



>THIS IS **THE WAY**

>THIS IS **NORTEL**

Nortel Multiservice Switch 7400/15000/20000

Frame Relay Technology Fundamentals

Document status: Standard
Document issue: 7.1S1
Document date: October 2005
Product release: PCR7.1 and up
Job function: Product Fundamentals
Type: NTP
Language type: U.S. English

Copyright © 2005 Nortel.
All Rights Reserved.

NORTEL, the globemark design, and the NORTEL corporate logo are trademarks of Nortel.



Contents

What's new	12
Frame relay overview	13
Basic elements of frame relay	14
Control plane (C-plane)	14
User data transfer plane (U-plane)	14
Frame relay frame	15
Data Link Connections (DLCs)	16
Frame relay interfaces	17
User-to-network interface (UNI)	17
Network-to-network interface (NNI)	18
Frame relay to ATM interface	18
Frame relay ISDN switched access interface on Multiservice Switch 7400 nodes	19
Multilink frame relay on Multiservice Switch 15000 and Multiservice Switch 20000 nodes	19
Committed information rate	20
CIR implementation	20
CIR algorithm	21
Excess information rate	22
EIR implementation	22
Relationship of CIR and EIR	25
Frame relay in a mixed Multiservice Switch 7400 and DPN network	26
Engineering considerations for frame relay and DPN interworking	27
Spared frame relay services on Multiservice Switch 15000 and Multiservice Switch 20000 nodes	28
Frame relay interactions with other Multiservice Switch spared services	28
Mixed frame relay and AAL1 CES on a single Multiservice Switch 7400 node MSA port	28
Interactions with other Multiservice Switch 15000 or Multiservice Switch 20000 node spared services	29
Hardware requirements	29
Frame relay feature list names	30



Frame relay connections	32
Permanent virtual circuits (PVCs)	32
Redundancy for PVCs	33
Switched permanent virtual circuits (SPVCs)	34
Redundancy for SPVCs	35
Frame relay switched virtual circuits (SVCs)	35
FR UNI switched connections	36
FR NNI switched connections	36
IP-optimized DLCIs	37
DLCI assignment	37
DLCI ranges for permanent and switched connections	38
FR NNI DLCI allocation policies	39
PVC status management	40
LMI protocols	40
FR UNI LMI application	41
FR NNI LMI application	41
Link integrity verification	41
Permanent connection status information	41
Asynchronous status reports	42
A-bit signaling	42
D-bit signaling	43
N-bit signaling	43
Multi-network PVC status	44
FR UNI automatic traffic parameter configuration	44
FR NNI LMI signaling exchange cycle	44
Frame relay call management	47
FR UNI switched virtual circuit (SVC) signaling	48
Q.933 signaling	48
FR UNI Q.922 signaling	48
Call setup procedures for SVC calls	49
FR NNI SVC and SPVC signaling	50
FR NNI Q.922 signaling	50
FRF.10 signaling	50
Call setup procedures for SPVC calls	50
Call setup procedures for SVC calls	51
DLCI data loopback mode	51
Frame relay virtual circuit disconnection and recovery	53
Disconnection triggered by application	53
Disconnection triggered by detection of network failure	53
Frame relay virtual circuit reconnection	54
PVC reconnection triggered by application	54
Loopback and A-bit signaling	54
Fast 1:1 FP equipment protection for Multiservice Switch 7400 FR UNI interfaces	55



Frame relay UNI management **57**

- Call redirection in a Multiservice Switch network 57
 - Call redirection for PVCs in a Multiservice Switch network 58
 - Call redirection for SPVCs in a Multiservice Switch network 59
 - Call redirection for SVCs in Multiservice Switch network 60
 - Situations for call redirection 61
 - Accounting records for call redirection 62
- Hunt groups 62
 - Hunt group members 63
 - Multiservice Switch Hunt Group Server 63
 - Hunt group call setup 63
 - Hunt group member availability status 64
 - Backup hunt groups 65
- Closed user groups (CUGs) 65
 - Simple CUGs 66
 - Selectable CUGs 66
 - CUG subscription at the FR UNI 66
 - Interlock codes 70
 - CUG signaling 71

Frame relay NNI management **72**

- Call redirection in a Multiservice Switch network 72
 - Call redirection for PVCs in a Multiservice Switch network 73
 - Call redirection for SPVCs in a Multiservice Switch network 74
 - Manual call redirection 75
 - Situations for call redirection 75
 - Accounting records for call redirection 76
- Gateway call routing to external networks 76
 - Hunt groups 77
 - Multiservice Switch Hunt Group Server 77
 - Hunt group call setup 78
 - Hunt group member availability status 79
 - Backup hunt groups 80

Bandwidth management **81**

- Bandwidth pool management 81
 - Bandwidth pools 82
 - Over- and under-subscribing port capacity 83
 - Full sharing of port capacity 84
 - Transfer priority and bandwidth pool assignment 85
- Connection admission control 86
 - Assigned bandwidth pools 88
 - Available link capacity 89
 - Bandwidth requirements for new connections 89



Ingress CAC behavior (outgoing calls)	89
Egress CAC behavior (incoming calls)	90
SVC call negotiation	91
Negotiated traffic parameters	92
FR UNI CIR negotiation	93
FR UNI EIR negotiation	94
FR UNI maximum frame size negotiation	94
FR UNI default traffic parameters for SVCs	94
FR NNI ingress and egress call negotiation	95
FR UNI CAC migration	95
Assigned bandwidth pool mapping	96
CAC bandwidth pool mapping	96

