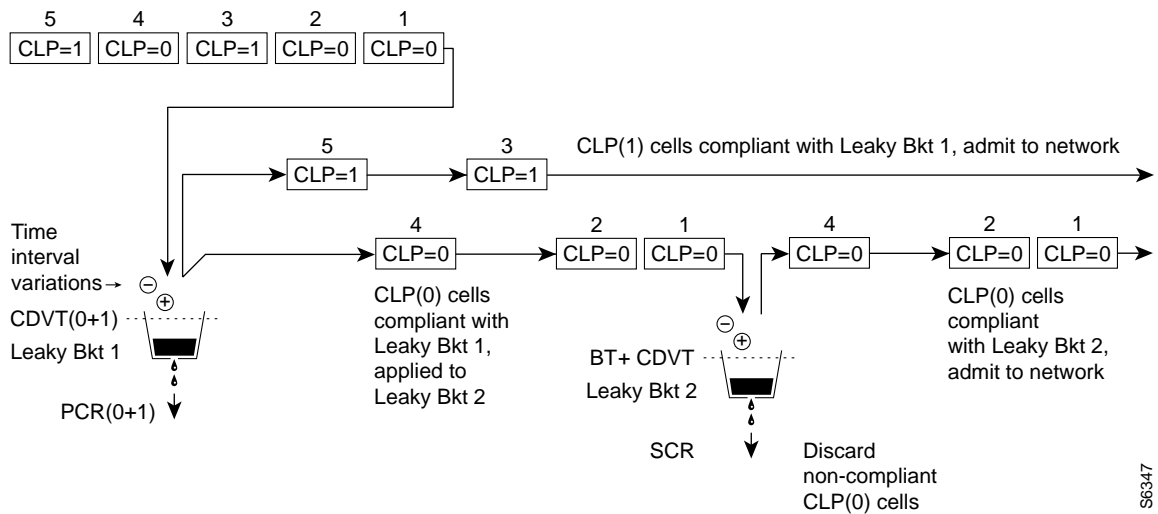
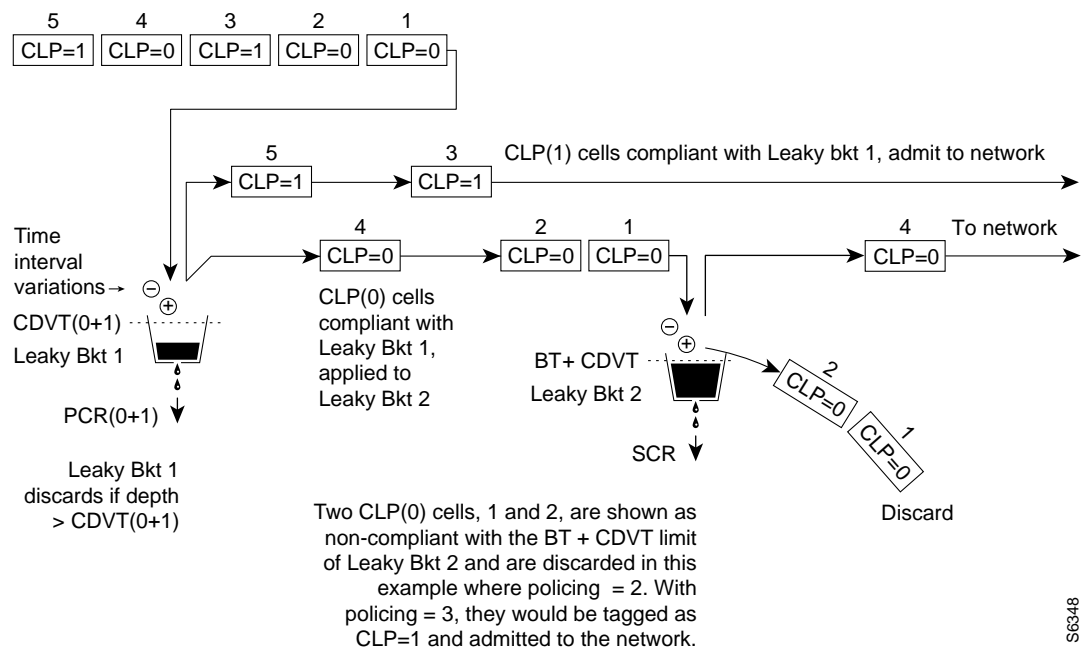


Figure 10-18 shows a VBR.2 connection policing example, with policing set to 2, and leaky bucket 2 non-compliant. Leaky bucket 2 is shown policing the CLP(0) cell stream for conformance with maximum burst size MBS (as specified by BT), and for compliance with SCR (sustained cell rate). In this example (policing set to 2), CLP tagging is not enabled, so the cells that have exceeded the BT + CDVT limit are discarded. In the example, either the sustained cell rate could have been exceeded for an excessive interval, or a data burst could have exceeded the maximum allowed burst size.

Fig e10-18 VBR.2 Connection, Leaky Bucket 2 Discarding CLP (0) Cells

Connection setup and compliance status:

VBR.2
 Policing = 2
 Bkt 1 depth < CDVT(0+1)
 Bkt 2 depth > BT + CDVT



81938

Figure 10-19 shows a VBR.1 connection policing example, with policing set to 1, and both bucket compliant. Leaky bucket 1 is policing the CLP (0+1) cell stream for conformance with the PCR limit. Leaky bucket 2 is policing the CLP (0+1) cell stream for conformance CDVT plus maximum burst size MBS (as specified by BT), and for compliance with SCR sustained cell rate.

Figur e10-19 VBR.1 Connection, Policing = 1, with Buckets 1 and 2 Compliant

Connection setup
and compliance status:

VBR.1
Policing = 1
Bkt 1 depth < CDVT(0+1)
Bkt 2 depth < BT + CDVT

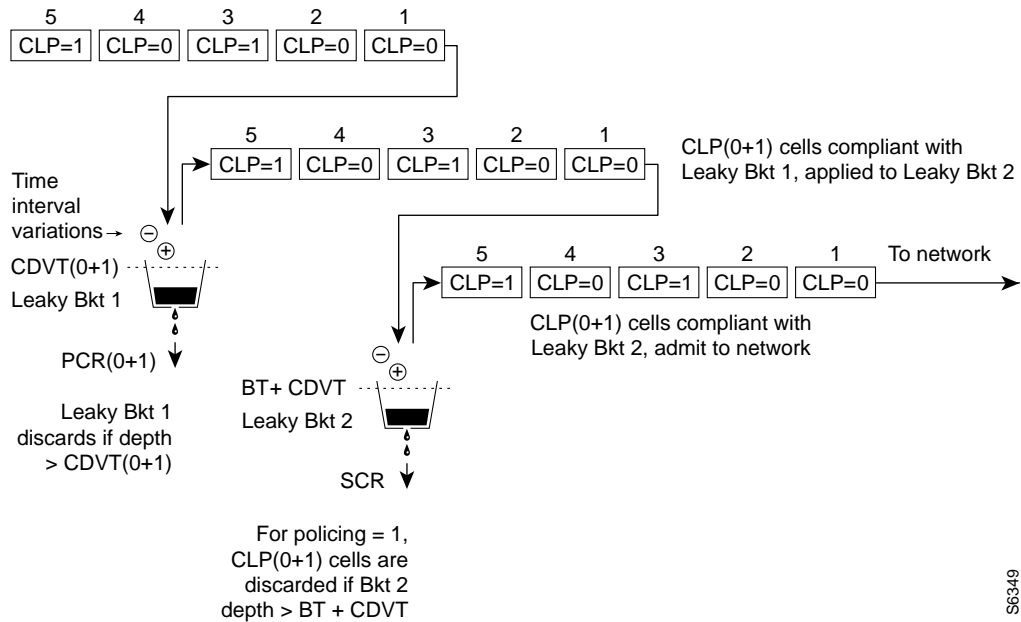
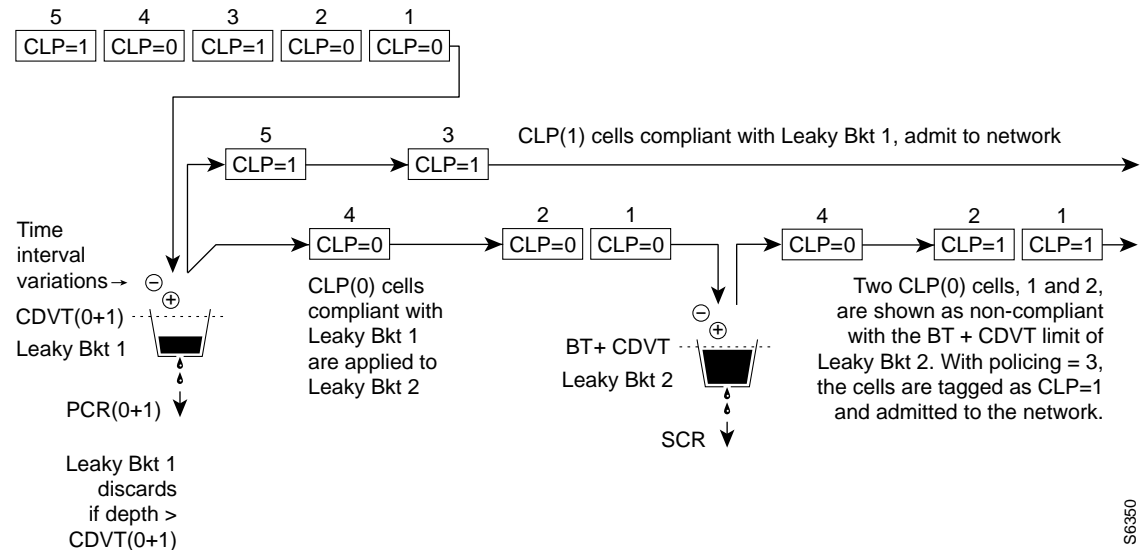


Figure 10-20 shows a VBR.3 connection policing example, with policing set to 3, and Leaky bucket shown as non-compliant. Leaky bucket 2 is shown policing the CLP(0) cell stream for conformance with maximum burst size MBS (as specified by BT), and for compliance with SCR sustained cell rate. For the policing = 3 selection, CLP tagging is enabled, so the cells that have exceeded the BT + CDVT(0+1) limit are tagged as CLP=1 cells and admitted to the network. In this example, either the sustained cell rate could have been exceeded for an excessive interval, or a data burst could have exceeded the maximum burst size allowed.

Figur e10-20 VBR.3 Connection, Policing = 3, with Bucket 2 non-compliant

Connection setup
and compliance status:

VBR.3
Policing = 3
Bkt 1 depth < CDVT(0+1)
Bkt 2 depth > BT + CDVT



S6350

ABR Connection Policing

Available Bit Rate (ABR) connections are policed the same as the VBR connections, but in addition use either the ABR Standard with VSVD congestion flow control method or the ForeSight option to take advantage of unused bandwidth when it is available.

UBR Connection Policing

The contract for a unspecified bit rate connection is similar to the ABR connection service for bursty data. However, UBR traffic is delivered only when there is spare bandwidth in the network.

When a connection is added, a VPI.VCI address is assigned, and UPC parameters are configured for the connection. For each cell in an ATM stream, the VPI.VCI addresses are verified and each cell is checked for compliance with the UPC parameters as shown in .

Leaky Bucket 1


Leaky bucket 1 polices the UBR connection for PCR compliance. When CLP=No (UBR.1), all cells that are compliant with leaky bucket 1 are applied to the network. However, these cells are treated with low priority in the network with % utilization default of 1%.

Leaky Bucket 2



When CLP=Yes (UBR.2), CLP(0) cells that are compliant with leaky bucket 1 are sent to leaky bucket 2. Since SCR=0 for leaky bucket 2, the bucket is essentially always full, and all the CLP(0) cells sent to leaky bucket 2 are therefore tagged with CLP being set to 1. This allows the network to recognize these UBR cells as lower priority cells and available for discard in the event of network congestion.

Figur e10-21 UBR Connection, UPC Overview

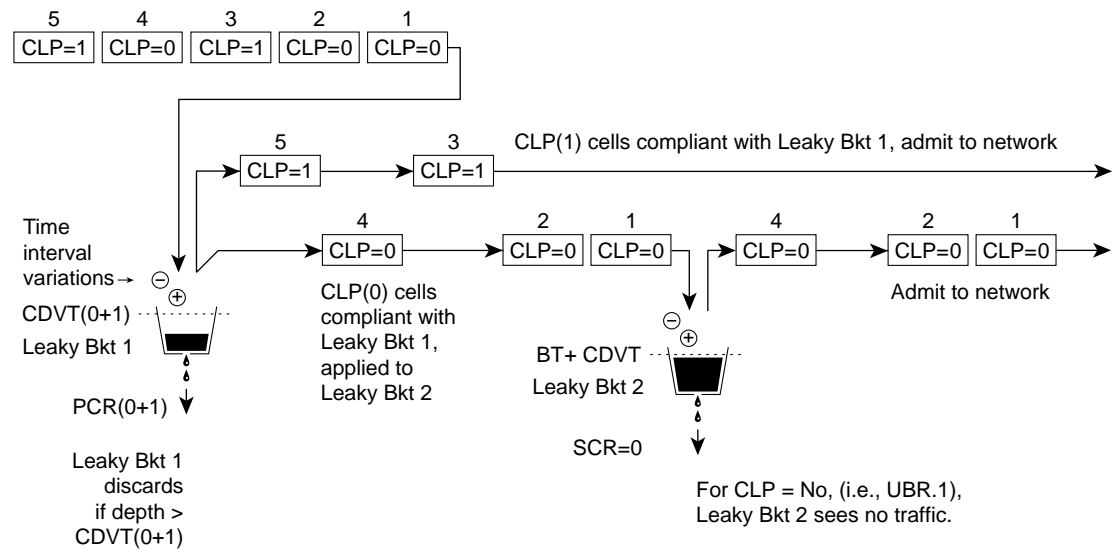
UBR Traffic



For UBR connections, the first bucket polices PCR compliance within the CDVT(0+1) limits. The second bucket, used when CLP is set to Yes, tags all CLP(0) cells.



CLP(0+1) cells to Leaky Bkt 1



Note: The notation 0, 1, and 0+1 refers to the types of cell being specified: cells with CLP set to 0, CLP set to 1, or both types of cells, respectively. For example, CLP(0), CLP(1), and CLP(0+1)

For CLP = No, (i.e., UBR.1), Leaky Bkt 2 sees no traffic.

For CLP = Yes, (i.e., UBR.2), CLP(0) cells that were compliant with Leaky Bkt 1 are sent to Leaky Bkt 2. Since SCR = 0 for Leaky Bkt 2, the bucket is essentially always full, and all cells are therefore tagged with CLP being set to 1. This allows the network to recognize these UBR cells as lower priority and available for discard in the event of network congestion.

S6351

LMI and ILMI Parameters

The following is a listing of the LMI and ILMI parameters for the ASI and BXM:

For ILMI information, refer to Table 10-8:

Table 10-8 LMI Parameters

Parameter	Description
VPI.VCI	VCCI for ILMI signaling channel equal 0.16
Polling Enabled	Keepalive polling
Trap Enabled	VCC change of state traps
Polling Interval	Time between GetRequest polls
Error Threshold	Number of failed entries before ILMI link failure is declared.
Event Threshold	Number of successful polls before ILMI link failure is cancelled.
Addr Reg Enab	SVC Address Registration procedures enabled.

For the LMI information, refer to Table 10-9.

Table 10-9 LMI Parameters

Parameter	Description
VPI.VCI	VCCI for LMI signaling channel equal 0.31
Polling Enable	Keepalive polling
T393	Status Enquiry timeout value
T394	Update Status timeout value
T396	Status Enquiry polling timer
N394	Status Enquiry retry count
N395	Update Status retry count

SVCs, ATM and Frame Relay

This chapter provides a summary of switched virtual circuits with respect to the BPX Service Node and collocated Extended Services Processor. For additional information, refer to the *BPX Service Node Extended Services Processor Installation and Operations* document.

This chapter contains the following:

- ATM and Frame Relay SVCs
- BPX Service Node Interfaces
- Signaling Plane
- Network Interworking Between Frame Relay and ATM
- Extended Services Processor
- Network Management
- Resource Partitioning

ATM and Frame Relay SVCs

With a collocated Extended Services Processor (ESP), the BPX Service Node adds the capability to support ATM and Frame Relay Switched Virtual Circuits (SVCs) in addition to support for Permanent Virtual Circuits (PVCs) as shown in Figure 11-1.

The Private Network to Network Interface (PNNI) protocol is used to route SVCs across the network. PNNI provides a dynamic routing protocol which is responsive to changes in network availability and will scale to large networks.

BPX Service Node resources, such as port VPI range and trunk bandwidth are partitioned between SVCs and PVCs. This provides a firewall between the two types of connections so that any SVCs that come on-line and off-line do not affect the availability of existing PVC services.

The following SVC connections are supported with Release 8.4:

- Frame Relay SVC connections between Frame Relay end users over an ATM network (Network Interworking)
- ATM SVC connections between ATM end users

The ESP provides the BPX Service Node with the ATM or Frame Relay signaling function. It interprets industry-standard signaling messages from ATM or Frame Relay CPE to provide the call setup and tear down for switched virtual circuits across the ATM network. In addition to SVC signaling, the ESP also performs PNNI routing, collects statistics, and processes alarms and billing records for SVC connections through the BPX Service Node.

PVCs and SVCs

Both permanent virtual circuits and switched virtual circuits are defined by ATM and Frame Relay standards groups.

PVCs

After being added to a network, permanent virtual circuits (PVCs) remain relatively static. The PVC only allocates a physical circuit and consumes bandwidth when there is data to send. However, the permanent virtual circuit remains in place, always available for use, and is similar to a dedicated private line in this respect.

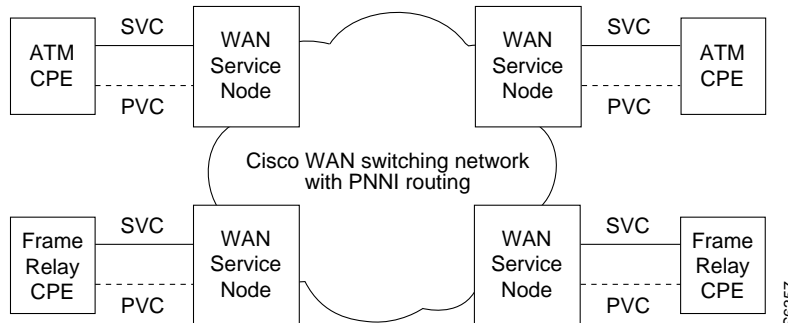
SVCs

A switched virtual circuit (SVC) only exists when there is data to send and a calling process has been initiated. With a switched virtual circuit, there must be some signaling mechanism to build a connection each time the user (ATM or Frame Relay device in this case) needs it. In addition, when the call is disconnected, there must be a mechanism for the orderly disconnection of the call, and the network's resources must be relinquished. During a disconnect, the Cisco WAN switching network sweeps through its connection tables and removes the connection.

ATM SVCs are ATM connections setup and maintained by a standardized signaling mechanism between ATM CPE (ATM user end systems) across a Cisco WAN switching network. ATM SVCs are created on user demand and removed when the call is over, thus freeing up network resources.

Frame Relay SVCs are Frame Relay connections setup and maintained by a standardized signaling mechanism between Frame Relay CPE (Frame Relay user end systems) across a Cisco WAN switching network. Frame Relays SVCs are created on user demand and removed when the call is over, thus freeing up network resources.

Figur e11-1 Wide Area Network with BPX Service Nodes



BPX Service Node Interfaces

The BPX Service Node supports the UNI and NNI interfaces for SVC operations as described in the following:

- UNI, that is the User Network Interface, is the interface for either ATM or Frame Relay customer premise equipment (CPE) to the BPX Service Node. The UNI is defined as any interface between a user device and an ATM network (i.e., an ATM switch). The UNI defines the signaling method which the CPE must use to request and setup SVCs through the wide-area ATM network. In addition, the UNI is used to send messages from the network to the CPE (i.e., user device) on the status of the circuit and rate control information to prevent network congestion.

For ATM SVCs, the UNI supports either the ATM Forum 3.0 or 3.1 signaling standards as well as traditional ATM PVCs. (Remember the BPX switch also supports high-speed ATM UNI ports.)

For Frame Relay, the UNI supports Frame Relay Forum Frame Relay User-to-Network SVC Implementation Agreement (FRF.4), which specifies the Frame Relay SVC signaling protocols. BPX Service Node Frame Relay UNIs (FRSMs) also support traditional Frame Relay PVCs.

- Network-to-Network Interface (NNI). The NNI is the interface to other BPX Service Nodes or foreign ATM Switches. The BPX Service Node supports Interim Inter-switch Protocol (IISP) 3.0 /3.1 or the Private Network to Network Interface (PNNI). These NNI interfaces provide the switching and routing functions between Cisco WAN switching networks and other networks. Information passing across a NNI is related to circuit routing and status of the circuit in the adjacent network. (Note that NNI could refer to both a connection between BPX Service Nodes and a connection between a BPX Service node and a foreign switch.)

Interim Inter-switch Protocol Routing

Interim Inter-switch Protocol (IISP) is an interim static routing protocol defined by the ATM Forum to provide base level capability until the Private Network to Network Interface (PNNI) was specified. The IISP provides users with some level of multi-vendor switch interoperability based on the existing ATM Forum UNI 3.1 specifications. IISP assumes no exchange of routing information between switching systems. It uses a fixed routing algorithm with static routes. Routing is done on a hop-by-hop basis by making a best match of the destination address in the call setup with address entries in the next hop routing table at a given switching system. Entries in the next hop routing table are configured by the user.

PNNI

The Private Network to Network Interface standards essentially define two protocols:

- Topology

The Private Network to Network Interface (PNNI) defines a protocol for distributing topology information between switches and clusters of switches. This information is used to compute paths through the network. A key feature of the PNNI mechanism is its ability to automatically configure itself in networks in which the address structure reflects the topology. PNNI topology and routing are based on the well-known link-state routing technique.


- Signaling

PNNI also defines a second protocol for signaling, that is message flows used to establish point-to-point connections across the ATM network. This protocol is based on the ATM Forum UNI signaling, with mechanisms added to support source routing, crankback, and alternate routing of call setup requests in case of connection setup failure.

Signaling Plane

To support ATM and Frame Relay SVCs, the BPX Service Nodes essentially overlay a signaling network over a traditional (that is PVC-based) network. This signaling network, indicated by the dashed lines in Figure 11-2, connects all of the BPX Service Nodes and extends to the CPE. The signaling plane establishes and maintains SVCs between the CPE, that is, end users, across a CiscoWAN switching wide-area ATM network.

Figur e11-2 BPX Service Node Network Signaling Plane



The signaling plane is created out of two basic types of signaling channels:

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

The signaling VCCs are normally configured during the provisioning of UNI ports and NNI trunks.

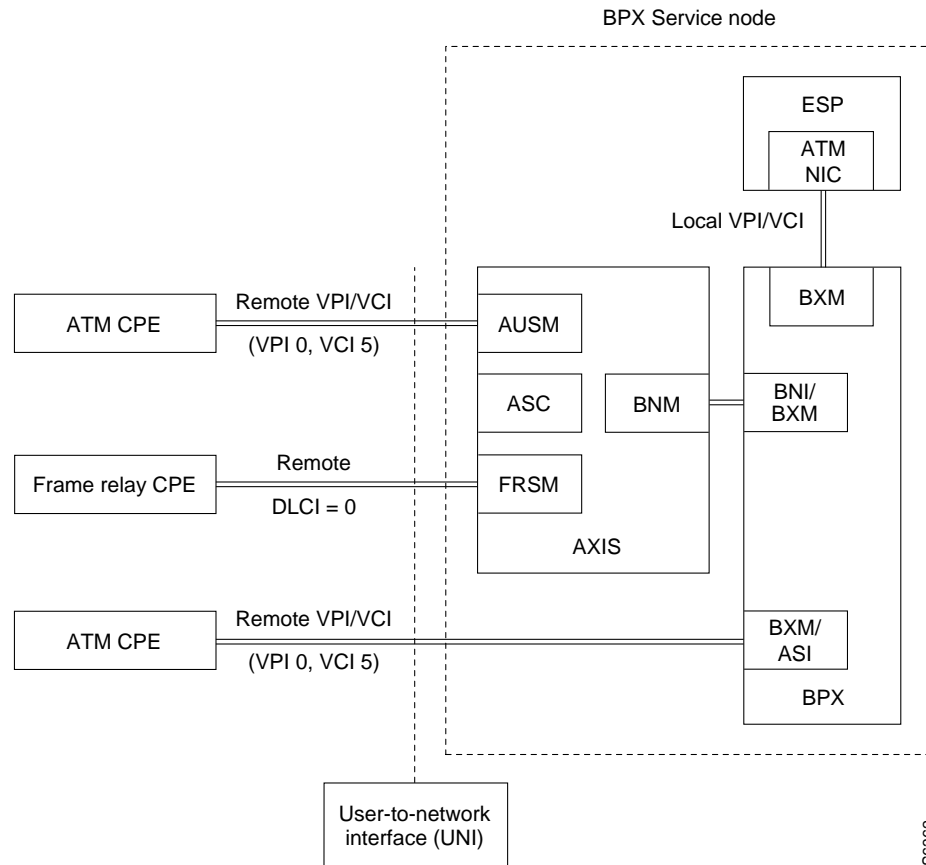
UNI Signaling Channel

There is an internal signaling VCC established between every UNI port on the BPX Service Node which will support ATM or Frame Relay SVCs and the ESP in the BPX Service Node. There are two types of UNI signaling channels supported by the BPX Service Node as shown in Figure 11-3.

- ATM UNI--For ATM CPE, these UNI VCCs extend from an ATM UNI port to the ESP. This is either a one segment cross-connect between the BXM (or ASI) attached to the ATM CPE and the BXM attached to the ESP within the BPX Service Node, or a two segment VCC from the Cisco MGX 8220 edge concentrator AUSM port connected to ATM CPE and the BXM attached to the ESP within the BPX Service Node. (Note that VPI 0 and VCI 5 are reserved on the ATM UNI port for ATM SVC signaling channels. The ILMI signaling channel will use VPI 0 and VCI 16, and the PNNI signaling channel will use VPI 0 and VCI 18.)

- Frame Relay UNI--For Frame Relay CPE, there will always be a two segment VCC between the Cisco MGX 8220 edge concentrator FRSM port connected to the ATM CPE and the BXM attached to the ESP within the BPX Service Node. (Note that DLCI 0 is reserved on the Frame Relay UNI port for Frame Relay SVC signaling channels.)

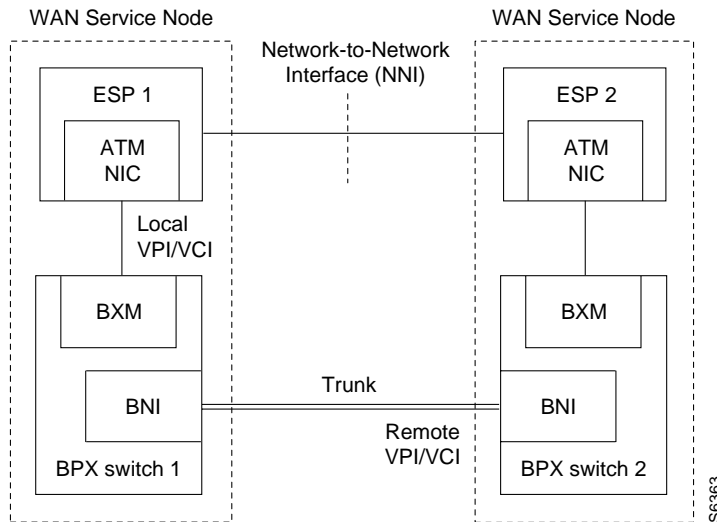
Figur e11-3 UNI Signaling Channels



NNI Signaling Channel

There is also a signaling channel established between each adjacent pair of BPX Service Nodes. This NNI signaling channel shown in Figure 11-4 is configured for either IISP or PNNI protocol. During IISP configuration, one side of the NNI signaling connection is configured as the user side and a weight is assigned. In the figure, the direct line between the ATM NICs indicates a logical connection; the physical connection is configured through BPX 1 and BPX 2.

Figur e11-4 ESP Signaling PVC



Network Interworking Between Frame Relay and ATM

Because the BPX is an ATM switch, Frame Relay SVCs that are setup and established across the Cisco WAN switching network must be translated into an ATM format to be carried across the network. At the far end, where typically the connection is terminated on another Frame Relay CPE, the ATM cell will have to be converted back to Frame Relay format. This is referred to as Network Interworking. Network Interworking can be performed between Frame Relay CPE and ATM CPE when the ATM CPE recognizes that it is connected to an interworking function (Frame Relay, in this case). The ATM CPE must then exercise the appropriate service specific convergence sublayer (SSCS). The SSCS will then convert the ATM cells to Frame Relay traffic.

In this release of the BPX Service Node, all Frame Relay SVC connections must be between Frame Relay CPE (that is, Frame Relay end users) or ATM CPE that is aware that it is performing Network Interworking, and all ATM SVC connections must be between ATM CPE (that is ATM end users). In other words, Service Interworking between ATM and Frame Relay SVCs is not supported in this release. (ATM and Frame Relay Service Interworking for PVCs is supported by the BPX Service Node as part of Release 8.4.)

Extended Services Processor

The Extended Services Processor (ESP) is an adjunct processor shelf integrated into the BPX Service Node.

The basic ESP features include:

- 140 MIPS CPU, with a 143Mhz clock
- 128 Megabytes of memory
- 4 Gigabyte hard disk

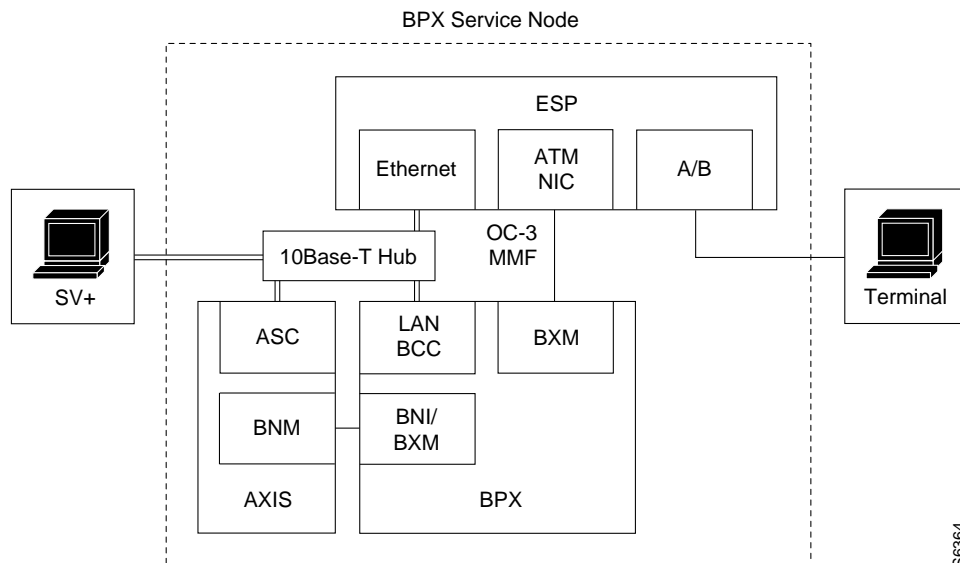
Available in either AC- or DC-powered models (ESP-AC or ESP-DC), the ESP is an orderable option for the BPX switch. The ESP can be configured in both non-redundant and redundant configurations. For the redundant configuration, two ESPs are installed in the BPX Service Node.

ESP Interfaces

The ESP uses three main physical interfaces, as shown in Figure e11-5:

- **Terminal port** for the direct connection of a terminal, such as a VT-100, to provide access for local configuration and to act as a console.
- **10Base-T Ethernet port** for connection to the Cisco WAN Manager workstation and to the BPX. Telnet or XTERM sessions can be established through the Ethernet port, and perform the same functions as can be performed with a directly connected terminal.
- **ATM Network Interface Card (ATM NIC)** for connection to the BPX. The ATM NIC is typically connected to a BPX BXM card using OC-3 multimode fiber connection with SC connectors. There are optional cables with built-in optical attenuation that allow BXM single mode fiber (SMF) backcards to be connected to the ESP ATM NIC.

Figur e11-5 ESP Physical Interfaces



The ESP also provides the following application interfaces:

- **SNMP** to configure and monitor the ESP.
- **TFTP** (trivial file transfer protocol) for uploading statistics, Call Detail Records (CDRs), and downloading configuration files and new software releases and revisions.
- **Telnet** for accessing the ESP remotely, such as from the Cisco WAN Manager workstation.

Redundant ESPs

ESPs can be installed in redundant pairs in the BPX Service Node. In a redundant pair, one ESP is active, that is it controls the switched services in the BPX Service Node, and the other ESP is standby. The redundant ESPs are known as peers. The ESPs will switch roles from active to standby and vice versa under the following conditions:

- Controlled switchover invoked from the active ESP Configuration Interface.
- The active ESP detects a major service affecting failure, such as a BPX polling failure or an ATM NIC card failure, and relinquishes the active role, by going Out of Service.
- The standby ESP detects that the inter-ESP paths have failed and assumes the active role.



Note

During a switchover, all SVC connections will be torn down and the ATM or Frame Relay CPE will have to initiate another SVC call to reestablish them.

Network Management

As shown in Figure 11-5, the BPX Service Node could have an Ethernet LAN connection to a Cisco WAN Manager workstation. Cisco WAN Manager discovers and monitors the ESP similarly to the way it does an Cisco MGX 8220 edge concentrator shelf. Cisco WAN Manager discovers the existence of the ESP when it is added to the BPX shelf with the **addshelf** command. After discovery, the ESP will be displayed on the Cisco WAN Manager topology map as a shelf attached to the BPX switch.

Cisco WAN Manager manages the BPX Service Node by providing:

- Telnet access to the ESP Configuration Interface.
- Configuration backup and restore
- Image download allows the user to select an ESP from the topology map and then select Image Download from a pull-down menu. A dialog screen displays the list of available image files and prompts the user to select a file. The selected file is sent to the ESP.

Resource Partitioning

During provisioning, resources on all UNI ports (both ATM and Frame Relay) are partitioned between SVCs and PVCs. Partitioning is performed using the BPX and Cisco MGX 8220 edge concentrator command line interfaces. This partitioning information is retrieved from the BPX by reading its port and trunk tables and from the Cisco MGX 8220 edge concentrator by reading the resource partitioning tables in the AUSM and FRSM MIBs.

The BPX line and routing or feeder trunk resources to be partitioned are:

- LCR range
- VPI range
- Port Queues
- Egress Queue pool size
- Bandwidth

Cisco MGX 8220 edge concentrator Feeder Trunk (BXM/BNI) resources to be partitioned are:

- LCN range
- Bandwidth

Cisco MGX 8220 edge concentrator AUSM port resources to be partitioned are:

- LCN range
- VPI range

Cisco MGX 8220 edge concentrator FRSM port resources to be partitioned are:

- LCN range
- DLCI range.

Frame Relay Connections

This chapter is provided for users who wish to have an in-depth knowledge of network frame relay connections and related functions and also describes the Port Concentrator Shelf (PCS) which extends the port capacity of an FRP on a Cisco IPX narrowband switch or of an FRM on a Cisco IGX 8400 series multiband switch from 4 high-speed ports to 44 low-speed ports.

The chapter contains the following:

- Connection Types
- Network Interfaces
- Congestion Prevention and Notification
- Connection Parameters
- Port Concentrator Shelf Frame Relay Connections

Cisco IPX Narrowband Switch, Cisco IGX 8400 Series Multiband Switch, and Cisco MGX 8220 Edge Concentrator

The examples in this chapter are based on Cisco IPX narrowband frame relay connections, but the general information up through the “Connection Parameters” section of this chapter is applicable to the Cisco IGX 8400 series multiband switch and to the Cisco MGX 8220 edge concentrator shelf. See the respective reference publications for frame relay information specific to the Cisco IGX 8400 series multiband switch and the Cisco MGX 8220 edge concentrator shelves.

Port Concentrator Shelf

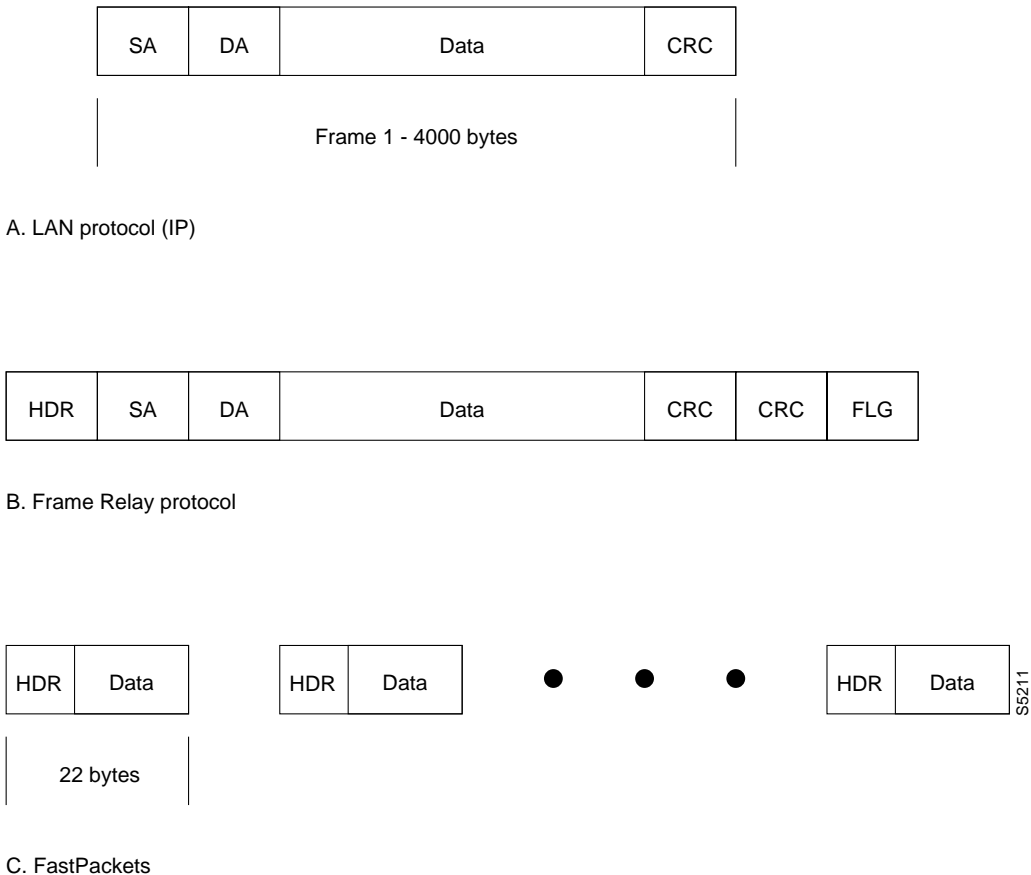
The use of the Port Concentrator Shelf (PCS) which extends the port capacity of an FRP on a Cisco IPX narrowband switch or of an FRM on a Cisco IGX 8400 series multiband switch from 4 high-speed ports to 44 low-speed ports is described in the “Port Concentrator Shelf Frame Relay Connections” section later in this chapter. The addition and configuration of frame relay connections via the ports on the PCS is essentially the same as that for connections added directly at an FRP port. The general information provided up through the “Connection Parameters” section of this chapter is applicable to the PCS, except as described in the “Port Concentrator Shelf Frame Relay Connections” section.

**Note**

In the following discussion, the FRM and NPM cards perform the same functions in the Cisco IGX 8400 series multiband switch as the FRP and NPC do in the Cisco IPX narrowband switch.

Frame relay can be used to transport virtually any higher layer data protocol. Higher layer protocol frames are encapsulated in Frame Relay frames, usually according to the scheme defined by Internet Standards RFC-1294. Figure 12-1 illustrates this for a typical LAN format (TCP-IP).

Figur e12-1 Example of Carrying Frame Relay Data with FastPackets



Frame relay surrounds the LAN data frame with its own protocol elements allowing it to carry LAN data transparently. The flags are used to identify the start and end of frame relay frames, which can be up to 4506 bytes long. Since each FastPacket contains only 20 bytes of data, a number of packets may be needed to transmit one frame relay frame (Items A and C in Figure 12-1). A Cisco WAN switching network recognizes three core frame relay elements:

- the frame delimiters (flags)
- a two-byte header
- the CRC byte

The frame relay network destination address is given by the Data Link Connection Identifier (DLCI) located in the header of each frame relay data frame. A connection's DLCI is simply a number used to distinguish the connection from all the other connections that share the same physical port between the user equipment and the frame relay port. It is assigned by the network operator when the connection is added to the network. DLCI values range from 0 to 1023 with 0 to 15 and 1007 to 1023 reserved for special use.