BNX interworking overview	98

BNX-lwf overview	98
BNX-lwf DLCI overview	100
BNX-lwf features	101
Virtual router	103
Frame relay IP server	103
Routing protocol	104
Choosing IP addresses	104
Cross-net signaling and LMI	104
LMI signaling	105
Cross-net signaling	105
State change of a DLCI	106
PVC tear down	108
Cross-net message format	109
PVC redirect	111
Manual and switch-on-demand redirect mode	112
Cross-net signaling for PVC redirect	112
Switch-on-demand prerequisites and enhancement	113
Redirection DTE notification	113

Frame relay ISDN switched access overview	114

Frame relay ISDN switched access overview	114
Summary of capabilities	115
Benefits of frame relay ISDN switched access	116
Frame relay ISDN switched access protocols	117
Frame relay ISDN switched access application examples	118
Bandwidth on demand	118
Disaster recovery	118
Occasional access	119
Switched access to frame relay	119
ISDN connection	119
Virtual circuit (VC) connection	120



ISDN protocol stack	121
B-channel connection	122
B-channel selection	122
Frame relay interface	122
Detailed description	123
ISDN call handling	124
Software modules	124
The data ISDN module	125
The Q.931 protocol module	125
The Call control module	126
The access service module	131
Security	131
Origin of the ISDN call	131
Channel selection	131
Release of an idle B-channel	132
Frame relay virtual circuit call handling	133
Permanent virtual circuit (PVC)	133
Switched virtual circuits (SVC)	133
Data network address	133
Accounting	136
Interaction between ISDN and virtual circuit connections	136
ISDN connection versus VCs	136
Clearing of virtual connections	136
End-to-end call establishment for VCs	137
Compatibility issues	138

Traffic management

140

Congestion control	140
Discard priority	141
Congestion notification	142
Congestion avoidance policy	142
Rate enforcement	143
Rate adaptation	144
Rate adaptation options	145
Rate adaptation sensitivity	148
Using rate adaptation across networks	149
Overflow routing	149
Traffic subscription options	151
Transfer priority	151
FR UNI routing class of service	152
FR NNI routing class of service	153
Emission priority	153
Flexible QoS mapping	154
IP class of service (CoS) mapping	155



IP packet fragmentation	156
Large frame transmission	156
End-to-end FRF.12 frame fragmentation	157
FR UNI weighted fair queue scheme	159
Multimedia traffic	161
Multimedia traffic classes	161
Packet forwarding of multimedia traffic	161
Frame-cell trunk interrupting mode engineering guidelines	162
Class of service types	163
Unrestricted CIR	163
Class of service A: Discard eligibility only (best effort)	163
Class of service B: CIR only (guaranteed)	164
Class of service C: CIR plus DE (premium)	164
Traffic management for BNX interworking	164
Data flow	165
IP datagram in FrUni/FrNni interworking with BNX switches	166
BNX header field settings for incoming frame relay traffic	168
Processing of BNX header fields in outgoing traffic	170
Quality of service	171
Congestion management	173

Frame relay accounting **177**

Benefits of frame relay accounting	178
Frame relay accounting meter	179
Accounting meter synchronization	179
In-band synchronization	179
Frame relay accounting record	181
CircuitId attribute and the accounting record	185
Egress protocol	186
Call clear	187
Peak water mark	187
Peak water mark interval	190
Example of peak water mark	192
Call management for switched virtual circuits	193
Generating reports	194
Accounting record generation mode - singleEnd	196
Accounting record generation mode - bothEnds	198
Frame relay billing policies	200
Flat-rate billing	201
Usage-sensitive billing	202
Quality of service parameters	202
Accounting for BNX-lwf	204
Mapping of attributes in the PVC accounting record	204
Attribute mapping for BNX-lwf DLCI accounting records	209



Troubleshooting	221
OSI state information for FR UNI components	222
OSI state information for FR NNI components	223
Alarms	225
Pinging a virtual circuit	225
Tracing data frames	226
Troubleshooting connection failure	226
Troubleshooting traffic problems	230
Troubleshooting accounting problems	231
Troubleshooting hardware resources	232
Clear cause codes	233
Diagnostic codes	237
FR UNI frame link monitor for E1 frame relay	241

BNX-lwf performance monitoring and fault management	242
Displaying BNX protocol operational attributes	243
Displaying BNX-lwf operational attributes	244
Use of the REDIRECT operational verb	245
Multiservice Switch-BNX frame relay alarms	246
Frame relay statistics for Multiservice Switch-BNX	247
Statistics	252
Multiservice Switch 15000 and Multiservice Switch 20000 node frame relay statistics and BNX-lwf	253
Behavior on control processor switchover	254
Frame relay ISDN switched access monitoring and troubleshooting	254
Monitoring the service	254
Examples	255
Operational attributes	255
OSI state	255
SNMP management	256
Maintaining the service	256
Detecting frame relay ISDN switched access problems	257

FR UNI compliance with standards	265
Compliance with Vendor Forum specification	265
Detailed compliance statement	265
Compliance with joint specification	266
Compliance with Annex D of T1.617	266
Compliance with Annex A of ITU-T Q.933	267
Compliance with frame relay UNI Interface Implementation Agreement	268
Compliance with FRF.4	268
Compliance with ITU-T Q.922	269
Compliance with ITU-T X36	271
Compliance with FRF.12	272