A frame length measurement is inserted at the end of the FastPacket carrying the last of the data for the frame. It is used to perform error detection on the whole frame relay frame. A ZBTISI algorithm is used to remove any bytes consisting of all zeros. The ZS bit indicates one or more data bytes were removed before being transmitted. One or more of the first data bytes acts as a pointer to the removed byte so it can be replaced at the receiving end.

At the receiving end, the Cisco IPX narrowband switch reassembles the packet data into a complete frame relay frame, checks for correct CRC and frame length, and outputs the frame to the destination device only if the frame relay CRC is correct. Frame relay ignores any CRC that may be associated with the LAN protocol.

Since the frame relay dataframe may be up to 4096 bytes long, it will likely require a number of FastPackets to transmit all of the frame relay data. The FRP card at the far end node will wait until the whole frame relay data frame has been received and checked before it will begin to output it to the far end user device.

All FastPackets for a frame relay connection travel along the same network route; they cannot get out of sequence as can be the case for other packet networks. But the frame relay packets are not necessarily contiguous over a packet line as there are other packet types that are being transmitted at the same time. The Cisco IPX narrowband switch inserts start of frame and end of frame codes into the frame relay message packet to assist the far end node in preserving the proper sequence of packets (Figure 12-2). An encapsulated frame code indicates the packet message frame contains a complete frame relay data frame.

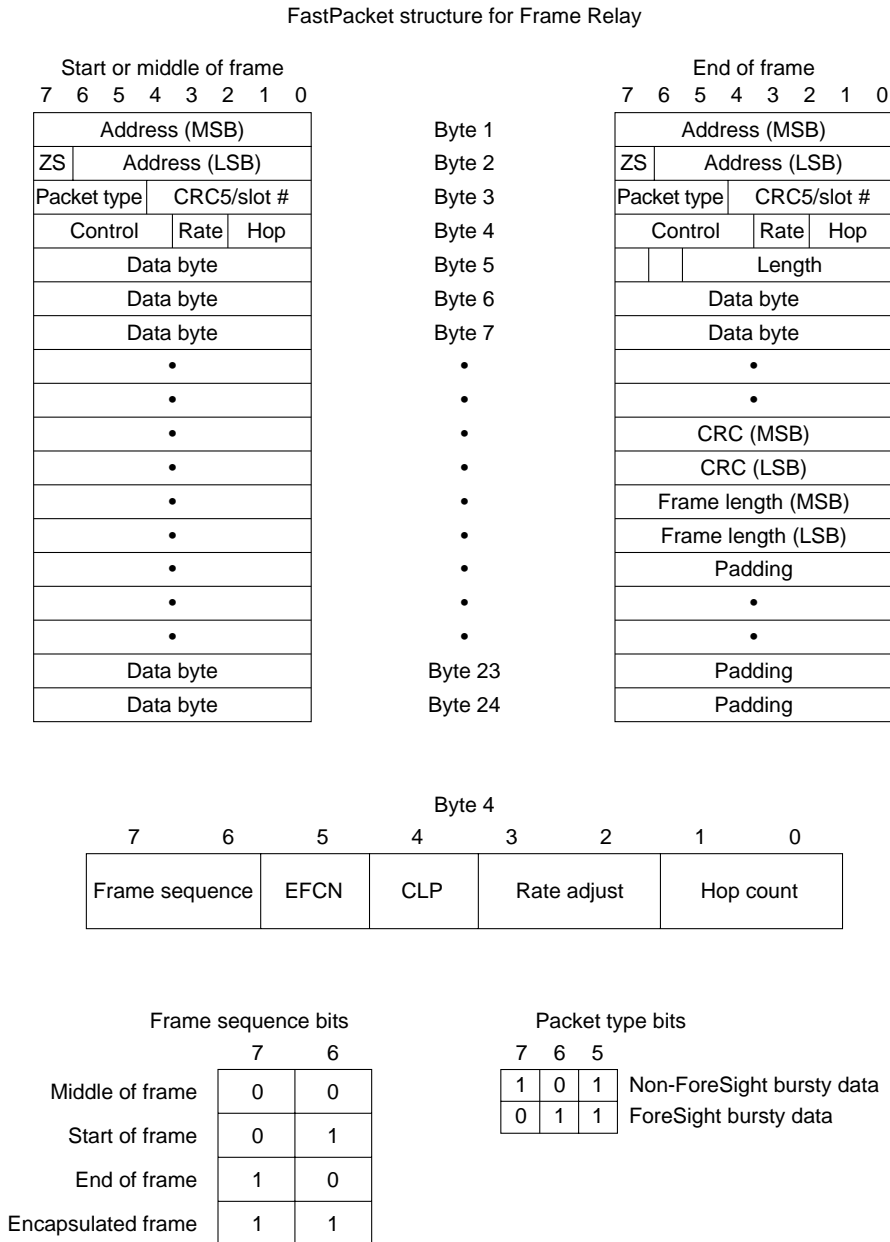
The hop count is incremented at each intermediate node to indicate how many hops in the network the packet has traveled. It is used at intermediate nodes to give higher priority in various packet line queues to frame relay data that has traveled over more hops (with associated longer delays) to attempt to even out the delay for various frame relay circuits.

The packet frame carrying the last frame relay frame (Figure 12-2) has a format different from all the rest. As it is unlikely that the last frame will be completely filled, the remaining bytes are used to carry a 2-byte CRC for error checking on the complete frame relay frame. The next two bytes indicate the length of the frame relay frame as a check against a dropped packet in case the CRC is still good. Any remaining packet message bytes are filled with padding (hex 7E). If the CRC check at the receiving end detects an error in the frame relay frame, the whole frame relay frame is discarded.

At the transmitting node, the frame relay Data Link Connection Identifier (DLCI) is replaced by the Cisco IPX narrowband packet header routing address for the destination node. At the destination node, the FRP card replaces this routing address with the source DLCI code before transmitting it to the user.

The maximum number of frame relay connections possible in a node is 1024. This is accomplished by bundling frame relay connections into groups. All bundled frame relay circuits have the same destination and are routed over the same route. Frame relay circuits originating at different FRPs in the same node may be bundled into the same group.

Figur e12-2 FastPackets for Frame Relay Data

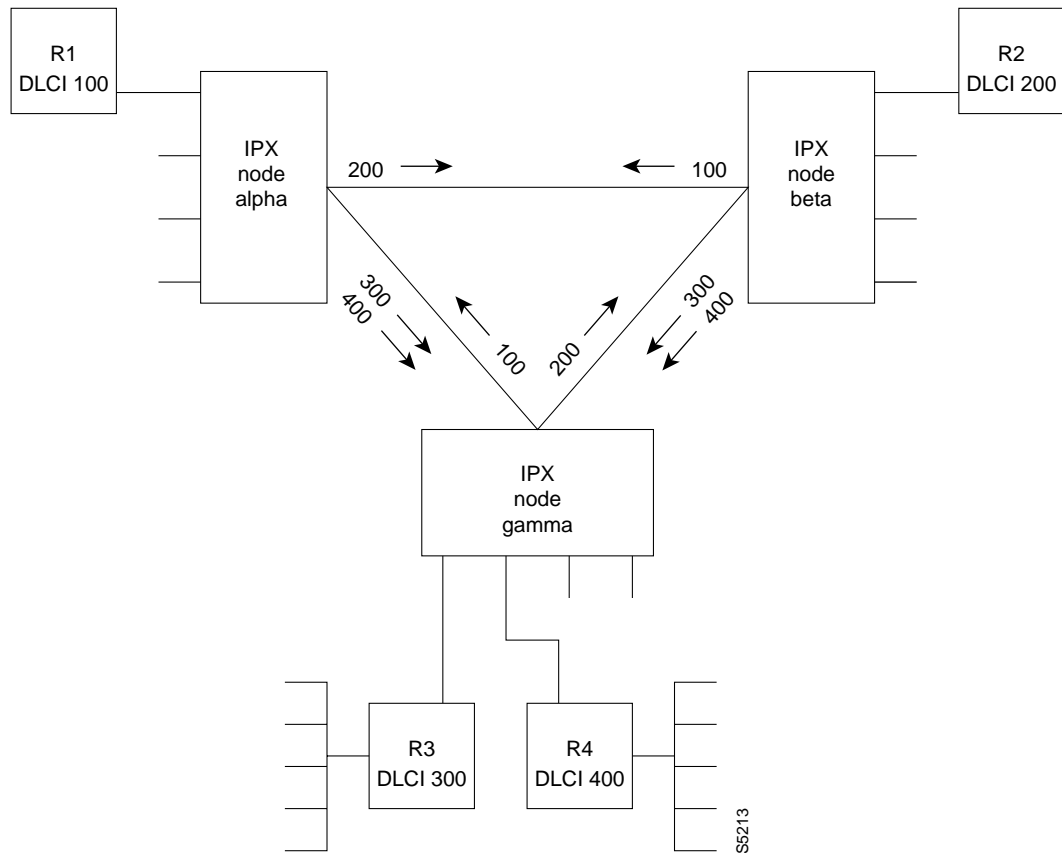


S5212

Frame Relay Addressing

An example of frame relay addressing is presented in Figur e12-3. It illustrates a simple three-node network with a single router at nodes alpha and beta (R1 and R2) connected to the first of four ports of an FRI card, and two routers (R3 and R4) connected to ports 1 and 2 of a FRI at node gamma.

Figur e12-3 Example of Frame Relay Addressing

**Note**

The Port ID field in the frame relay port record is not used by the Cisco IPX narrowband switch except for adding bundled connections and to help administer global addressing.

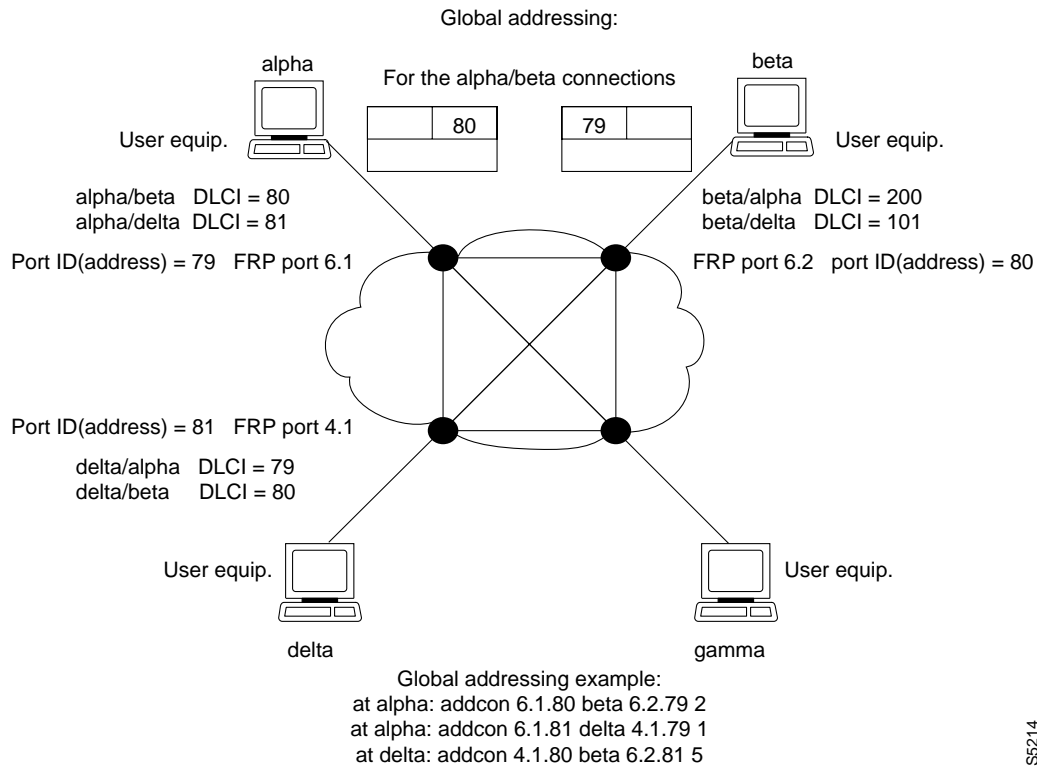
A Port ID can be assigned to each Cisco IPX narrowband frame relay port using the `configure FR port` command. As an option, the Port ID can be used as a starting point for assigning DLCI's. The frame relay connection between alpha and beta has been added with a DLCI of 100 at alpha and 200 at beta. Likewise, the connection between nodes alpha and gamma have been assigned a DLCI of 300 for router R3 and 400 for router R4.

When router 1 at alpha wants to send a frame to router 2 at beta, it inserts the DLCI 200 in the frame relay header. If it wants to send a frame to router 4 at gamma, it inserts 400 in the DLCI field. Each router will have a list of all available destinations (routers) connected to the network and their corresponding DLCI number.

In this example, each router in the network has a different DLCI. A circuit from anywhere in the network to router 2, for example, will use a DLCI of 200 to identify the destination as node beta, router R2. This makes it easy for anyone in the network to associate destination locations with a simple numbering scheme. When a frame is broadcast to multiple destinations, there will be no confusion as each destination address is unique.

This addressing scheme is often called Global Addressing. Global Addressing uses the DLCIs to identify a specific end device, which makes the frame relay network look more like a LAN. Figure 12-4 illustrates a network using global addressing.

Figur e12-4 Example of Frame Relay Global Addressing



Another user community that is not connected with the first could have its own global addressing scheme using some, or all, of the same DLCIs. This might be the case when a public network provider has a number of customers, each with their own frame relay network, running on the same Cisco IPX narrowband hardware.

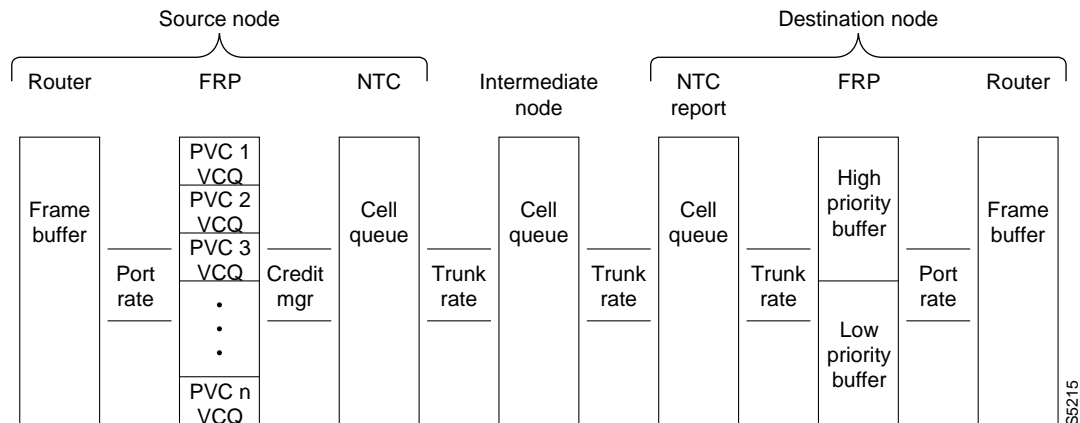
Since the DLCIs have only “local significance” the only real restriction on the use of DLCIs is that they are not used for more than one destination from the same port. This is not to say that the same DLCI numbers could not be reused at another port at the same or another Cisco IPX narrowband node. And, indeed, another addressing scheme might assign the same DLCI to both ends of each virtual circuit between nodes.

For example, router R2 at beta could have a DLCI of 100 and 100 could also be used as the DLCI for router R1 at alpha. Now, a frame originating at either alpha or beta would use 100 as its destination address and the same PVC to transmit the frame. But this addressing scheme can be confusing for the same number refers to two different destinations. This type of addressing convention is sometimes called connection addressing since the address is only unique at the local port. Another addressing scheme, Local Addressing, reuses DLCIs at each node and only at that node are they unique.

Frame Relay Signal Flow

Figure 12-5 illustrates a block diagram of a simple three-node network and the buffers encountered by the frame relay data. Frames are assembled by the router and are transmitted to the FRP port at the Port Rate determined by the hardware. The frame header contains the DLCI from the User Device.

Figur e12-5 Frame Relay Signal Flow



This DLCI, along with the port number, is used to determine a unique permanent virtual connection within the Cisco IPX narrowband network. Each PVC has its own individual queue, VC Q, at the input to the FRP. Each VC Q is user specified and can be up to 65 Kbytes deep. A FRP card can have up to 252 PVCs associated with it.

The FRP uses this DLCI in a lookup table to find a FastPacket address for routing through the network. This address is placed in the header of one or more Cisco IPX narrowband packets into which the frame is loaded and the packets are forwarded to the appropriate trunk card where they are queued in a frame relay Cell Queue for transmission out to the network. There are separate cell queues for ForeSight and non-ForeSight traffic. The Cisco IPX narrowband Credit Manager determines the rate at which the frames are unloaded from the cell queue and applied to the network.

The packets are then routed across the Cisco IPX narrowband network. At each intermediate node, the packets containing frame relay data are queued in their Cell Queues for transmission over the next hop. Packets containing frame relay data share the same network facilities as packets containing voice, low-speed data, synchronous data, PCC, and video traffic.

At the receiving end FRP, packets containing frame relay data are loaded into one of two Port Buffers. The original frame is reconstructed from the packets and then the source DLCI is replaced with the destination DLCI. This identifies the source of the data.

One of the two buffers is for connections marked as high priority, the other is for normal frame relay data. Data in the high priority queue is transmitted to the destination router first. One final stage of buffering is found in the router as it receives data from the FRP in bursts.

Each FRP VC Q buffer can be allocated from 1 to 65 Kbytes. The larger buffer sizes increase the overall connection delay but larger buffers also minimize the possibility of buffer overflow. Buffer space that is not used by frame relay connections can be allocated to other types of connections.

Connection Types

There are three types of frame relay connections available with the Cisco IPX narrowband switch:

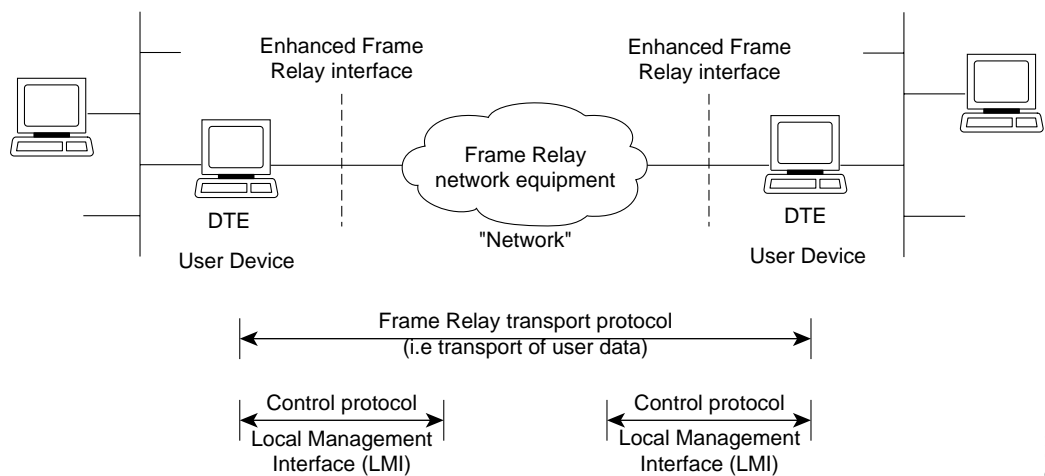
- normal connections
- bundled connections
- point-to-point (frame forwarding) connections

Normal Frame Relay Connections

The Cisco IPX narrowband switch provides permanent virtual circuits (PVC) for interconnecting user data devices (routers, bridges, and packet switches). The PVCs are created internally in the Cisco IPX narrowband switch, using routing tables and FastPacket switching. The user device is connected to the Frame Relay Interface (FRI) card set installed in the Cisco IPX narrowband switch, which provides the adaptation layer function to convert between the frame relay format and the Cisco IPX narrowband FastPacket format (Figure 12-6).

In Cisco IPX narrowband networks, all packets belonging to a particular PVC travel along the same route. This means all packets carrying the frame relay data for a particular destination are by definition transmitted in sequence and experience approximately the same delay, unlike previous low-speed packet networks. Cisco IPX narrowband frame relay packets may travel over multiple network hops (10 maximum) to reach their destination.

Figur e12-6 Basic Frame Relay Service



S5216

The data transfer protocol used by frame relay is very simple. The user device transmits data frames, based upon the core functions of Q.922 (LAPD), to the frame relay network. The frame relay network looks at the first two octets of the frame Data Link Connection Identifier (DLCI) and proceeds to forward the frame to the destination user device.

The only processing performed by the Cisco IPX narrowband switch on the frame is bit insertion and verification and a Frame Check Sequence (FCS) for error checking. Termination of the data link is not performed. The destination DLCI is replaced with a DLCI that identifies the source of the data frame.

Because so little processing is performed, the frame relay service can offer much higher speed interconnection between user devices than can be achieved with conventional packet switches that relay layer 3 protocol data units. Refer to ITU-T and ANSI Standards for further definitions on data transfer protocol and procedures.

Bundled FR Connections

Bundled frame relay connections simplify the specification of large numbers of FR connections (up to 1000 circuits per node). This feature is used to connect a set of consecutively numbered frame relay ports together.

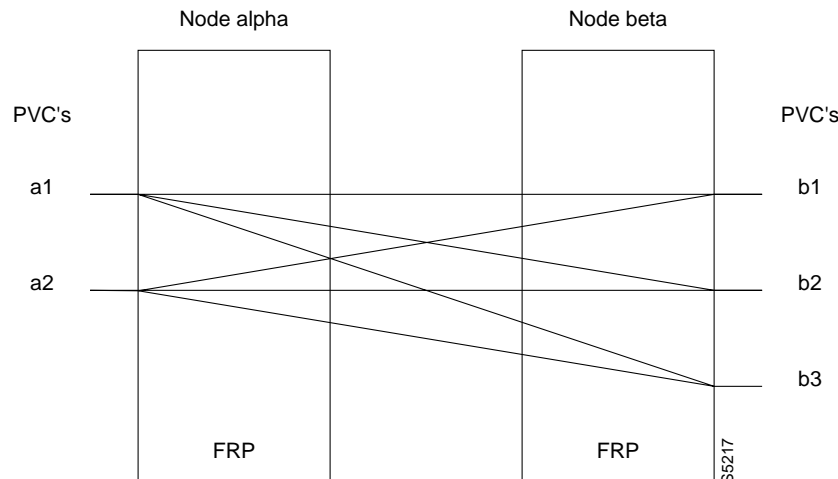
As with grouped connections, a bundled connection can consist of up to 16 virtual circuits with the same routing and same destination node. However, all circuits in a bundled connection must be located on the same FRP card. They do not necessarily have to originate and terminate on the same port at each end but the ports must be on the same card. For example, each of the four ports can have four connections per port.

An example of a mesh bundled connection is illustrated in Figure 12-7 as follows. Suppose we have two ports at the local node (node alpha) that we want to connect to three ports at the remote node (node beta). This will result in six PVCs (2 X 3) fully interconnecting the five ports. The Add Connection (addcon) command does not need to individually specify each connection, only the ports.

This will result in six connections as follows:

```
alpha 1 to beta 1      and      alpha 2 to beta 1
alpha 1 to beta 2      and      alpha 2 to beta 2
alpha 1 to beta 3      and      alpha 2 to beta 3
```

Figur e12-7 Bundled Connections



Unlike grouped connections, the circuits in a bundled connection must all be added at the same time and use the port ID of each FRP port for the destination address as the DLCI. The Add Connection (addcon) command for bundled connections uses only the slot and port number for specifying and consecutively numbers the PVCs, starting with the port number, for the first port to make it simple to add bundled connections.



Note

Bundled connections cannot be grouped.

Frame Forwarding Connections

Frame forwarding supports a data connection like that described for the SDP or LDP. It uses the capabilities of the frame relay card to provide bursty data, DFM-like data compression for high speed data connections. Frame forwarding is configured on an individual FRP port basis.

This feature is used to interconnect various data applications that do not conform to the frame relay Interface Specification, such as bridges, switches, front end processors, and other equipment, supporting SDLC, HDLC, or LAP-B interfaces.

1. Frame length must be between 5 to 4506 bytes.
2. Valid ITU-T CRC/FCS required.
3. Flags must be 7E hex.

In this configuration all frames received on a local FRP port are transmitted via a single PVC to a remote FRP port, and all frames received on a remote FRP port are transmitted via a single PVC to a local FRP port. Refer to the discussion of frame relay in Chapter 2 for port hardware interface description and operating bit rates. Note that a frame forwarding connection is still a frame-oriented interface, as in the case of normal frame relay connections.

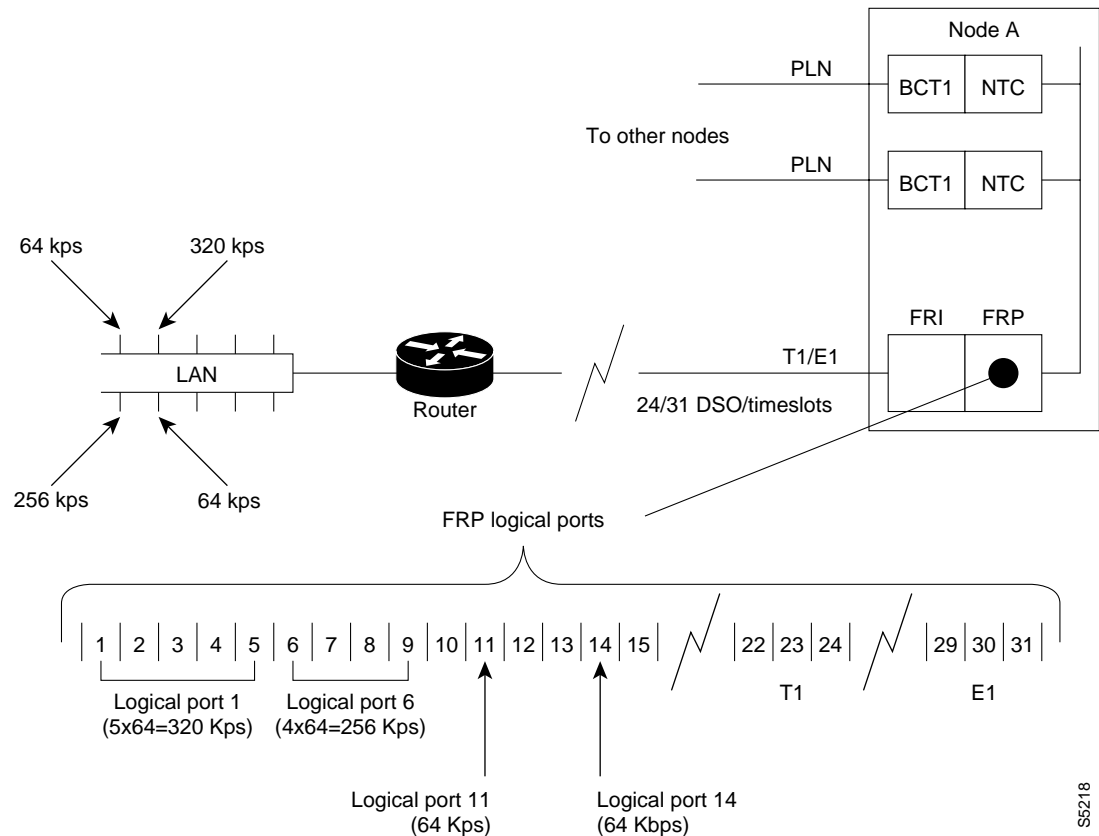
T1/E1 Frame Relay

The Frame Relay T1/E1 application allows the user to group FRP DS0/timeslots into “logical ports”. These logical ports may be a single DS0/timeslot or groups of contiguous DS0 timeslots. Logical ports that consist of multiple DS0/timeslots are at the full rate of 64 Kbps per timeslot. Frame Relay LMI is simultaneously supported on a maximum of 31 T1/E1 logical ports.

Logical ports that consist of single DS0 timeslots may be configured in 56 Kbps or 64 Kbps. If configured for 56 Kbps the Cisco IPX narrowband switch will strip off the signalling bit in the incoming octet and stuff a “1” in the outgoing octet. This 56 Kbps rate is typically used for groomed DDS circuits that appear on a T1/E1 line. Figure 12-8 is a simplified illustration of multiple and single DS0/timeslots comprising logical ports.

Logical ports are created with the Add Frame Relay Port (**addfrport**) command, which associates a line number (circuit line) and DS0/timeslots to form a logical port. The lowest timeslot number of the created group becomes the logical port number. The created logical port number is used to *up* the port, add connections, and display statistics. Logical ports are deleted using the Delete Frame Relay Port (**delfrport**) command, which ungroups any multiple DS0/timeslots and/or unassign a single DS0/timeslot logical port.

Figur e12-8 Multiple and Single DS0s Forming a Logical Port



S5218

Network Interfaces

There are two types of network interfaces possible at frame relay ports, a User-to-Network Interface (UNI) and a Network-to-Network Interface (NNI). The User-to-Network Interface is defined as the port where a user device, such as a router, interfaces with a CiscoWAN switching wide area network carrying the frame relay traffic. However, the functions performed by each network interface is quite different as will be discussed in the following paragraphs.

A Network-to-Network Interface is a port that forms a boundary between two independent wide area networks, e.g. a Cisco WAN switching network and another network and may or may not consist of Cisco WAN switching equipment. There is no user device connected, only another network port. Each network interface in a Cisco WAN switching network consists of a port on a FRP card.

User to Network Interface (UNI)

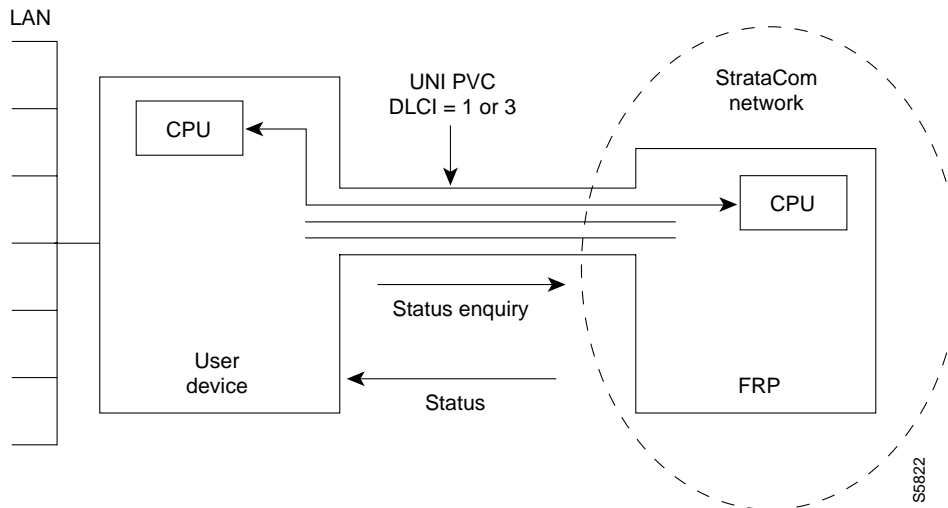
The User-to-Network Interface for frame relay permanent virtual circuits (PVC) is a defined set of protocols and procedures. Currently, the Cisco IPX narrowband switch supports UNI via the following protocols: StrataCom LMI, ITU-T Q.933 Annex A, and ANSI T1.617 Annex D. Each of the three protocols is quite similar and only the StrataCom Local Management Interface (LMI) will be discussed here.

LMI transmits on a logical connection between the Cisco IPX narrowband switch and the user device separate from the data path and uses DLCI 1023 (Figure 12-9). This connection is a special PVC, carrying messages between the Cisco IPX narrowband frame relay port 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 PVC
- Notification of the priority of each virtual circuit.
- The StrataCom LMI can provide a simple 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 Cisco IPX narrowband network.

Figur e12-9 Frame Relay User-to-Network Interface



Network-Network Interfaces

Frame Relay networks utilizing Cisco WAN switching nodes can be seamlessly connected together and to other frame relay networks adhering to standards set forth by the Frame Relay Forum. Internetwork connections originate within one network and terminate inside another independent network. For example, a circuit might originate within a private network, pass through a public switched network, and terminate within a third, private network.

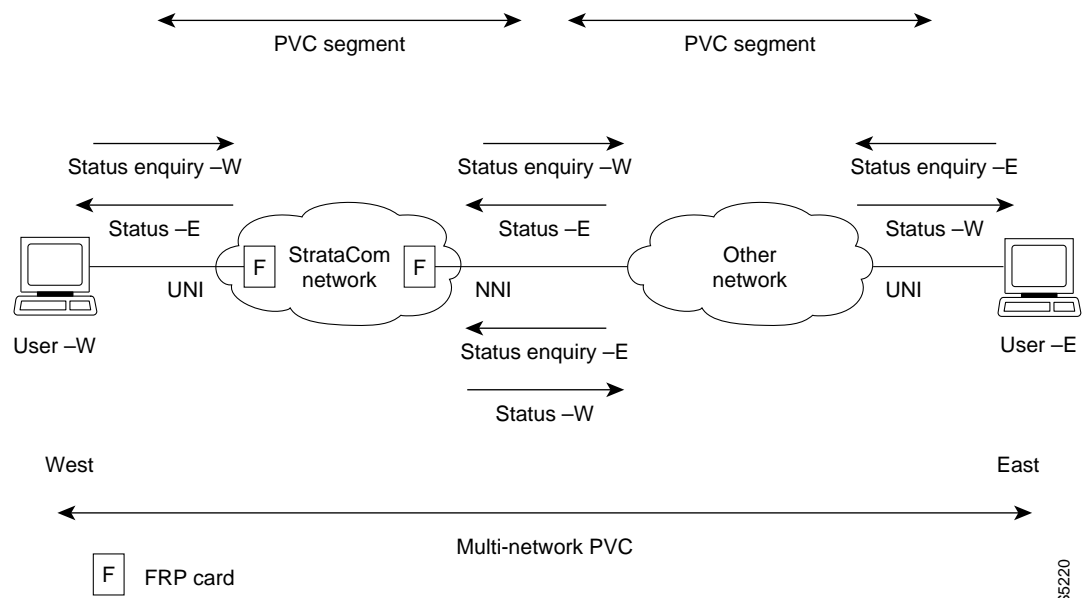
The Frame Relay Network-to-Network Interface (NNI) is used to forward the connection status through to the User-to-Network Interface (UNI) ports within the different networks. Inter-network messages flow between Network-to-Network Interface (NNI) frame relay ports to report internetwork connection status to user devices wherever they are located. In order for this feature to be completely implemented, both networks must support NNI.

Within a Cisco WAN switching flat 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. There are three possible status to report

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

This is illustrated in Figure 12-10 with a multi-network PVC connecting a user at the West end of the connection (User-W) connecting to a Cisco WAN switching network and a user at the East end (User-E) to an adjacent, independent, network. This connection is segmented with that portion of the connection traversing one of the networks called a PVC segment.

Figur e12-10 Example of NNI in a Multi-Network PVC



At the boundary of the two networks are two nodes, one in each network, with one or more frame relay NNI ports in one node connecting to a like number of ports in the node in the other network. These ports carry the internode connection. Each of the NNI ports constantly monitors the connection interface

between the two networks. Each of the frame relay ports will periodically poll the corresponding port in the other network using a Status Enquiry message (heartbeat). The interval of this polling is configurable, set by the Link Integrity Timer (T391), and normally polls every six seconds.

This same port expects to receive a Status message back from the other network port within a certain interval, set by the Polling Verification Timer (T392), indicating the NNI ports are active and communicating properly.

If a Status message is not received before the Polling Verification Timer times out, an error is recorded. When a preset number of errors, set by the Error Threshold and Monitored Events Count, a Port Communications Fail message is generated and returned to the UNI end of the connection. This is displayed on the Cisco IPX narrowband alarm screen.

When the heartbeat signal indicates the ports are functioning properly, the port sends out a Full Status Enquiry message to the corresponding port in the other network requesting a status of all of its connections. This occurs approximately once a minute, set by the Full Status Polling Cycle (N391). The port responds by returning a Full Status Report indicating the status (active, failed, disconnected) of all connections carried by this port.

The connection status is transmitted using a bit, the active bit (A-bit), in the frame relay frame as defined by the various frame relay standards committees. Since the connection is bidirectional, the NNI protocol must also be bidirectional with both directions of transmission operating similarly, but independently.

In Figure 12-10, the physical layer network-to-network interface for a CiscoWAN switching network is a port on a FRP card. The FRP in the Cisco WAN switching network at the network-to-network interface must have at least one of its ports configured as a NNI port. The two NNI ports send Status Enquiries and receive Status messages back and forth to confirm proper port communication.

After a specified number of heartbeat cycles, the FRP on the east side of the local (CiscoWAN switching) network sends a Full Status Request message and the corresponding NNI port in the other network replies with a Full Status Report indicating the condition of all of the PVCs over that port for the connection segment in the other network.

The FRP in the east side of the CiscoWAN switching network builds a table using information from the other network received in the A status bit. This table stores the status for each PVC. If there is a change in status, this FRP generates a special packet (Operation, Administration, and Maintenance FastPacket) that is used to send this change of status information back to the west side FRP.

A similar table is built in the west side FRP. The A-bit status, reflecting the status of the PVC in the "other" network segment, is logically ANDed with the status of the PVC in the Cisco WAN switching segment and the resulting status now reflects the end-to-end status of the PVC. This status is available to the User Device at the east side of the Cisco WAN switching network. If, and when, the User-W sends a Status Enquiry to the FRP over the UNI port, the connection status is transmitted to the User-W device. The process is repeated in the opposite direction to propagate the PVC status to the User-E device at the other end of the connection.

A Cisco WAN switching network may be upgraded to support NNI in a gradual manner. It is not required that all FRPs in a network be upgraded to enable NNI at one or more ports. Connections may be established between NNI ports and UNI ports using pre-Model F or H FRPs. However, NNI will not send the A-bit status to old UNI ports to prevent these FRPs from logging errors.

Congestion Prevention and Notification

Each frame relay virtual circuit has an assigned information rate. Because of the bursty nature of most data protocols, not all devices will use all of their assigned information rate all of the time. This has two consequences:

- Some devices may be allocated more than their assigned bandwidth at moments in time when other devices are not using theirs.
- Network trunks may be provisioned with less bandwidth, anticipating that the actual bandwidth used by all connections will be less than their assigned information rates.

Because any frame relay device can transmit data at up to its physical access rate for extended periods of time, congestion control mechanisms must be implemented within the frame relay network to ensure the fair allocation of finite network bandwidth among all users. The network must also provide feedback to user devices on the availability of network bandwidth to permit them to adjust their transmissions accordingly.

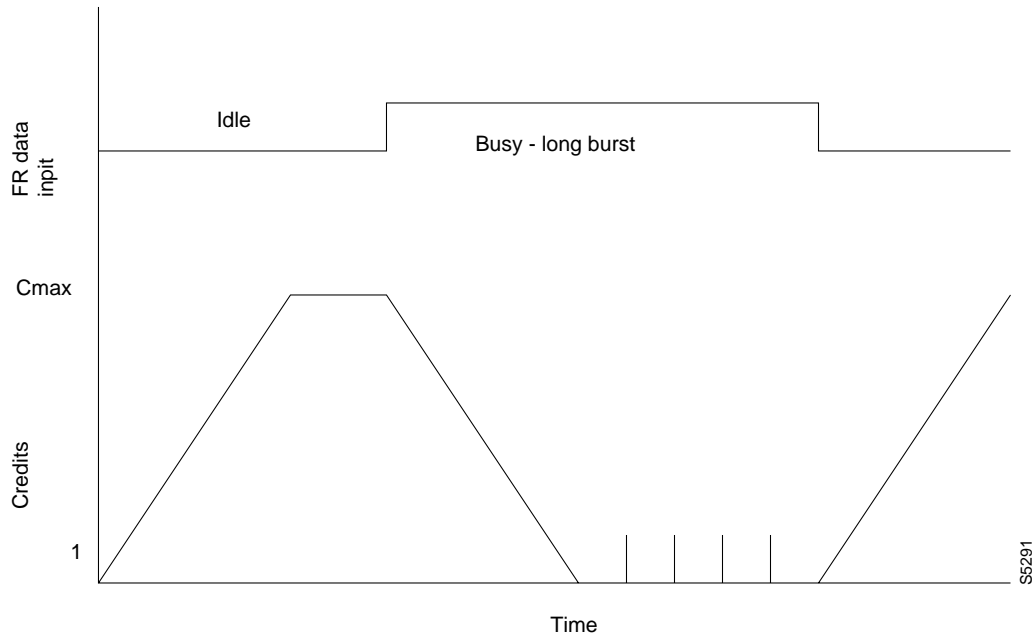
Credit Manager

The basis of congestion avoidance and notification within Cisco WAN switching frame relay service is the Credit Manager. The Credit Manager actively regulates the flow of data from each frame relay virtual circuit into the FastPacket network. The rate at which data is admitted to the network depends on parameters assigned to the virtual circuit by the network administrator, and may also depend on the current state of resources within the network.

A “credit manager” software control limits the size of these initial bursts of frame relay data. Each connection is allowed to send one FastPacket to the network in exchange for one “credit”. Credits are accrued by each connection at a rate sufficient to allow it to send the required number of packets to achieve its configured minimum bandwidth. If a connection does not need its credits to send packets immediately, it is allowed to accumulate a limited number of them for future use.

Cmax provides a maximum credit value in packets to a connection. A connection accumulates credits continuously at a rate up to a maximum accrual of Cmax. A connection spends credits when it transmits packets of data. When credits are available, a burst of Packets is transmitted as long as credits are available; and when credits are exhausted, Packets are transmitted at the minimum information rate (refer to Figure 12-11).

Figur e12-11 Cisco IPX Narrowband Switch Credit Manager Operation



Credits are accumulated at a fixed rate with the normal frame relay feature based on the connections specified minimum information rate/committed information rate and the average frame size. With ForeSight frame relay, to be discussed in detail later in this chapter, credits are accumulated at a variable rate based on CIR as well as the instantaneous available bandwidth on the packet trunks in the network.

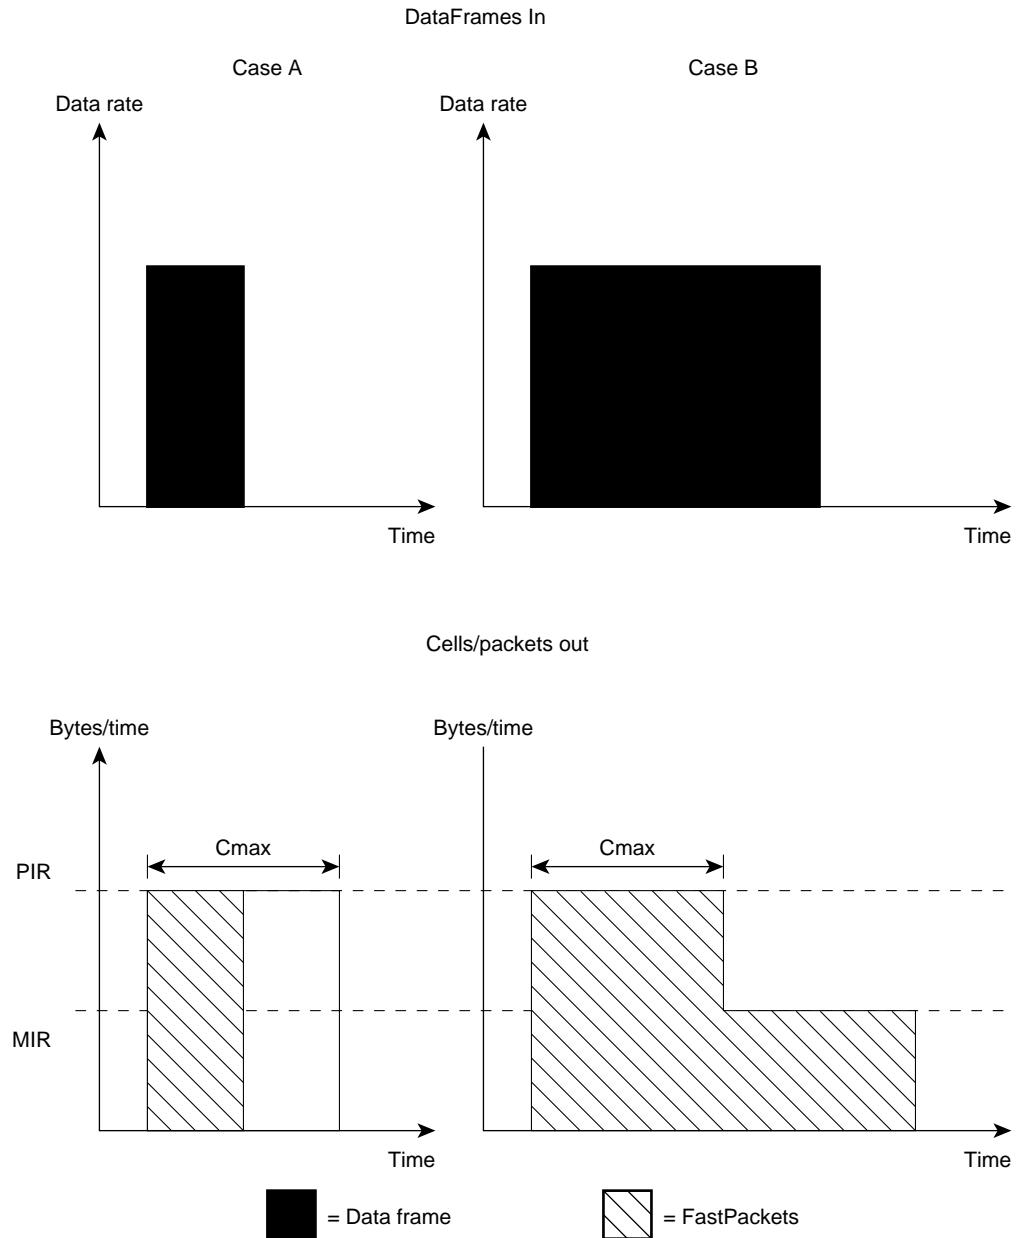
Cmax is the maximum number of credits that may be saved. It also represents, therefore, the maximum number of packets that a connection may send in rapid succession. Once the connection has used all of its available credits, it is required to cease sending packets until it has been awarded another credit.

Since frames received from the user equipment typically are broken into multiple packets, Cmax is typically set to the number of packets resulting from the average frame size. This allows a frame to be received in the FRP, packetized, and sent without incurring any unnecessary delay in the FRP. Conversely, setting Cmax to 1 limits the connection to its configured minimum bandwidth unless ForeSight is enabled on the connection.

If a connection is forced to withhold packets until it receives additional credits, space is needed to store those packets. The amount of buffer space available for this purpose is specified, in bytes, by VC_Q. The default buffer space is 64 Kbytes per connection.

Figure 12-12 shows two cases. Case A shows an example where there is sufficient bandwidth and credits available to transmit the frame of data immediately at the peak bandwidth. Case B shows an example where the frame of data is transmitted at the peak rate until credits are exhausted. Then the remaining frame of data is transmitted at the minimum guaranteed.

Figur e12-12 Bandwidth and Credits



S5221

Explicit Congestion Notification

Explicit Congestion Notification is a form of flow control to signal the onset of network congestion to external devices in an Cisco IPX narrowband frame relay network. ECN detects congestion primarily at either the source or at the destination of network permanent virtual circuits.

To be effective, external frame relay devices should respond correctly to control the rate at which the send information to the network. This feature results in data transmission at the optimum rate for the channel and reduced possibility of packet loss due to excess bursts by the external device.

With Explicit Congestion Notification, the Cisco IPX narrowband FRP card(s) signals network congestion to external devices by setting bits in the frame relay data frame header. These bits, BECN and FECN, are reserved for rate control in the forward (FECN) and backward (BECN) directions of transmission. The external devices detect the number of frames received with these BECN and/or FECN bits set and should adjust their transmitting rate up or down accordingly.

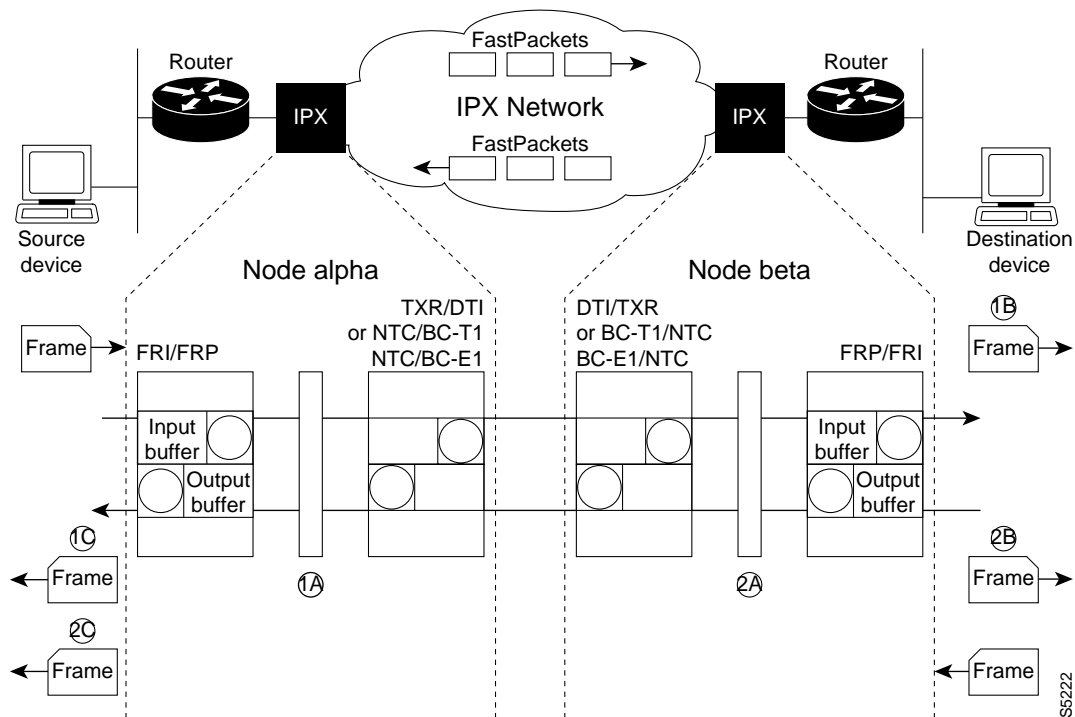
**Note**

The ECN feature requires the user device to take action to reduce the transmitted information rate applied to the network, not the network device (Cisco IPX narrowband switch).

The Explicit Congestion Notification bits may be set as a result of congestion detected at either the transmitting end of the network at the FRP in the source node or at the receiving end of the network at the FRP in the destination node. ECN does not necessarily detect congestion that may occur at intermediate nodes in the queues of the NTC or AIT trunk cards. The ForeSight feature, discussed in a later section, is an additional form of control that addresses this problem.

Refer to Figure 12-13 illustrating a simple two-node frame relay network for the following discussion. A user device, typically a router, connects to one of the four physical ports on an FRI/FRP card set in the Cisco IPX narrowband switch. At the source node, Node alpha, the FRP queues the data in a separate input buffer for each PVC. At the destination node, Node beta, the FRP queues the data to be transmitted to the user device in a single output buffer for all PVCs. (Actually, there are two buffers for each port, one for high priority data and another for low priority data as will be discussed later in this section).

Figur e12-13 Explicit Congestion Notification Example



Source ECN

Typically, network congestion occurs at the source of traffic. This can occur when the traffic being generated by the source device momentarily exceeds the network bandwidth allocated. ECN can be used at the source to relieve congestion at this point.

For this example, let's examine frames originating at the left hand side of Figure 12-13 and arriving at the destination user device on the right side of the figure. As frames are received from the source user device they are queued at the input buffer in the FRP at Node alpha. The FRP monitors the depth of the data in the input buffer for each PVC. When the frame relay data in this queue exceeds a preset threshold, VC Q ECN threshold, the FRP declares that this particular PVC is congested.

When congestion is declared, the forward direction is notified by setting the FECN bit in the data header of all outgoing data frames [1B] towards the network (alpha to beta). This will be detected by the destination user device (not FRP) and may or may not take some action to limit the data being applied to the network.

At the same time, the BECN bit is set in all outgoing data frames for this PVC towards the source device (connected to alpha) [1C] to notify equipment in the reverse (backwards) direction. The source device may be able to restrict the rate at which it transmits frames in an attempt to reduce the congestion.

Destination ECN

In a similar manner, the two ECN bits may also be set as a result of congestion detected at the destination side of an Cisco IPX narrowband network. This may result when a large number of PVCs from all over the network all terminate on a single FRP port at the destination node. For this example, let's look at what may happen at the FRP in Node beta.

As frames are received from the source user device they are queued at the output buffer in the FRP at Node beta. The FRP monitors the depth of the output buffer for each port on the FRP. When the frame relay data in this queue exceeds a preset threshold, PORT Q ECN threshold, the FRP declares that this particular port is congested.

When congestion is detected, the FRP sets all FECN bits in frames for all PVCs transmitted in the forward direction to the destination user device [1B] as well as all BECN bits in frames [2B] for all PVCs terminating at this port in the network. The net effect is approximately the same except the ECN mechanism affects all PVCs on a port at the destination whereas source ECN affects only individual PVCs that are congested.

ForeSight


ForeSight is an optional software feature that provides a closed-loop feedback mechanism for controlling the rate at which users can apply data to the network. It improves the efficiency of the Cisco IPX narrowband network for carrying frame relay data, especially during periods of light usage and maintains consistent performance during peak loading without dropping frames. ForeSight controls the data flow at the FRP, and is not dependent on the user device as is congestion prevention using FECN and BECN.

ForeSight provides congestion avoidance by monitoring the transmission of FastPackets carrying frame relay data throughout the network and adjusting the rate at which data is allowed to enter the network. ForeSight allows the FRP card to send packets at a rate that varies between a minimum and a maximum based on the state of congestion in the network along the route.

When the Cisco IPX narrowband switch receives an initial burst of data from a frame relay user device, it sends this data out on the network at a rate set by the Quiescent Information Rate (QIR) parameter as shown in Figure 12-14. This rate is usually set higher than the Committed Information Rate guaranteed the user.

The FastPacket and ATM trunk buffers used by ForeSight connections are separate from the buffers used by normal frame relay connections. If the initial FastPackets do not experience any congestion on the network, the information rate is stepped up in small increments towards a maximum set by the Peak Information Rate (PIR).

Figur e12-14 ForeSight Operation



If the FastPackets are processed by a node where there is congestion (trunk card buffers close to being full), an explicit forward congestion notification (EFCN) bit in the FastPacket header is set. When this packet is received at the destination node, the EFCN bit is examined by the destination FRP card. The far end FRP card then sends a message in the reverse direction (using the RA bit) to indicate the congestion to the near end node.

When the congestion message is received at the source FRP, the data rate is reduced in larger increments towards a Minimum Information Rate (MIR), the minimum guaranteed data rate for the connection. The FRP restricts the bandwidth allocated to the connection by reducing the rate that it issues credits to the PVC.

The connection transmit and receive data rates are controlled separately as the congestion may be present in one direction but not the other. The data rate allocated to the connection is directly controlled by the rate at which the Cisco IPX narrowband switch Credit Manager allocates credits. Without ForeSight, the credit allocation is based on MIR and credits are allocated statically.

With ForeSight the credits are allocated dynamically providing a much better utilization of network bandwidth. The net effect is to provide greater throughput for bursty connections when the network has available bandwidth yet preventing congestion on the network when the extra rate cannot be accommodated.

ForeSight can be enabled for a frame relay class. If a class is configured with ForeSight enabled, a new connection added using that class will automatically have ForeSight enabled. Once the connection is added, ForeSight can be disabled. For maximum benefit, all frame relay connections should be configured to use ForeSight when it is installed.

As part of the calculation of the ForeSight algorithm, the Cisco IPX narrowband switch measures the round trip delay for each ForeSight connection in the network. This delay is measured automatically by sending a special test packet over the connection. The far end FRP returns the packet and the round trip delay is measured this packet is received at the near end. This delay essentially consists of transmission delay as the test packet is a high priority packet type and experiences minimal processing and queuing delays. Since the network topology can change without notice (due to reroutes, etc.), this delay is measured periodically.

A Test Delay (**tstdelay**) command provides additional information about connection delay to the user. This test is performed only on command and, since normal frame relay FastPackets are used, includes the delay caused by processing and queuing throughout the connection. The delay factor is stored and available for display using either the Test Delay (**tstdelay**) command or Display Connection (**dspon**) command.

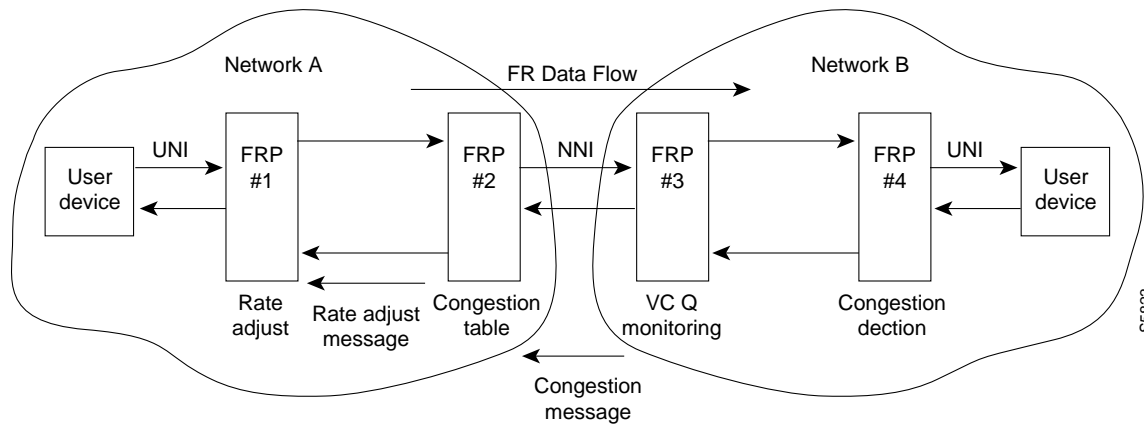
ForeSight Across Multiple Networks

ForeSight requires communicating the connection congestion status from the terminating end of a PVC back to the originating end of a PVC. But when a connection extends across two or more independent Cisco WAN switching networks, a means must be found to communicate the connection status between the networks for ForeSight to operate properly. This is accomplished in a manner similar to that just described for the Frame Relay NNI in the “Network-Network Interfaces” section earlier, except that ForeSight uses a Consolidated Link Layer Management (CLLM) message protocol at the junction port of the two independent networks.

An example, illustrated in Figure 12-15, has a frame relay connection originating in Cisco WAN switching Network A and terminating in Cisco WAN switching Network B. Congestion occurring in Network A is detected in FRP #2, which sends a congestion message back to FRP #1. Congestion occurring in Network B, however, is detected in FRP #4, which returns a congestion message back to FRP #3 at the network-to-network interface to Network B.

Periodically, FRP #3 monitors its VC queue congestion status for all the PVCs at the ingress of Network B. A message, using the CLLM protocol, carrying network status is returned to the NNI port at FRP #2. If no connection is congested, the message will still be sent across the port but there will be no connection DLCI entries listed.

Figur e12-15 ForeSight Operation Across Multiple Networks



FRP #2 builds a software congestion table listing only internetwork connections experiencing congestion in Network B. Each entry in the congestion table initiates a rate adjustment message that is transmitted across Network A to FRP #1 using a special PCC control packet, the Operation, Administration, and Maintenance (OAM) packet type. If no message is received from FRP #3 after one second, it is assumed there is no congestion on any of the PVC segments in Network B and the congestion table is cleared.

A message indicating congestion on a connection results in a downward rate adjustment and a “Other Network Congestion” indication for each PVC segment experiencing congestion in Network B. The rate adjustment is as follows:

1. If Frame Loss has occurred or the FRP #2 transmit queue is full, a fast down rate adjust signal is sent.
2. If the FRP #2 transmit queue depth exceeds the ECN Threshold, or if the last EFCN bit is set, or if an Other Network Congestion message has been received, a down rate adjust signal is sent.
3. If no congestion is experienced in either network, an up rate adjust signal is sent.

Discard Eligibility

Under conditions of network overload, it may become necessary for the frame relay service to discard frames of user data. Discard Eligibility is a mechanism defined in frame relay standards that enables a user device to mark certain frames to be discarded preferentially if this becomes necessary. This frame discard mechanism can be disabled on a port by port basis by software command.

One bit in the frame relay frame header (DE) is used to indicate frames eligible for discard. User devices may optionally set this bit in some of the frames that they generate. This indicates to the network that if it is necessary to discard frames to prevent congestion, to discard frames with the DE bit set first.

User frame relay data is buffered at the entry point to the network in a separate buffer for each frame relay port. The size of this buffer is specified when the port is defined (using the Configure Frame Relay Port (**cnffrport**) command).

When configuring a port, the network operator enters a ‘DE Threshold’ parameter, specified as a percentage of the input buffer. During normal operation, when the PVC buffer is filled above this threshold, any new frames received with DE set will be immediately discarded. Only frames not marked

DE will be accepted and queued for transmission. This has the effect of discarding frames before they are applied to the network where they may cause congestion. This function, however, is effective only if the user sets the DE bit.

Cisco WAN switching frame relay connections also employ another feature, called Internal Discard Eligibility, in an attempt to prevent network congestion. It does not depend on the user to set the DE bit.

Frames received above the CIR with the DE bit not set are monitored in a sliding window, T_c , and a counter in the FRP port. When the counter reaches a certain point, B_c , the IDE bit is set by the port before the frame is queued for transmission. Any node in the network these frames encounter congestion, they will be discarded.

Internal DE is a reserved bit in a Cisco WAN switching trailer label at the end of the last FastPacket or ATM cell for the frame. At the receiving FRP, the system can be instructed to map the IDE bit to the DE bit or not as a user option.

Cell Loss Priority

Cell Loss Priority (CLP) is a feature that allows FastPackets to be discarded in a selective manner. If the network is congested and packets must be discarded, packets with the CLP bit set will be discarded before other packets. The CLP bit is located in the control byte of the FastPacket header (refer to) or ATM cell.

The CLP bit is set by the NTC card for either of two conditions:

- For all cells/packets associated with frames where the user device has set the DE bit.
- For cells/packets transmitted above the CIR after an initial burst (same requirement for setting the IDE bit).

There are two CLP thresholds associated with NTC or AIT bursty data queues, a high threshold and a low threshold. If the high threshold in the queue is exceeded, FastPackets or cells will be discarded to prevent network congestion. They will continue to be discarded until the queue fill drops below the low threshold.

Circuit Priority

Frame relay virtual circuits can be designated as either high priority or low priority. This is used to reduce the circuit delay for delay-sensitive protocols like SNA. The Configure Frame Relay Channel Priority (**cnfchpri**) command is used to set the priority for each PVC. The priority set is communicated to the user device via LMI.

Each FRP has two output buffers for frame relay data, one for frames received with high priority, one for low priority frames. Data in the high priority frame relay buffer is unloaded and transmitted to the user device before any data in the low priority output buffer is unloaded. Since the priority basically affects the unloading of the far end FRP port queue, changing the priority requires that the Cisco IPX narrowband switch be able to communicate to the remote end of the connection.

Connection Parameters

Parameters for each frame relay connection are set by assigning a class of service (COS) to the connection. These parameters are set by either the Configure Frame Relay Connection (**cnffrcon**) or Configure Frame Relay Class of Service (**cnfcos**) commands. Parameters that define frame relay connections are listed below:

- Frame relay port (local and remote).
- Data Link Connection Identifier, or DLCI (local and remote).
- Minimum guaranteed bandwidth (transmit and receive).
- Committed Information Rate
- Peak bandwidth (transmit and receive).
- Burstiness coefficient, Cmax (transmit and receive).
- FRP buffer allocation parameter, VC_Q.
- Explicit congestion notification queue size, ECN Q.
- Percent utilization factor for the FR class, %utl.
- Quiescent information rate, QIR.
- Connection type, standard frame relay or with optional ForeSight.

There are two connection parameters that may be entered in one of two formats: StrataCom format and standard Frame Relay format. You enter which parameters are being used with the **cnfsysparm** command. Refer to Table 12-1 for the parameters that can be chosen. Note that these parameters are not numerically equal.

Table 12-1 StrataCom vs. Standard Frame Relay Parameters

StrataCom Parameters	Standard Frame Relay Parameters
Peak Information Rate (PIR)	Excess burst (Be)
VC Queue Depth (VC_Q)	Committed burst (Bc)

Minimum Information Rate, MIR—A connection's minimum bandwidth (information rate) is specified in Kbps. This is used by ForeSight algorithm and represents the lowest information rate that will be assigned when there is congestion on the network. This rate will be reached only during times there is congestion over an extended time. MIR can be set from 2.4 to 2048 Kbps.

Committed Information Rate, CIR—A connection's minimum bandwidth (information rate) is specified in Kbps. It represents the minimum bandwidth that is guaranteed to be available to the user. If a user transmits data at a rate exceeding CIR, the DE bit will be set after a time to indicate those cells may be selectively discarded. CIR can be set from 2.4 to 2048 Kbps.

Peak Information Rate, PIR—This parameter is also used by ForeSight. It is a connection's peak bandwidth (information rate) that the connection may use during data bursts when there is excess bandwidth available and no congestion on the network connection. PIR can be set from 2.4 to 2048 Kbps.

Excess burst bandwidth, Be—If frame relay standard parameters are used, excess burst bandwidth (Be) is specified in place of PIR. Be is specified in the range from 0 to 65,535 bytes. PIR is related to Be by:

$$\text{PIR} = \text{MIR} * (1 + \text{Be}/\text{Bc})$$

Burstiness Coefficient, Cmax—A connection's minimum bandwidth specifies, indirectly, the minimum number of packets per second that a connection is allowed to generate. However, a connection is allowed to exceed that minimum for short bursts. Cmax specifies, in packets, the allowed size of those bursts. This value is also used in allocating packet trunk buffer space. Cmax is a value in the range 1 to 255. If MIR is less than port speed, the larger Cmax is set, the smaller the delay in the source FRP. Increasing Cmax can have the same effect as increasing MIR. However, large Cmax can cause occasional congestion on the Cisco IPX narrowband trunks.

Buffer Allocation Parameter, VC_Q—VC_Q specifies the maximum queue size, in bytes, reserved in the FRP for the connection. As such, it sets the maximum allowable delay in the source FRP. It can have a value of from 1 to 64 Kbytes. This is a buffer where the frame is held before transmission through the credit manager. It is necessary because the external device can transmit frames to the port at the port speed while the credit manager imposes a maximum FastPacket rate on the connection.

Committed Burst Bandwidth, Bc—If frame relay standard parameters are used, the committed burst bandwidth (Bc) is specified instead of VC_Q. Bc is specified in the range from 1 to 65,535. VC_Q is related to Bc by the following:

$$VC\ Q = Bc * (1 - MIR/AR) \text{ where } AR \text{ is the port access rate.}$$

Explicit Congestion Notification queue depth, ECN_Q—This parameter sets a threshold in the input VC_Q buffer on the FRP for the connection. When exceeded, both the FECN and BECN bits are set in the frame relay frame. FECN bits are set in any frames sent into the network and BECN bits are set in any frame whose DLCI matches a connection with a congested VC_Q. ECN_Q can be set separately for both the transmit and receive directions.

Quiescent Information Rate, QIR—QIR is the initial port transmission rate utilized for ForeSight connections after a period of inactivity. The initial QIR can be set by the network administrator at some value between the Minimum Information Rate and the Peak Information Rate.

Percent Utilization factor, %util—Indicates what percent of the CIR that will actually be used. If the port is expected to be fully utilized, the Cisco IPX narrowband switch will reserve enough depth in the VC_Q to handle the traffic. If a lesser value of utilization is set, the Cisco IPX narrowband switch will reduce the VC_Q depth reserved for the connection. This, correspondingly, reduces the packets per second reserved for the port.

Typically if the port is expected to handle few, high speed inputs or if ForeSight is used on the connection, leave the % utilization factor at 100%. If there are many low speed inputs, you may want to reduce the utilization factor based on the assumption that not all connections are going to be active at any one time.

Connection Type—The connection type, (FST/fr), specifies whether the connection is to utilize the ForeSight option (FST) or if it is only a standard frame relay connection (fr). Both frame relay and ForeSight are purchased options.

Frame relay can be enabled without ForeSight in which case only congestion in the FRP queues will be detected and the bandwidth allocation for the connection is static. Additionally, if ForeSight option is enabled, congestion along the network connection will be detected and the connection's bandwidth allocation dynamic.

Configuring Connection Bandwidth



Note

The information presented here is to assist users who wish to optimize their frame relay network. This data is based on lab experiments and initial live network observation.

All parameters can use the default values from the frame relay connection class specified. The user may, however, explicitly specify any bandwidth parameters when establishing connections. The values for the minimum bandwidth (MIR), quiescent bandwidth (QIR), and maximum bandwidth (PIR) in both directions are specified in Kbps. The user can override the default for any, or all, of the bandwidth parameters in the Add Connection (addcon) command or the user can adjust any of these values after the connection is made using the Configure Frame Relay Connection (**cnffrcon**) command.

The software checks the validity of the values entered. It checks in the Configure Frame Relay Connection and Add Connection commands that the sum of the MIR is less than or equal to the line speed of the port. A warning message is generated if it is exceeded. The PIR value is checked against the line speed of the port for all connections. If this is exceeded, the user receives a warning message but the change will be made.

In selecting routes for the connections, the system software checks the bursty data buffers to see if the network will support the proposed connection. If there is not enough bandwidth, then another route will be selected or the connection will not be routed.

Data is received by the FRP from the user equipment in frames that may vary in size from 5 to 4510 bytes. Each of these frames must be broken into pieces to be encapsulated in FastPackets. The number of packets from one frame depends on the size of the frame.

When a connection is added, the Cisco IPX narrowband switch verifies that there is enough packet bandwidth to support at least the minimum information rate. To do this, the Cisco IPX narrowband switch converts the minimum information rate from Kbps to packets per second for loading considerations.

The number of packets that result from a frame is calculated as:

Integer value of [MIR/160]

For example, assume a connection with a minimum information rate of 256 Kbps.

Minimum bandwidth = 256Kbps x 1000/160 bits per FastPacket = 1600 packets/sec.

Setting MIR, QIR, and PIR for Foresight Connection

The MIR should initially be set to the required CIR for the connection and the percent utilization factor for ForeSight connections should be set to 100%. This is a very conservative setting and will give full availability to the connection(s) as well as assuring the user always receives the full CIR or higher.

In this case, the Cisco IPX narrowband switch will assign new connections to a packet trunk until the sum of the minimums equals the available bandwidth of the packet trunk. If, after time, it appears that the packet trunks are underutilized, the utilization factor may be reduced. This allows the Cisco IPX narrowband switch to assign more connections to the packet trunk.

There is an inverse relationship between overall connection delay and network congestion. An increase in MIR results in a decrease in end user delay but increases the probability of packet trunk congestion.

If minimum bandwidth is adjusted for a working connection, the system may reroute the connection. However, if the speed is changed by only a small amount, the connection will not be rerouted. Since the bandwidth used can be asymmetrical (greater in one direction than the other), the user can specify the MIR differently for each direction of a connection.

The Peak Information Rate (PIR) parameter is used to set an upper limit to the transmitted data rate when ForeSight is being used. When there is unused packet trunk bandwidth, the transmitted rate is allowed to climb to the PIR. If there are few connections on the packet trunk or if the MIRs are set at

or close to the CIR, there is less likelihood of the packet line getting oversubscribed and the PIR can be set at or near the access rate of each connection. However, if this is not the case, the PIR should be set at some level above QIR but less than the access rate.

The minimum that PIR can be set by the user is the CIR required by the circuit. The maximum is 2048 Kbps, which is the maximum port speed a FRP card permits. It is suggested that PIR be initially set to the access rate (AR) of the user device connected to the FRP port. It does little good to set PIR greater than AR as the user device will limit the maximum data rate to AR. PIR must be set to a rate greater than QIR.

Quiescent Information Rate, QIR, is used on ForeSight connections to set the initial burst rate. It must be set somewhere between MIR and PIR. ForeSight then modifies the transmit rate up or down depending on the bandwidth available. If the application consists of short bursts as with transaction processing or database inquiry, setting QIR high is useful in getting a quick user response. It is less effective when the data transmission consists of long intervals of activity.



Note

MIR, QIR, and PIR are not used for non-Foresight connections.

Setting CIR

CIR is generally specified by the usage subscribed to, or purchased, by the user. This is the guaranteed minimum data rate that is guaranteed to the user under all network conditions. If there is no congestion on the network, the user will experience higher throughput. The system uses CIR to determine the setting of the DE bit in the frame relay frame and the CLP bit in the ATM or FastPacket header.

The system calculates the Committed Burst (Bc) and Burst Duration (Tc) from the CIR, VC_Q, and Access rate for the port as follows:

$$Bc = VC_Q / [1 - (CIR/AR)] \quad \text{and} \quad Tc = Bc / CIR$$

If the user does not know what CIR to specify initially, the network administrator needs to know the port speed, the bandwidth needed by the user of the connection, and the average frame size. Determining these parameters may be difficult, especially for new systems. Generally, CIR for a non-Foresight connection should be set to a value at least 1/3 greater than the heaviest expected continuous load on the connection to prevent excessive queuing delay and data loss in the FRP. CIR should not be set greater than the port speed as this wastes network bandwidth.

After a frame relay connection has been up and running for a time, the statistics gathered by Cisco WAN Manager can assist in fine tuning frame relay connections. Statistics that should be observed to monitor this include:

- Bytes received at FRP port with DE bit set
- Frames and bytes received in excess of CIR at both FRP ingress and egress
- FastPackets/Cells transmitted to network with CLP bit set

If a user has a low percentage of traffic sent above the CIR, and a low or zero number of BECN frames the CIR may be reduced without sacrificing performance. If a high percentage of traffic sent above the CIR, the user may be able to improve application performance by subscribing to a higher CIR.

Prior to release 7.0, the system calculated CIR from MIR as follows:

$$CIR = MIR * \%utilization$$

Setting VC Q Depth

It is important to remember that the maximum data rate supported by a FRP is about 2 Mbps and can be allocated to only one port or spread over the four FRP ports. In general, setting VC Q depth to the default of 65535 is recommended for new installations. However, this may result in an excessive amount of delay if the CIR is low.

If there are many connections with relatively small CIR values originating from a port, then there is a possibility that the entire memory pool becomes allocated. A few connections could unfairly utilize memory and cause data loss on all connections. In this case, values should be set according to some suggested recommendations listed below.

1. CIR equals port speed, frame size is constant—set VC_Q = frame size.
2. CIR equals port speed, frame size varies—set VC_Q = 5X avg. frame size.
3. CIR less than port speed, frame size is constant—set VC_Q = 25X avg. frame size up to the max. of 65536 bytes.
4. Average frame size is unknown—set VC_Q = 65536 bytes.

Setting Cmax

Cmax can be set to a value of from 1 to 255 for each connection. The higher the number, the more data will be allowed on the network on an initial burst. For example, since a packet carries 20 bytes, with traffic with an average frame size of 100 bytes, a Cmax of 10 to 15 will allow one complete frame to be transmitted to the network with no delay. Settings beyond this do not reduce delay significantly.

Without ForeSight, the initial data rate is fixed at the CIR. With ForeSight, this can be adjusted by setting QIR considerably higher than CIR (MIR). If MIR is equal to port speed, Cmax can be small. If MIR is much less than the port speed, a value that can result in reducing delay is to set Cmax appropriate to generate sufficient packets to send two or three average frames in a burst.

If there are relatively few frame relay connections and the average frame size is small for most of them, it might be feasible to set high Cmax values for the connections. However, high values on too many connections originating at many different ports can cause congestion on the trunks. It is recommended that the total of the Cmax values for all connections using a given trunk not be allowed to exceed the BData Q Depth for that trunk. A suggested procedure in choosing Cmax are:

1. MIR equals port speed, frame size is constant—set Cmax = 1.
2. MIR equals port speed, frame size varies—set Cmax = 5
3. MIR less than port speed, frame size is constant—set Cmax = [(avg. frame+22) divided by 20]. This equals the number packets needed for one average frame.
4. Average frame size is unknown—set Cmax = 30 if utl is greater than 75%, 20 if utl is 50 to 75%, 10 if utl is 25 to 50%, and 5 if utl is less than 25%.

Setting ECN Q Depth

Some suggestions in choosing ECN Q depth is to set it to 1 or 2 times the mean frame size. If an approximate mean size is not available, refer to the following suggestions.

1. Mean frame size is known—set ECN Q Depth = 2X mean frame size.
2. Mean frame size can be estimated—set ECN Q Depth = 500, 1000, or 2000 for small (bytes or hundreds of bytes) frame size, medium (500 to 1000 bytes) frame size, or large (1000 to 2000 bytes) frame size, respectively.

- Mean frame size is unknown—set ECN Q Depth = 2000 if the traffic is primarily batch or 1000 if primarily transactions.

Frame Relay Port Parameters

A port on a FRP frame relay card is defined by the following parameters. These parameters can be observed using the Display Frame Relay Port (**dsfrport**) command and changed using the **cnfrport** command.

- Port location in the Cisco IPX narrowband switch (slot and port).
- Port clock speed and type.
- Port ID (DLCI) associated with the port.
- Transmit port queue depth.
- ECN queue threshold.
- DE Threshold.
- LMI mode.
- LMI protocol parameters

Port Location—This is the physical card slot where the FRP card is located in the Cisco IPX narrowband switch and the port number on the card (1 to 4) being assigned.

Clock Speed and Type—This is the data clock rate of the port (i.e. port speed) and how the port clocking is to be configured. Allowable port speed ranges from 2.4 to 2048 Kbps. Each port can be configured as either DCE or DTE. In addition the port clock can be configured normal clocking or loop timing.

Port ID—This is a Data Link Connection Identifier number assigned to the port. It may be left to the default of 0 if not using bundled connections as it is not used by the Cisco IPX narrowband switch. For bundled connections, where the user does not have to enter a specific DLCI, this number is used as the beginning DLCI of the mesh bundle.

Port Queue Depth—Is the number of bytes to allocate to the transmit port buffer for this frame relay port. The maximum size that can be allocated is 65,535 bytes. This should be sized set to some multiple of the average frame length being transmitted by the user device and the length of the bursts expected.

DE Threshold—Is the port queue discard eligibility threshold above which frames received from the user device with the DE bit set will be discarded to prevent queue overflow. Valid entries for this parameter range from 0 to 100% with reference to the capacity of the virtual circuit queue. A setting of 100% effectively disables DE for the port.

Signalling Protocol (LMI mode)—This indicates the Local Management Interface mode and protocol to be used for this port. The basic LMI can be enabled or disabled, as well as the asynchronous update process and the GMT time feature. Refer to Table 12-2 for valid entries and their description.

Table 12-2 LMI Protocols Supported

LMI =	LMI Status	Protocol Used
0	Disabled	None
1	Enabled	StrataCom LMI, asynchronous update process (unsolicited sending of update status messages) and GMT time features are enabled.
2	Disabled	None

Table 12-2 LMI Protocols Supported (continued)

LMI =	LMI Status	Protocol Used
3	Enabled	StrataCom LMI but asynchronous update process disabled.
4	Enabled	UNI uses ITU-T Q.933 Annex A parameters.
5	Enabled	UNI uses ANSI T1.617 Annex D parameters.
6	Enabled	NNI uses ITU-T Q.933 Annex A parameters.
7	Enabled	NNI uses ANSI T1.617 Annex D parameters.

LMI Protocol parameters—If LMI is enabled for the port, some or all of these various parameters may be modified to tailor the LMI to the user device attached to the port. These parameters described in Table 12-3.

Table 12-3 LMI Protocol Parameters

Parameter	Valid Entries	Descriptio
Asynchronous Status	y or n (n)	Defines if status reports should be sent asynchronously to the connecting device over the Cisco IPX narrowband switch frame relay LMI channel. If no, the Cisco IPX narrowband switch waits for a status request from the connecting device.
Polling Verification Timer	5–30 sec. (15)	Sets the interval for the keepalive timer. This should be set to 5 seconds more than the heartbeat timer in the user device.
Error Threshold	1–10 (3)	Sets the threshold for errors in the signalling protocol before an alarm is generated.
Monitored Events Count	1–10 (4)	Indicates how many events in the signalling protocol should be monitored for keepalive.
Communicate Priority	y or n (n)	Indicates if the Cisco IPX narrowband switch should communicate the PVC port SNA priority (high or low) to the user device.
Upper RNR Threshold	0–255 (75)	Sets the upper threshold (as a % of the max. VC Q depth) above which a congestion indication i sent to the user device.
Lower RNR Threshold	0–255 (25)	Sets the lower threshold (as a % of the max. VC Q depth) above which an end to congestion indication is sent to the user device.
Min. flags per frame	Any value > 1 (1)	Indicates the number of flags to expect between frame relay frames.