Compliance with FRF.16 272

FR NNI compliance with standards 273

Detailed compliance statement 273
Standards guiding frame relay development 274
Compliance with Annex D of T1.617 274
Compliance with Annex A of ITU-T Q.933 275
Compliance with Frame Relay NNI Phase 1 Implementation Agreement 276
Compliance with FRF.10 (subset of ITU-T X.76) 277
Compliance with ITU-T X.76 (year 2000 revision) 278
Compliance with FRF.12 286
Compliance with FRF.16 287

Australian TS014 compliance statement 288

Terminology 288
General information and definitions 289
General requirements 289
Layer 1 requirements 290
Layer 2 requirements 290
Layer 3 requirements 296
Compatibility requirements 303
Appendices requirements 304

European ISDN compliance statement 305

Terminology 305
European ISDN standards for PRI access 306
 Layer 1 requirements 307
 Layer 2 requirements 307
 Layer 3 requirements 316
French Euro-numeris ISDN network standards 333
German ISDN network standards 336

Japanese ISDN compliance statement 338

Terminology 338
Layer 1 requirements 338
Layer 2 requirements 339
Layer 3 requirements 343

North American NI-2 compliance statement 353

Terminology 353
Layer 1 requirements 354
Layer 2 requirements 354
Layer 3 requirements 361

Standards 370

Frame relay technical references 370
ITU-T references 371



ANSI references 371



What's new


There were no new features added to this document.

Attention: To ensure that you are using the most current version of an NTP, check the current NTP list in NN10600-000 *Nortel Multiservice Switch 7400/15000/20000 What's New*.

Frame relay overview

Frame relay is a fast access service that provides high performance connectivity for applications such as IP network interconnection. The frame relay service transfers data frames between two networks or user devices (for example, routers) over a virtual connection. See the figure [Example of a frame relay network \(page 13\)](#).

Example of a frame relay network



Navigation

- [Basic elements of frame relay \(page 14\)](#)
- [Frame relay interfaces \(page 17\)](#)
- [Committed information rate \(page 20\)](#)
- [Excess information rate \(page 22\)](#)



- [Relationship of CIR and EIR \(page 25\)](#)
- [Frame relay in a mixed Multiservice Switch 7400 and DPN network \(page 26\)](#)
- [Spared frame relay services on Multiservice Switch 15000 and Multiservice Switch 20000 nodes \(page 28\)](#)
- [Frame relay interactions with other Multiservice Switch spared services \(page 28\)](#)
- [Hardware requirements \(page 29\)](#)
- [Frame relay feature list names \(page 30\)](#)

Basic elements of frame relay

The frame relay protocol uses the physical layer and a portion of the link layer of the OSI model. The result is reduced error checking and more processing capacity to deliver additional network throughput. In particular, frame relay

- provides bidirectional frame transfer
- preserves frame order during transfer
- detects transmission, format, and operational errors
- transports frames transparently. The network can modify only the data-link connection identifier (DLCI), congestion bits, DE bit, and frame check sequence.

For more information about frame relay service characteristics, see the following sections:

- [Control plane \(C-plane\) \(page 14\)](#)
- [User data transfer plane \(U-plane\) \(page 14\)](#)
- [Frame relay frame \(page 15\)](#)
- [Data Link Connections \(DLCs\) \(page 16\)](#)

Control plane (C-plane)

The control plane (C-plane) carries signaling information for both permanent and switched calls. C-plane procedures include implementation of the Local Management Interface (LMI) for monitoring PVCs, as well as implementation of Q.933 signaling for setting up switched calls for FR UNI and FRF.10 signaling for setting up switched calls for FR NNI.

User data transfer plane (U-plane)


The ITU-T and ANSI standards define the data transfer protocol for the user data transfer plane (U-plane). ANSI T1.602 (LAPD) and its subset T1.618 form the basis for this protocol. See the appendix [FR UNI compliance with standards \(page 265\)](#) and [FR NNI compliance with standards \(page 273\)](#).



Frame relay frame

A frame relay frame consists of an opening flag, an address field, a user data field, a frame check sequence (FCS) field, and a closing flag. See the figure [Frame relay frame \(page 15\)](#).

Frame relay frame



The minimum size of the user data field is one octet, and the maximum size is 8187 octets. The maximum number of octets between (but not including) a frame's opening and closing flags is $8187 + 2$ (address) + 2 (FCS) = 8191.

Each address field in a frame relay frame consists of two octets subdivided as follows (see the figure [Frame relay address field \(page 16\)](#)):

- DLCI bits

The DLCI includes ten bits within the two address octets. DLCIs can have values in the range 16-1007 inclusive. Network administrators can impose other restrictions on DLCI allocation.

- Command/Response indication (C/R) bit

The frame relay data link protocol does not use the C/R bit. The user device can set the C/R bit to any value; the network carries the C/R bit unchanged.

- Forward explicit congestion indication (FECN) bit

The FECN bit tells the recipient to start congestion avoidance procedures for traffic that transmits in the same direction as the FECN. Both the network and the user device can set the FECN bit; the network does not clear a FECN that the user sets.

- Backward explicit congestion indication (BECN) bit

The BECN bit tells the recipient to start congestion avoidance procedures for traffic traveling in the opposite direction from the BECN. Both the network and the user device can set the BECN bit; the network does not clear a BECN bit that the user sets.



- Discard Eligibility (DE) bit

DE bit indicates whether the frame is discarded during congestion. The user device can set the discard eligibility (DE) bit to any value; the network can set, but cannot clear, the DE bit.

Frame relay address field

8	7	6	5	4	3	2	1
(msb) DLCI						C/R	0
DLCI (lsb)				FECN	BECN	DE	1

msb Most Significant Bit
lsb Least Significant Bit
FECN Forward Explicit Congestion Notification
BECN Backward Explicit Congestion Notification
DE Discard Eligibility Indicator
C/R Command/Response indication

PPT 3383 037 AA

Data Link Connections (DLCs)

Frame relay uses data-link layer addressing to multiplex and demultiplex different user data streams within the same access channel. Each user data stream within the physical access channel is called a data-link connection (DLC).

Each DLC has a local address called the data link connection identifier (DLCI), which identifies different DLCs within the same channel. During data transfer, all frames that belong to a particular connection have the same DLCI. DLCIs are unique within a single access channel, so they have local significance only.

You can configure each DLCI with a unique traffic contract that supports the user's subscribed service level agreement (SLA). Each DLCI provides a common set of frame statistics used to monitor the performance of each connection. These statistics consist of discard counts for bytes and frames, as well as counts for frames successfully delivered in each direction.

The dynamic packet routing system (DPRS) transmits frame relay traffic across a Nortel Multiservice Switch network. Frame relay traffic can also use multiservice cut-through switching (MCS). MCS provides many-to-one multiplexing for data service connections, while it uses DPRS facilities for call



setup. For more information about MCS, see NN10600-440 *Nortel Multiservice Switch 7400 Operations: Frame Relay Managed Cut Through Switching*.

Frame relay interfaces

The Nortel Multiservice Switch frame relay service supports the following frame relay interfaces:


- [User-to-network interface \(UNI\) \(page 17\)](#)
- [Network-to-network interface \(NNI\) \(page 18\)](#)
- [Frame relay to ATM interface \(page 18\)](#)
- [Frame relay ISDN switched access interface on Multiservice Switch 7400 nodes \(page 19\)](#)
- [Multilink frame relay on Multiservice Switch 15000 and Multiservice Switch 20000 nodes \(page 19\)](#)

User-to-network interface (UNI)

The frame relay user-to-network interface (UNI) service provides a standard interface between the user device and the network.

The figure [Example of a frame relay UNI connection \(page 17\)](#) shows the connection between two frame relay UNI users, Device A and Device B. In this figure, devices can communicate using different data addressing plans (that is, User A uses E.164 addressing, and User B uses X.121 addressing).

Example of a frame relay UNI connection






Network-to-network interface (NNI)

The Nortel Multiservice Switch frame relay network-to-network interface (NNI) service provides a communication interface between two frame relay networks. This service also transmits the network status signaling information throughout the network. The frame relay NNI service provides the end user with an accurate picture of the network connection, which can span several different frame relay networks. The frame relay NNI service also ensures multi-vendor network compatibility with networks that conform with the FRF.2 Implementation Agreement.

The figure [Example of a frame relay NNI connection \(page 18\)](#) shows how two frame relay users connect Device A and Device B over a hybrid frame relay network. Device A uses Multiservice Switch frame relay UNI service to access a Multiservice Switch frame relay network. Multiservice Switch frame relay networks use Multiservice Switch frame relay NNI service to connect to an external frame relay network. Device B uses a non-Multiservice Switch frame relay UNI service to connect to the external network.

Example of a frame relay NNI connection



Frame relay to ATM interface

Nortel Multiservice Switch frame relay to ATM (FR-ATM) interworking service maps frame relay PVCs to and from asynchronous transfer mode (ATM) PVCs. the FR-ATM interface supports both service interworking (in accordance with the FRF.8 standard) and network interworking (in accordance with the FRF.5 standard).

See NN10600-920 *Nortel Multiservice Switch 7400/15000/20000 Operations: Frame Relay to ATM Interworking* for a description of the FR-ATM service and implementation procedures.



Frame relay ISDN switched access interface on Multiservice Switch 7400 nodes

Frame relay ISDN switched access provides an interface to the frame relay network through a Nortel Multiservice Switch 7400 node over a switched access path (B-channel). The ISDN signaling channel (D-channel) establishes this switched access path. A PRI link with the D-channel on one timeslot and the B-channels on the other timeslots support the service.

See [Frame relay ISDN switched access overview \(page 114\)](#) for more information about the frame relay ISDN switched access service.

Multilink frame relay on Multiservice Switch 15000 and Multiservice Switch 20000 nodes

Nortel Multiservice Switch 15000 and Multiservice Switch 20000 nodes support multilink frame relay (MLFR) in accordance with the FRF.16 Implementation Agreement. The MLFR service provides increased bandwidth in FrUni and FrNni connections through physical interface emulation for frame relay devices. The emulated physical interface consists of one or more DS1 links, called bundle links, configured as a single bundle. By combining multiple DS1 logical interfaces into a single bundle, you can design a frame relay interface that supports bandwidth capacity from the equivalent of two DS1s to the equivalent of a full DS3.

The traffic transmitted on each bundle link is timed to ensure that the maximum measured differential delay between links is within a provisionable range. If the maximum differential delay is provisioned to be too small, frames may be lost. If this value is too high, more memory is required for buffering and the overall performance of the multilink feature decreases. The Timestamp information field is used to measure differential delay.