Setting Port Spee

In some cases, determining a reasonable port speed may be difficult. For example, if a user has eight 64 Kbps lines between a hub site and several remote sites, a first estimate for port speed would be 512 Kbps. However, if only 30% of the line capacities are used then 256 Kbps would be a more than adequate port speed.

If the approximate total bandwidth needed is known, port speeds should be set to at least 1.5 times the total bandwidth needed. Higher settings may be used to reduce delay or allow room for future growth. In cases where older technology (e.g. modems or DDS lines) are being replaced with frame relay, port speeds should be set to the total of the line bandwidths currently being used until the bandwidth needed has been determined from statistics.

For new networks where the bandwidth needs are unknown, start out with a high port speed and adjust downward based on collected Cisco IPX narrowband statistics. If the speed of the line being replaced is known, use it. If not, set the MIR and port speed the same and use a low %utl factor, i.e 100 divided by the number of connections to the port.

Setting Port Queue Parameter

The setting for port queue depth sets a maximum allowable delay in the terminating FRP. Larger settings can reduce the probability of discarded frames. In most cases, end-user delay caused by queuing in the FRP will be much less than end-user delay caused by lost data. Therefore, larger settings can reduce delay by reducing discarded frames. Unless there is good reason to change it, leave this at the default of 65,535 bytes.

If the mean frame size is known, set the ECN Q threshold to twice the mean frame size. If it is unknown, start with a value of 2000. Once FECN and BECN support has been implemented in end-user devices, the setting for ECN Q depth can be used to reduce network delay, congestion, and data loss.

Port Concentrator Shelf Frame Relay Connections

Setting up Ports and Connections

Port Command

Prior to adding connections, Port Concentrator Shelf (PCS) ports are configured from the Cisco IPX narrowband or Cisco IGX 8400 series multiband user interface with the **cnffrport** command. The parameters associated with the **cnffrport** command, along with the complete set of commands available for frame relay ports, are described in the *Command Reference Manual*.

PCS ports are specified by *<slot.port>*, where *slot* is the slot number in which the PCS-connected FRM-2 or FRP-2 card resides, and *port* is the PCS port in the range 1-44.

Connection Command

Connections are activated via commands. The commands are concerned with the activating, configuring, and reporting of statistics for frame relay connections. Each of the commands described in the frame relay chapter of the *Command Reference Manual* is supported for PCS frame relay connections.

Maximum Number of Frame Relay Connections

The maximum number of PCS frame relay connections through an FRM-2 or FRP-2 card is 252. This is an average of five connections per port, but any mix is acceptable as long as the total for all 44 PCS ports does not exceed 252.

Cisco WAN Manager Connections Manager

PCS Frame Relay connections can also be managed by the Cisco WAN Manager Connection Manager, as described in the *Cisco WAN Manager Operations publication*.

Connection Command Sequence

When not using the CiscoWAN Manager Connection Manager, the following commands are required to set up a frame relay connection:

-
- Step 1** Activate a frame relay port with the **upfrport <slot.port>** command.
 - Step 2** Use the **cnffrport** command to establish parameters for the port.
 - Step 3** Use the **dspcls** command to view existing frame classes. Select a class if a suitable class exists, otherwise create a class with the **cnffrcs** command.
 - Step 4** Access the node at the remote end of the proposed connection with the **vt** command, and again use the **upfrport** and **cnffrport** commands as in Step 1 and Step 2 above.
 - Step 5** Create the connection with the **addcon** command, specifying the class selected in Step 3 above.
-

Frame Forwarding Connections

Frame forwarding connections are specified with the **addcon** command, using an asterisk (*) for the DLCI number.

Configuring Connection Bandwidth

Existing frame relay connection bandwidth algorithms supported for FRM or FRP frame relay connections extend to PCS frame relay connections.

The software checks the validity of the values entered. It checks in the configure frame relay connections (**confrecon**) and add connection (**addcon**) commands to ensure that the sum of the MIR is less than or equal to the Port Concentrator port speed. A warning message is generated if it is exceeded. The PIR value is checked against the Port Concentrator port speed for all connections. If this is exceeded, an warning message is generated but the change will be made. There is no validity checking of the MIR or the PIR values against the concentrated link speed.

Optimizing Traffic Routing and Bandwidth

Refer to the **cnfchutl**, **cnfcos**, and **cnfpref** commands in Chapter 11 of the Command Reference manual for information on optimizing frame relay connection routing and bandwidth utilization.

PCS Network Functions

Frame Relay Interworking

Complete interworking allows connections to terminate on different types of endpoints within a network. A frame relay connection entering the network from a PCS port can terminate on:

- Another PCS Port
- An FRM or FRP Port
- A FastPAD Port
- An ASI Port
- An AIT Trunk
- A BNI Trunk

Frame Forwarding

Frame forwarding connections provide a mechanism for connecting non-frame relay (e.g. HDLC and SDLC) data. Similar to frame relay handling, HDLC frames are encapsulated into Cisco IPX narrowband frames with routing information permitting them to be terminated at any other Cisco IPX narrowband destination as listed above.

Maximum Throughput

Maximum throughput through the FRM-2 or FRP-2 card when connected to PCS is 2000 100-byte frames per second.

Maximum throughput through the PCS is 800 frames per second for each module of 11 logical ports.

Maximum Number of Connections

The maximum number of connections available per FRM-2 or FRP-2 card is 252 connections. Thus, total number of connections on all 44 ports cannot exceed this limit. This is an average of 5 connections per port, but any mix is acceptable as long as the total of 44 ports does not exceed 252 connections.

Delay

As described under “Frame Processing by the PCS”, below, additional processing of data frames is required for connections utilizing the PCS.

The amount of delay introduced depends on a large number of network characteristics. In an average application of a PCS port operating at 64K bps, utilizing 200-byte frames with the remainder of ports on the concentrated link operating at 50% capacity, the additional delay caused by a PCS at either end of a connection is in the range of 25ms more than the same connection directly through FRM-2 or FRP-2 ports.

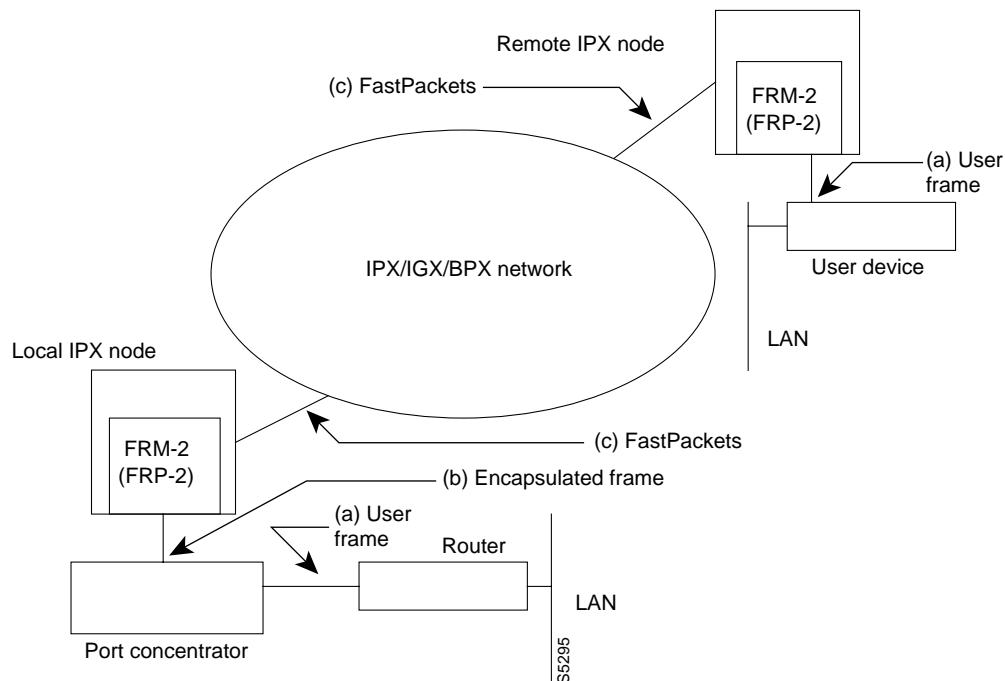
Frame Processing by the PCS

As the PCS passes frames arriving on its ports to the FRM-2 or FRP-2, the FRM-2 or FRP-2 disassembles the frame into FastPackets and passes them across a Cisco IPX narrowband /Cisco IGX 8400 series multiband network to a destination. When FastPacket frames from the network destined for a port on a PCS arrive on an FRM-2 or FRP-2, the FRM-2 or FRP-2 re-assembles the FastPackets into a complete frame to pass to the PCS. The PCS then delivers the frame to the user device.

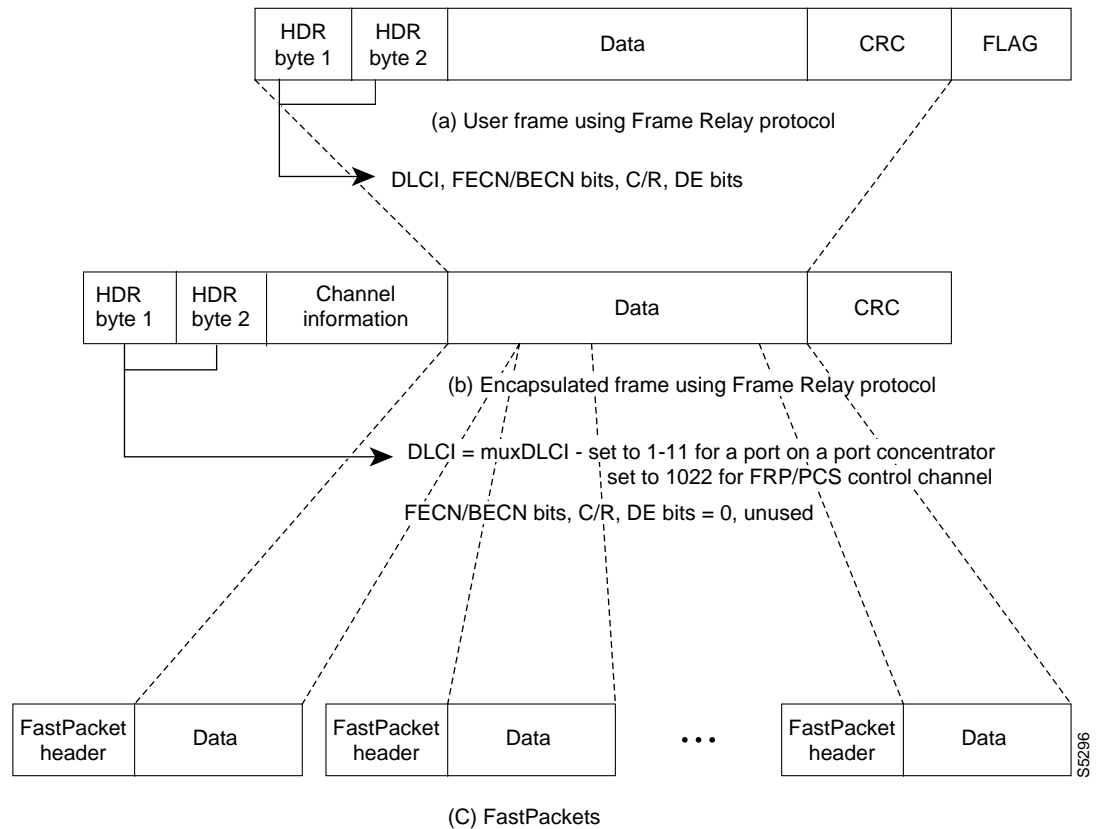
The PCS performs the minimum amount of work necessary to process incoming data frames. Besides multiplexing/demultiplexing data between the concentrated link and its 11 ports, the PCS does not handle any other frame relay features. The FRM-2 or FRP-2 performs all frame relay functions for the 44 PCS logical ports as if they were physical ports. This includes functions of permanent virtual circuit management, ingress queuing management, and signaling protocol support. Other functions, such as egress queuing management and statistics gathering are performed by the PCS.

Figure 12-16 below shows an example of a PCS frame relay connection within a Cisco IPX narrowband/Cisco IGX 8400 series multiband/Cisco BPX 8600 series broadband network. Figure 12-17 illustrates the breakdown of PCS frame relay data entry into Cisco IPX narrowband/Cisco IGX 8400 series multiband FastPackets. The frame formats at point (a), (b) and (c) shown in Figure 12-16 are described in greater detail in Figure 12-17.

Figur e12-16 PCS Frame Relay Access to an Cisco IPX Narrowband/Cisco IGX 8400 Series Multiband/Cisco BPX 8600 Series Broadband Network



Figur e12-17 Breakdown of PCS Frame Relay Data into FastPackets



Frame Relay Format

Frames arriving from a PCS port are encapsulated in a Cisco IPX narrowband frame consisting of a 4-byte header and a 2-byte CRC checksum. In the opposite direction, a frame delivered to a PCS from an FRM-2 or FRP-2 will be in a Cisco IPX narrowband frame. The PCS strips off the Cisco IPX narrowband header and CRC and passes the data out to the connected device. On the Cisco IPX narrowband header, only the DLCI field will be used. This field is referred to as "muxDLCI".

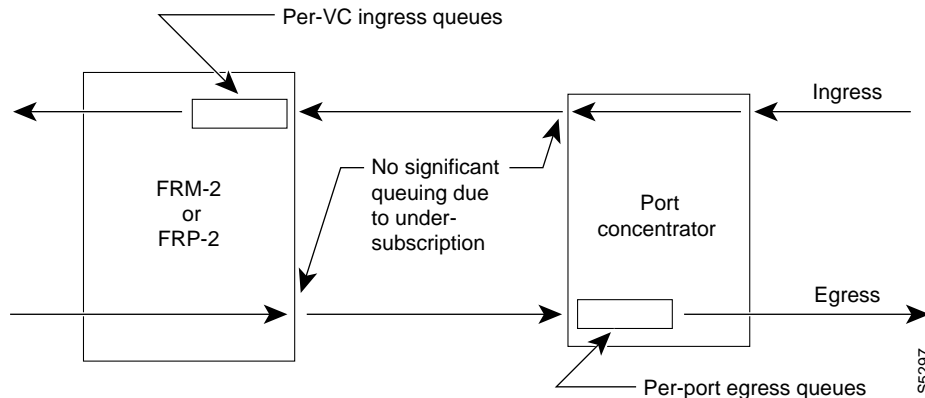
The muxDLCI specifies the port number on a PCS which is either a source or a destination of that frame. A frame containing the control message between PCS and FRM-2 or FRP-2 will have the muxDLCI set to 1022. The control information includes code download, port configuration, port status and statistics, and failure event reporting. This default DLCI of the control frame is not configurable in release 1.0.

The Cisco IPX narrowband frame format is used only on the concentrated link. The PCS appends the Cisco IPX narrowband frame header and CRC checksum when it passes frame to the FRM-2 or FRP-2. The FRM-2/FRP-2 removes the Cisco IPX narrowband frame header and CRC before constructing the FastPackets. Once the FastPackets arrive at the destination FRM-2 or FRP-2, the FRM-2 or FRP-2 reconstructs the original user data frame. The Cisco IPX narrowband header and CRC is then added before a frame is sent to the PCS. The PCS removes the Cisco IPX narrowband header and CRC and sends the frame out on the port as specified in the muxDLCI.

Frame Queuing with the Port Concentrator

Ingress queues are managed on a per-PVC basis, and are implemented on the FRM-2 or FRP-2 card. From PCS to FRM-2 or FRP-2, there are no ingress queues since concentrated link speeds are fixed at 512K bps and the maximum for each set of 11 ports must always be 448K bps or less. The Foresight egress queue is managed on a per-port basis (rather than per-connection as is done on ingress). This queue is implemented on the PCS identically to existing FRM-2 or FRP-2 egress queues (Figure 12-18).

Figure 12-18 PVC Ingress and Egress Queues



The frame header for incoming frames to the PCS contains the DLCI from the user device. This DLCI, along with the Port Concentrator port number (muxDLCI), is used to determine a unique permanent virtual connection. The FRM-2 or FRP-2 implements a separate ingress queue for each PVC. Once a frame is received by the FRM-2 or FRP-2, the frame is queued up in its appropriate PVC queue on the FRM-2 or FRP-2 based on the DLCI. The frame is then routed to the destination node.

The destination of a frame may be another FRM-2 or FRP-2 port in which case the data will be passed to the user device. If the destination of a frame is another PCS, the FRM-2 or FRP-2 card will also receive the packet. In either case, the FRM-2 or FRP-2 will determine the destination port number from the virtual circuit identifier found in the FastPacket header. The frame is reconstructed and placed in the corresponding port queue. For a frame destined for an FRM-2 or FRP-2 port, the frame is placed directly on the physical port queue and transmitted out to the user device. For the frame destined for a PCS port, the packet is placed in the logical port queue and transmitted out on a concentrated link. The PCS then passes the data out to the user device.

If the packet is destined for an FRM-2, FRP-2, or PCS port, after reconstructing the frame from the FastPackets, the FRM-2 or FRP-2 replaces the source DLCI with destination DLCI before placing the packet into the port queue.

Signalling Protoco

The following signaling protocols currently available on a normal frame relay port will be supported on the 44 frame relay ports:

- Unidirectional StrataCom LMI
- Unidirectional/Bidirectional CCITT Annex A
- Unidirectional/Bidirectional ANSI Annex D

The frame containing LMI and CLLM information and the control information between FRM-2 or FRP-2 and Port Concentrator has priority over user data frames.

ForeSight

ForeSight provides congestion avoidance by monitoring the transmission of FastPackets carrying frame relay data throughout the network and adjusting the rate at which the data is allowed to enter the network. ForeSight allows the FRM-2 or FRP-2 card to send packets at a rate that varies dynamically between the minimum information rate (MIR) and the peak information rate (PIR) based on the state of congestion.

ForeSight will be supported on a frame relay connection terminating on a PCS. The FRM-2 or FRP-2 handles the ForeSight algorithm for all 44 ports in the ingress direction. The egress direction is handled by the PCS. On ingress, since the concentrated link speed will always be higher than the total port speed of all 11 ports, there will always be bandwidth available on the link. However, when the network is congested, the PVC receive queue on the FRM-2 or FRP-2 can be full in which case some data receiving from the concentrated link will be dropped by the FRM-2 or FRP-2.

On egress, the FRM-2 or FRP-2 passes the frame to the Port Concentrator at the concentrated link speed. The Port Concentrator maintains the per port egress queue for ForeSight. The Port Concentrator reports egress congestion status to the FRM-2 or FRP-2. The FRM-2 or FRP-2 forwards the congestion status to the remote end of the connection using standard ForeSight mechanisms. The Port Concentrator sets egress FECN and ingress BECN bits as necessary based on egress congestion levels. The Port Concentrator also drops DE frames and queue overflow frames as necessary based on egress congestion levels. The related channel and port statistics are also reported to the FRM-2 or FRP-2.

Frame Forwarding

Frame forwarding connections provide a mechanism for connecting non-frame relay frames (HDLC and SDLC). The frame forwarding on a Port Concentrator is implemented in the same way as a normal frame relay connection. When an HDLC frame arrives on a port, the Port Concentrator simply encapsulates the frame in a Cisco IPX narrowband header and passes it to the FRM-2 or FRP-2 along with the muxDLCI. The FRM-2 or FRP-2 uses the muxDLCI to determine the PVC and route the packet to the destination node. The FRM-2 or FRP-2 at the destination node determines the destination port and places the frame in the appropriated logical port queue.

ForeSight is supported on frame forwarding connections. In addition to supporting frame forwarding on PCS connections, frame forwarding is also supported on PCS-to-FRM/FRP and PCS-to-FastPAD interworking connections.

Bandwidth and Routin

Network routing decisions within the Cisco IPX narrowband/Cisco IGX 8400 series multiband/Cisco BPX 8600 series broadband network for a PCS frame relay connection are the same as those for other frame relay connections. The rate parameter specified for connections, together with the utilization parameter for the channels are used to determine the LU (Load Unit) for the Cisco IPX narrowband segment of a connection. The frame relay connection is mapped to the BData-A queue for the non-ForeSight case and to the BData-B queue for the Foresight case. Courtesy downing is also available on a PCS frame relay connection.

Local connections from one PCS port to another are routed to the FRM-2 or FRP-2 card and back. Local switching is not supported by the PCS.

Frame Relay to ATM Network and Service Interworking

This chapter describes frame relay to ATM interworking which enables frame relay traffic to be connected across high-speed ATM trunks using ATM standard Network and Service Interworking.

The chapter contains the following:

- Interworking
- Service Interworking
- Networking Interworking
- ATM Protocol Stack
- AIT/BTM Interworking and the ATM Protocol Stack
- AIT/BTM Control Mapping, Frames and Cells
- Management, OAM Cells
- Functional Description
- Management

Interworking

Interworking allows users to retain their existing Frame Relay services, and as their needs expand, migrate to the higher bandwidth capabilities provided by Cisco BPX 8600 series broadband 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.

Two types of Frame Relay to ATM interworking are supported, Network Interworking (Figure 13-1) and Service Interworking (Figure 13-2). The Network Interworking function is performed by the AIT card on the Cisco IPX narrowband switch and by the BTM card on the Cisco IGX 8400 series multiband switch, and the FRSM card on the Cisco MGX 8220 edge concentrator. The FRSM card on the Cisco MGX 8220 edge concentrator and the UFM cards on the Cisco IGX 8400 series multiband switch also support Service Interworking. See Figure 13-3 for some examples of ATM to Frame Relay Interworking.

Network Interworking: Part A of Figure 13-1 shows typical frame relay to network interworking. In this example, a frame relay connection is transported across an ATM network, and the interworking function is performed by both ends of the ATM network. The following are typical configurations:

- Cisco IPX narrowband frame relay (shelf/feeder) to Cisco IPX narrowband frame relay (either routing node or shelf/feeder)
- Cisco MGX 8220 edge concentrator frame relay to Cisco MGX 8220 edge concentrator frame relay
- Cisco MGX 8220 edge concentrator frame relay to Cisco IPX narrowband frame relay (either routing node or shelf/feeder)

Service Interworking: Part B of Figure 13-1 shows a form of network interworking where the interworking function is performed by only one end of the ATM network, and the CPE connected to the other end of the network must itself perform the appropriate service specific convergence sublayer function. The following are example configurations:

- Cisco IPX narrowband frame relay (either routing node or shelf/feeder) to Cisco BPX 8600 series broadband or Cisco MGX 8220 edge concentrator ATM port
- Cisco MGX 8220 edge concentrator frame relay to Cisco BPX 8600 series broadband or Cisco MGX 8220 edge concentrator ATM port.


**Note**

In the items listed above, a Cisco IGX 8400 series multiband switch can be substituted for each instance of an Cisco IPX narrowband switch.

Figur e13-1 Frame Relay to ATM Network Interworking


Part A

Network interworking connection from CPE Frame Relay port to CPE Frame Relay port across an ATM Network with the interworking function performed by both ends of the network.




Part B

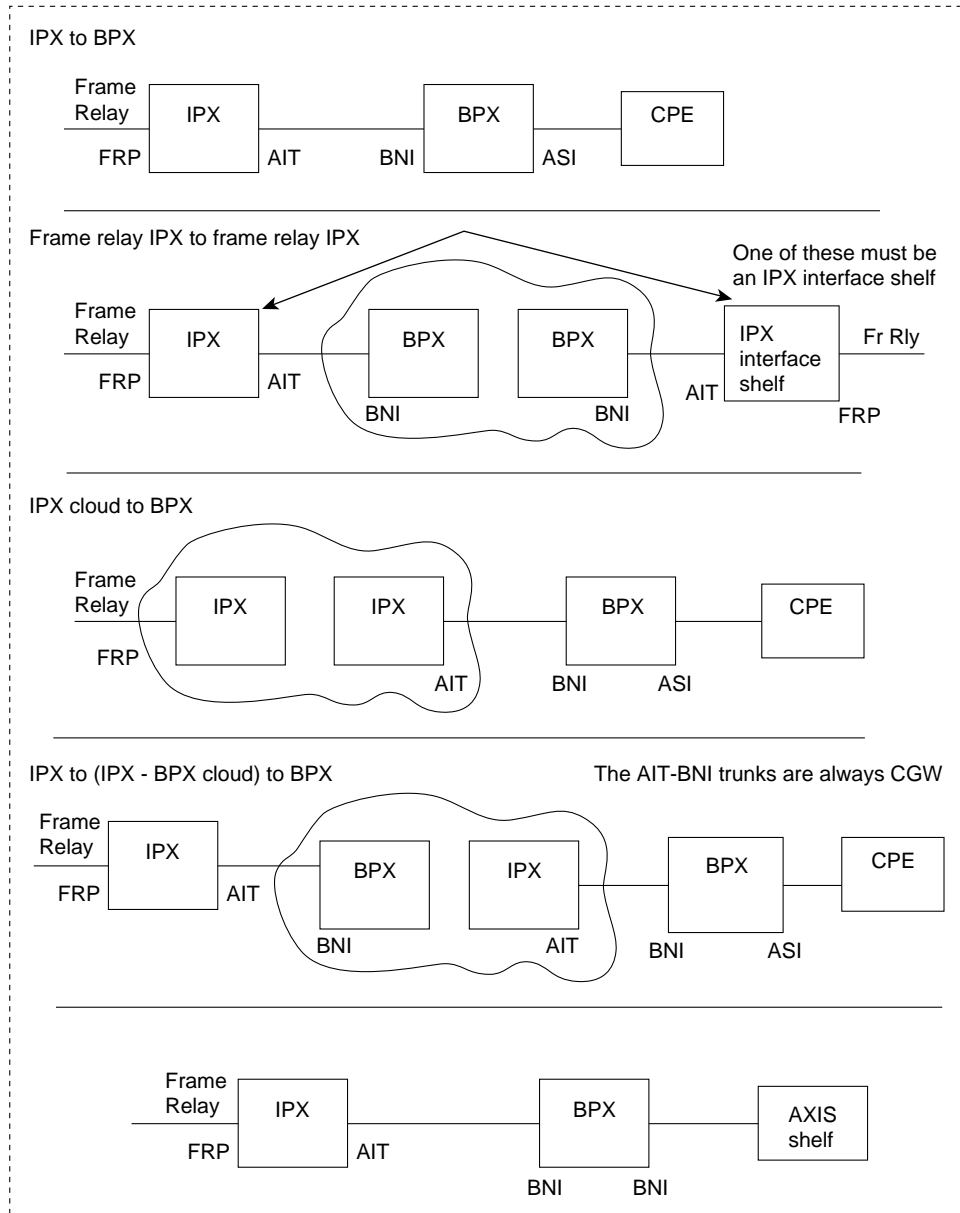
Network interworking connection from CPE Frame Relay port to CPE ATM port across an ATM network, where the network performs an interworking function only at the Frame Relay end of the network. The CPE receiving and transmitting ATM cells at its ATM port is responsible for exercising the applicable service specific convergence sublayer, in this case, (FR-SSCS).



Figur e13-2 Frame Relay to ATM Service Interworking



Figur e13-3 Frame Relay to ATM Interworking Examples with AIT Card on a Cisco IPX Narrowband Switch



AIT Interworking Examples

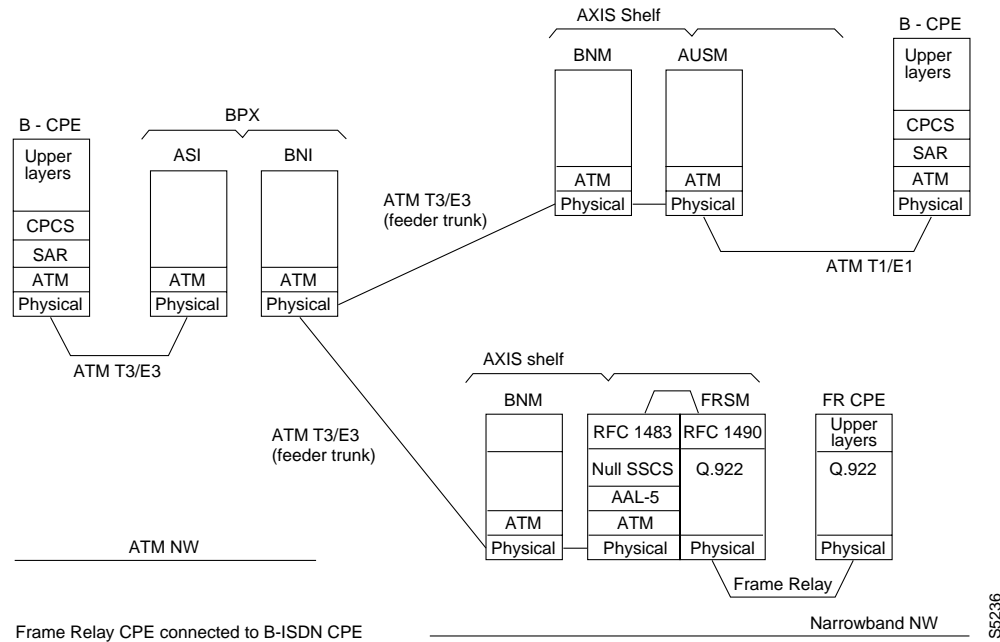
55239

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 13-4).

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 Cisco MGX 8220 edge concentrator shelf, e.g., ASI, AUSM. Translation between the Frame Relay and ATM protocols is performed in accordance with RFC 1490 and RFC 1483.

Figur e13-4 Frame Relay to ATM Service Interworking Detail



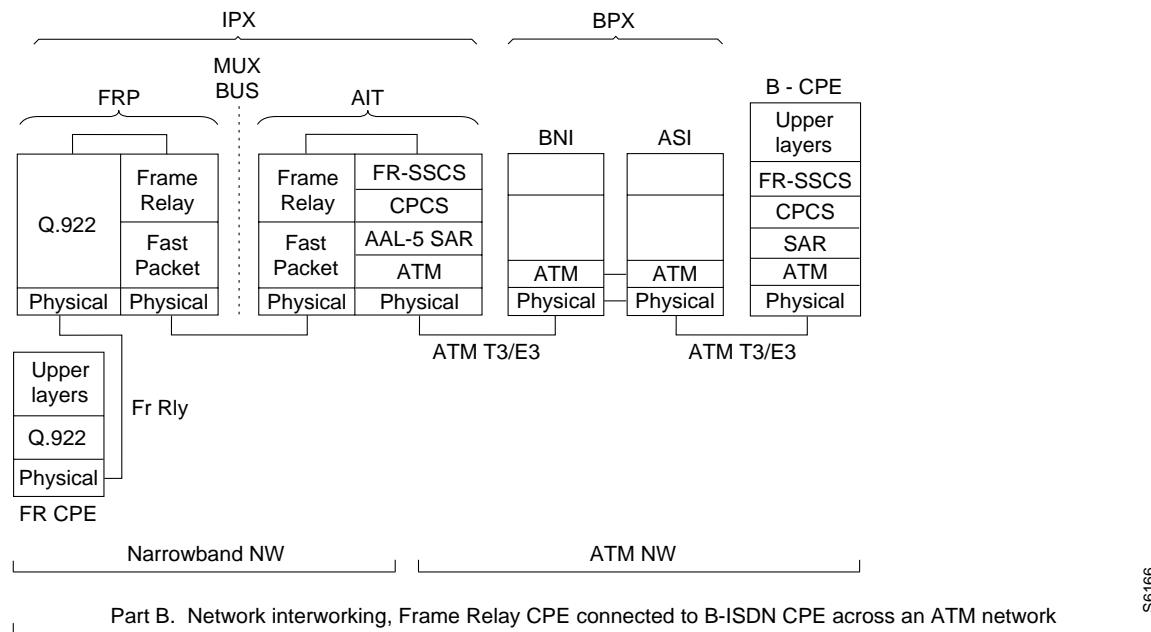
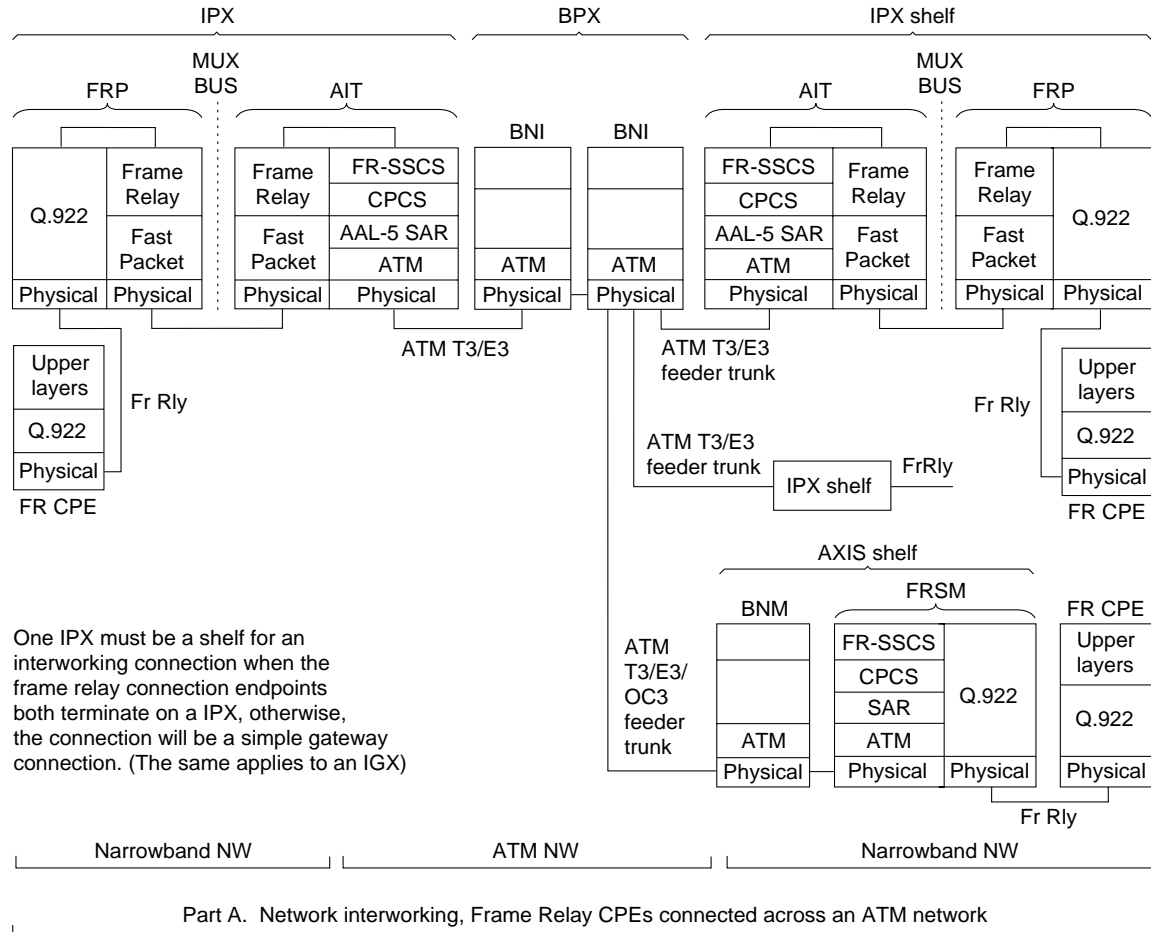
55236

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 13-5.

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 13-5.

Figur e13-5 Frame Relay to ATM NW Interworking Detail



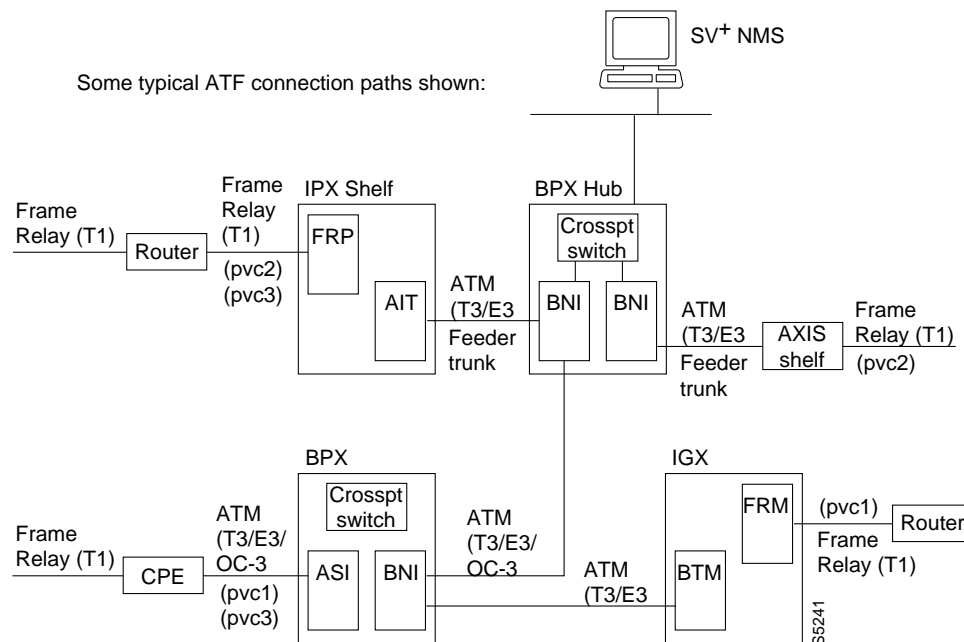
S6166

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

- Cisco IPX narrowband frame relay (shelf/feeder) to Cisco IPX narrowband frame relay (either routing node or shelf/feeder)
- Cisco MGX 8220 edge concentrator frame relay to Cisco MGX 8220 edge concentrator frame relay
- Cisco MGX 8220 edge concentrator frame relay to Cisco IPX narrowband frame relay (either routing node or shelf/feeder)
- Cisco IPX narrowband frame relay (either routing node or shelf/feeder) to Cisco BPX 8600 series broadband or Cisco MGX 8220 edge concentrator ATM port
- Cisco MGX 8220 edge concentrator frame relay to Cisco BPX 8600 series broadband or Cisco MGX 8220 edge concentrator ATM port.
- In the items listed above, a Cisco IGX 8400 series multiband switch can be substituted for each instance of a Cisco IPX narrowband switch.

On the Cisco IPX narrowband switch, interworking is performed by the AIT card, and on the Cisco IGX 8400 series multiband switch by the BTM card. A simplified example of the connection paths is shown in Figure 13-6. 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 Cisco BPX 8600 series broadband to Cisco IPX narrowband hop in a Frame Relay to ATM connection's path.

Figur e13-6 ATF Connections, Simplified Example



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 Cisco 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 a Cisco IPX narrowband switch/Cisco IGX 8400 series multiband switch, 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).

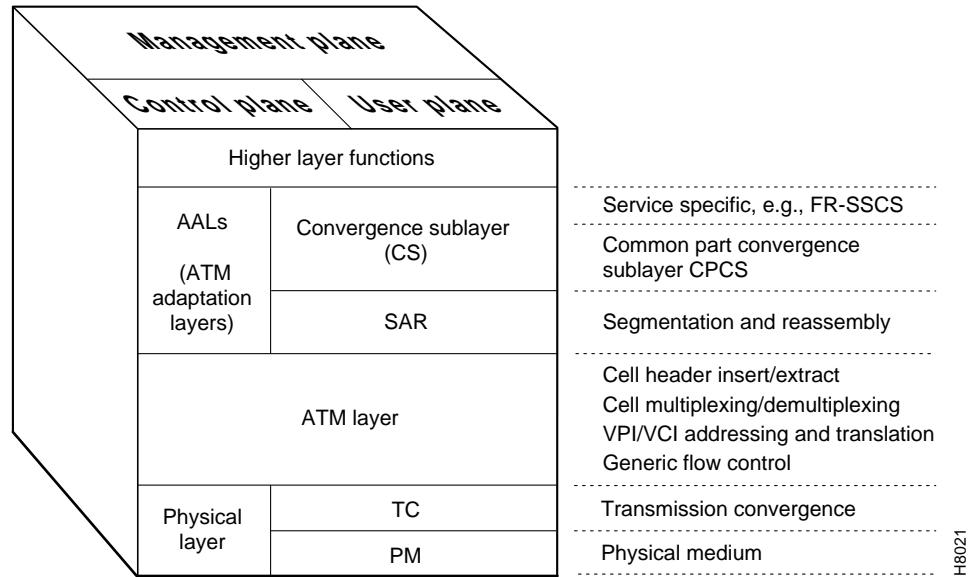
For each connection, only a single dlci can be played out for all traffic exiting the port, and is inserted into the frame headers. The standard LAPD framing format is played out the port on the FRP/FRM.

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 13-7. 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.