By default, frames are sent through bundles in one piece; in other words, fragmentation is disabled. To minimize queuing latency, you can enable fragmentation so that frames above the provisioned fragmentation threshold are broken down into smaller pieces (fragmented). This allows frame fragments to be transmitted on different bundle links, thus balancing the load between bundle links and maximizing the bandwidth allocated for that bundle. You can also provision the fragment size to be used by a bundle. Fragmented frames are transmitted in the same sequence in which they were received and then reassembled prior to being sent to the data link layer. If an error occurs and a fragment is lost, that frame is discarded.

MLFR service supports the dynamic addition and removal of bundle links, which results in a fluctuation in the amount of total bandwidth available on the interface. Adding one or more bundle links allows for easy service expansion when traffic levels increase as well as efficient utilization of the DS3 channelized facility.



The MLFR service also provides redundancy when multiple DS1 links are configured on a single bundle; if a bundle link fails, that link is disabled and the remaining links in the bundle continue to support network traffic with reduced throughput.

Committed information rate

The committed information rate (CIR) is the rate (in kbit/s) at which the network agrees to transfer information over a connection under typical conditions.

For more information, see the following sections:

- [CIR implementation \(page 20\)](#)
- [CIR algorithm \(page 21\)](#)

CIR implementation

The CIR applies to the rate of data entry into the network. CIR implementation is based on a fast-forward technique. Frames enter the network and transmit to the destination port at the maximum possible rate. This CIR implementation means that instantaneous rates are in excess of the CIR rate, in contrast with other implementations that delay or leak frames into the network at the CIR rate.

A frame relay port can have multiple connections, with two assigned, independent CIR values (one in each direction of a full-duplex PVC). Independent bi-directional CIRs allow better network engineering to match the asymmetrical data flows common in data applications.

The CIR allows the network to match itself to the application's burstiness. The following two service parameters determine the allowed amount of burst:

- committed burst size (Bc)

The Bc is the maximum amount of data (in bits) that a network agrees to transfer under normal conditions over a committed rate measurement interval.

- committed rate measurement interval (Tc)

The Tc is the time interval over which the network measures rates and burst sizes. Typically, the measurement interval is proportional to the traffic's burstiness.

CIR is a function of the committed burst and measurement interval, according to the following relationship:


$$\text{CIR} = \text{Bc}/\text{Tc}$$



You can configure the CIR, Bc, and Tc to optimize network resources by matching them to application characteristics. The maximum burst can exceed the Bc limit for a short time. Nortel Networks recommends that you configure the Bc to a minimum of the application's window size to reduce the chance of frame discard.

See the figure [Relationship of CIR, Bc, and Tc parameters \(page 21\)](#).

Relationship of CIR, Bc, and Tc parameters




CIR algorithm

The CIR algorithm is based on a credit system and uses a token pool. The token pool is a counter that indicates the available credit, based on how much of the allocated bandwidth is already used. Initially this pool contains a number of tokens equivalent to the number of bytes in Bc bits of data (the subscribed Bc level).

As data transmits, the pool decreases by an amount equal to the number of data bytes that enter the network, and increases in proportion to the CIR value (that is, Bc tokens per Tc seconds). The token pool needs only one available token to transmit frames. When the token pool is less than or equal to zero, the network either tags the frame as discard eligible before routing it through the network, or discards the frame immediately. See the figure [CIR algorithm \(page 22\)](#).



CIR algorithm



Excess information rate

The excess information rate (EIR) is the sustainable information rate in excess of CIR that the network delivers if there is available bandwidth. The total information rate is CIR + EIR. For more information, see [EIR implementation \(page 22\)](#).

EIR implementation


EIR allows you or the network to tag frames on an incoming data stream as discard eligible (DE). When either you or the network sets the DE bit in the frame relay frame, it indicates that the data bits in this frame belong to the excess burst (Be). See the figure [Discard eligibility \(page 23\)](#).

The network treats DE frames as excess and discards them first at the onset of level 2 or moderate congestion. The network discards frames not tagged discard eligible only when severe congestion occurs. The network does not clear a DE bit that you set. For more information, see [Congestion control \(page 140\)](#).



You can select a different class of traffic (DE and non-DE classes) for different applications on the same connection. This selection is network-independent (that is, you set the DE bits on frames of your choice). Alternatively, you can dedicate one connection to one class of traffic and another connection on the same port to another class of traffic.


Discard eligibility



If the DE bit is set, the Be pool decreases by an amount equal to the number of data bits that enter the network in that frame. See the figure [The network sets the DE bit \(page 24\)](#).




The network sets the DE bit



When a user sets the DE bit, the frames are counted immediately against the Bc token pool, and enter the network as discard eligible. See the figure [The user sets the DE bit \(page 25\)](#).



The user sets the DE bit



Traffic marked as discard eligible (that is, DE=1) by the user is discarded at the Nortel Multiservice Switch node frame relay port if the rate of DE traffic exceeds the provisioned EIR, regardless of the state of the CIR. If you want to prevent such discards, increase the provisioned EIR to accommodate user DE=1 traffic. If the amount of user DE=1 traffic is unknown, configure the EIR to be equal to the access line speed.


Relationship of CIR and EIR

The network attempts to deliver information at a rate equal to CIR + EIR. When a frame arrives at the network, the network checks the Bc token pool for token availability. If the Bc token pool is exhausted, but at least one Be token is available, the network tags the frame as discard eligible (that is, DE=1) and the frame fast forwards. The Be token pool decreases accordingly. The Be token pool, on the other hand, replenishes at a rate proportional to EIR (that is, Be tokens during Tc seconds).

The figure [Relationship of CIR and EIR parameters \(page 26\)](#) shows normal traffic flow, for CIR parameters, where the traffic bursts are within the Bc. For EIR parameters, the figure shows traffic that exceeds the Bc. This traffic is Be traffic, which the network tags as DE (discard eligible).



Relationship of CIR and EIR parameters



Frame relay in a mixed Multiservice Switch 7400 and DPN network

Nortel Multiservice Switch nodes' frame relay interworks seamlessly with DPN frame relay. The compatible subnet allows packet routing, PVC creation, congestion signaling, and LMI A-bit transmission. Subscribers to DPN frame relay service can connect to users on other external networks through the Multiservice Switch nodes' frame relay NNI service.

The frame relay service in a Multiservice Switch network interworks with the DPN frame relay service as long as there is a Multiservice Switch 7400 series node in the network.



When a Multiservice Switch 7400 node operates in an existing DPN network, the node can set the FCI bit in packets that travel to a DPN-based frame relay port. If the Multiservice Switch trunks are congested, the node sets the FCI, and the DPN statistics fields count these FCIs as if they were FECNs. The Multiservice Switch 7400 node can also set the BCI bits. The DPN statistics fields count these BCIs as BECNs.

Engineering considerations for frame relay and DPN interworking

Nortel Multiservice Switch 7400 nodes' performance is much higher than DPN performance, and improper engineering can degrade service. The configured CIR values and the node's frame relay port must not exceed the maximum recommended throughput for the destination DPN port.

See 241-1001-153 *DPN-100 System Engineering Guidelines* for these maximum recommended throughput values. For additional information about Multiservice Switch 7400 node and DPN frame relay interworking, see 241-1001-120 *DPN-100 Frame Relay Service User Guide and Specification* and the generic release reports for that DPN.

To interwork Multiservice Switch 7400 frame relay with the DPN frame relay service, follow the connection guidelines above and ensure that the following fields in the DPN frame relay PVC envelopes contain these default values:

- packet size sent = 4096
- packet size received = 4096
- throughput class sent = 13

Use the sensitivity setting and the call priority setting to select emission priority on DPN. Select one of the following four emission priorities and queues:

- high priority, delay sensitive
- high priority, throughput sensitive
- normal priority, delay sensitive
- normal priority, throughput sensitive

If the sensitivity is delay, traffic with that label transmits across the network on the path with the least delay. The transmission order depends on which of the four combinations you select. High priority, delay sensitive traffic transmits with the highest priority, while normal priority, throughput traffic transmits with the lowest priority. A quota scheme ensures that all combinations receive service.



Spared frame relay services on Multiservice Switch 15000 and Multiservice Switch 20000 nodes

On Nortel Multiservice Switch 15000 and Multiservice Switch 20000 nodes, applications and services are categorized as hot, warm or cold standby based on their sparing behavior. The following interfaces are warm standby features when provisioned on a spare LP:

- frame relay user-to-network interface (UNI)
- network-to-network interface (NNI)
- frame relay to ATM interface (FR-ATM)
- IP over frame relay (IP/FR)

The IP/FR interface is a warm standby feature during a hitless software migration only.

Warm standby features reduce service outages during an equipment switchover. During an equipment switchover, warm standby applications and features incur a longer outage of service than hot standby applications and features, but not as long as cold standby applications and features. As well, all connections must be re-established.

See NN10600-550 *Nortel Multiservice Switch 7400/15000/20000 Common Configuration Procedures* for a description of hitless services and hot, warm and cold standby applications and features.

Frame relay interactions with other Multiservice Switch spared services

The frame relay service can interact with other Nortel Multiservice Switch node services when it is running in combination with:

- other services on the same Multiservice Switch 7400 function processor. See [Mixed frame relay and AAL1 CES on a single Multiservice Switch 7400 node MSA port \(page 28\)](#).
- other spared services on Multiservice Switch 15000 or Multiservice Switch 20000 nodes. See [Interactions with other Multiservice Switch 15000 or Multiservice Switch 20000 node spared services \(page 29\)](#).

Mixed frame relay and AAL1 CES on a single Multiservice Switch 7400 node MSA port

Nortel Multiservice Switch 7400 nodes' 32-port DS1/E1 Multi-Service Access (MSA) function processors support multiple types of services, including frame relay, which otherwise would be provided by separate existing function processors.



On the 32-port DS1/E1 MSA function processor, you can configure a single port to support mixed structured AAL1 CES and frame relay UNI services. See NN10600-720 *Nortel Multiservice Switch 7400/15000/20000 Operations: AAL1 Circuit Emulation*.

Interactions with other Multiservice Switch 15000 or Multiservice Switch 20000 node spared services

As a warm standby application or feature, the frame relay service on Nortel Multiservice Switch 15000 or Multiservice Switch 20000 nodes can operate with a hot standby application or feature on the same FP without affecting the ability of the hot standby application or feature to provide hitless services. See NN10600-550 *Nortel Multiservice Switch 7400/15000/20000 Common Configuration Procedures* for a description of hitless services and hot, warm and cold standby applications and features.

Hardware requirements

The table [Frame relay capabilities on various function processors \(page 30\)](#) lists the interfaces you can configure to support frame relay services. For information about performance specifications for the frame relay services and FPs, see Nortel Multiservice Switch Release Notes.