Figur e13-7 ATM Layers



AIT/BTM Interworking and the ATM Protocol Stack

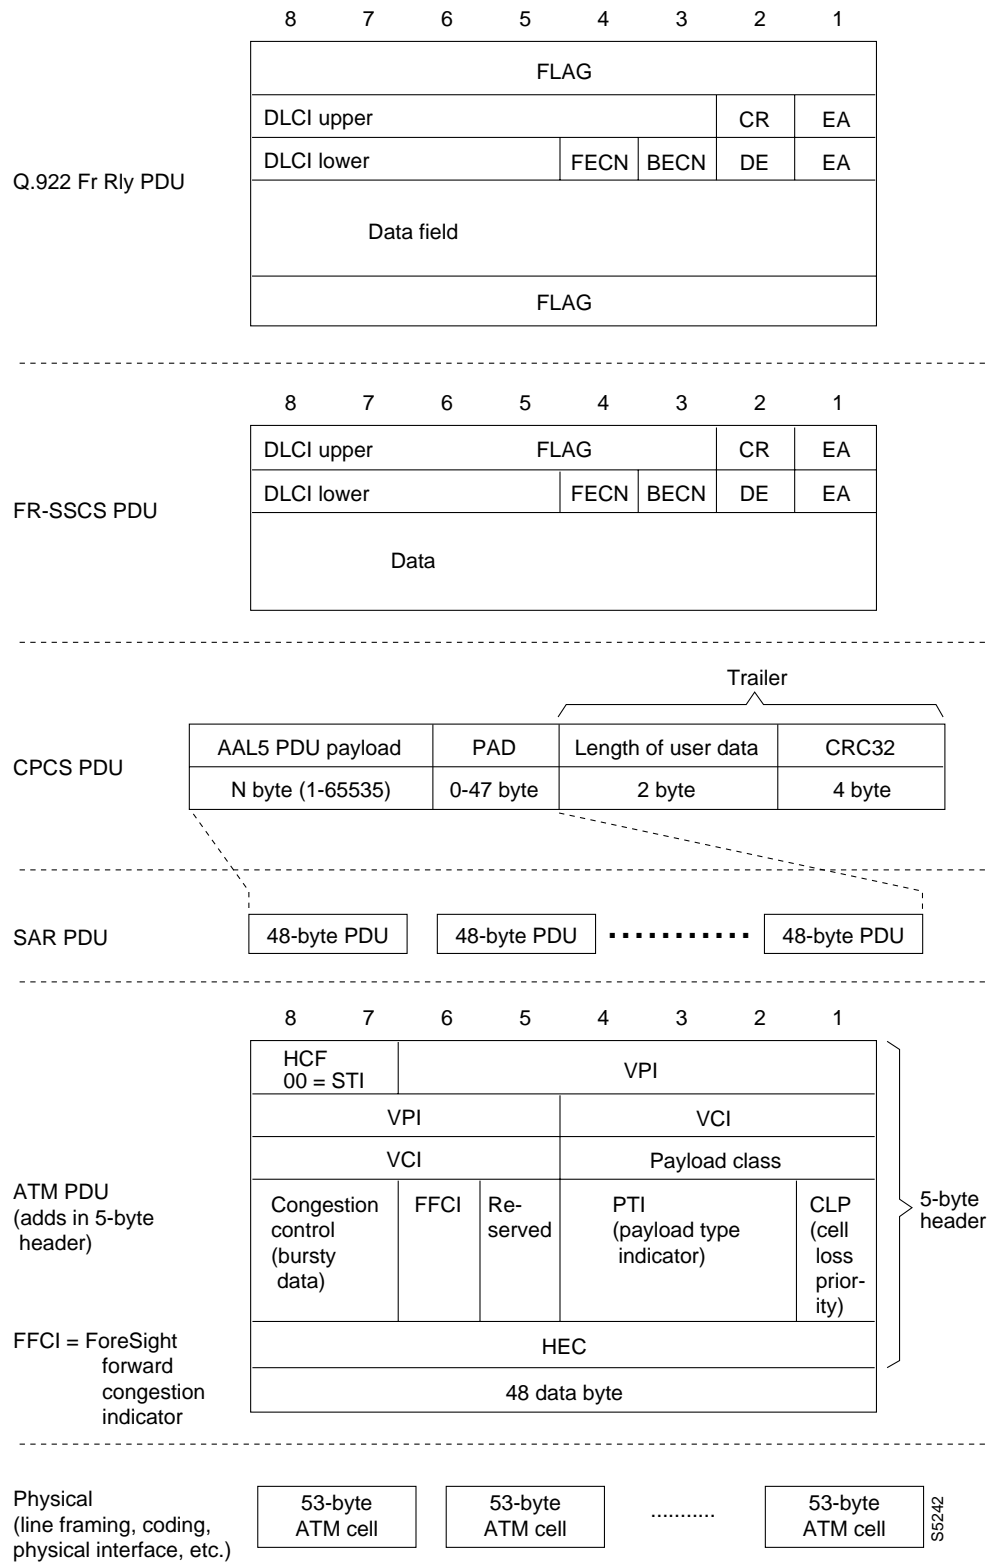
ATM to Frame Relay interworking (ATF) performs various tasks including the following:

- 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 event

Figure 13-8 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 Cisco IPX narrowband switch/Cisco IGX 8400 series multiband switch 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 Cisco IPX narrowband switch, incoming frames are converted to FastPackets by the FRP card. The FastPackets are then routed to the AIT card via the Cisco IPX narrowband switch 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 FRSSCS [or FRCS in shortened notation] layer (Frame Relay Service Specific Convergence Sublayer) 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.

Figur e13-8 Protocol Stack Operation



AIT/BTM Control Mapping, Frames and Cells

In addition to performing DLCI to VPI/VCI conversion, the network interworking feature provided by the AIT card in the Cisco IPX narrowband switch or by the BTM in the Cisco IGX 8400 series multiband switch maps cell loss priority, congestion information, and management information between frame relay and ATM formats as follows:

CELL LOSS PRIORITY, Frame Relay to ATM Directio

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 Cisco IPX narrowband/Cisco IGX 8400 series multiband NIW (network interworking):

- CLP is always 0
- CLP is always 1

CELL LOSS PRIORITY, ATM to Frame Relay Directio

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 directio

- EFCI is always set to 0.

CONGESTION INDICATION, ATM to Frame Relay Directi

- 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 Managemen

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
- Cisco 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.

Limitation

- 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 Cisco BPX 8600 series broadband endpoints; it is supported at Cisco IPX narrowband switch endpoints
- Gateway terminated inter-domain connections
- Via connections through Cisco IPX narrowband switches

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 Cisco IPX narrowband switch or the Cisco BPX 8600 series broadband switch 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 Cisco BPX 8600 series broadband switch, and bits per second (bps) on the Cisco IPX narrowband switch/Cisco IGX 8400 series multiband switch. 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**
- **dspon**
- **dspon**

Port Management



Note

ELMI is an enhancement to LMI. ELMI adds capabilities that are not currently supported in LMI so that network switches, e.g., Cisco BPX 8600 series broadband switch, Cisco IGX 8400 series multiband switch, etc., can inform a user (routers, bridges, etc.) about network parameters such as various quality of service (QoS) parameters. Depending on the implementation, these might be such parameters as Committed Information Rate (CIR), Committed Burst Size (Bc), Excess Burst Size (Be), maximum Frame Size, etc. ELMI is supported by the Cisco IGX 8400 series multiband UFM-U and UFM-C cards.

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 CiscoWAN Manager or by the Cisco BPX 8600 series broadband 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 Cisco BPX 8600 series broadband switch. 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 Cisco BPX 8600 series broadband switch is modified to include the ASI-1 card.
- The card background test that exists for the FRP card on the Cisco IPX narrowband switch 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 nondestructive, as compared with the Cisco IPX narrowband switch `tstdly`.

User Command

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:

- 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 Cisco 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 Cisco IPX narrowband switch. 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 Command

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

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

User Command

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 Cisco BPX 8600 series broadband switch, the Cisco IPX narrowband switch, the Cisco IGX 8400 series multiband switch, or the Cisco MGX 8220 edge concentrator. Intra- and inter-domain interworking connections are supported.

Connection configuration parameters are endpoint-specific. Thus, the ATM-only parameters are only configurable on the Cisco BPX 8600 series broadband switch end. The Cisco IPX narrowband switch does not know about these parameters, so they cannot be configured or displayed at the Cisco IPX narrowband switch end. Parameter units are endpoint-specific also. Units on the Cisco BPX 8600 series broadband switch are cells per second, units on the Cisco IPX narrowband switch 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 Cisco BPX 8600 series broadband switch-Cisco IPX narrowband switch hops these connections route over must provide the complex gateway function. Cisco IPX narrowband switch-Cisco IPX narrowband switch 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 cell 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 Cisco BPX 8600 series broadband switches. Cisco IPX narrowband switches 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 **'cnfstparm'** 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.

PART 5

DATA AND VOICE CONNECTIONS

Synchronous Data Connections

This chapter is provided for users who wish to have an in-depth knowledge of the Cisco IPX narrowband switch/Cisco IGX 8400 series multiband switch synchronous data connections and related functions. It describes the basic flow of information through the network for various data connection types. It also describes the data compression features, data clocking, and data channel conditioning available.

The chapter contains the following:

- Data Connection Types
- Data Block Diagram Signal Flow
- Data Control Leads
- Data Compression
- Data Clocking

Data Connection Types

The types of synchronous data connections available include:

- High-speed synchronous data channels.
- Low-speed synchronous/asynchronous data channels.
- Subrate and superrate channelized DS0 connections.

Synchronous Data Connections (SDP)

Point-to-point high-speed synchronous data connections interface to the Cisco IPX narrowband switch on Synchronous Data PAD (SDP) cards that transmit at rates from 1.2 Kbps to 1.024 Mbps. These connection types are used for direct computer-to-computer transfer of large files. Up to four ports per card are supplied interfacing data, clocking, and control leads.

Low-speed Data Connections (LDP)

For more efficient transmission of multiple low-speed data circuits (up to 19.2 Kbps), the Low-speed Data PAD card is used. It can interface both synchronous as well as asynchronous data, typically from data terminals, point of sale terminals, etc. A special version of the LDP is used to interface with common carrier 56 Kbps DDS circuits in the US. Four or eight ports per card are supplied interfacing data, clocking, and control leads.

Channelized Data Connections

Channelized data connections using T1 or E1 connections are interfaced to the Cisco IPX narrowband switch by using the Channelized Data PAD (CDP) card. These connections use individual or multiple DS0's to provide data connections operating at rates from 56 Kbps to 512 Kbps, and provide a direct interface with user's data submultiplexers, voice/data multiplexers, and other data devices that operate over T1 or E1 lines

Channelized data connections fall into two categories, **subrate** and **superrate**. Subrate data connections consist of multiple low-speed data logical connections sub-multiplexed into a single 64 Kbps DS0. Common multiplex standards include DS0A and DS0B for T1 and X.50 for CEPT E1. DS0 subrate data rates cover the range of 2.4, 4.8, 9.6 and 56 Kbps. Cisco IPX narrowband systems with a CDP card will accommodate one subrate data channel per DS0 channel and DS0A signalling.

A superrate data connection (sometimes called bundled connection) consists of a single data circuit that spans multiple DS0s, up to a maximum of eight. Superrate connections are available as aggregates of 56 Kbps or 64 Kbps up to 512 Kbps on a CDP E1 or T1 port. All bits transported are considered data bits, because no supervision or control bits are imbedded in the data stream. Superrate 56 Kbps bit streams are bit stuffed up to 64 Kbps by the external multiplexer by inserting a "1" in the first or last bit position (user selectable) of the data byte. The CDP extracts and discards this bit for packet transmission and reinserts it at the far end Nx56 connection.

With superrate data connections, the timeslot and port assignments are flexible. Channels are specified independently at both endpoints when adding these connections to the Cisco IPX narrowband switch and must be balanced in number (for example a superrate connection transmitting with 4 DS0s must terminate with four channels but the channels do not have to have the same destination). The channel range may be contiguous or alternating but not randomly assigned within the T1 or E1 frame.

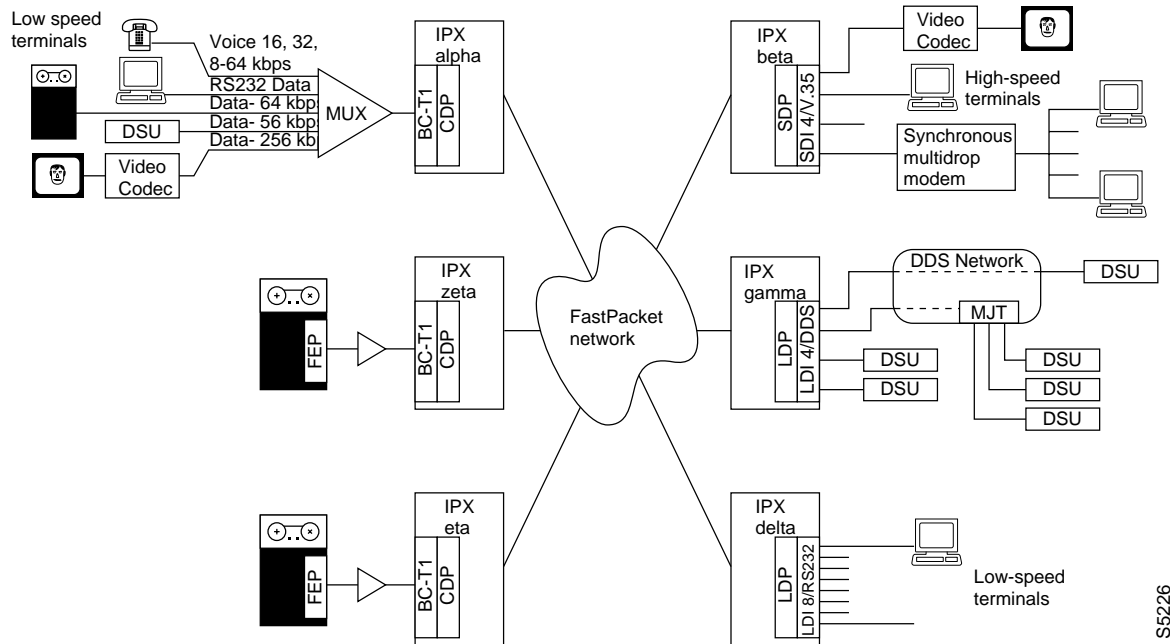
There are two transmission modes for subrate data connections: interpretive and transparent. In the transparent mode, the Cisco IPX narrowband switch does not look for any supervisory or control information in the data stream; it assumes all bits are data bits. The interpretive mode, on the other hand, expects that the data stream from the customer includes some supervisory bits and the CDP must do some processing to extract this supervisory information.

CDP to CDP connections always run in the transparent mode. CDP to SDP or LDP run in the interpretive mode as these channels provide separate control lead inputs and outputs. Supervisory signals in the DS0A mode are distinguished from data signals in interpretive mode connections. Some of the features available with the SDP/SDI and LDP/LDI cards are not available on the CDP with the subrate or superrate connections. These include isochronous clocking, fast EIA, and Data Frame Multiplexing.

Figure 14-1 illustrates a hypothetical Cisco IPX narrowband switch network with six nodes, all of which are equipped with a CDP, to illustrate the various types of data connections possible. At nodes zeta and eta, there are two mainframe computers that need to exchange bulk data. Each is equipped with a front end processor for high-speed data transfer with its own high-level protocol. Each has a DS1 interface to the BC-T1 back card associated with the CDP.

If the full bandwidth of the DS1 is not required for this computer-to-computer link, groups of DS0s can be bundled together to provide Nx56 or Nx64 Kbps virtual circuits on each CDP up to 512 Kbps each. This requirement for only part of the full DS1 could be accommodated by a fractional T1 circuit interface.

Figur e14-1 Examples of CDP Data Connections



At Node alpha, the CDP is serving as the Cisco IPX narrowband interface to a customer's multiplexer which, in turn, is fed by a variety of sources: low-speed (RS232) terminals, high-speed computer links, synchronous DDS data (from a DSU), and high-speed synchronous data from a video codec. Node alpha is capable of supporting connections with all of the other nodes.

The CDP supports channelized data connections. For example, node alpha routes 256 Kbps data from a video codec to node beta where the data is terminated by an SDP with an appropriate SDI back card. The 256 Kbps data connection is presented to the CDP by the multiplexer as a bundle of four DS0 timeslots, yet it terminates on an SDP as a serial, synchronous V.35 data port.

Note that to support clear channel Nx64 Kbps connections such as this, the multiplexer feeding node alpha must enforce the one's density requirement with B8ZS (T1) or HDB3 (E1) coding. Otherwise, the connection could be carried as Nx56 Kbps with the multiplexer providing the bit stuffing. Or one could use alternating 64 Kbps channels with ones density requirements met in unused channels.

Likewise, the packet trunks connecting node alpha and node beta must guarantee transparency to packets that require either B8ZS trunk lines or some form of transparency encoding. This can be either 7/8 ths or 8/8 ths-inverted encoding. The 7/8 ths coding stuffs an extra one bit into each byte. The CDP will discard these extra bits prior to packetization. Note that the 56 K bit stuffing at the port interface is separate and independent of the 7/8 bit stuffing on the packet trunk.

Another example of data connectivity is the connection of the computer at node alpha with the synchronous terminal at node beta. The remote end is terminated by another SDP data port and connected to a cluster of synchronous data terminals via a multidrop modem.

An example of interconnectivity with the DDS network is presented between node alpha and node gamma. The DDS connection is presented to the CDP as a single DS0 timeslot carrying 56 Kbps of useful information. This illustrates a typical “tail” DDS circuit. The LDP4/DDS at node gamma is used to connect to the DDS network with two of its four ports as well as provide other local DSU connections.

Another data connection illustrated is the low-speed (RS232) asynchronous data connection. This enters node alpha as Sub-Rate Data Multiplexed (SRDM) timeslots, in which a single DS0 can carry one or more low-speed channels. The CDP DS0A currently carry only a single data channel regardless of operating speed. The LDP at node delta terminates these bundled data connections as individual low-speed data ports. This obviates the need for a separate multiplexer at node delta as is used at node alpha.

Interface to DDS Network

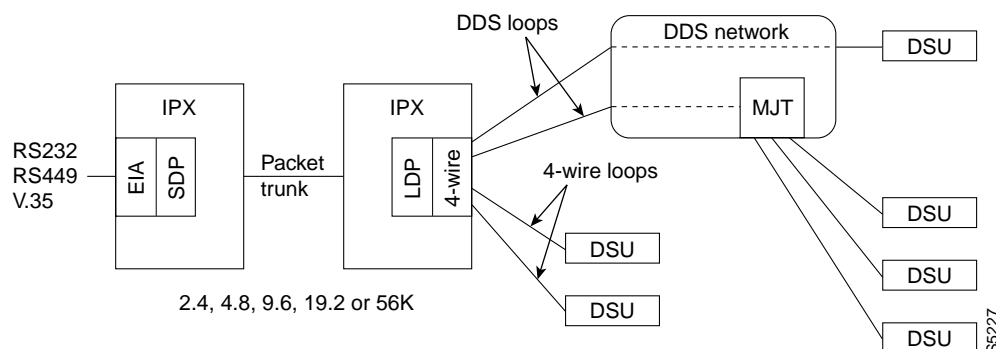
Digital Data Service (DDS) circuits can be obtained from AT&T and other common carriers to provide domestic customers with nationwide data connections over switched public T1 networks. These DDS trunks operate synchronously at rates up to 56 Kbps and have imbedded maintenance and configuration commands available. These DDS circuits can be used to extend data service to users who are not served directly by a Cisco IPX narrowband network.

The Cisco IPX narrowband switch provides a direct interface to the AT&T DDS network using the LDI4/DDS back card and LDP, Model B or later front card. Each LDI/DDS provides four DDS ports, each of which can be configured to act as either a Data Service Unit (DSU) or an Office Channel Unit (OCU). The LDI is software configured to operate at a selected synchronous data rate of between 2.4 and 56 Kbps.

When an LDI/DDS port is configured as a DSU, it provides a 4-wire DDS circuit connecting point. The port acts like a DTE using external clocking since the LDI must be synchronous with the DDS network timing. This circuit terminates on a OCU inside the serving central office providing the access point to the DDS network.

This is illustrated in the top two circuits in Figure 14-2. The DDS circuit may be a point-to-point connection (#1) from OCU to OCU or a multipoint connection, using a Multiple Joint Terminator (MJT), which connects several terminating circuits together.

Figure 14-2 Typical DDS Network Configurations



An LDI/DDS port can also be configured as an Office Channel Unit to terminate a local data user over a DDS-type data circuit. This circuit connects point-to-point to a remote DSU and associated user device over privately-owned non-loaded lines. This is illustrated in the bottom two circuits of

Figure 14-2. In the OCU mode, the port acts like a DCE and loops the timing back to the local DSU. The DSU permits the user to be located several miles from the Cisco IPX narrowband switch and provides DDS-type loopbacks for testing.

For example, in Figure 14-2, a remote user of node beta may come into the Cisco IPX narrowband switch on port #3 then be connected to the DDS network off port #1 to communicate with another user who may be thousands of miles away. The allowable distance from the Cisco IPX narrowband switch to the remote DSU ranges between approximately 3 to 30 miles depending on operating bit rate and cable gauge of the telephone cable.

An alternate method for a user located adjacent to an Cisco IPX narrowband switch to access the DDS network is by using a SDP front card and associated SDI back card at the local node. The user device must be located within the nominal distance restrictions of RS-232, RS-449, or V.35. The packet network, in turn, interfaces with the public switched network at this or any other node using a 56 Kbps circuit. The SDP must be configured to supply timing to the local user device as it must ultimately be synchronized to the DDS network.

This is illustrated in the circuit between the two Cisco IPX narrowband switches in Figure 14-2. An originating user device, colocated with a Cisco IPX narrowband switch, interfaces to node alpha using the SDP/SDI. At node beta, there is an interface to a DDS office using an LDP and LDI/DDS operating as a data service unit. The Cisco IPX narrowband switch network is configured with a 2.4, 4.8, 9.6, or 56 Kbps data circuit between the two nodes.

Data Block Diagram Signal Flow

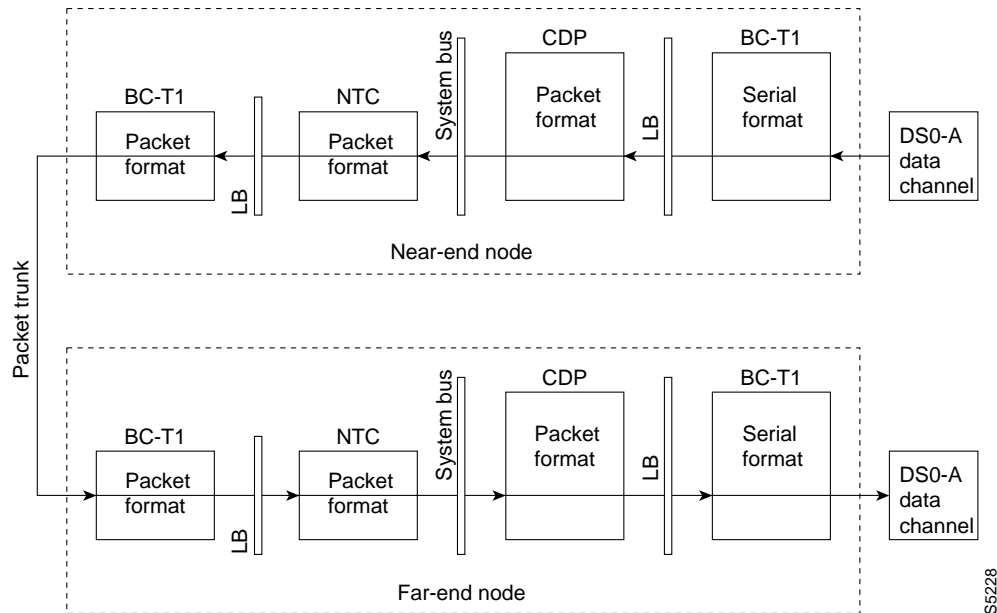
The following paragraphs describe the block diagram signal flow for the various synchronous data circuits through the Cisco IPX narrowband switch network.

CDP Data Signal Flow

Data connections processed by the CDP terminate on the BC-E1 or BC-T1 backcard (CDP does not operate with a BC-SR, SDI, or LDI backcard) which stores DS0 timeslots separately in a receive frame buffer. A Local Bus (LB) is required between the CDP card and the back card. When a complete frame of data has been received, it is forwarded to the CDP for processing (Figure 14-3).

The processing performed in the CDP depends on the circuit specifications entered by the user and may include bit stuffing, data inversion, or timeslot resequencing. For subrate channels, the data processing may also include extracting supervisory bits, signalling conversion, and detection and synchronization to the subrate frame pattern.

Figur e14-3 Data Circuit Flow with CDP



The CDP demultiplexes the subrate frame pattern and extracts individual channels. Each subrate data channel is processed and transported as separate logical channels in the Cisco IPX narrowband switch network. The CDP can handle a maximum of 32 individual logical connections.

The processed data is packetized, queued for transmission, and then written to the system bus. Supervisory data bits are packetized separately from user data and applied to the system bus. These packets are pulled from the system bus by either an NTC if the data destination resides on another Cisco IPX narrowband switch node or another CDP for an intranode connection.

At the receiving node, packets for superrate (bundled) connections are extracted from the system bus and buffered. Higher rate channels (64 Kbps and above) in general do not use timestamped packets and, as such, are buffered and pass directly to signal processing. The processing is the inverse to that done to the data on the transmitting end. After processing, the data is immediately assembled into a frame, along with samples from other connection types, and sent to the associated back card where the frames are played out into corresponding timeslots on the circuit line to the user device.

Subrate channels are lower speed and are often timestamped. These packets are transferred through the receiving end CDP according to the timestamped packet synchronizing algorithm with buffering to compensate for variable network delays. After processing, the individual subrate channels are multiplexed into a single DS0 timeslot. A data frame is accumulated then and to the back card.

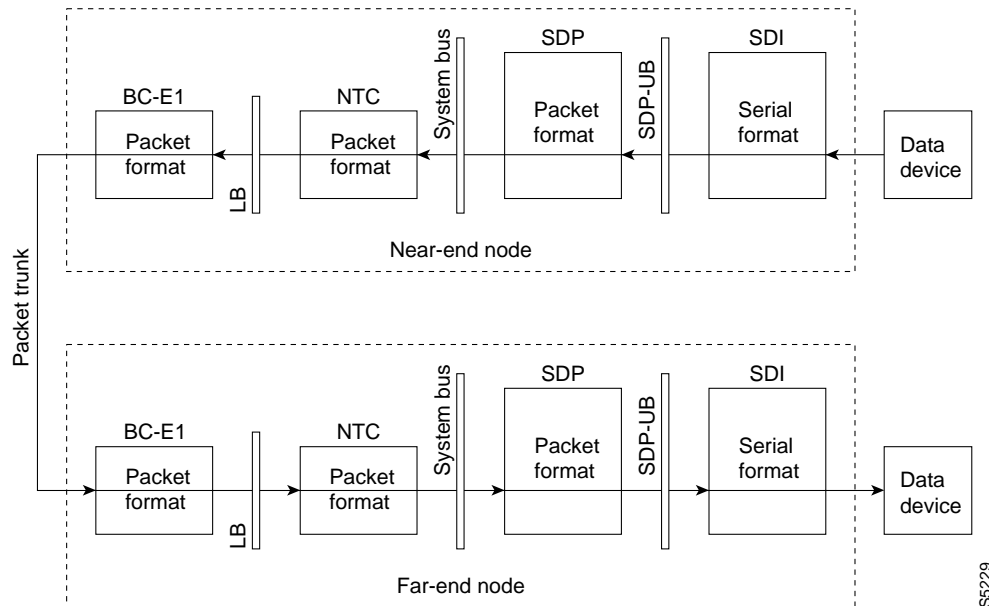
Data connections may also be read off the system bus by a SDP or LDP depending on the connection type. The signal flow is essentially the same except the supervisory packets can be translated into control lead status at the corresponding SDI or LDI back card.

SDP/LDP Synchronous Data Channel Signal Flow

Interface to the Cisco IPX narrowband switch for circuits from customer synchronous data devices use the SDP or LDP card and associated SDI and LDI interface card. Customer data and associated control lead information from various devices are transmitted to the Cisco IPX narrowband switch on data lines

with standard data interfaces (for example RS232, V.35, RS449/X.21). The SDP card is connected to the back card via a SDP Utility Bus (SDP-UB). The LDP card is connected to the back card via a Local Bus (LB). Figure 14-4 illustrates the signal flow diagram for synchronous data.

Figur e14-4 Synchronous Data Flow



Depending on the user data device and its operating speed, data will be received by a Synchronous Data Interface (SDI) card, as shown in Figure 14-4, a Low-Speed Data Interface (LDI) card, or a Frame Relay Interface (FRI) card. The data back card provides the physical interface.

A plug-in daughter board allows the back card to present DCE or DTE interfaces independently on each channel. Clocks may either be taken from the user's equipment (at either end of the connection) or provided by the Cisco IPX narrowband switch as a submultiple of its system clock.

Data interface cards all perform a similar function of monitoring the transmit data and control leads, serial to parallel conversion of the data, level shifting, and, in the case of balanced data leads, converting to unbalanced logic signals. The data is forwarded to one of the Data Packet Assembler/Disassembler (PAD) or service interface cards over a utility bus connecting the interface card to the PAD card.

Again, depending on data type, the Data PAD card may be one of the following: Synchronous Data PAD (SDP) or Low-Speed Data PAD (LDP). All input data is treated as synchronous data. If it is low-speed, asynchronous data, the LDP oversamples it to make it synchronous (the bit rate is generally set to five times the asynchronous data rate).

The Data PAD cards provide transmit clock, either normal, looped, or split, to the external data device to clock in the data. The SDP buffers the data to permit isochronous operation. Repetitive patterns in the incoming data are suppressed with a feature called Data Frame Multiplexing (DFM) and assembled into packets of up to 168 bits with a destination header.

The packets are applied to specific timeslots on the TDM MUXBUS portion of the system bus where they are picked off by a NTC or AIT trunk card associated with the proper destination code as previously described for voice packets. From this point on, the data packets are treated just like voice packets. The trunk cards do not mix voice and data in a packet queue. There are separate queues for voice and the various types of data.

Generally, a packet is not sent until its payload is full. But since low-speed data does not fill up its queue very rapidly, it is quite likely that there might be long delays in sending the data. Because of this, low-speed data circuits can be assigned a timestamp that is used to limit the delay in sending these packet types.

At the far end, the flow is the reverse. The packet is received by an NTC card, then sent via the system bus to the receiving data card. Here, it is read off the bus into a receive buffer, allowing for clock and transmission fluctuations, and then clocked out as serial data on the back card connector.

The Data PAD card is a physical-layer device, and relies on the user's equipment to handle errors and retransmission. Packet headers are checked for CRC in the NTC card, and corrupted packets discarded, but the data stream out of the Cisco IPX narrowband switch may contain errors or missing bits. The data PAD does not detect these conditions or retransmit to correct them.

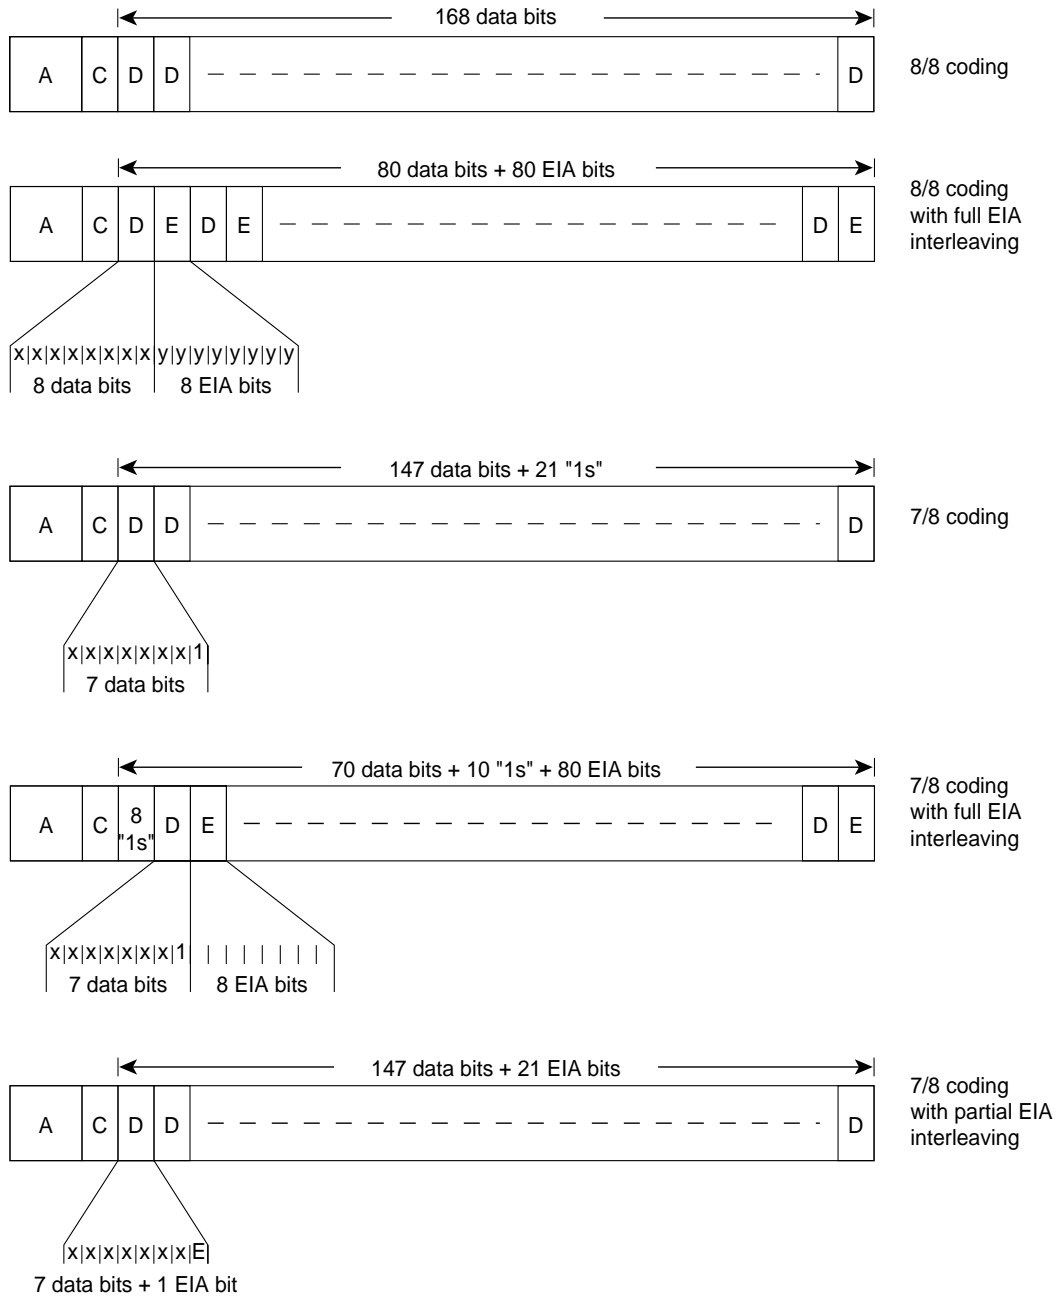
Data Control Leads

Data channel control lead status is transmitted through Cisco IPX narrowband switch networks in one of three data channel signalling connection types:

- Normal (non -interleaved) connections—a separate packet is generated for the control leads and is transmitted in a FastPacket separate from the FastPacket carrying the data.
- Fast EIA (interleaved) connections—the data and seven control leads are carried in the same FastPacket but in alternating, 8-bit bytes. The remaining control leads are carried as in the non-interleaved connection.
- Partially interleaved connections—seven data bits and one select control lead bit are carried in the same 8-bit byte in the FastPacket. The remaining control leads are carried as in the non-interleaved connection.

Refer to description of FastPacket data packets in the “Non-Timestamped Data Packets” section of **Chapter 9, FastPackets and Narrowband Trunks**, and Figure 14-5 illustrating these three modes. The data back cards allow up to 12 data control leads per data circuit to be sampled and transmitted to the far end of the connection. The choice is dependent on DCE-DTE options and software configuration.

Figur e14-5 Various non-Time Stamped Data Formats



S6201

Non-Interleaved Control Leads

For **non-interleaved connection** the Cisco IPX narrowband switch treats up to twelve control leads impartially, any signal received on a control lead pin at one end may be transmitted to any pin at the other end either directly or with translation. Therefore, it is possible to have different interfaces for the same connection (for example V.35 to RS449). Table 14-1 shows equivalent names and pins for RS-232, RS449 and V.35.

Table 14-1 Data Control Lead Equivalency

Src	Name	V.35	RS 9	Pin	RS 232C	Pin	RS 232D	Pin	Functi	Fast EIA
DTE	TXD	P/S	SD	4/22	BA	2	BA	2	XMT Data from DTE	no
DCE	RxD	R/T	RD	6/24	BB	3	BB	3	RCV Data to DTE	no
DTE	RTS	C	RS	7/25	CA	4	CA	4	Request to Send	F4
DCE	CTS	D	CS	9/27	CB	5	CB	5	Clear to Send	F4
DCE	DSR	E	DM	11/29	CC	6	CC	6	Data Set Ready	F3
DCE	DCD	F	RR	13/31	CF	8	CF	8	Carrier Detect (RLSD)	F7
DCE						9		9	Positive Test Voltage	no
DCE						10		10	Negative Test Voltage	no
DCE	QM					11		11	Equalizer Mode	no
DTE	pin 11					11		11	Sometimes used for data	no
DCE	SDCD		SRR		SCF	12	SCF	12	Sec. Carrier Detect	no
DCE	SCTS		SCS		SCB	13	SCB	13	Sec. Clear to Send	no
DTE	STxD		SSD		SBA	14	SBA	14	Sec. XMT Data	F5
DTE	NS		NS	34		14		14	New Sync	F7
DCE	TxC	Y/a	ST	5/23	DB	15	DB	15	XMT Clock	no
DCE	SRxD		SRD		SBB	16	SBB	16	Sec. RCV Dat	F5
DCE	RxC	V/X	RT	8/26	DD	17	DD	17	RCV Clock	no
DCE	DCR					18			Divided RCV Clock	no
DTE	SRTS		SRS		SCA	19	SCA	19	Sec. Request To Send	no
DTE	DTR	H	TR	12/30	CD	20	CD	20	Data Terminal Ready	F3
DCE	SQ		SQ	33	CG	21	CG	21	Signal Quality Detect	no
DCE	RI	J	IC	15	CE	22	CD	22	Ring Indicator	F2
DTE	SF		SR,S F	16	CH	23	CH	23	Signal Rate Select	no
DCE	SI		SI	2	CI	23	CI	23	Signal Rate Indication	no
DTE	XTC	U/W	TT	17/35	DA	24	DA	24	External XMT Clock	no
DTE	BSY		IS	28		25			Busy (In Service)	F1
DTE	LL		LL	10					Local Loopback	F2
DTE	RL		RL	14			RL	18	Remote Loopback	F6
DTE	TM	K	TM	18					Test Mode	F6

Table 14-1 Data Control Lead Equivalency (continued)

Src	Name	V.35	RS 9	Pin	RS 232C	Pin	RS 232D	Pin	Function	Fast EIA
DTE	SS		SS	32					Select Standby	no
DCE	SB		SB	36			TST	25	Standby Indicator	F1

In the non-interleaved mode, the leads are sampled at a rate set by the operator. If there has been a change since the last sample, the front card assembles a supervisory packet of the same type as the connection's data packets and transmits it to the far end. There the packet is processed by the data front card and the new states of the leads latched onto the outputs.

Because the EIA packetize path is separate from the data path and control leads are sampled every 50 msec. maximum (20 times/sec), the synchronization between control lead and data changes is accurate only to the update interval +50 msec. (100 to 1050 msec).

The non-interleaved mode is the most efficient as far as packet trunk bandwidth utilization is concerned. FastPackets carrying control lead information are sent only sporadically and the payload has plenty of room for control lead status from a number of data channels. However, the delay between far end receiving data and the corresponding control lead change of state at the transmitting end is quite large.

Interleaved (Fast EIA) Control Leads

When **interleaved ("Fast") EI** connections, are specified, up to seven leads may be updated at the interleaved rate. The control lead status change is tied very closely to the data (within 1 byte interval). Combinations of control leads for both Cisco IPX narrowband switch data channels as DCE or DTE are listed in Table 14-1 in the Fast EIA column. The remaining control leads may be sent as non-interleaved at 0–20 updates/sec. With interleaved EIA connections, there is a trade-off between the control lead update rate and packet trunk bandwidth as 50 percent of the FastPacket payload is taken up with control lead bytes.

Any input pin designated F(n) in Table 14-1 may be mapped to an output pin with the same F(n) up to a maximum of 8 pins in the "Fast EIA" column. With fast EIA connections, only some combinations of EIA lead correspondence are carried by hardware and synchronized to the data:

- Local output RI can follow remote input DTR.
- Local output CTS can follow remote input DTR.
- Local output DCD can follow remote input DTR.
- Local output DCD can follow remote input RTS.
- Other combinations shown in Table 14-1.

A sixth control lead (RL) may be configured on an LDP in order to support the use of three control lead outputs on a DTE connection. This is required for V.24 and RS232 circuits that use two leads, local loopback (LL) and remote loopback (RL), for data circuit loopback control. To support these leads, either RTS or DTR must be given up, in addition to using the currently unused pin 11 DTE output. With the sixth EIA lead available, the LDP supports three control leads in DCE mode (RTS, DTR, and RL) and three leads in DTE mode (CTS, DSR, and DCD).

The SDP/LDP supports secondary channel features by treating the signals as EIA leads. These signals are available for interleaved EIA connections, so the maximum rate is from 5 percent to 2.5 percent of the primary connection, depending on configuration.

Partially Interleaved (Embedded) Control Leads

The **partially interleaved** or **embedded EIA** mode is a compromise between the non-interleaved and interleaved modes and is used for DFM data. The partially interleaved mode allows one bidirectional control lead (RTS if DCE or CTS if DTE) to be encoded as the eighth bit in each data word, providing quick control lead response without significantly affecting bandwidth requirements. The remaining control leads go as separate packets as in non-interleaved mode.

This mode is used for data 19.2 Kbps and under and is not available for DDS trunks. The data must be specified as 7/8 coding to reserve the eighth bit for the selected control lead and requires zero suppressed trunks using B8ZS or HDB3 coding. Table 14-2 lists the functions that the LDP Embedded EIA mode supports.

Table 14-2 LDP Embedded EIA Mode Features

Features Support	Features Not Supported
Bidirectional RTS/CTS leads.	Delayed RTS/CTS loopback feature is not available.
Control lead tracks the data within 24 bit intervals.	Connections with 7 or 16 bit pattern matching is not supported.
Data rates up to 19.2 Kbps.	8/8 bit coding is not supported.
DFM bit pattern matching.	DDS trunks.
7/8 coding.	

Control Lead Conditioning

Control leads may be held high or low or change to follow any control lead input from the far end of the connection. Also, conditioning of the control leads by the Cisco IPX narrowband switch can be specified for when the connection has failed or is looped for maintenance. Conditioning is specified by the Interface Control Template (ICT) in software and is set by the user.

It is also possible to configure local lead changes depending on local inputs. This is done in hardware, so there is almost no delay. The valid configurations for local conditioning are any control lead to an output control lead. When remote CTS is controlled by the local RTS the user may program a delay of from 1 to 255 milliseconds.

Data Compression

Data compression, sometimes called Data Frame Multiplexing (DFM), is used in the SDP, LDP, and FRP to reduce the bandwidth requirements for most types of low-speed data connections. It is not currently employed by the CDP for subrate or superrate data.



Note

Data Frame Multiplexing is an optional feature that must be purchased and enabled on each node.

For DFM to be applied to a connection, four conditions must hold as follows:

- The cards at each end must support DFM (SDP/LDP's Revision BA or later, all new LDPs).
- The data rate must not exceed 128 Kbps or 64 Kbps (Rel. 7.0 and earlier) and fast EIA mode not selected.
- The channel information at each end of the connection must have DFM enabled.
- The connection type must allow DFM (timestamped packets, no fast EIA).

DFM relies on a repetitive pattern suppression algorithm. Data packets where DFM is specified contain blocks of 152 consecutive bits of user data. The algorithm searches for packets where a pattern is repeated and suppresses transmission of these packets. Since the receiver must be able to reconstruct the data of the suppressed packets, it is programmed to repeat the last 7, 8 or 16 bits of the previous packet received. Therefore, DFM requires a pattern of 7, 8 or 16 bits to be repeated for at least 140 bits (dependent on connection type) before a packet can be suppressed and bandwidth savings realized.

The expected pattern length must be set in software for DFM to work correctly. It should be set to the number of bits in a character of the user's protocol, since repetitive idle codes are the patterns normally suppressed. Normally, both directions of a connection use the same protocol, so the pattern length in each direction should be the same. In nearly all cases, the pattern length should be 8 bits.

The transmitting front card uses the sequence field in the packet to indicate the position the packet should be placed in the receive buffer. The receiver can tell from the sequence how many packets have been suppressed since the last received, and therefore reconstruct the data stream.

Even if the data stream is entirely repetitive, every sixteenth packet is transmitted to allow the receiver to synchronize to the sequence number and keep the connection alive. Therefore, the minimum utilization of a DFM connection is 6 to 7 percent.

DFM does not introduce more delay to data connections. However, the receive buffer is deeper than without DFM, so if tail circuit clocking is incorrectly configured, the buffer can fill up resulting in an increase in the end-to-end data delay. The increased delay can cause high-level protocols to fail, where previously they had masked the wrong clock configuration by retransmissions.

Data Clocking

Synchronous data transmission requires that the data timing be transmitted along with the data bits. Synchronous data transmission is used almost exclusively in the data communications industry for almost all connections except the slowest data rates.

Synchronous Data Clock Source

In general, there are two sources of data channel clocking referenced to the Cisco IPX narrowband switch:

- internal—the Cisco IPX narrowband switch supplies timing to the external user data device.
- external—data timing is supplied to the port by the external user data device.

When the clock source is internal, the timing signal (clock) that the Cisco IPX narrowband switch supplies to the customer data terminal is obtained by dividing the network clock down to the desired data rate. The data device must be able to accept clock from the communications device (most newer equipment is designed to do this). This is the normal synchronous data channel transmission and is generally used for all data rates of 56 Kbps and above.

When the clock source is external to the Cisco IPX narrowband switch, the customer data device supplies transmit timing along with the data to the Cisco IPX narrowband switch. This timing is sent through the network. The receive clock from the Cisco IPX narrowband switch at the far end of the connection is a reproduction of the transmit clock from the external data device at the near end.

In the data world this type of timing is often called isochronous clocking. Isochronous clocking is generally used for data rates of 112 Kbps or less. Data channel timing in this configuration is carried in the control lead packets. If a connection is configured to accept a clock and for some reason there are no clock transitions, no packets will be sent.

With isochronous timing, the front card local to the clock source calculates the actual frequency and sends a supervisory packet to the far end, containing adjustments for the baud rate generator. This allows the clock to be propagated through the Cisco IPX narrowband switch (though with some delay). If the clock changes dramatically, such as after a modem fallback, the data will be held to mark for a few seconds while the far end clock tracks to the new rate. The transmitting card generates an “overspeed clock” alarm when the incoming clock is 2 percent greater than the nominal rate. At this stage, the connection will fail and the clock at the far end will lose synchronization.

**Note**

Many field problems arise because external equipment that is supposed to synchronize with the Cisco IPX narrowband switch-provided clock does not do so, but free-runs. This can result in data corruption and buffer overruns, which may go undetected for some time because the high-level protocol masks errors by retransmission.

Data Clock Configurations

The Cisco IPX narrowband switch supports a number of data channel clocking configurations. Each port is configured independently. The SDI and LDI data back cards can be configured to operate as either a DCE or as a DTE.

Normally, the Cisco IPX narrowband switch data card serves as the DCE and the user data device as the DTE. In this mode, the Cisco IPX narrowband switch provides data circuit transmit and receive timing to the user data device to synchronize the flow of data. This was described previously as internal timing.

If necessary, the SDI and LDI cards can be configured to operate as a DTE for when the data device is a modem instead of a terminal device. In this mode, the Cisco IPX narrowband switch receives timing from the data device. This was described previously as external timing.

After specifying whether the Cisco IPX narrowband switch data card is to be a DCE or DTE, the data card is software configured to operate in one of three data clocking configurations as follows:

- normal clock.
- split clock.
- loop clock.

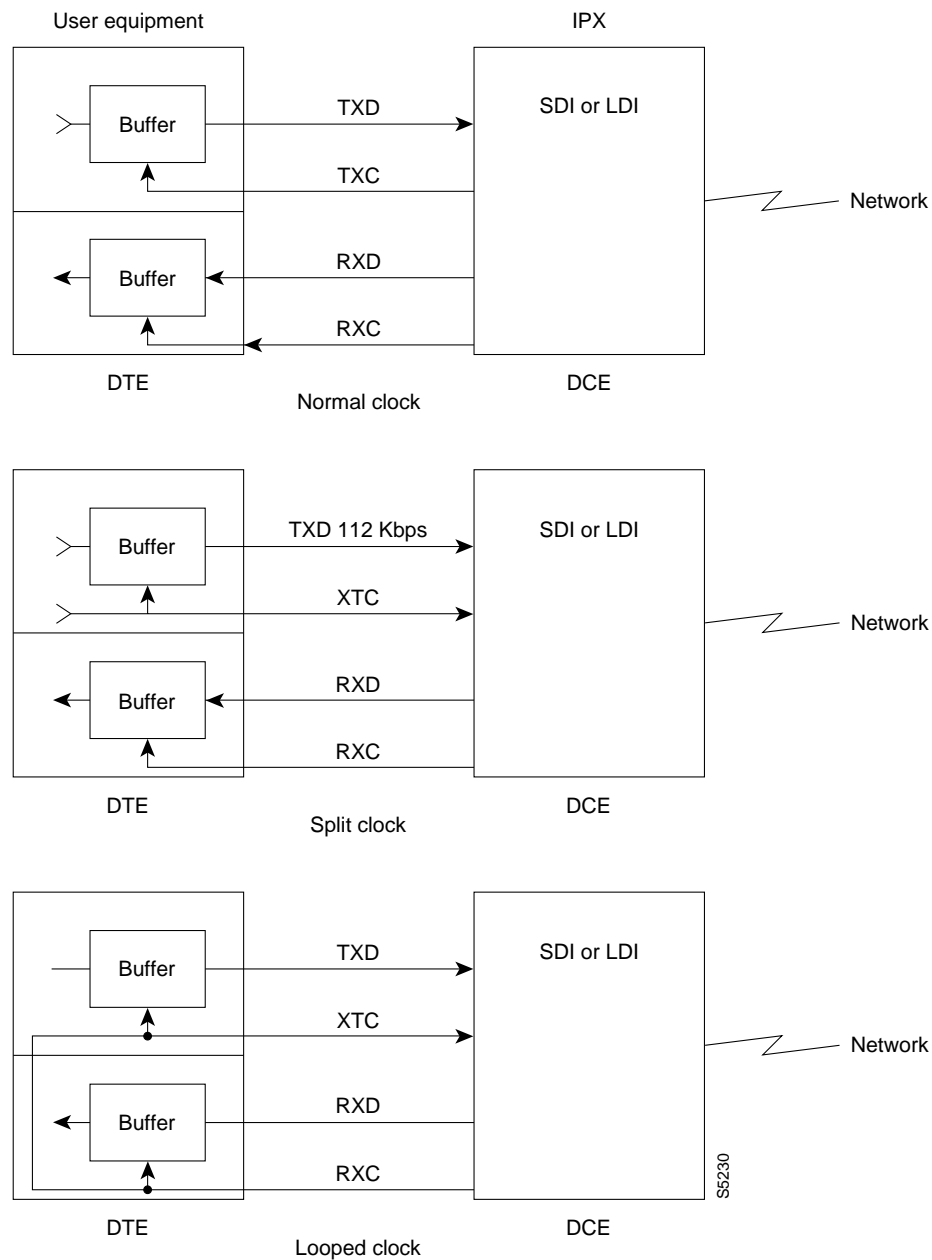
Refer to Figure 14-6 and the following discussion that illustrates these clocking configurations with the Cisco IPX narrowband switch data channel set for DCE operation. If the Cisco IPX narrowband switch were set for DTE, the directions of the clock lines will be reversed. Generally, the data channels are configured the same at both ends of the circuit but this is not a requirement.

In the **normal clock** mode, the Cisco IPX narrowband switch data interface card transmits both RxC and TxC to the DTE, which uses it to synchronize both transmit and receive data streams to and from the Cisco IPX narrowband switch. For “DCE normal” the TxC and RxC clocks transmitted by the Cisco

IPX narrowband switches are synchronous and the external equipment **must** synchronize both Tx and Rx to these clocks. This is a truly synchronous mode and is the recommended mode whenever possible for maximum throughput.

If the Cisco IPX narrowband switch is set for DTE and normal clock, the external data device will be sending both transmit and receive timing to the Cisco IPX narrowband switch. The two clocks (Tx and Rx) **must** be frequency-locked, but not necessarily phase-locked to each other, since the Cisco IPX narrowband switch only counts and synchronizes the far end to Rx. “DTE normal” implies an isochronous clock operation.

Figur e14-6 Data Clocking Configurations



In the **split clock** mode, the Cisco IPX narrowband switch data interface card sends RxD synchronized to RxC to the DTE and receives both TxC and TxD synchronized to the DTE. Since the clock is not synchronized to the Cisco IPX narrowband switch, it is referred to as anisochronous clocking. This timing is carried through the Cisco IPX narrowband switch network to the far end where it synchronizes the far end RxC and RxD. In this mode, the external equipment **must** use the clock output by the Cisco IPX narrowband switch to clock in the appropriate data stream (RxD for DCE, TxD for DTE).

In the **looped clock** mode, the Cisco IPX narrowband switch data interface card sends a receive clock (RxC) to the DTE along with the receive data. The DTE takes this RxC and uses it to time both receive data and transmit data. It also loops it back to the Cisco IPX narrowband switch as transmit clock (TxC) or external transmit clock (XTC) and expects to receive data synchronous with that clock. In this mode, the external equipment **must** synchronize to the clock output by the Cisco IPX narrowband switch and use that clock for its transmit data stream.

Isochronous clocking is implied by “DTE—normal” “DCE—split” and “DTE—split” modes (although the clock mode at the other end of the connection must be compatible). Isochronous connections are buffered at the receive end for an additional 10 msec to allow for clock fluctuations, and automatically use a 20/sec EIA update rate to maintain accuracy.

**Note**

Modem tail circuits in particular may need to be reconfigured to work properly with the Cisco IPX narrowband switch.

The clock types may be set differently at each end of a data circuit with some care in selecting the options. The restrictions to keep in mind are:

- Make sure that Cisco IPX narrowband switch channel is set for the proper configuration, DCE or DTE. It must be opposite from the connecting device for example if the user device is or acts like a terminal, it will be a DTE and the Cisco IPX narrowband switch data interface must be set to act like a DCE.
- The Cisco IPX narrowband switch cannot support two isochronous clocks in the same direction on a connection, though one isochronous clock in each direction is possible.
- Isochronous clocking can only be used for data rates of 112 Kbps or less.
- A data circuit cannot have data timed from two different sources of timing unless the sources are ultimately synchronized from some master clock source.

If both ends of the data circuit are set for DTE normal, indicating the external device is supplying timing to the Cisco IPX narrowband switches at both ends, a source of external timing for synchronization is required. The Cisco IPX narrowband switch cannot synchronize a different source of clocking at different ends of the connection unless they are both synchronized to the same master clock. If the external equipment is not synchronized to a common clock, buffer over-flows and under-flows will result, and data will be lost.

Voice Connections

This chapter is provided for users who wish to have an in-depth knowledge of the Cisco IPX narrowband/Cisco IGX 8400 series multiband voice connections and related functions. This chapter describes voice and signaling flow as well as digital processing topics such as voice compression, speech detection, modem detection, and echo cancelling.

The chapter contains the following:

- Voice Circuit Types
- Voice Circuit Signal Flow
- Voice Processing Features

Voice Circuit Types

An Cisco IPX narrowband switch/Cisco IGX 8400 series multiband switch voice circuit is transmitted using one or more DS0 time slots in the T1 bit stream. The circuit is defined within the node by one of nine voice frequency (VF) connection types depending on the compression characteristics applied as indicated in Table 15-1. These characteristics are with respect to the CDP card (Cisco IPX narrowband switch) and the CVM card (Cisco IGX 8400 series multiband switch) and are set in software for each connection. See the “Voice Processing Features” section for more information.

Table 15-1 CDP/CVM Circuit Types

Type	Descriptio
p	A p-type connection carries 64 Kbps PCM voice and supports A-law or μ -law encoding and conversion, level adjustment (gain/loss), and signaling.
t	VF t-type connections carry 64 Kbps clear channel data traffics.
c	When using the CVM card for voice circuits, the level of ADPCM compression must be specified by the (x) parameter and can be 32, 24, or 16 for 2:1, 3:1, or 4:1 compression. For example “c16” indicates full 4:1 compression withVAD.
a16z c16z	The c16z and a16z use a 4-level 16 Kbps ADPCM to ensure ones-density but at a decrease in voice quality. The "z" acts as a route -avoid specification whose effect is to avoid lines with zcs (zero code suppression).

Table 15-1 CDP/CVM Circuit Types (continued)

Type	Descriptio
a32d c32d	The a32d and c32d circuit types are used for Enhanced Instafax. These types support high speed circuits but stay at 32 Kbps when a high-speed circuit is detected and thus support compression that would otherwise be unavailable. Unless a32d or c32d is selected, the circuit switches to 64 Kbps when a high-speed modem or FAX is detected on a voice circuit. This selection is made using the cnfvchparam command. The selection of a32d or c32d depends on the performance of the modem/FAX using the circuit and should be selected for error-free operation.
a32 a24 a16	Uses ADPCM only. Can select 32 Kbps, 24 Kbps, or 16 Kbps compression. Compressed code avoids all zeros and can be used on lines with no other zero code suppression techniques. Modified 16 Kbps compression.
c32 c24 c16	Uses both ADPCM and VAD. Can select 32 Kbps, 24 Kbps, or 16 Kbps ADPCM compression. Compressed code avoids all zeros and can be used on lines with no other zero code suppression techniques. 16 Kbps compression is non-standard.
p	A p-type connection carries 64 Kbps PCM voice and supports A-law or μ -law encoding and conversion, level adjustment (gain/loss), and signaling.
t	VF t-type connections carry 64 Kbps clear channel data traffics.

The p type assumes the connection is carrying 64 Kbps PCM voice and permits specifying A or μ -law encoding and conversion, level adjustment (gain/loss), and signaling. The VF “t” type connections, however, are assumed to be carrying 64 Kbps clear channel data traffic and restrict the use of the voice parameters.

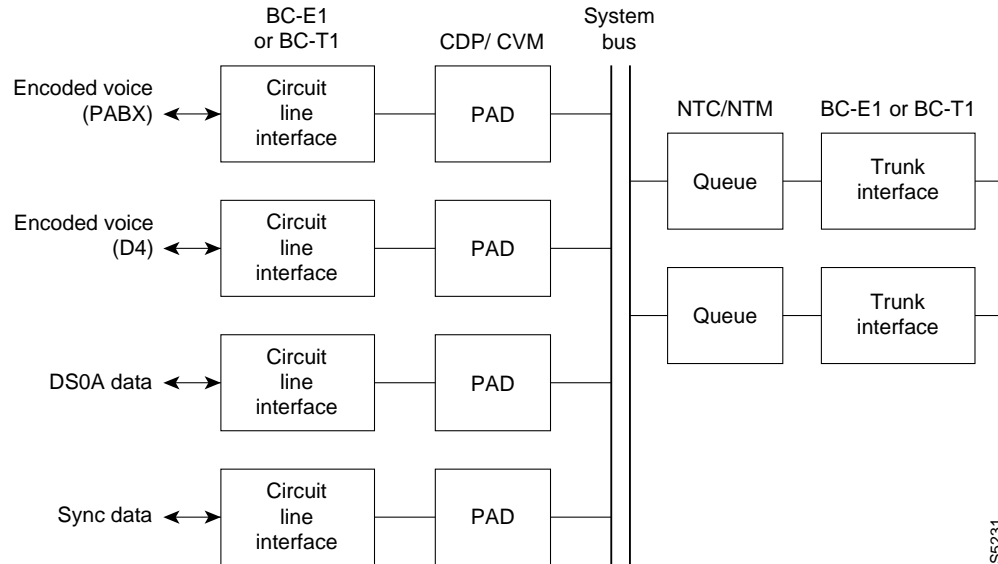
When using the CDP/CVM card for voice circuits, the level of ADPCM compression must be specified by the (x) parameter and can be 32, 24, or 16 for 2:1, 3:1, or 4:1 compression. For example “c16” indicates full 4:1 compression with VAD. The c16z and a16z use a 3-level 16 Kbps ADPCM to ensure one’s density but at a decrease in voice quality.

The a32d and c32d circuit types are used for Enhanced Instafax. When a high-speed modem or FAX is detected on a voice circuit, the circuit reverts to 64 Kbps unless a32d or c32d is selected. This provides compression otherwise unavailable. The selection of a32d or c32d depends on the performance of the modem/FAX using the circuit and should be selected for error-free operation.

Voice Circuit Signal Flow

Figure 15-1 illustrates the general voice circuit signal flow through a Cisco IPX narrowband switch or a Cisco IGX 8400 series multiband switch. Encoded voice circuits, originating with PABX’s and D4 channel banks, are applied via a T1 or E1 circuit line to an interface card on the Cisco IPX narrowband switch/Cisco IGX 8400 series multiband switch that converts the input data to a standardized DS1 or E1 format. The encoded VF data is then packetized by a CDP (Cisco IPX narrowband switch) or CVM (Cisco IGX 8400 series multiband switch) card.

Figur e15-1 Generalized Cisco IPX Narrowband Switch/Cisco IGX 8400 Series Multiband Switch Voice Signal Flow



The packetized VF data is then applied to the Cisco IPX narrowband switch/Cisco IGX 8400 series multiband switch system bus where it is made available to any other card in the node. Packets of data destined for a remote node are gathered from the system bus and queued up for transmission by a trunk card such as the NTC (Cisco IPX narrowband switch) or NTM (Cisco IGX 8400 series multiband switch) and sent to the remote node through an appropriate trunk interface card such as BC-T1 or BC-E1. The queues are arranged by each of the six FastPacket types described previously.

The sequence of events is reversed at the receiving node. Data that is destined for another circuit terminating on the local node is received from the system bus by another PAD card where it is depacketized and sent off to the local user device.

Voice Signal Flow with CDP Card



Note

This description is with respect to the CDP card in the Cisco IPX narrowband switch, but the CVM card in the Cisco IGX 8400 series multiband switch functions similarly.

Voice traffic originates from a PABX or channel bank and is transmitted over a circuit line to the Cisco IPX narrowband switch/Cisco IGX 8400 series multiband switch where it is terminated on an E1 or T1 back card, depending on the type of transmission facility used for the circuit line (Figur e15-2). The back card is primarily an interface card that monitors the received signal for alarm conditions or codes, converts the bipolar signal to unipolar logic level data, and extracts clock from the input signal. The back card makes this clock available to the node for frame synchronization.

The back card receives this input data and stores it in a frame buffer to reduce the effects of jitter on the circuit line. The BC-E1 can be programmed to operate with either of the two E1 line formats, AMI or HDB3. The BC-T1 can be programmed for either of the three T1 line formats, AMI, B8ZS, or ZCS.

For all voice circuit types except the “t” type, the data stream from the user will be assumed to be carrying voice circuit signaling information in specific bit positions (interpretive mode). With the transparent mode for “t” type connections, all bits are assumed to be data bits and no signaling

processing is performed on these circuits. For E1 circuits, the signaling bits may be associated with each frame, Channel Associated Signaling (CAS), or in a separate frame by itself, Common Channel Signaling (CCS) and is specified in the system software.

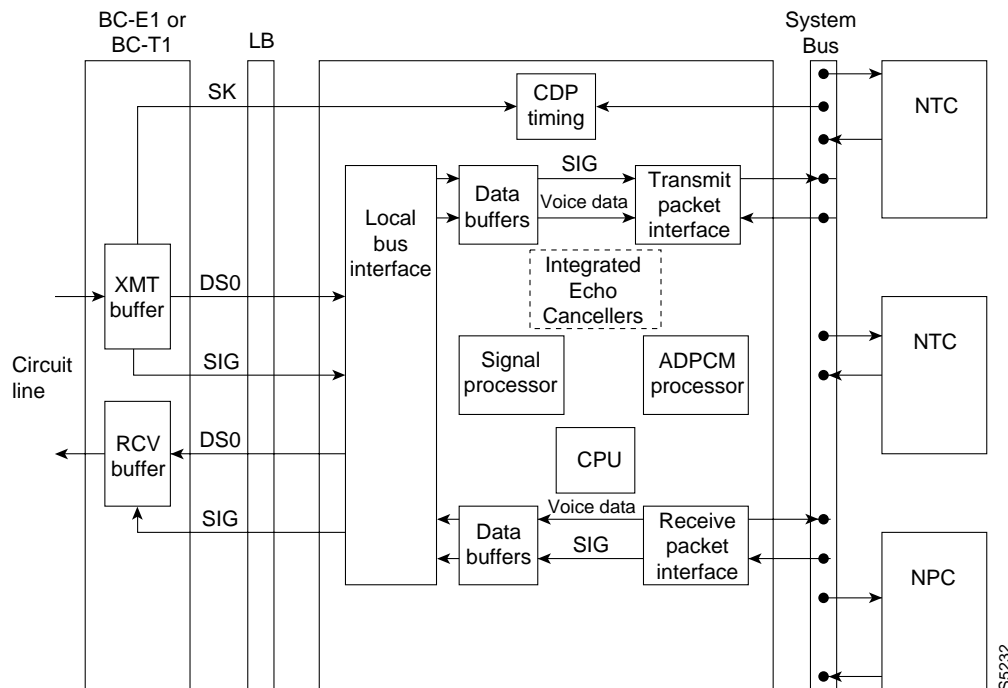
For circuits defined as one of the interpretive modes, the signaling bits are stripped off by the back card and forwarded separately to the CDP for further signaling bit processing (see next section). The voice data bits are then passed to the CDP card at 125 μ s. intervals via the Local Bus.

The CDP (Figure 15-2) performs specified signal processing functions on standard PCM voice data on a per-channel basis. The circuit data, temporarily stored in the CDP's buffers, is analyzed for silent intervals for VAD as well as for modem detection. Gain or loss may be added to the circuit data, an A-law to μ -law conversion may be provided, and voice compression in the form of VAD and/or ADPCM may be performed. If a modem tone is detected on a channel, any compression or gain control is disabled for the duration of the modem activity.

Each CDP may be equipped with optional Integral Echo Cancellers (IECs). The IECs come in two models, a 24-channel IEC for use with T1 circuit lines and a 32-channel for E1 circuit lines. This provides quality voice connections while eliminating the need for any external echo cancellers.

When so equipped, the data is modified according to input parameters from the user to reduce echo in the reverse direction. The gain adjustment precedes the echo cancelling function so that high-level input signals may be padded down, if necessary, for best echo cancelling results. When echo cancellers are employed on channels that require μ -law to A-law conversion, the channels must be numerically paired (for example channels 1 and 2 must be configured for the same signal processing.)

Figur e15-2 Voice Block Signal Flow with CDP



The ADPCM compression used by the CDP is fully compatible with the algorithm adopted by ITU-T G.721, G.723, and G.726. It offers either 32 Kbps, 24 Kbps, or 16 Kbps ADPCM selected by the user. It is not, however, compatible with a similar algorithm used by the VDC card used in earlier version Cisco IPX narrowband switch systems.

The applicable processing functions are entered by the system operator where they are decoded by the controller card (NPC). The controller card communicates this information, as well as other background information, to the CDP and associated BC-E1 or BC-T1 back card over the System Bus.

CDP operations are timed using the internal Cisco IPX narrowband switch timing. Timing information recovered from the incoming signal at the back card is compared to the internal Cisco IPX narrowband switch timing and, if they are not synchronous, the CDP performs controlled frame slips to prevent buffer over or underflow. A frame slip is performed by repeating or skipping a frame's worth of data. Frame slips are not desirable and can be avoided with proper network clocking.

After all signal processing functions have been performed on the data, it is queued up, assembled into packets with a destination address header, and placed on the TDM bus of the System Bus structure. From there, the packets are forwarded to an NTC or AIT trunk card to be routed through the network. If the destination is another circuit off this same node (internode switching), the CDP does not packetize the data but buffers then applies each byte of data onto the TDM bus where it is received from the bus by another CDP.

The NTC packet transmitter looks at the destination address of every packet on the TDM bus and compares it to a routing table to determine whether or not to take the packet off the bus. Cisco IPX narrowband switch packets destined for transmission to the far end node are pulled from the TDM bus by the NTC. If a destination address match is found, the NTC places the voice packet from the CDP in its queue to await transmission on the associated packet line.

NTC FastPacket buffers for a particular E1 or T1 packet line are loaded up with all packets having the same destination header. The NTC calculates a CRC and the FastPackets are read out at the 1.544, 2.048 Mbps or subrate line rate.

The NTC implements packet line frame synchronization and alignment based on the 5-bit CRC field in the FastPacket data format. This makes the FastPacket frame independent from the E1 transmission line frame and multiframe alignment (FA and MFA) signals and allows the full utilization of the available bandwidth. At the remote node, the process just described occurs in reverse order using the same card configuration.

Voice Channel Signaling



Note

This description is with respect to the CDP card in the Cisco IPX narrowband switch, but the CVM card in the Cisco IGX 8400 series multiband switch functions similarly.

Signaling on digitally encoded voice channels consists of one or two bits that accompany each voice channel to represent the status of the channel (on-hook, off-hook, dialing, etc.). Changes in the status of a particular voice channel are represented by changes in the signaling bit state for the particular channel. Timing of signaling bit transitions with respect to voice information on the same connection is not precisely maintained through the Cisco IPX narrowband switch/Cisco IGX 8400 series multiband switch network as the transmission paths are not the same.

On a T1 circuit line unless “no signaling” is selected with the **cnfvcht** command, signaling is carried using up to four bits for each channel, the A, B, C, and D bits. These bits are inserted in the least significant bit position of the PCM sample in frames 6 and 12 of every superframe by the PABX or D4 Channel Bank. On an E1 circuit line, signaling is performed using these same four bits for each channel. The A, B, C, and D bits for a particular channel are carried in timeslot 16 in every sixteenth frame (each frame carries the four signaling bits for two channels in channel 16).

Regardless of the type of circuit line (T1 or E1), the Cisco IPX narrowband switch/Cisco IGX 8400 series multiband switch recognizes two types of dialing formats: inband (for DTMF) and dial pulsing. These are software configured channel-by-channel. The Cisco IPX narrowband switch/Cisco IGX 8400 series multiband switch supports the transmission of both of these signaling types with appropriate time delays.

When a channel is configured for inband signaling, the only signaling bit transitions expected are for off-hook/on-hook information and winks. Dialing information is transmitted using tones. In this mode, the signaling information transitions more rapidly than with dial pulse signaling.

When a channel is configured as pulse dial type, signaling bits are used to convey dialed digits. The signaling bits will change every 40 to 60 msec within a dialed digit, with pauses between digits of greater than 100 msec. resulting from the slowness of the rotary dialing.

The designation of a channel as “inband dial” allows the signaling delay parameter to be set low, reducing the delay through the network. For “pulse dial” channels, the signaling delay is set greater than the largest expected pulse width within a digit, allowing the whole digit to be sent in one packet. Default parameters work well with 10 pps signaling. Note, however, that pulse signaling will not be transmitted properly on a channel designated “inband” because after each transition is received the next will be delayed longer than is proper.

The channel signaling type (DPO, PLAR etc.) that is selected for each voice channel in the Cisco IPX narrowband switch/Cisco IGX 8400 series multiband switch is only used to provide a window on the instantaneous signaling state of a channel (and supply correct conditioning if a connection fails). So long as the inband/pulse choice is made correctly, signaling bit transitions will be transmitted through the Cisco IPX narrowband switch/Cisco IGX 8400 series multiband switch network whatever the configured signaling type. The software has several screens that allow the operator to view the status of the various signaling bits on a specified channel.

Signaling Bit Flow with CD



Note

This description is with respect to the CDP card in the Cisco IPX narrowband switch, but the CVM card in the Cisco IGX 8400 series multiband switch functions similarly.

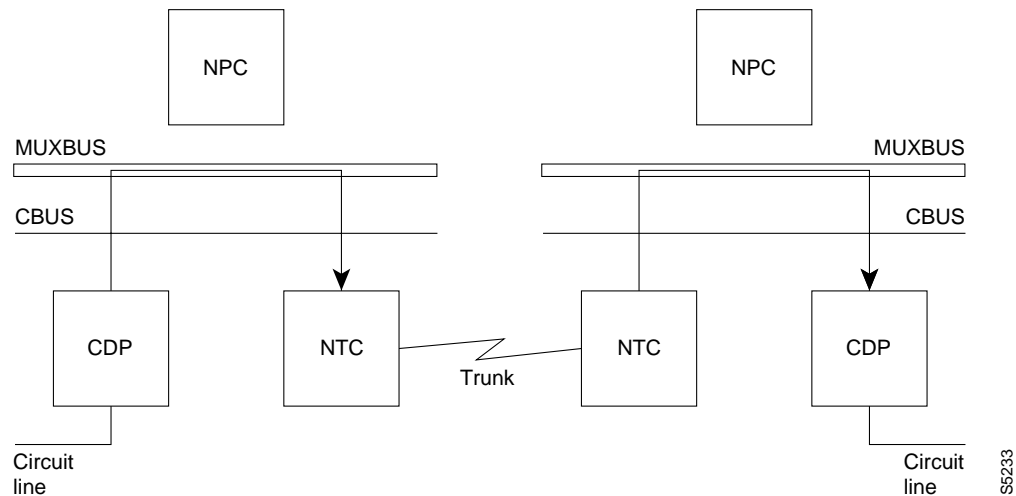
Signaling bits are extracted by the BC-T1 or BC-E1 back card. The CDP demultiplexes the T1 or E1 signal and separates the voice data bits from the signaling bits. Both voice bits and signaling bits are temporarily stored in RAM buffers in the back card. An Interdigit Signaling timer debounces the signaling bit transitions to eliminate noise. Upon demand, the back card sends the signaling bits to the CDP over a local bus. A separate local bus is used for the voice bits for maximum data transfer. Signaling bits, occur at a slower rate and are transferred using a second local bus.

The CDP detects signaling bit transitions and integrates the changes of state to minimize false signaling. The CDP data processing modules perform any protocol conversions that may be specified for a channel and packetizes the signaling transition. Each change of state of a signaling bit is packetized in a timestamp packet. In addition to the packet timestamp, the CDP adds a delay field in the packet payload that is used to ensure the signaling packets are received in the proper sequence and with acceptable end-to-end delay to support a maximum of 11 pulses/second dial pulsing.

The signaling packets are sent over the MUXBUS where they are picked off by the NTC and forwarded to their destination. If there are no signaling bit transitions detected, heartbeat packets are sent every 2 to 30 seconds to maintain the far end signaling state.

At the far end, the signaling packets are received by the NTC, which forwards them directly to the appropriate CDP over the MUXBUS (Figur e15-3). The signaling state is detected and the signaling data is forwarded to the associated back card where it is inserted in the circuit's receive data frame and sent to the user device.

Figur e15-3 CDP-to-CDP Signalling



Voice Processing Features

Digital signal processing provides many of the features available for voice connections through the Cisco IPX narrowband switch. The following paragraphs describe in detail these features.

ADPCM Voice Compression

The Cisco IPX narrowband switch/Cisco IGX 8400 series multiband switch uses a standard algorithm (ITU-T G.721, G.723, and G.726) to implement Adaptive Differential Pulse Code Modulation and encodes only the difference between one encoded voice sample and the following sample, not the actual value of the sample amplitude. Since the range of values for the difference is much smaller than the actual value of the sample, it requires fewer bits for encoding this difference.

With ADPCM, a formula is applied to the previous outgoing sample that predicts the range of values for the next sample. The transmitter then encodes the difference in level between the previous sample and the present sample, on a scale set by the prediction, and transmits this information to the far end receiver. Since the receiver has the same data to work with as the transmitter, it is able to reconstruct the PCM data input to the transmitter.

Conversion of 8-bit PCM samples to ADPCM, as well as packetizing the ADPCM, is performed together in the CDP (Cisco IPX narrowband switch) or CVM (Cisco IGX 8400 series multiband switch) card. Compressed voice code words from the CDP/CVM can be selected by the user to be either 4-bit (using 32 Kbps of bandwidth), 3-bit (24 Kbps) or 2-bit (16 Kbps) rather than 8-bit PCM samples. The resulting bandwidth savings with ADPCM generally compromises voice quality.

ADPCM compression relies on the signal being close to its model of speech to make such bandwidth reduction. High-speed modem and FAX signals are distorted by ADPCM to the extent that operation above 4800 baud generally will not work reliably on ADPCM connections. When PCM channels are used for either of these two functions, all compression should be disabled.

Speech Detection

If the Voice Activity Detector (VAD) option is active, a speech detector is used to identify periods of silence or idle VF channels to provide further compression in addition to ADPCM. This function is performed by the CDP card.

The PCM signal is first passed through a high-pass filter to remove DC bias and 50/60 Hz. component that might disturb the speech or modem detectors. The energy of each sample is measured, and the average speech power is calculated. If the power is greater than -31 dBm0 (dBm0 equals a digital milliwatt), the signal is recognized as speech. This is a simplification of the VAD function. The full implementation of VAD in the Cisco IPX narrowband switch requires more calculations than simple power level.

When there are short periods of silence between syllables or words, the speech detector switches on and off quickly, chopping the previously smooth conversation. To prevent this, whenever the speech detector has been on but speech is no longer detected, a timer is started. While the timer is running, packets are still sent, bridging short periods of silence. The trade-off here is that packets are still sent for a short time when the channel is silent.

Since the speech detection test takes some time to make, speech was actually present some time before the speech detector was switched on. To accommodate this, a buffer is used and the first few packets of speech are sent out using high priority packets to minimize additional delay factors. This practice prevents front-end clipping of talkspurts but this increases overall connection delay.

There are various software parameters that may be changed by the system operator to optimize the compromise between clipping and utilization. If VAD is used for a voice circuit, echo cancellers will be required to allow better voice quality without echo. The CDP (Cisco IPX narrowband switch) or CVM (Cisco IGX 8400 series multiband switch) card has echo cancelling built into the card and external echo cancellers are not required when using this card.

In the receiving CDP/CVM, packets are buffered and played out as they are received. If there are no packets to transmit to the user device, the card will insert pink noise to simulate the background noise when the far end is off-hook, minimizing the perceived choppiness introduced by VAD. Prior to Rel. 7.1, the level of this pink noise was fixed by the system. In Rel. 7.1 and afterwards, the background noise level is sampled at the transmitting CDP/CVM and reproduced at the receiving CDP/CVM to better approximate the actual background noise level.

The Adaptive Voice feature automatically disables VAD during periods when there is unused network bandwidth available. This allows transmission of the compressed voice without the effects of VAD. As the extra bandwidth is allocated to other connections, VAD is dynamically enabled on all voice circuits to free up the needed bandwidth.

Modem Detection

The V.25 modem detector recognizes the steady 2100 Hz tone output by V.25 compliant modems and FAX machines to disable echo cancellers at the beginning of transmission.

When a channel is declared V.25 modem, a message is sent to the controller card, and software coordinates upgrade the connection at both ends from ADPCM to PCM (unless already PCM, or otherwise configured in the “**cnfvchparm** V.25 Detect” field). Therefore, the connection effectively

becomes “p” until the modem removal criteria (bidirectional power below “MDM Low Pwr Thrsh” for a duration greater than “MDM Detect Silence Max”) are met, at both ends, when it returns to the configured state. Software attempts to synchronize changes at both ends to within 100 msec, but there is a short silence on the channel while reconfiguration takes place.