You can bind the FR UNI service to either a physical or logical interface. Logical links do not use physical ports: they provide a cost effective way to reduce the number of physical ports in use and to free those ports for other services. You can create a logical link between the following

- two frame relay interfaces
- a FR UNI interface and a FR NNI interface
- a frame relay interface and a FR ATM interface

Specific feature combinations are subject to existing exclusivity rules and are limited to product roll-out considerations. For more information, see Nortel Multiservice Switch Release Notes.

Not every function processor listed in the table [Frame relay capabilities on various function processors \(page 30\)](#) is compatible with the Nortel Multiservice Switch 7400, Multiservice Switch 15000, and Multiservice Switch 20000 platforms. For the details on every frame relay FP in the table below, see the corresponding section of NN10600-551 *Nortel Multiservice Switch 7400/15000/20000 FP Configuration Reference*.



Frame relay capabilities on various function processors

Function processor	Link speeds (kbit/s)	maximum number of frame relay interfaces (single link mode)	maximum number of frame relay interfaces (fractional link mode, n*64 or n*56)
4-port DS1	56 - 1544	4	8
4-port DS1Ch	56 - 1544	4	96
8-port DS1	56 - 1544	8	8
32-port DS1 MSA	56 - 1544	32	500
1-port DS3	44736	1	Not applicable
1-port DS3Ch	44736	28	28
4-port DS3Ch ATM Fr	44736	112 (DS1links)	1024
4-port DS3Ch ATM Fr (when unchannelized)	44736	4 (DS3 links)	Not applicable
4-port E1	56 - 2048	4	8
4-port E1Ch	56 - 2048	4	124
1-port E3	34368	1	Not applicable
32-port E1 MSA	56 - 2048	32	500
1-port HSSI	1000 - 50000	1	Not applicable
1-port STM -1Ch FR	63 - 2040	63 (E1 links)	768
8-port V.11	9.6 - 7680	8	Not applicable
8-port V.35	9.6 - 3840	8	Not applicable

Frame relay feature list names

The tables [FR UNI feature list names \(page 30\)](#) and [FR NNI feature list names \(page 31\)](#) list the corresponding feature names used for provisioning frame relay services.

FR UNI feature list names

FR UNI service	Feature name
FR UNI supporting PVC, SVC, and P-SPVC connections	frUniAllDprs
FR UNI supporting PVC connections	frameRelayUni
FR UNI supporting PVC and SVC connections	frameRelayUniPvcSvc
(1 of 2)	



FR UNI feature list names (continued)

FR UNI service	Feature name
FR UNI supporting PVC and P-SPVC connections	frameRelayUniPvcSpvc
FR UNI supporting Fast 1:1 equipment protected PVC connections	frUniPvcDprs frUniPvcSparingDprs
(2 of 2)	

FR NNI feature list names

FR NNI service	Feature name
FR NNI supporting PVC, SVC, P-SPVC and S-SPVC connections	frNniAllDprs
FR NNI supporting PVC connections	frameRelayNni
FR NNI supporting SPVC connections	frNniSpvcDprs
FR NNI supporting SVC connections	frNniSvcDprs
FR NNI supporting PVC connections	frNniPvcDprs



Frame relay connections

Use the following sections to learn more about frame relay connections on Nortel Multiservice Switch nodes.

Navigation

- [Permanent virtual circuits \(PVCs\) \(page 32\)](#)
- [Switched permanent virtual circuits \(SPVCs\) \(page 34\)](#)
- [Frame relay switched virtual circuits \(SVCs\) \(page 35\)](#)
- [DLCI assignment \(page 37\)](#)
- [PVC status management \(page 40\)](#)
- [Frame relay call management \(page 47\)](#)
- [FR UNI switched virtual circuit \(SVC\) signaling \(page 48\)](#)
- [FR NNI SVC and SPVC signaling \(page 50\)](#)
- [DLCI data loopback mode \(page 51\)](#)
- [Frame relay virtual circuit disconnection and recovery \(page 53\)](#)
- [Fast 1:1 FP equipment protection for Multiservice Switch 7400 FR UNI interfaces \(page 55\)](#)


Permanent virtual circuits (PVCs)

A frame relay permanent virtual circuit (PVC) is a logical connection that you configure at both ends. Once a PVC is configured, it remains in place even when it is not being used. The PVC connection originates at a master endpoint and terminates at a slave endpoint. The PVC connection may terminate on either a FR UNI or a FR NNI. See the figure [Frame relay PVC \(page 33\)](#).


For multi-segment PVCs (that is, PVCs that span more than one network), you must configure each segment of the connection manually. See the figure [Frame relay PVC \(page 33\)](#).



Frame relay PVC



FR NNI PVC spanning multiple networks



Redundancy for PVCs

You can configure PVC redundancy in Nortel Multiservice Switch networks through a call redirection server. The call redirection server directs failed call attempts to alternate destinations.

For more information on FRUNI PVC redundancy, see [Call redirection for PVCs in a Multiservice Switch network \(page 58\)](#). For more information on FRNNI PVC redundancy, see [Call redirection in a Multiservice Switch network \(page 57\)](#).



Switched permanent virtual circuits (SPVCs)

A frame relay switched permanent virtual circuit (SPVC) is a logical connection that allows configured SPVC endpoints in a single network or two different networks to establish a switched connection across the NNIs that connect the networks.

A configured SPVC that uses the switched connection between networks is a permanent SPVC (P-SPVC). The switched connection at the NNI that enables connectivity between networks is a switched SPVC (S-SPVC).