Echo Cancelling

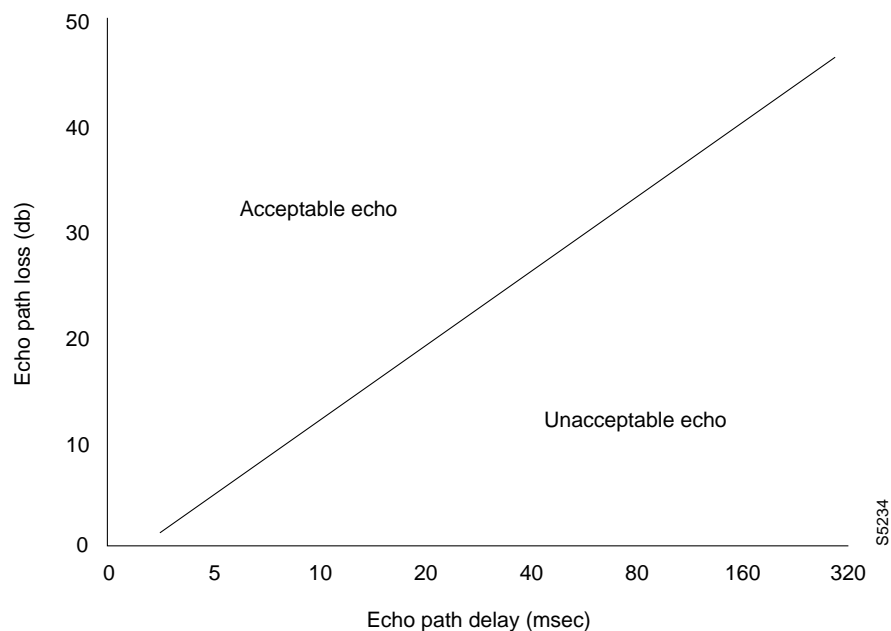
It is practically impossible to obtain good echo performance from an Cisco IPX narrowband switch/Cisco IGX 8400 series multiband switch network without using echo cancellers. The echo cancellers are provided internally as an option with the CDP/CVM card. The performance of echo cancellers depends on a number of factors that can be optimized in the Cisco IPX narrowband switch/Cisco IGX 8400 series multiband switch by the advanced user for particular conditions.

Echo is a function of how much energy is reflected back to the talker or listener, and the time delay between the original signal and the reflection. The graph in Figure 15-4 illustrates the generally accepted boundary between acceptable and unacceptable echo.

The reflections of speech are produced by impedance mismatches in 2-wire circuits that are found to some extent in all 2-wire networks. The reflected power is generally in the order of 10 to 20 dB less than the forward power level as indicated in the figure. User perception of echo is aggravated by circuit delay. In circuit-switched environments, with a typical delay less than 5 msec, the perceived echo is below the threshold of annoyance.

However, when Cisco IPX narrowband switch/Cisco IGX 8400 series multiband switch equipment is introduced into the network, the delay increases by 30-100 msec. The increased delay is primarily due to voice activity detection, packetizing, and queuing processes. Figure 15-4 shows that for this delay, the reflected signal must be 20 to 30 dB less than the original signal. If nothing were done, the echo would be intolerable.

Figur e15-4 Echo Effects Diagram



The CDP/CVM echo canceller is a digital signal processor that continuously monitors the digitized speech going through the Cisco IPX narrowband switch/Cisco IGX 8400 series multiband switch towards the network (transmit direction) and inserts loss in the reverse direction. It calculates the expected echo power level and subtracts it from the received signal. Therefore, echo path loss can be increased by 30 dB to 40 to 50 dB, which, at a delay of 50 to 100msec, results in an acceptable echo.

Echo cancellers have a tone detection feature for FAX/fast modem detection. If a 2100 Hz tone is detected, echo cancelling is disabled until the signal power in the connection has fallen, indicating the end of the call. This feature should normally be enabled.

The following measures can improve echo performance on voice circuits:

1. Make sure all the parameters that can be set for the echo canceller are correct.
2. Avoid any satellite circuits. They have a tremendous amount of delay as compared to a terrestrial circuit.
3. If possible, substitute 4-wire circuits for 2-wire, eliminating echo from that source.
4. Reduce the circuit Echo Return Loss (ERL) as much as possible. Adding loss in the tail circuit is always helpful but this may reduce VAD performance (increase clipping and choppiness). Finding particular 2-wire terminations causing echo is sometimes possible. These can be improved by line build-out or impedance options on trunk cards.
5. Echo cancellers usually have a limit to the delay they can accommodate (often 32 msec). If the tail circuit delay (not packet delay) is close to this limit, an extended version of the echo canceller may be needed.
6. Echo cancellers have difficulty with double-talk when the signal levels in each direction of a call are different by more than 10 dB. It may be possible to change the network loss plan to allow for this.
7. The Cisco IPX narrowband switch/Cisco IGX 8400 series multiband switch-introduced delay may be reduced by changing some software parameters if only on-net calls are made. Contact Customer Service for assistance on this.
8. Try reducing the hop count for troublesome circuits. This is likely to make only a small improvement, but may be useful

Voice and Signaling Conditioning

With almost all digital multiplexers and PABXs, the signaling bits are forced to a predetermined state when the transmission link has failed. This is used to drop all calls in progress and block any new access to the voice circuits. Usually, the predetermined state is a “busy” but other conditioning sequences may be used (for example idle for a short interval to drop all calls in progress followed by a permanent busy until the fault clears).

With a point-to-point transmission link, where all 24 channels end up in the same destination, the channel bank or PABX conditioning detects the fault on the digital trunk line and provides the conditioning. However, with transmission through the Cisco IPX narrowband switch/Cisco IGX 8400 series multiband switch, each DS0 circuit may have a different destination from all others from the same source. Therefore, it is up to the Cisco IPX narrowband switch/Cisco IGX 8400 series multiband switch to condition the signaling when it detects a network trunk failure (such as packet out of frame).

The signaling conditioning for each individual voice circuit is specified by a conditioning template that the user selects using the Configure Voice Channel Interface Type (**cnfvcht**) command. The command specifies the channel on-hook (idle) state and the signaling state forced by the CDP (Cisco

IPX narrowband switch) or CVM (Cisco IGX 8400 series multiband switch) when a connection fails. After detecting a connection failure, the channel voice and signaling conditioning is instantaneously applied.

Most of the common VF circuit types are conditioned by using a preset template. This template also permits the user to specify the voice code bits transmitted to the user during connection failure. These bits are usually set to the idle code (all voice bits = 1 for D4 μ -law encoding) but can be anything the user chooses to configure.

It is important to select the proper channel type at installation of any new circuit. For instance, if the convention selected for E&M channels (idle for 2 seconds, then busy) is selected for PLAR circuits (auto ringdown) it causes PLAR circuits to ring when connections are failed. Changing the conditioned signaling to “on-hook” corrects this problem.

Level Adjustment

Voice connections in a network must concern themselves with providing sufficient talk volume at the receive end with a low level of circuit noise and distortion for acceptable voice conversations to be held. In worldwide telephone networks this is accomplished with the use of loss plans. Loss plans assign acceptable circuit loss for each element of the network that a telephone call may traverse.

In general, subscriber loops (from the telephone to the central office) are allowed to have anywhere from 0 dB to 6 or 8 dB of voice frequency (VF) loss. This loss is primarily determined by the size and length of the cable pairs that connect the local phone or to the office switching equipment. Rarely are there any gain devices employed in the local loop. Since this local loop is found at both ends of the talk path, the total loss may be as high as 15 or 16 dB, which would be quite noticeable to both parties.

Most local calls and all long distance calls travel over one or more connections (hops) between central switching offices on circuits called VF trunks. These differ from subscriber lines in that trunks are not dedicated to a particular caller but are used as needed to route calls between central offices. It is in these trunks that VF level control can be accomplished.

Many businesses employ Private Branch Exchanges (PBXs) so the local loop may be short (between the phones in an office and the PBX) but the distance from the PBX to the central office may be some distance. PBX's usually connect to telephone offices using trunks.

Bell System loss objectives for various types of trunks are listed in Table 15-2. These trunks may or may not have gain associated with them. Most trunks in the telephone networks of industrialized nations use some form of digital transmission (T1 or E1 or higher bit rates) where circuit gain or loss can be tightly controlled.

Table 15-2 Loss Objectives for Common Carrier Trunks

Trunk Type	Loss Objective
Toll Connecting	VNL + 2.5 dB
Intertoll	VNL
Tandem	0.5 dB balanced offices 1.5 dB unbalanced offices
Direct and Tandem	3 dB

Via Net Loss (VNL) ranges from 0.5 dB for short trunks up to 165 miles to 2.9 dB for extremely long trunks up to 1850 miles. For any trunk that employs echo suppressors, VNL should be 0 dB. In addition to insuring adequate talk volume at each end of the circuit, the loss plan is used to reduce the effects of

echo on the voice path. A certain amount of loss is required to prevent echo levels from becoming objectionable. If the loss plan for each trunk type as stated in Table 15-2 is followed, there should be sufficient loss to keep echo from being a problem.

It is strongly recommended that the Cisco IPX narrowband switch/Cisco IGX 8400 series multiband switch voice paths be equipped with echo cancellers because of the delay that voice packets experience over an Cisco IPX narrowband switch/Cisco IGX 8400 series multiband switch network. Note that it is not desirable to have more than one echo suppressor device on a talk path in attempt to further reduce echo as they tend to interfere with each other sometimes resulting in worse echo performance than with no echo cancellers at all.

Digital interoffice trunks do not have the problem of loss so much as analog trunks as once the voice signal is encoded, the level remains constant and the primary concern is to keep the bit error rate low enough to avoid noticeable clicks on the talk path. Digital circuits are inherently 4-wire trunks until terminated on a 2-wire channel unit thereby avoiding the repeated 2-wire to 4-wire conversions that contribute to echo problems.

In the Cisco IPX narrowband switch/Cisco IGX 8400 series multiband switch, the per-channel voice frequency gain can be digitally adjusted over a range of -8 to + 6dB in the CDP/CVM. The default for all voice channels is 0 dB. It is recommended that any gain adjustment be done at the circuit ends on an analog basis (at the channel bank or PABX channel) rather than within the network to minimize the distortion introduced.

Most VF encoders in channel banks, digital PABX's, digital switching offices, and even the Cisco IPX narrowband switch have a limited range over which they will encode linearly. In addition, too high a signal level input to the encoder will overload the encoder resulting in clipping of the VF signal. The overload point is approximately +3 dBm. And any signal level below approximately -45 dBm0 will not be encoded accurately.

Table 15-3 lists the typical quantization distortion added for a single-hop Cisco IPX narrowband switch/Cisco IGX 8400 series multiband switch trunk with 32 Kbps ADPCM employed. The impairment indicated is approximately the same for all digital PAD values in the range of 1 to 8 dB. One exception is the 6 dB (or multiple of 6 dB) A-Law PAD. It introduces negligible impairment for signals ranging to approximately -30 dBm0, and therefore attracts zero units of quantizing distortion.

Table 15-3 Quantization Distortion Pertaining to the Cisco IPX Narrowband Switch/Cisco IGX 8400 Series Multiband Switch

Digital Proces	Quantization Distortio
32 Kbits per second ADPCM (with adaptive predictor) combination of a PCM-ADPCM-PCM tandem conversion.	2.5 units
Digital lossPAD (8-bit, A-Law or μ -Law)	0.7 units

P A R T 6

**CONNECTION MANAGEMENT
NOTES**

Connection Management

This chapter is provided for users who wish to have an in-depth knowledge of the Cisco IPX narrowband connection management functions. It describes packet queuing and the various queue types. It also discusses circuit routing and rerouting, delay for various types of connections, and circuit bandwidth requirements and utilization.

The chapter contains the following:

- Packet Queuing
- Delay in a Cell Network
- Routing and Rerouting
- Bandwidth Allocation
- Traffic Statistics

Packet Queuing

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

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

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

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

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

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

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

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

Then, as frames go by, each of these packet types is eligible to accrue credits as indicated in Table 16-1.

Table 16-1 Credit Accumulation

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

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

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

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

Delay in a Cell Network

The overall delay of information through the network includes:

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

Delay in Packet Frame Relay Networks

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

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

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

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

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

Delay at the Source

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

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

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

Delay in the Network

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

- intermediate node delay
- propagation delay

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

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

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

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

Delay at the Sink

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

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

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

Synchronous Data Connection Delay

For all time-stamped and non-timestamped data packets, the number of information bytes in a packet varies from 4 to 21 bytes depending on the type and speed of the connection. The packetization delay for the two types of data packets can be calculated by using Table 16-2 or Table 16-3 or looked up in the tables at the end of this chapter.

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

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

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

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

Ways to Reduce Data Connection Delays

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

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

Table 16-2 Calculation of Non-Timestamped Data Packet Overall Delay

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

1. Includes trunk queuing delays

2. Isochronous connections only

Table 16-3 Calculation of Timestamped Data Packet Overall Delay

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

1. Isochronous connections only

2. Ages timestamp to 40 (default).

Voice Connection Delay

Non-VAD voice packets

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

VAD voice packet

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

Transmission delay across a trunk is generally a function of the distance travelled. For a terrestrial trunk, signals travel an average of about 100 miles per millisecond, or 0.01 msec/mile. For a satellite trunk, the propagation delay of the signal to the satellite and back adds approximately 300 msec per satellite hop. Table 16-4 can be used to calculate delay for the four types of voice connections.

Table 16-4 Sources of Delay in Voice Connections

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

Table 16-4 Sources of Delay in Voice Connections (continued)

Delay Source	t & p	v	a16	a24	a32	c16	c24	c32
Transmission delay (satellite, per hop)	300	300	300	300	300	300	300	300
Miscellaneous dejitter delays (per hop)	0.5	0.5	0.5	0.5	0.5	0.5	0.5	0.5

Table 16-5 illustrates some typical expected delays for various number of terrestrial hops. The calculations do not include any PBX or channel bank delay.

Table 16-5 Typical Voice Connection Delays

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

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

Maximum Hops—Voice Connections

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

Routing and Rerouting

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

Load Model

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

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

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

Load Model and Routing for Frame Relay

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

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

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

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

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

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

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

Routing Algorithm

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

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

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

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

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

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

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

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

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

Causes of Rerouting

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

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

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

Reroute Priority and Order

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

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

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

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

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

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

After the COS timer, priority is given to connections with the highest bandwidth (packets/second) of the group awaiting rerouting. This is because, as available bandwidth diminishes, it is more difficult to find routes for the higher bandwidth connections. The data block for each connection contains the packet/second requirement, so prioritizing is easy. The general rerouting priority order is given in Table 16-6.

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

Table 16-6 Priority for Rerouting

Priorit	ConnectionType
1	high speed data connections (>64 Kbps)
2	"p" or "d" connections
3	"a" connections

Table 16-6 Priority for Rerouting (continued)

Priorit	ConnectionType
4	"v" connections
5	"c" connections
6	low speed data connections (< 9.6 Kbps)

Courtesy Downing

Courtesy downing is the process of monitoring voice connections and downing those connections that are configured for downing when they go inactive. This frees up network bandwidth for other uses, generally for bandwidth reservation.

Only voice type connections can be monitored for inactivity, and then only when the on-hook status is configured by the user. All other connection types and voice types with no defined on-hook status are treated as active and cannot be courtesy downed.

System Message Traffic Routing

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

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

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

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

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

Bandwidth Allocation

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

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

Network Trunk Bandwidth

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

T3 Framed

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

E3 Framed

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

T1 Framed

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

E1 Framed

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

E1 Unframed

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

Subrate:

Depends on the number of DS0's available in the subrate trunk. See Table 16-7.

Table 16-7 Subrate Packet Line Bandwidth

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

**Note**

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

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

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

Each 64 Kbps time slot, or DS0, provides $1/3 \times 1000$ or approximately 333 packets per second of available bandwidth on a trunk. Table 16-7 shows the packet bandwidth available on a subrate trunk as a function of the number of DS0s regardless of the trunk type.

Voice Compression Bandwidth Requirements

The bandwidth required on a trunk to carry the information on a DS0 circuit depends on which one of the five compression types is selected for the circuit as indicated in Table 16-8. The equivalent bit rate after compression is also listed in this table.

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

Table 16-8 Cisco IPX Narrowband Switch Voice Grades

Type	Equivalent Bit Rat	Required BW
p	64 Kbps	381 pkts/sec.
t	64 Kbps	381 pkts/sec.

Table 16-8 Cisco IPX Narrowband Switch Voice Grades (continued)

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

1. Assumes 50% VAD.

VAD and DFM Effects

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

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

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

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

Data Channel Packet Generation Rates

The synchronous data channels use widely varying amounts of trunk bandwidth depending on whether they use timestamped data packets or not and how the control lead information is carried. Refer to Table 16-9 through Table 16-13 for bandwidth requirements or calculate as follows.

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

- Fast EIA packets, 7/8 coded, have 70 data bits per packet.
- DFM packets, 8/8 coded have 152 data bits per packet.
- DFM packets, 7/8 coded have 133 data bits per packet.

Exceptions are the low-speed connections listed in Table 16-11, Table 16-12, and Table 16-13, where partially-filled packets are used to reduce packetization delay. Divide the bit rate of the connection by the number of user bits per packet and the result is the number of packets/second.

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

Traffic Statistics

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

Table 16-9 Data LoadTable with Standard EIA and No DFM

Bit Rat	7/8 Coding			8/8 Coding		
	Kbps	Pkt/s	Byte/pk	Delay, ms	Pkt/se	Byte/pkt
1.2	43	4	24	38	4	27
1.8	65	4	16	57	4	18
2.4	35	10	29	30	10	33
3.2	46	10	22	40	10	25
3.6	52	10	20	45	10	22
4.8	35	20	29	30	20	33
6.4	46	20	22	40	20	25
7.2	52	20	20	45	20	22
8	58	20	18	50	20	20
9.6	69	20	15	60	20	17
12	86	20	12	75	20	14
12.8	92	20	11	80	20	13
14.4	103	20	10	90	20	11
16	115	20	9	100	20	10
16.8	120	20	9	105	20	10
19.2	138	20	8	120	20	9
24	172	20	6	150	20	7

Table 16-9 Data LoadTable with Standard EIA and No DFM (continued)

Bit Rat	7/8 Coding			8/8 Coding		
	Kbps	Pkt/s	Byte/pk	Delay, ms	Pkt/se	Byte/pkt
28.8	206	20	5	180	20	6
32	229	20	5	200	20	5
38.4	275	20	4	240	20	5
48	343 ¹	20 ¹	3 ¹	300	20	4
56	381	21	3	350	20	3
57.6	392	21	3	360 ¹	20 ¹	3 ¹
64	436	21	3	381	21	3
72	490	21	3	429	21	3
76.8	523	21	2	458	21	3
84	572	21	2	500	21	2
96	654	21	2	572	21	2
112	762	21	2	667	21	2
115.2	784	21	2	686	21	2
128	871	21	2	762	21	2
144	980	21	2	858	21	2
168	1143	21	1	1000	21	1
192	1307	21	1	1143	21	1
224	1524	21	1	1334	21	1
230.4	1568	21	1	1372	21	1
256	1742	21	1	1524	21	1
288	1960	21	1	1715	21	1
336	2286	21	1	2000	21	1
384	2613	21	1	2286	21	1
448	3048	21	1	2667	21	1
512	3483	21	1	3048	21	1
672	4572	21	1	4000	21	1
768	5225	21	1	4572	21	1
772	5252	21	1	4596	21	1
896	6096	21	1	5334	21	1
1024	6966	21	1	6096	21	1
1152	7837	21	1	6858	21	1
1344	9144	21	1	8000	21	1

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

Table 16-10 Data LoadTable with Standard EIA and DFM

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

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

Table 16-11 Data LoadTable with Partially-Filled Packet and No DFM

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

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

Table 16-12 Data LoadTable with Partially-Filled Packet and DFM

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

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

Table 16-13 Data LoadTable with Fast EIA

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

Table 16-13 Data LoadTable with Fast EIA (continued)

Bit Rat	7/8 Coding			8/8 Coding		
	Kbps	Pkt/se	Byte/pkt	Delay, ms	Pkt/sec	Byte/pkt
230.4f	3292	10	1	2880	10	1
256f	3658	10	1	3200	10	1
288f	4115	10	1	3600	10	1
336f	4800	10	1	4200	10	1
384f	5486	10	1	4800	10	1
448f	6400	10	1	5600	10	1
512f	7315	10	1	6400	10	1

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

Table 16-14 Data LoadTable—Fast EIA with Partially-Filled Packet

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

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

PART 7

REFERENCE

Network Specifications

This appendix lists preliminary information for Release 9.1. (Refer to on-line documents for latest information.)

Broadband Trunk Interfaces

Trunk Interfaces:	T3 and E3
Trunk Protocol:	Asynchronous Transfer Mode (ATM)
Trunks per Node:	32 DS3 or E3 trunks per BPX of which 16 can be shelf trunks 1 active DS3 or E3 trunk per IPX
Trunk Bandwidth:	96,000 cells/sec. (T3) 80,000 cells/sec. (E3)
Trunk Interfaces:	OC3
Trunk Protocol:	Asynchronous Transfer Mode (ATM)
Trunks per Node:	
Trunk Bandwidth:	353,208 cells/sec. (OC3)
Trunk Interfaces:	OC12
Trunk Protocol:	Asynchronous Transfer Mode (ATM)
Trunks per Node:	
Trunk Bandwidth:	1,412,830 cells/sec.

Narrowband Trunk Interfaces

North American:	T1 Fractional T1 (minimum of 4 DS0s) T2
International:	CEPT E1 Subrate, 64 Kbps to 1.920 Mbps Japanese “Y1”
Trunk Bandwidth:	1.536 Mbps, 8,000 FastPackets/sec. (T1) 1.984 Mbps, 10,333 FastPackets/sec. (framed E1) 2.048 Mbps, 10,666 FastPackets/sec. (unframed E1) Subrate and fractional T1 depend on data rate selected

Narrowband Channel Interfaces

Frame Relay:	Channelized T1: 56 Kbps, N x 64 Kbps, up to 1.536 Mbps. Channelized E1: G.703, N x 64 Kbps, up to 1.984 Mbps. V.35: configurable from 56 Kbps to 2.048 Mbps
Data:	V.35: 56 Kbps to 1.344 Mbps. RS-449: 56 Kbps to 1.344 Mbps.
Voice:	Channelized T1 Channelized E1 Channelized Japanese “TTC” DC5A and E&M channel associated signaling FAX/high-speed modem

ATM Trunk Interface (BXM-T3/E3 Cards)

Characteristic	T3 (DS3)	E3
Line Rate	44.736 Mbps +/- 20 ppm	34.368 Mbps +/- 20 ppm
Line Code	B3ZS	HDB3
Cell Transfer Rate	96,000 cells per second	80,000 cells per second
Framing	ANSI T1.107, T1.107a	ITU T G804, G.832
Signal Level	TA-TSY-000773 (PLCP)	ITU-T G.703
Transmission Convergence Sublayer	DS3 PLCP frame format DS3 HEC mapped format	G.832 E3 frame format
T3 (DS3) and E3		
Port Interface	Trunk or UNI	
ATM Layer Protocol	LMI, ILMI	
Port Alarm Processing	LOS, LOF	
Connector	SMB	
Indicators	Card status, Port status	

ATM Trunk Interface (BXM-155 Cards)

Line Rate	155.52 Mbps	
Line Code	NRZ	
Signal Level	Min dBm	Max dBm
MMF LED TX	-22	-15
MMF LED RX	-34	-10
SMF IR TX	-15	-8
SMF IR RX	-34	-10
SMF LR TX	-5	0
SMF LR RX	-34	-10
Framing Format:	STS-3c, STM-1	
Port Interface:	LMI, ILMI	

ATM Cell Rate:	353,208 cells/sec.
Jitter:	ATM Forum UNI 3.1
ATM Layer Protocol:	LMI, ILMI
Port Alarm Processing:	LOS, LOF, LOP, Path AIS, Path Yellow
Line Errors Counted:	
Connector:	SC for MMF, FC-PC for SMF (IR) and SMF (LR)
Max. Cable Lengths:	MMF ~2 KM SMF IR ~20 KM SMF LR ~40 KM
Indicators:	Card status, Port status.

ATM Trunk Interface (BXM-622 Cards)

Line Rate	622.08 Mbps	
Line Code	NRZ	
Signal Level	Min dBm	Max dBm
SMF IR TX	-15	-8
SMF IR RX	-28	-8
SMF LR TX	-2	+2
SMF LR RX	-28	-8
Framing Format:	STS-12c, STM-4	
Port Interface:	LMI, ILMI	
ATM Cell Rate:	1,412,830 cells/sec.	
Jitter:	ATM Forum UNI 3.1	
ATM Layer Protocol:	LMI, ILMI	
Port Alarm Processing:	LOS, LOF, LOP, Path AIS, Path Yellow	
Line Errors Counted:		
Connector:	SMF-FC	
Max. Cable Lengths:	MMF ~2 KM	
	SMF IR ~20 KM	
	SMF LR ~40 KM	
Indicators:	Card status, Port status.	

ATM T3 Trunk Interface (BNI-T3, LM-3T3)

Line Rate:	44.736 Mbps \pm 20 ppm, asynchronous
Line Code:	B3ZS
Signal Level:	DSX-3

Framing Format:	C-bit parity is monitored. No other framing or control bits in the DS3 frame are either altered or monitored
Protocol:	Physical Layer Convergence Protocol per AT&T Publication TA-TSY-000772 and 000773
ATM Cell Rate:	96,000 cells/sec. Limited to 80,000 cells/sec. when interfacing with the IPX
Alarms Sent:	AIS Remote
Alarms Received:	AIS Loss of Signal Remote Loss of Framing
Line Errors Counted:	BPV. Parity Bit Errors
Jitter:	Meets ACCUNET T45 specification (Pub 54014)
Connector:	75 ohm BNC
Recommended Cable Lengths:	900 feet (275 m.) max. using specified cable 450 feet (150 m.) to a DS3 crossconnect
Indicators:	Card status Port status

ATM E3 Trunk Interface (BNI-E3, LM-3E3)

Line Rate:	34.368 Mbps \pm 20 ppm, asynchronous
Line Code:	HDB3
Signal Level:	CCITT G.703
Framing Format:	CCITT G.804, G.832
Port Interface:	75 ohm unbalanced
Barrier:	Fully barriered per EN 41003
ATM Cell Rate:	80,000 cells/sec.
Jitter:	per CCITT G.823
ATM Layer Protocol:	per CCITT I.361 with HEC
Port Alarm Processing:	AIS Loss of Signal Remote Alarm Indication Loss of Framing
Line Errors Counted:	BPV. Parity Bit Errors
Connector:	75 ohm BNC
Max. E3 Cable Lengths:	900 feet (275 m.) using specified cable
Indicators:	Card status Port status

ATM OC3 Trunk Interface (BNI-OC3, LM-OC3)

Line Rate:	155.52 Mbps	
Line Code:	NRZ	
Signal Level:	Max	Min
MMF TX	-8 dBm	-15 dBm
MMF RX	-8 dBm	-28 dBm
SMF LR TX	0 dBm	-5 dBm
SMF LR RX	-10 dBm	-34 dBm
Framing Format:	STS-3c, STM1	
Port Interface:	LMI, ILMI	
ATM Cell Rate:	353,208 cells/sec.	
Jitter:	< 0.01 UI p-p, < 0.1 UI rms	
ATM Layer Protocol:	LMI, ILMI	
Port Alarm Processing:	LOS, LOF, LOP, Path AIS, Path Yellow	
Line Errors Counted:	Section BIP8, Line BIP24, Line FEBE, Path BIP8, Path FEBE	
Connector:	MMF SC SMF FC/PC	
Max. Cable Lengths:	MMF ~ 2 KM SMF IR ~20 KM SMF LR ~40 KM	
Indicators:	Card status Port status	

ATM Service Interface (BXM-T3/E3 Cards)

Capacity:	8 or 12 ports per card.
Interface:	DS3/T3 or E3
Line Rate:	DS3 44.736 Mbs, E3 34.368 Mbs
No. of channels per card:	16,000
No. of channels per node:	
VPI Addressing Range:	ATM UNI 3.1 compliant
VCI Addressing Range:	ATM UNI 3.1 compliant
Queues:	16 COS with 32 Virtual Interface (VI) queues

ATM Service Interface (BXM-155 Cards)

Capacity:	4 or 8 ports per card.
Interface:	OC-3c/STM-1
Line Rate:	155.52 Mbs
No. of channels per card:	16,000
No. of channels per node:	
VPI Addressing Range:	ATM UNI 3.1 compliant
VCI Addressing Range:	ATM UNI 3.1 compliant
Queues:	16 COS with 32 Virtual Interface (VI) queues

ATM Service Interface (BXM-622 Cards)

Capacity:	2 ports per card.
Interface:	OC-12c/STM-4
Line Rate:	622.08 Mbps
No. of channels per card:	16,000/32,000
No. of channels per node:	
VPI Addressing Range:	ATM UNI 3.1 compliant
VCI Addressing Range:	ATM UNI 3.1 compliant
Queues:	16 COS with 32 Virtual Interface (VI) queues

ATM Service Interface (ASI-1, LM-2T3

)

Capacity:	2 ports per card
Interface:	T3
Line Rate:	96,000 cells/sec.
No. of channels per card:	1000
No. of channels per node:	1000 or 5000 (grouped)
VPI Addressing Range:	0–255 (UNI), 0-1023 (NNI)
VCI Addressing Range:	1–4095
Queues:	32, 16 per line (port) includes CBR, VBR, and ABR queues.

ATM Service Interface (ASI-1, LM-2E3)

Capacity:	2 ports per card.
Interface:	E3
Line Rate:	80,000 cells/sec.
No. of channels per card:	1000
No. of channels per node:	1000 or 5000 (grouped)
VPI Addressing Range:	0–255 (UNI), 0-1023 (NNI)
VCI Addressing Range:	1–4095
Queues:	32, 16 per line (port) includes CBR, VBR, and ABR queues.

ATM Service Interface (ASI-2, LM-OC3)

Capacity:	2 ports per card.
Interface:	OC3
Line Rate:	353,208 cells/sec.
No. of channels per card:	1000
No. of channels per node:	1000 or 5000 (grouped)
VPI Addressing Range:	0–255 (UNI), 0-1023 (NNI)
VCI Addressing Range:	1–4095
Queues:	24, 12 per line (port) includes CBR and VBR queues

Network Synchronization

Network Clock:	One clock source per domain must be defined. Network clock cannot propagate across domain boundary.
Node Clock:	44.736 MHz, ± 50 Hz (Stratum 3 per AT&T Pub 62411). Internal 8 KHz and 192 KHz clocks for synchronizing network interfaces.
Clock Sources:	Internal, free-running oscillator. Phase-locked to any appropriate network interface. External input at T1 or E1 rate.
Clock Output:	Single clock output at T1 or E1 rate for synchronizing co-located IPX node(s) or CPE.
Clocking Hierarchy:	Network-wide or domain-wide (for structured networks) clock selection.

Network Management

NMS Terminal:	SUN Microsystems and Cisco WAN Manager, Informix, Wingz, and HP OpenView software required for network management, connection provisioning, graphical display of network status, statistics gathering and display, automatic downloading of switch software, and save and restore of node configurations.
Control Terminal:	DEC VT100 or equivalent terminal for basic system configuring and alarm monitoring.
Remote Alarm Reporting:	Autodial modem may be connected to the auxiliary port on one or more IPX or IGX nodes for automatic reporting of network alarms.
Remote Diagnostics:	Autoanswer modem may be connected to the control port one or more IPX, IGX, or BPX nodes for remote diagnostic access by Cisco Customer Service or other authorized personnel.
Network Control Ports:	Three ports per node, two with RS 232C interface, one with Ethernet LAN interface. (Nodes with PCC controller have only the two RS-232 ports).
Alarm Notification:	Status of all trunks and nodes in network distributed to and stored by each individual node. Reported to Cisco WAN Manager terminal at connecting node.
External Alarms:	Meets Bellcore Compatibility Bulletin #143 and AT&T Technical Reference PUB 43801 DS1 (T1) facility alarm requirements.

A

A

A-bit (active bit)	The bit in the frame relay frame header that indicates the status of the far end user device and the status of the PVC segment in the foreign network.
A-law	An analog to digital encoding scheme used to convert voice samples to an 8-bit data word used in CEPT E1 multiplex equipment. (See also μ -law.)
ABR (Available Bit Rate)	ATM connection type for bursty traffic, such as data. Provides closed loop control of service rate that allows connections to use additional bandwidth when available. ABR may be used with ATM Traffic Management 4.0 standards VSVD flow congestion control, or with the proprietary ForeSight flow congestion control. See also, CBR and VBR.
ACO (Alarm Cut Off)	A switch to turn off the audible alarm outputs from a node while leaving the visual alarm outputs unchanged.
adaptive voice	An optional feature of the IPX that disables VAD from connections using it whenever there is excess bandwidth available to allow the normal encoded voice to be carried on the packet line. See also VAD.
ADPCM (Adaptive Differential Pulse Code Modulation)	A compression method that samples voice 8,000 times per second, and uses the changes between samples as the basis for compression. Increases the capacity of a T1 line from 24 to 48 channels.
ADTF (Allowed Cell Rate Decrease Factor)	Time permitted between sending RM cells before the rate is decreased to ICR.
AIT (ATM Interworking Trunk Card)	The AIT front card provides an ATM trunk interface for the IPX. The AIT operates in conjunction with a backcard, AIT-T3 or AIT-E3.
AIT-E3 (ATM Interworking Trunk E3 Interface Card)	The AIT-E3 backcard provides an E3 interface for the AIT (IPX) or BTM (IGX) ATM trunk cards.
AIT-T3 (ATM Interworking Trunk T3 Interface Card)	The AIT-T3 backcard provides a T3 interface for the AIT (IPX) or BTM (IGX) ATM trunk cards.
alternate routing	An automatic rerouting of a failed connection by a node to a new route through the network to maintain service.
AMI (Alternate Mark Inversion)	The line code used for T1 and E1 lines where the “1s” or “marks” on the line alternate between positive polarity and negative polarity.

A

arbiter	A BPX administration processor that polls each network port to control the data flow in and out of the crosspoint switch matrix.
ARC (Alarm Relay Card)	An alarm front card for the IPX.
ARI (Alarm Relay Interface Card)	An alarm interface back card for the IPX and IGX.
ARM (Alarm Relay Module)	An alarm front card for the IGX.
ASM(Alarm/Status Monitor Cards)	An alarm front card and back card set for the BPX.
ATM (Asynchronous Transfer Mode)	Data transmission that uses a very flexible method of carrying information, including voice, data, multimedia, and video between devices on a local or wide area network using 53-byte cells on virtual circuits. The 53 byte cell consists of data and a small header. See also cell relay.
ATM Switched Virtual Circuits (SVCs)	A member of the INS product family that uses ATM SVC Server Shelves and software to enhance a Cisco WAN switching network with ATM switched virtual circuits.
ATM SVC Server Shelf	An adjunct processor used in the INS ATM SVC application to enhance traditional Cisco WAN switching networks with ATM switched virtual circuits. The ATM SVC Server Shelf is co-located with and connected to a BPX.
auxiliary port	An RS-232 port on the front panel of the SCC card used for connecting a printer or an out-dial modem. This port is a one-way, outgoing port.

B

B3ZS (Bipolar with Three Zero Suppression)	A protocol for T3 lines that converts a channel word with three consecutive zeros into a code which, at the far end, is converted back to three zeros.
B8ZS (Bipolar with Eight Zero Suppression).	A T1 line protocol that converts a channel word with eight consecutive zeros into a code which, at the far end, is converted back to eight zeros. Allows 64 Kbps clear channel operation while assuring the ones density required on the T1 line.
bandwidth reservation	An IPX software feature that allows circuits to automatically become active (or “upped”) at a specified time and date and downed at some later time and date. For circuits that do not need to be available 100% of the time.
B channel	In ISDN, a full-duplex, 64-kbps channel used to send user data. Also known as the bearer channel. Compare with D channel.
BCC	The control card for the BPX
BC-E1 (Backcard E1)	E1 interface card used on IPX and IGX
BC-E3 (Backcard E3)	E3 interface card used on IPX and IGX
BC-J1 (Backcard J1)	J1 interface card used on IPX and IGX
BC-SR (Backcard Subrate)	Subrate interface card used on IPX and IGX
BC-T1 (Backcard T1)	T1 interface card used on IPX and IGX
BC-T3 (Backcard T3)	T3 interface card used on IPX and IGX
BC-Y1 (Backcard Y1)	Y1 interface card used on IPX and IGX
BDA (Bframe Destination Address)	The address of the slot.port.channel for which the Bframe is destined. This address is part of the Bframe header and is only used across the switch fabric locally in the node.
Bframe	The BPX frame is the 64-byte format for messages used to encapsulate ATM cells which are sent across the switch fabric.
bipolar violations	Presence or absence of extra “1” bits on a T1 transmission facility caused by interference or a failing line repeater. These extra or missing bits interrupts one of the rules for bipolar pairs of a digital transmission line.
BISDN (broadband ISDN)	ITU-T communication standards designed to handle high-bandwidth applications. Compare with ISDN.

B

BNI (BPX Network Interface Card)	The front card used to network BPX switches together and to connect to MGX 8220 shelves, and IPX and IGX nodes configured as shelves. Supports T-3, E-3, and OC3 trunks carrying ATM cells.
BPX	A high-speed broadband, high-capacity ATM cell relay network switch from Cisco for private and public networks.
BRI (Basic Rate Interface)	ISDN interface composed of two B channels and one D channel for circuit-switched communication of voice, video, and data. Compare with PRI.
bundled connections	Frame relay connections grouping a number of ports into one permanent virtual circuit.
BTM (Broadband Trunk Module)	The BTM provides an ATM trunk interface for the IGX. The BTM operates in conjunction with a backcard, AIT-T3 or AIT-E3.
BXM	A series of BPX cards, BXM-T3/E3, BXM-155, or BXM-622 which can be configured for either trunk or line (service access) modes. These cards support ATM Traffic Management 4.0, including VSVD congestion flow control.

C

CAC (Connection Admission Control)	A function exercised by switches during connection setup to determine whether or not a connections requested QoS will violate QoS for existing connections.
CAS (Channel Associated Signalling)	A signalling mode in E1 transmission where the signalling bits for all 30 E1 channels are carried in timeslot 16. Timeslots 1 to 15 and 17 to 31 carry encoded voice bits only.
CBR (Constant Bit Rate)	ATM Connection type for constant bit rate traffic such as voice or synchronized data requiring a low variation in delay. See also, VBR and ABR
CCDV (Compliant Cell Delay Variation)	A parameter utilized in defining ATM Constant Bit Rate service. The amount of delay that is acceptable between ATM cells for them to be accepted as compliant (usable).
CCITT (Consultative Committee for International Telephone and Telegraph)	An International telecommunications advisory committee established under the United Nations to recommend worldwide standards for data and voice communications.
CCS (Common Channel Signalling)	A carrier signalling mode in E1 transmission where signalling bits are not used. CCS typically separates user data from signalling information. A signalling channel is used to provide signalling for all other user data channels in the system.
CDP (Channelized Data PAD)	An IPX dual-purpose front card that can carry voice traffic, a combination of voice and data, or just data. The CVM card is used in conjunction with a BC-T1, BC-E1, or BC-J1 backcard.
CDVT (Cell Delay Variation Tolerance)	Controls time scale over which the PCR is policed.
Cell	A unit of data with a fixed number of bytes. For ATM the cell size is 53 bytes.
cell relay	A form of digital communications using fixed length cells consisting of data and a small header IPX FastPacket was an early implementation of cell relay. The 53 byte ATM cell consists of data and a small header.
CEPT	CEPT is the European Conference of Posts and Telecommunications Administrations. This association is comprised of European Telecommunications service providers that participate in relevant areas of the work of CEN/CENELEC.
CGA (Carrier Group Alarm)	A major alarm condition for a T1 multiplexer or PABX that results in all channels being taken out of service.
channel	The logical end point for a connection.
circuit line	A T1 or E1 line that connects a user device, such as a PABX or channel bank to the IPX. Carries customer DS0 voice and data circuits. See also line.
clear channel capability	When all eight bits of a channel word in the T1 line signal are available for transmitting customer data with no restrictions on content. Also referred to as 64 Kbps clear channel.

C

Cmax	A frame relay connection parameter that specifies the number of packets allowed in the initial burst of data after which the data bandwidth is reduced to the connection's minimum specified bandwidth.
CLLM	Consolidated Link Layer Management. A protocol used to transmit ForeSight messages across the frame relay NNI port.
CLP (Cell loss priority)	Cell loss Priority. CLP Hi and CLP Lo thresholds are configurable.
Complex Gateway	Refers to interworking of a connection with respect to the IPX and IGX nodes. For example, in a Frame Relay to ATM interworking, the Frame Relay data is extracted from FastPackets and transformed to ATM cells with redundant overhead bits discarded.
composite data rate	The sum of the data rates for all circuits transmitting on the same synchronous or frame relay data card.
control port	An RS-232 port on the face plate of a back card for a controller card, (BCC, NPC, NPM) that may be used for connecting a control terminal. This port is bi-directional.
COS (Class of Service)	The priority assigned each user connection. Defines which circuits get rerouted first during a network failure.
courtesy downing	A software feature that is used to conserve network bandwidth by automatically “downing” a voice connection when the signalling status indicates an inactive (on-hook) circuit. The circuit is automatically “upped” when the circuit becomes active.
CRC (Cyclical Redundancy Check)	A method of error checking that detects errors in a block of data. Unlike parity checks, the CRC can detect multiple data errors within the block and thus equipment using a CRC error check can derive an error rate.
crosspoint switch	A two-dimensional data switch type that is arranged in a matrix of all input connections along one axis and all output connections along the other axis. Each input and output line has a switch point where the two axes intersect that can be enabled (switch closed) or disabled (switch open). The central matrix switch providing the switching matrix for traffic routing by the BPX node.
CSU (Channel Service Unit)	A network protection unit that terminates any T1 span line connected to the carrier's central office, providing receive direction regeneration and maintenance loopback for the 1.544 Mbps signal.

D

D4-format	A digital signal format with 24 eight-bit channels plus one synchronizing bit per T1 line. Channels are assigned in a straight, numeric sequence.
DACS (Digital Access and Control System)	Equipment, usually found in the telephone company central office, that is used to groom and retime the 24 channels in a DS1 signal. Individual DS0 channels can be cross-connected from one DS1 source and inserted in another DS1 source either with the same or with a different channel number.
DAS Server Shelf	The adjunct processor used in INS Dial-Up Frame Relay applications to provide frame relay dial-up and dial-backup circuits. The DAS Server Shelf is co-located with and connected to an IPX or IGX.
DCE (Data Communications Equipment)	As defined by the RS-232 standard, any device that transmits or receives information. Usually used with data terminal equipment (DTE, like a computer or network node).
D channel	A message-oriented ISDN signalling channel, typically carried in DS24 of a PRI on T1 facilities or TS16 of a PRI on E1 facilities. Compare to B channel.
DDS (Digital Data Service).	An AT&T dial-up data service offering for 2.4 to 56 Kbps over subscriber loop cable. Requires a Data Service Unit, DSU, at customer premise for interface to the DDS trunk.
Device Code	The first 8 bits of a FastPacket Address.
DFM (Data Frame Multiplexing).	An optional feature that saves data channel bandwidth by analyzing data channel content and suppressing repetitive data patterns.
Dial Access Switching	Another name for the INS Dial-Up Frame Relay application.
Dial-Up Frame Relay	An INS application that uses a DAS Server Shelf and software to enhance CiscoWAN switching networks with frame relay soft permanent virtual circuits (SPVCs) for dial-up dial-backup connections.
DLCI (Data Link Connection Identifier)	A field in a frame relay data packet that identifies the destination for the data.
domain	A grouping of nodes sharing common interests or attributes.
domain name	A unique name consisting of the letter “D” immediately followed by a number (1–8) delineated by a “.” (period) from the node name (1–8 characters maximum). Example: D1.alpha
domain number	A number from 1–8 assigned with the cnfdmn command. The number assigned is part of the domain name.
DPNSS	Digital Private Network Signalling System. A common-channel message-oriented signalling protocol commonly used by private branch exchanges (PBXes). The INS Voice Network Switching application supports DPNSS signalling.
DS0 (Digital Signal 0)	A 64 Kbps channel used to transmit encoded voice and/or data. There are 24 DS0 channels in a circuit T1 (DS1) line. DS0 data is transmitted using one or more DS0 circuits in a T1 or E1 circuit line.

D

- DS0A** An extension of DS0 that defines the format for assembling various low-speed data circuits (1.2 to 19.6 Kbps) into a single 64 Kbps DS0 channel.
- DS1 (Digital Signal 1)** A digital transmission standard that carries 24 individual channels in a bipolar, high-speed line signal at 1.544 Mbps. DS1 signal level is $\pm 3V$.
- DSI (Digital Speech Interpolation)** An algorithm that analyzes DS0 voice bits for non-speech codes. Suppresses these bits to conserve packet line bandwidth and inserts a code to indicate to the far end that these bits have been removed. Similar to DFM for data channels. Also, referred to as VAD (Voice Activity Detection).
- DTE (Data Terminal Equipment)** As defined by the RS-232 standard, any device that generates or utilizes information. See also, DCE.

E

E

E1	European transmission service at the rate of 2.048 Mbps
E3	Transmission service at a rate of 34.368 Mbps
ECN (Explicit Congestion Notification)	A frame relay feature to signal the onset of network congestion to external devices. Sets FECN and BECN bits in frame relay header to indicate forward and backward congestion.
EFCl (Explicit Forward Congestion Indication)	
ICR (Initial Cell Rate)	The rate at which a source should initially send after an idle period

F

Fast EIA	Same as interleaved EIA. Seven data circuit control leads in each direction are transmitted in alternating bytes with data. For fast control lead response to data being turned on and off but with a sacrifice in packet line bandwidth.
FBTC (Frame Based Traffic Control)	An AAL5 frame based traffic control that provides the possibility of discarding the whole frame, not just one compliant cell. This avoids wasting bandwidth by continuing to send the cells in a frame once a cell has been found to be non-compliant.)
FGCRA (Frame Based Generic Cell Rate Algorithm)	An enhancement option to GCRA that allows an entire frame to be discarded if any of its cells are non-compliant, rather than transmitting a partial frame over the network.
flat network	A non-structured network, a network in which there are no junction nodes or domains.
foreign network	An adjacent network that is owned and managed by a different party than the one that owns the local network.
ForeSight	A proprietary optional feature that uses feedback techniques to dynamically allocate extra bandwidth to frame relay and ATM connections when the network bandwidth is available and not used by other connections. See also VSVD.
frame forwarding	A software feature allowing point-to-point frame relay type connection for various data applications that do not conform to the Frame Relay Interface Specification.
FPC (FastPAD Back Card)	The FPC is used with an FTC (IPX) or FTM (IGX) card. The FPC provides either a T1, E1, V.35, or X.21 interface.
frame relay connection class	A tag for a frame relay circuit which indicates the class of service to be provided for this connection. Parameters associated with a connection class include minimum information rate guaranteed, peak information rate expected, maximum network delay, etc.

F

FRI (Frame Relay Interface Card)	The backcard for an FRP (IPX) or FRM (IGX) card. The FRI provides V.35, X.21, T1, or E1 interfaces.
FRP (Frame Relay PAD)	An IPX frame relay front card that supports 1-4 data ports, and in single-port mode, operates up to 2.048 Mbps. The card is used in conjunction with FRI-V.35, X.21, T1, or E1 backcards.
FRM (Frame Relay Module)	An IGX frame relay front card that supports 1-4 data ports, and in single-port mode, operates up to 2.048 Mbps. The card is used in conjunction with FRI-V.35, X.21, T1, or E1 backcards.
FRM-2 (Frame Relay Module)	An IGX frame relay front card that provides an interface to the frame relay Port Concentrator Shelf (PCS). The card is used with the FRI-2-X.21 backcard which connects to the PCS.
FRP-2 (Frame Relay Module)	An IPX frame relay front card that provides an interface to the frame relay Port Concentrator Shelf (PCS). The card is used with the FRI-2-X.21 backcard which connects to the PCS.
FRP-2 (Frame Relay Module)	An IPX frame relay front card that provides an interface to the frame relay Port Concentrator Shelf (PCS). The card is used with the FRI-2-X.21 backcard which connects to the PCS.
frame relay Service	A packet interface data transmission protocol used for connecting widely-separated LANs. Characterized by long intervals of no data to be sent interspersed with bursts of large volumes of data; sometimes referred to as “bursty data”.
frame slip	A T1 error condition caused by a timing problem between the network and the IPX. When this happens, the IPX inserts a blank DS1 frame or drops an idle DS1 frame so there is no loss of customer data.
FRTT (Fixed Round Trip Time)	The sum of the fixed and propagation delays from the source to a destination and back.
Full Status Report	A message sent across the NNI indicating the A-bit status of all connections routed across this NNI frame relay port.
FTC (FastPAD Trunk Card)	An IPX frame relay front card that provides an interface to a FastPAD. The FTC is used with an FPC backcard. that provides either a T1, E1, V.35, or X.21 interface.
FTM (FastPAD Trunk Module)	An IPX frame relay front card that provides an interface to a FastPAD.

G

gateway	An IPX node that is configured to handle both T1 and E1 packet and circuit lines for direct interface international circuits. (See also Seamless International IPX Network.)
GCRA (Generic Cell Rate Algorithm)	GCRA is a “continuous leaky-bucket” process that monitors the cell depth in the input queue for each PVC to determine whether to admit a new cell to the network without setting the Cell Loss Priority bit.
global addressing	A frame relay addressing convention that uses the DLCI to identify a specific end device somewhere else in the frame relay network. In a global addressing scheme, the DLCI is a unique number for each port in the network.
grouped connections	Frame relay connections grouping a number of ports onto one permanent virtual circuit. Similar to bundled connections except the grouped connections do not have to be contiguous, nor do they all have to be added simultaneously.

H

HDB3 (High Density Bipolar Three)	A new line interface for E1, similar to B8ZS for T1, which eliminates patterns with eight or more consecutive zeros. Allows for 64 Kbps clear channel operation and still assure the ones density required on the E1 line.
HDP (High Speed Data PAD)	An IGX front card that supports one to four medium speed, synchronous data channels.

I

IGX	A multi-service, multi-band ATM cell relay network switch from Cisco for private and public networks.
Intelligent Network Server (INS)	INS is the broad name for a range of products that enhance traditional CiscoWAN switching networks. These products include Dial-Up Frame Relay, Voice Network Switching, and ATM Switched Virtual Circuits.
intra-domain	Connections within a domain including but not going beyond the junction nodes.
inter-domain	Connections between domains through junction nodes.
interleaved EIA	Same as “Fast EIA”.
IPX	A narrowband cell relay network switch from Cisco for private and public networks.
ISC (International Support Center)	The customer support group at Cisco that provides assistance in solving network or equipment problems over the telephone
ISDN (Integrated Services Digital Network)	A service provided by the telephone company or OCC that supports combined customer voice and data connections over the twisted pair subscriber loop. Requires special equipment at the customer premises and a connecting central office switch that is capable of providing ISDN.

J

- J1** A multiplexed 24-channel circuit line to a PBX conforming to the Japanese TTC-JJ-20 circuit standard. Similar to E1, it operates at 2.048 Mbps.
- junction node** A node handling inter-networking of domains.
- junction trunk** A packet line connecting junction nodes.

L

- LCON** The logical connection used to represent an individual routing entity.
- LDM (Low Speed Data Module)** An IGX data front card that supports up to 8 synchronous or asynchronous data ports. When used with an LDI4/DDS, an LDP can provide 56-Kbps Digital Data Service (DDS) interfaces to the IGX.
- LDP (Low Speed Data PAD)** An IPX data front card that supports up to 8 synchronous or asynchronous data ports. When used with an LDI4/DDS, an LDP can provide 56-Kbps Digital Data Service (DDS) interfaces to the IPX.
- LEC (Lower Expansion Card)** An expansion back card for the IPX32 that connects upper shelf bus to lower shelf bus and the active NPC to standby NPC.
- line** Connects a user device to a service interface, for example, a router to an ASI or AUSM card, a data line to a data card, a frame relay line to an FRP or a port concentrator, or a T1 or E1 line to a CDP card.
- link** The network connection between two nodes.
- LMI (Local Management Interface)** The protocol and procedures for control of IPX frame relay connections. Used for configuration, flow control, and maintenance of these connections.
- local addressing** A frame relay addressing convention that uses the DLCI to identify the IPX frame relay port at the interface between the user device and the frame relay network. In local addressing, a particular DLCI is used only at the local FR connection. The DLCI may be reused at any other IPX node in the network.
- local alarm** An IPX alarm indicating that the associated T1 line is down due to a local failure of the its receive path.
- local bus** An IPX utility bus (LB/0 or LB/1), located on the midplane, which provides the electrical connections between various front and back cards. For example, the front and back cards of the Low Speed Data PAD group (LDP and LDI) plug into this utility bus.
- logical port** A frame relay circuit consisting of either 1, 6, 24 (T1), or 31 (E1) contiguous DSO's on a T1 or E1 physical port.

M

major alarm	A local or remote failure that is affecting operation of the network.
MBS (Maximum Burst Size)	Maximum number of cells which may burst at the PCR but still be compliant.
MCR (Minimum Cell Rate)	The minimum cell rate that is supported by an ATM connection for an ABR connection.
MIR (Minimum Information Rate)	The minimum information rate that is supported by a frame relay connection.
minor alarm	A local or remote failure that is not affecting operation of the network, but nonetheless should be investigated.
MUXBUS	A high-speed IPX backplane bus that carries data and timing between card slots for both circuit line and packet line data. Consists of the TDM bus carrying the data and the system clock bus that is used to synchronize all data flowing on and off the TDM bus.

N

n+1 redundancy	A redundancy method in which a group of cards share the same standby redundant card.
Network-to-Network Interface (NNI)	The protocol at a frame relay port that serves as a bidirectional interface between a local Cisco WAN switching network and a separate and independent “other” network.
node	An IPX/IGX/BPX serving as a connection point to the network. At a node, connections from service lines are routed to trunks for transmission to other nodes in the network.
NPC (Network Processor Card)	Micro-processor based system controller front card that contains the software used to operate the IPX.
NPM (Network Processor Module)	Micro-processor based system controller front card that contains the software used to operate the IGX.
Nrm	Maximum number of cells a source may send for each forward RM cell, i.e., an RM cell must be sent for every Nrm-1 data cells.
NTC (Network Trunk Card)	IPX front card that coordinates fastpacket trunk traffic to another node via a number of backcards: T1, E1, Y1, and subrate (RS449, X.21, and V.35)
NTM (Network Trunk Module)	IGX front card that coordinates fastpacket trunk traffic to another node via a number of backcards: T1, E1, Y1, and subrate (RS449, X.21, and V.35)

O

- OC-3** Standard optical transmission facility rate of 155.20 Mbps.
- OCC (Other Common Carrier).** In the United States, reference to all the other telecommunications companies providing various transmission services other than AT&T.

P

- packet line** Packet line referred to a line used to carry FastPackets between IPX nodes in a network. The term in these documents is replaced by the more general “trunk” which is defined as a physical link from node to node, node to shelf, or node to network. The trunk may be one that supports 24-byte FastPacket (packet trunk), or one that supports 53 byte ATM cells (cell trunk).
- packet switching** A system that breaks data strings into small units (packets), then individually addresses and routes them through the network.
- PAD (Packet Assembler/Disassembler)** A device that converts a serial data stream into discrete packets in the transmit direction and converts the received packets back into a serial data stream. Adds header information in the transmit packet to allow it to be routed to the proper destination.
- partially-interleaved EIA** One control lead in each direction, generally RTS-CTS, is transmitted in same byte as seven data bits. For fast control lead response to data being turned on and off.
- PBX (private branch exchange)** Digital or analog telephone switchboard, classified as customer premise equipment (CPE), used to connect private and public telephone networks
- PCM (Pulse Code Modulation)** The system for transmitting telephone signals digitally. Voice is sampled 8000 times per second and converted to an 8-bit digital word.
- PCR (Peak Cell Rate)** The maximum rate for an ATM connection at which cells are allowed into the network.
- PCS (Port Concentrator Shelf)** The PCS is an external shelf that expands the capacity of the FRP card. The PCS is used with the FRP-2 (IPX) or FRM-2 (IGX) card to 44 frame relay connections. The PCS connects to the FRI-2.X.21 backcard.
- PIR (Peak Information Rate)** The peak level in bits per second allowed for a frame relay connection.
- PLCP (Physical Layer Convergence Protocol)** A protocol defined for use with Switched Megabit Data Service. Used on DS3 ATM trunks in the BPX.
- PLPP (Physical Layer Protocol Processor)** A custom VLSI processor used in the T3 ATM port interface of the BPX BNI card to handle the coding and decoding of the PLCP bit structure. Functions handled by the PLPP include header check sequence generation and checking, DS3 framing, and optional payload scrambling/descrambling.
- plesiochronous network** A network where there is more than one source of network timing. The multiples sources must be operating at the same frequency but are not phase locked (synchronous) with each other.

P

PNNI (Private Network to Network Interface)	<p>PNNI consists of two parts :</p> <ol style="list-style-type: none">1. Protocol for distributing topology information between switches and groups of switches.2. Protocol for signaling process for establishing connections across a network.
port	<p>Refers to a signal connection on a data back card that interfaces to a customer circuit or data device. The number of ports on a card ranges from 1 to 8 depending on the particular card type.</p>
PRI (Primary Rate Interface)	<p>An ISDN interface to primary rate access. Primary rate access consists of a single D channel for signalling and 23 (T1) or 30 (E1) B (bearer) channels for user data. A PRI is typically carried on T1 or E1 facilities.</p>
privilege level	<p>A level between 1 and 6 that is assigned to each IPX command. Each operator is assigned a privilege level by the system administrator. The operator may only access and execute commands equal to or lower than his or her own privilege level. Level 1 is the highest and level 6 is the lowest</p>
PVCs	<p>Permanent Virtual Connections (circuits). Connections that are assigned but not connected until data is sent, thereby not using bandwidth when idle.</p>

Q

- Q.921/Q.931** ITU-T specifications for the ISDN user network interface (UNI) data link layer.
- QSIG** A common-channel message-oriented signalling protocol, defined by the European Telecommunications Standard Institute (ETSI), commonly used by private branch exchanges (PBXes). The INS Dynamic Network Switching application supports QSIG signalling to the Cisco WAN switching network.
- queue** A buffer that is used to temporarily hold data while it waits to be transmitted to the network or to the user.

R

- RIF (Rate increase factor)** Controls the amount by which the cell transmission rate may increase upon receipt of an RM cell.
- RDF (Rate decrease factor)** Controls the amount by which the cell transmission rate may decrease upon receipt of an RM cell.
- red alarm** Another name for local alarm as the local alarm lamp on most digital transmission equipment is red in color.
- remote alarm** An IPX alarm indicating that the associated T1 line is down due to a receive line failure on another node. See also yellow alarm.
- RPS (repetitive pattern suppression)** Also called data frame multiplexing (DFM). An option for data circuits where repeating strings of data are replaced on the packet line by a single occurrence of the data string and a code that indicates to the far end how many repetitions of the string were being transmitted. Used to conserve network bandwidth.
- robbed bit signaling** A type of signaling used on T1 lines where the signaling bits for each channel are substituted for the least significant voice bit in each channel word during frames 6 and 12.
- RS-232** A physical and electrical interface standard for a low-speed, unbalanced, serial, data interface adopted by the EIA committee on data communications. Generally used for data circuits operating at data rates below 56 Kbps.
- RS-422/423** Another EIA standard electrical interface for serial data circuits operating at higher data rates than RS232. RS422 is a balanced interface; RS423 is unbalanced. Uses RS-449 for the physical interface (connector).
- RS-449** The physical interface for the RS422 and RS423 electrical interfaces. Contains the Processor Controller Card and the PCC utility bus, and provides system timing and control via the system bus.

S

SAR (Segmentation and Reassembly)	The process of breaking a dataframe containing data from a number of virtual paths or circuits apart so that the individual paths/circuits can be switched by reassembling the data into a new frame with a different sequence.
SCC (System Clock Card)	An IPX backcard that works in conjunction with the NPC. The SCC provides a centralized clock generation function and provides serial and LAN port interfaces.
SCM (System Clock Module)	An IGX backcard that works in conjunction with the NPM. The SCM provides a centralized clock generation function and provides serial and LAN port interfaces.
SCR (Sustainable Cell Rate)	Rate above which incoming cells are either tagged or discarded.
SDP (Synchronous Data PAD)	An IPX front card that supports one to four medium speed, synchronous data channels.
SDI (Synchronous Data Interface)	The back card for the SDP (IPX) or HDM (IGX) cards. The SDI is available with V.24, X.21, and V.35 interfaces.
seamless international network	An IPX network that is configured to carry traffic over international borders (E1-T1 or T1-E1)—see also gateway.
Simple Gateway	Refers to FastPacket to ATM interworking with respect to the IPX and IGX nodes. In the simple gateway mode, FastPackets are encapsulated in their entirety into cells. See also, complex gateway.
SIU (Serial Interface Unit)	A set of circuits common to all BPX cards used for transmitting and receiving via the crosspoint switch.
Soft PVC	A PVC in the INS Dial-Up Frame Relay application that is dormant in the networks database until is activated by a call into the network by a user.
spanning tree	An IPX network topology in which there is only one path available between any two sources in a frame relay multicast group. Spanning trees are required to prevent frames broadcast from a single source to multiple receptors from circulating endlessly around the network a result of frame relay circuits not having properly closed loops.
speech detection	Determining the presence or absence of speech for Digital Speech Interpolation. Performed in either the CDP card or VDP card in an IPX node.
split clock	A data clocking configuration where the timing for the transmit data is obtained from one source (e.g. user device) and the timing for the receive data is obtained from another source (e.g. IPX).
Status Enquiry	A message transmitted by a FR NNI port requesting an updated status from the attached foreign network. This message is used as a heartbeat to detect when a port has failed.
StrataBus	On the BPX, contains crosspoint wiring used to carry ATM trunk data between both the network interface and service interface modules and the crosspoint switch as well as providing control, clock, and communications.

S

StrataView Plus	A Unix-based workstation and software used as a network management system (NMS) for Cisco WAN switching networks. It is part of the StrataSphere group. Provides a graphical user interface for configuration, maintenance, administration of the network. Collects and displays network statistics. Now called Cisco WAN Manager.
StrataSphere	A standards based multi-protocol management architecture that includes CiscoWAN Manager, StrataSphere Connection Manager, StrataSphere BILLder, StrataSphere Modeler, and StrataSphere Optimizer.
subrate data	Multiple low-speed data circuits carried in a single DS0 timeslot.
superrate data	Single high-speed data circuit carried in multiple DS0 timeslots.
SCR (Sustained Cell Rate)	Long term limit on the rate a connection can sustain.
SVC (switched virtual circuit)	A virtual circuit that is dynamically established on demand and torn down when transmission is complete. SVS do not need to reserve any network resources when they are not in use. Called a switched virtual connection inATM terminology. Compare with PVC.
system bus	A two-part IPX data bus. One part carries system commands between the PCC all other IPX cards; the other carries time division multiplexed data.

T

TT1	The standard US. multiplexed 24-channel voice/data digital span line. Operates at a data rate of 1.544 Mbps.
T3	Transmission service at DS3 rate of 44.736 Mbps
TBE (Transient Buffer Exposure)	The negotiated number of cells that the network would prefer to limit the source to send during the start-up period.
TDM (time division multiplexing)	The process of combining several communication channels by dividing a channel into time increments and assigning each channel to a timeslot.
timestamp	A field in certain FastPacket formats that indicates the amount of time the packet has spent waiting in queues during the transmission between its source and destination nodes. Used to control the delay experienced by the packet.
Trm	An upper bound on the time between RM cells for an active source, i.e., RM cell must be sent at least once every Trm msec.
trunk	A physical link between two nodes. The trunk may be one that supports 24-byte FastPackets (packet trunk), or one that supports 53 byte ATM cells (cell trunk)
trunk conditioning	A set of signalling and information bits that indicate a DS1 line failure.
trunk queues	The buffers in packet line cards (NTC, TXR) where the various FastPackets are queued up for transmission over the packet line(s). The buffers attempt to prioritize each packet so it experiences minimum delay.

S

μ-law	An analog to digital encoding scheme used to convert voice samples to an 8-bit data word used in D3/D4 T1 multiplex equipment.
UBR	Unspecified Bit Rate
UNI (User to Network Interface)	The user to network interface, as for ATM connection to CPE. See also NNI.
UPC (Usage Parameter Control)	A general procedure for controlling the rate of user data applied to an ATM network. There are a number of different algorithms for performing UPC. See also GRCA.
USART (Universal Synchronous/Asynchronous Receiver Transmitter)	A single-chip device used in certain applications that allows microprocessors to communicate with input/output (I/O) devices.
User to Network Interface (UNI)	The protocol at a frame relay port that passes information between the network and the user device attached to the port.

V

- V.21** A CCITT interface standard often used for data transmission over modems.
- V.35** A data communications interface standard adopted by the CCITT. Often used for data circuit operating at 56 Kbps and above
- VAD (Voice Activity Detection).** Used to statistically compress voice by not sending packets in the absence of speech.
- VBR (Variable Bit Rate)** Connection type for variable bit rate traffic such as bursty data. See also, CBR and ABR.
- VC QDepth** VCQFrame relay buffer allocation parameter that specifies the maximum queue size reserved in the FRP card for the FR connection.
- virtual circuit** A circuit that acts like it is an individual transmission path but is actually shared with other circuits over a single transmission path. See also PVCs.
- VNS** The adjunct processor used in the INS Voice Network Switching application. The VNS is co-located with and connected to an IGX or IPX.
- Voice Network Switching** An INS application used to provide voice or data switched virtual circuits over a CiscoWAN switching network for PBXes using either QSIG or DPNSS signalling.
- VS/VD (Virtual Source/Virtual Destination)** ATM Forum Traffic Management 4.0 method of providing congestion flow control for ABR connection types. Resource Management (RM) cells are used to convey management information between sources and destinations.
- vt (virtual terminal)** An IPX control terminal that is the active control terminal at one node but is physically attached to another node.

W

- WAN (Wide Area Network)** A network of transmission circuits generally spanning a large region or territory for transmission of voice and data between widespread end users. An IPX network is an example of a WAN.

X

- X.21** A CCITT standard for data interfaces transmitting at rates up to approximately 2 Mbps.
- X.25** A commonly-used standard that defines the protocol for low-speed data packet networks
- XON/XOFF** A simple communications protocol for controlling the flow of data from one device to another. An XON sent from a receiving device indicates it is ready to accept data and the transmitting device may begin to output data. An XOFF from the receiving device indicates that it can no longer store any more data and the transmitting device should temporarily cease transmitting.

Y

- Y-cable(s)** A short adapter cable forming an electrical branch (thus the term Y) for connecting a single customer data or trunk connection to two identical back cards to provide hardware redundancy on the IPX.
- Y-cable redundancy** A redundancy type used in the IPX when a 1:1 card redundancy is implemented using a split or Y-cable for the data connection between the user device and the primary and standby IPX interface card.
- Y1** A digital trunk conforming to the Japanese “Y” circuit standard, for use as a packet line. Similar to T1, it operates at 1.544 Mbps.
- yellow alarm** Another name for remote alarm as the remote alarm lamp on digital transmission equipment is always yellow in color.

Symbols

%Util **10-9**

A

AAL **8-8**

classes of traffic **8-8**

AAL5 Frame Base Traffic Control **10-9**

ABR (Available Bit Rate) **10-2**

ABR and ATFST Connections **10-1**

ABR Connection Policing **10-3**

ABR STD **10-1**

active bit (A-bit) **12-1**

adaptation

FastPacket to ATM **8-11**

adaptation layers

and traffic classes **8-8**

Adaptive Differential Pulse Code Modulation

See ADPCM

Adaptive Voice feature **15-8**

address information **2-2**

addshelf command **11-8**

ADPCM **2-6, 4-8, 4-17**

voice compression

ADPCM **15-7**

ADTF **10-1**

AIT card **8-11**

AIT card set **5-4**

alarm outputs **4-23**

alarm reporting, automatic **7-1**

and ATM **8-16**

ASI-1 card

Usage Parameter Control **4-3**

ASI-1 card set **5-10**

Asynchronous Transfer Mode

general overview **8-1**

network connection types **5-8**

services supported **5-5**

asynchronous transfer mode

definition **1-25**

AT3-6M Interface Adapter **5-8**

ATM

and T3 trunks **8-15**

cell headers **8-3**

E3 rate **8-16**

ingress queues **4-3**

payload length **8-3**

ATM (see Asynchronous Transfer Mode)

ATM, theory of **8-1**

ATM adaptation layer **8-8**

ATM connections

definition **4-3, 10-**

ATM Forum Traffic Management 4.0 **10-**

ATM layers **8-2**

ATM NIC **11-**

ATM physical layer **8-2**

ATM Service Interface **5-10**

ATM switch

and WAN networks **8-12**

definition **8-12**

ATM switching

advantages **8-12**

ATT 4ESS **4-5**

ATT 5 ESS **4-5**

AUSTEL **4-5**

AutoRoute **1-23, 1-34**
 AXIS Command Reference **xxix**
 AXIS Reference Manual **xxix**
 See Related documentation

B

BAM **8-10**
 bandwidth **12-3**
 Bandwidth Reservation **4-19**
 BC-J1 back card **5-15**
 BC-Y1 back card **5-15**
 BDS **5-10**
 B-ISDN
 definition **8-1**
 BPX **2-1**
 description **xxvi, 2-1**
 features **1-6**
 BPX Addressing Mode **8-10**
 BPX Service Node **11-1**
 BPX Service Node Extended Services Processor
 Installation and Operation **xxix**
 BPX Service Node Reference **xxviii**
 BPX switch **8-12**
 Broadband Integrated Services Digital Network **8-1**
 bundled circuits, frame relay
 bundled circuits **12-**
 Bursty Data Services **5-10**
 BXM, ATM Connections **10-**
 BXM-15 **10-1**
 BXM-622 **10-**
 BXM-T3/E3 **10-1**

C

CAM **8-9**
 CAS **1-8, 1-32**
 CBR (Constant Bit Rate) **10-2**
 CBR Connections **10-1**

CBR Single Leaky Bucket **10-2**
 CDP card **5-23**
 CDVT(0+1) **10-9**
 cell headers
 ATM **8-3**
 Cell Loss Priority
 definition **12-2**
 cell loss priority **12-2**
 cell relay technology **1-25**
 CES **5-10**
 Channel Associated Signalling **4-17**
 Channelized Data Pad card **5-23**
 CI **10-3**
 CIR
 setting **12-2**
 Circuit Emulation Service **10-1**
 Circuit Emulation Services **5-10**
 circuit priority **12-23**
 classes of traffic
 associated AAL Layers **8-8**
 class of service **1-23, 16-1**
 clocking, pass synch **5-26**
 clocking configuration
 looped clock mode **14-16**
 split clock mode **14-16**
 clocking configurations
 normal clock mode **14-1**
 clock sources **5-24**
 closed-loop **1-24 10-5**
 Cloud Addressing Mode **8-9**
 CLP **10-9**
 CLP Hi **10-1**
 CLP Lo **10-10**
 Cmax
 setting **12-2**
 Command Reference Manual **xxviii**
 See Related documentation
 configuration save/restore **4-22**
 congestion management **1-24, 10-**

congestion prevention **12-1**

Connection Parameters **10-**

connections

- bandwidth
- port concentrator **12-3**

CWM

- port concentrator **12-3**

frame relay

- port concentrator **12-3**

maximum number for port concentrator **12-3**

port concentrator

- delay **12-3**
- ForeSight **12-3**
- frame processing **12-3**
- frame relay format **12-3**
- frame relay signal flow **12-36**
- maximum throughput **12-33**
- signalling protocol **12-3**

connection testing **7-13**

conventions, documentation **xxix**

Courtesy Downing **4-18**

CRC check **12-3**

Credit Manager

- general description **16-2**

credit manager **4-14**

D

data flow with CDP **14-**

data flow with SDP **14-6**

data frame multiplexing

- description **4-15**

Data Link Connection Identifier **5-19**

dataport, DS0A **4-15**

dataport, synchronous **4-14**

DDS interface **14-4**

DFM **2-6, 14-1**

dialing formats

- dial pulsing **15-**

- inband **15-6**

Digital Signal Level 1 (DS1) rate

- definition **9-8**

Discard Eligibility **12-2**

discard eligibility **12-2**

displaying connection status **7-8**

displaying network alarms **7-3**

displaying node status **7-6**

displaying trunk loading **7-10**

displaying trunk status **7-5**

disruptive tests **7-1**

DLCI, see Data Link Connection Identifier

DLCI 0 **11-**

DNS Installation and Operation manual

- overview **xxv**

documentation

- conventions **xxix**

DPNSS **4-8**

DS1 rate

- definition **9-8**

E

E1 network

- supported interfaces **5-13**

E1 network, example **5-13**

E3 **8-16**

echo cancellers **15-10**

echo on voice circuits

- definition **15-9**

ECN, See Explicit Congestion Notification

EFCI **10-10**

EIA

- lead correspondence **14-1**

ELMI **1-8, 1-25**

Enhanced Instafax **4-18**

ER **10-**

ESP **11-1, 11-**

- Application interfaces **11-7**

redundant pairs **11-8**
 Ethernet port **11-7**
 ETSI **4-5**
 event log, maintenance log **7-9**
 Explicit Congestion Notification **12-1**
 explicit rate **10-3**
 Extended Services Processor **11-1**
 See ESP
 external device window **7-14**

F

fast EIA data control **9-5**
 FastPacket **1-25**
 cell length **1-27**
 FastPacket, definition **9-8**
 FastPacket to ATM **8-11**
 FastPAD
 definition **5-20**
 FastPAD network
 example **5-21**
 FBTC **10-**
 FCES **10-9**
 features **1-6**
 firmware download **4-20**
 Flow Control External Segment **10-9**
 ForeSight **1-24, 10-5 12-1**
 across multiple networks **12-2**
 definition **4-10**
 port concentrator **12-3**
 setting MIR, QIR, and PIR **12-26**
 Fractional T1 line
 characteristics **9-9**
 Frame **12-3**
 frame addressing **5-17**
 frame forwarding **4-13 12-32**
 definition **4-13**
 frame processing
 port concentrator **12-3**

frame relay
 call processing **4-6**
 connections
 maximum for port concentrator **12-32**
 frame addressing **12-**
 IPX-supported data port physical interfaces **4-11**
 port parameters **12-2**
 frame relay, congestion prevention **12-1**
 Frame Relay, description **12-8**
 Frame Relay, T1/E1 port **4-11**
 frame relay addressing **5-19**
 frame relay connections
 parameters **12-24**
 port concentrator **12-3**
 frame relay connection status **5-19**
 frame relay data packet types **9-6**
 Frame Relay feature
 advantages **4-10**
 frame relay feature **4-10**
 Frame Relay PAD **5-17**
 frame relay signal flow
 port concentrator **12-3**
 frame relay standard **4-10**
 framing mode, CCITT G.832 **8-16**
 framing mode, PLCP **8-15**
 FRF.4 **11-**
 FRP **5-17**
 FRTT **10-1**

H

hardware administration **4-24**
 high priority packets **9-2**

I

IBS **10-1**
 ICR **10-1**

IGX Reference Manual **xxvii**

See Related documentation

IISP **11-3, 11-5**

ILMI **11-**

ingress queues

ATM **4-3**

Instafax **4-18**

integral echo cancellers **15-4**

Interfaces

ESP **11-7**

Interim Inter-switch Protocol

See IISP

international gateway **5-22**

international networks

example **5-22**

interpretive mode **14-2**

interworking

port concentrator **12-3**

IPX **2-5**

connecting to BPX **5-4**

voice circuit **15-1**

IPX8 applications **5-14**

IPX features **2-5**

IPX Reference Manual **xxvii**

See Related documentation

isochronous clocking **14-1**

L

LMI Interface **12-1**

LMI protocol **5-18 12-1**

LMI protocol parameters **12-3**

logical ports **4-11**

loopback tests available **7-11**

looped clock mode **14-16**

Loss of Packet alarm **9-13**

M

maximum throughput

port concentrator **12-3**

MBS **10-**

MCR **10-9**

N

Network congestion **11-3**

network delay **16-**

network functions

port concentrator

interworking **12-33**

network interfaces

NNI **12-1**

UNI **12-1**

network management **4-23**

network statistics **7-14**

network synchronization **5-24**

Network-to-Network Interface **5-19**

network troubleshooting **7-2**

network-trunk administration **4-24**

NI **10-3**

NMS Interfaces **4-25**

NN **11-3**

NNI **11-**

Signaling channel **11-5**

NNI (see Network-Network Interface)

NNI header **8-3**

NNI status **12-13**

Non-interleaved data connections **14-**

non-timestamped packets **9-4**

normal clock mode **14-1**

Nrm **10-1**

nrt-VBR (Non-Real Time Variable Bit Rate) **10-2**

NTC packet framing **9-13**

NT DMS100 **4-5**

NT DMS250 **4-5**

NTT 4-5

O

OptiClass 1-23, 1-35

P

packet lines, description 9-7

Packet trunk format

E1 9-10

subrate 9-12

Packet trunk format, fractional T1 9-9

packet trunk format, T1 9-8

PAD

Frame Relay 5-17

parameters

frame relay connection 12-2

Partially interleaved connections 14-8

Partially interleaved data connections 14-8

payload length

ATM 8-3

PCR 10-9

Permanent Virtual Circuits

definition 1-26

permanent virtual circuits 1-11, 8-9

pleisiochronous network 5-26

PNNI 11-

PNNI signaling 11-

Policing 10-1

port concentrator

connection

frame processing 12-3

connections

bandwidth 12-3

delay 12-3

ForeSight 12-3

frame relay format 12-3

frame relay signal flow 12-36

maximum throughput 12-33

signalling protocol 12-3

through CWM 12-3

frame relay connections 12-3

maximum number of frame relay connections 12-3

network functions

interworking 12-33

port queue parameters

Frame Relay 12-31

port speed

setting Frame Relay 12-3

Private Network to Network Interface 11-

PVCs 8-9

Q

QIR, See Quiescent Information Rate

QSIG 4-9

Quantization Distortion 15-12

queues

ingress 4-3

queuing 16-

queuing, later releases 16-

Quiescent Information Rate 12-2

R

RDF 10-1

redundancy, power supplies 4-21

Related documentation xxv

IGX Installation Manual xxv

IPX Installation Manual xxix

Resource Management 10-3

RIF 10-1

RM 10-

RM cells 10-3

robbed bit signalling 4-17

routing
 system message **16-1**

RPS **2-6**

rt-VBR (Real-Time Variable Bit Rate) **10-2**

S

SAM **8-9**

SC connector **11-**

SCR **10-9**

Service Node

Network Management **11-**

setting CIR **12-2**

setting MIR, QIR, and PIR

ForeSight connections **12-26**

Signaling **11-3**

signaling **11-4**

signalling channel **4-8**

signalling conditioning **15-10**

signalling protocol

port concentrator **12-3**

signalling VCCs **11-**

Simple Addressing Mode **8-9**

sixth control lead **14-11**

SNA Credit Manager **12-1**

SNMP **11-7**

software downloading **4-19**

software restore **4-20**

software updates **3-3**

SONET **8-2**

Sonet/SDH Interfaces **1-20**

specifications

ATM E3 trunk interface **A-7, A-8**

ATM Service Interface, T3/E3 **A-3, A-10, A-11**

ATM T3 trunk interface **A-5**

speech detection **15-8**

split clock mode **14-16**

Standard ABR **10-3**

status enquiry **12-14**

status message **12-14**

STI header **8-3**

StrataView Plus Operations Guide **xxvii**

subrate data connections **14-**

superrate data connections **14-**

SVC **11-**

SVCs **8-9**

switch arbiter **2-4**

switched virtual circuits **8-9**

Synchronous Optical Network **8-2**

system message

routing **16-11**

system software **1-35**

T

T1/E1 conversion **5-23**

T1 network

example **5-11**

T3 trunks

and ATM **8-15**

TBE **10-1**

Telnet **11-**

Terminal port **11-**

TFTP **11-7**

timestamped data, packet format **9-5**

Topolog **11-**

Traffic **10-2**

traffic classes

and adaptation layers **8-8**

Traffic Policing **10-2**

Traffic Policing Definitions **10-**

transparent mode **14-2**

Trm **10-1**

troubleshooting commands **7-10**

trunks

bandwidth allocation **16-1**

E1 framed bandwidth **16-1**

E1 unframed bandwidth **16-1**

subrate bandwidth **16-13**
 T1 framed bandwidth **16-12**
 trunks, broadband **4-1, 4-3**
 trunks, narrowband **4-2**

U

UBR (Unspecified Bit Rate) **10-2**
 UBR Connection Policing **10-3**
 UBR Connections **10-2**
 UNI **11-**
 UNI (see User-Network Interface)
 UNI header **8-3**
 Usage Parameter Control **4-3**
 User-to-Network Interface **4-12**
 frame relay **5-18**

V

VAD **2-6, 4-16 15-8**
 See Voice Activity Detection
 VAD (see Voice Activity Detection)
 VBR and ATFR Connections **10-13**
 VBR Dual-Leaky Bucket Policing **10-2**
 VCI/VPI, description **8-6**
 VC Q depth
 setting **12-2**
 VC Qdepth **10-1**
 virtual circuit identifier **8-4**
 virtual path identifier **8-4**
 virtual trunk **1-20**
 Voice Activity Detection **4-8**
 Voice Activity Detector **15-8**
 voice channel signalling **15-5**
 voice circuit types **15-1**
 voice flow with CDP **15-3**
 voice packet type **9-3**
 VPI 0 and VCI 1 **11-4**

VPI 0 and VCI 18 **11-**
 VPI 0 and VCI 5 **11-4**
 VSVD **10-9**

W

WANs
 advantages over dial-up and leased lines **1-10**
 implementing **1-9**