SPVCs eliminate the need to configure each network segment of a multi-segment PVC connection.

In the two different networks joined by a switched connection across the NNIs, the originating network includes the master endpoint of the SPVC connection while the terminating network includes the slave endpoint. Single-ended configuration for an SPVC is not supported.

See the figure [Frame relay SPVC spanning multiple networks \(page 35\)](#) for an illustration showing the location of both endpoints. The SPVC connection may terminate on a FR UNI or a FR NNI.


With FR NNI, you only need to configure the end-points that terminate the end-to-end SPVC call, and to ensure that intermediate network nodes support FRF.10 call signaling. Route selection is automatic, and the dynamic characteristic of an SPVC connection enables the re-establishment of a call around a failed NNI.

The SPVC connection may originate and terminate on either a FR UNI or a FR NNI. In addition, P-SPVC status is signaled in the LMI full status report and asynchronous status report.

The FR UNI service supports P-SPVCs only; for information about switched connections between frame relay networks, see [FR UNI switched connections \(page 36\)](#).



Frame relay SPVC spanning multiple networks



Redundancy for SPVCs

You can configure SPVC redundancy in a Nortel Multiservice Switch network through a Multiservice Switch node's call redirection server. The call redirection server directs failed call attempts to alternate destinations.

For more information on FRUNI SPVC redundancy, see [Call redirection for SPVCs in a Multiservice Switch network \(page 59\)](#). For more information on FRNNI SPVC redundancy, see [Call redirection for SPVCs in a Multiservice Switch network \(page 59\)](#).

You can also configure SPVC redundancy in a Multiservice Switch network for both the originating and destination networks.

In the originating network, you can configure SPVC redundancy through the Hunt Group Server (HGS). The HGS receives call requests from the network and forwards them to an appropriate hunt group member. Hunt groups provide resiliency and allow the network to route around a failed or congested FR NNI. For more information, see [Gateway call routing to external networks \(page 76\)](#).

In the destination network, you can configure SPVC redundancy through a Call Redirection Server (CRS). If the primary SPVC endpoint is unable to accept an SPVC call request, the setup request goes to the CRS. The CRS directs the failed call attempt to a backup SPVC endpoint.

Frame relay switched virtual circuits (SVCs)

Frame relay services support [FR UNI switched connections \(page 36\)](#) and [FR NNI switched connections \(page 36\)](#).




FR UNI switched connections

The FR UNI service supports switched virtual circuit (SVC) connections. An SVC is a logical connection that is established on an as-required basis. SVCs are set up dynamically and torn down as required by subscriber applications. User equipment initiates SVC call establishment and terminates the connection when it is no longer needed. See the figure [Frame relay SVC \(page 36\)](#).

SVCs do not require user configuration, but network nodes must be configured to support SVC signaling. In addition, The FR UNI SVC service supports traffic parameter negotiation. For more information, see [SVC call negotiation \(page 91\)](#).

You can configure SVC resiliency in a Nortel Multiservice Switch network through a call redirection server. The call redirection server directs failed call attempts to alternate destinations. For more information about FR UNI SVC resiliency see [Call redirection for SVCs in Multiservice Switch network \(page 60\)](#).

Frame relay SVC



FR NNI switched connections

The FR NNI service offers resiliency for SVC and S-SPVC connections through the Hunt Group Server (HGS). If a FR NNI fails, all new calls destined for the failed NNI are routed to an available NNI. In addition, if an NNI becomes heavily utilized, new calls may also be redirected to an alternate NNI, to provide loadsharing.

The FR NNI service supports both [Switched permanent virtual circuits \(SPVCs\) \(page 34\)](#) and S-SPVCs (see) [FR NNI switched connections \(page 36\)](#).




FR NNI switched virtual circuit (SVC) connections enable switched connectivity between frame relay networks. They allow NNIs in adjacent networks to establish a switched connection on demand.

SVCs are set up dynamically and torn down as required by subscriber applications. They do not require user configuration, but network nodes must be configured to support SVC signaling. User equipment initiates and tears down the SVC calls on an as-needed basis. See the figure [FR NNI SVC connection across networks \(page 37\)](#).

The FR NNI SVC service supports traffic parameter negotiation. For more information, see [SVC call negotiation \(page 91\)](#).

FR NNI SVC connection across networks



IP-optimized DLCIs

An IP-optimized data link connection identifier (DLCI) can directly bind to a virtual router protocol port. This type of DLCI is linked to the Nortel Multiservice Switch frame relay user-to-network interface (FRUNI), which eliminates the need for a frame relay DTE and simplifies provisioning of IP over frame relay.

For more information on IP-optimized DLCIs, see NN10600-800 *Nortel Multiservice Switch 7400/15000/20000 IP Technology Fundamentals*.

DLCI assignment

All frame relay connections have an assigned DLCI value. The user assigns a DLCI value to PVCs and P-SPVCs when configuring the connections. The network assigns a DLCI value to SVCs when processing a call setup request.

To avoid possible conflicts for DLCI assignment on the same interface, you can define different ranges for permanent and switched calls. If you do not configure the range properly, collisions occur and calls are rejected.



For more information, see the following sections:

- [DLCI ranges for permanent and switched connections \(page 38\)](#)
- [FR NNI DLCI allocation policies \(page 39\)](#)

DLCI ranges for permanent and switched connections

Defining DLCI ranges for permanent and switched calls ensures that DLCI collisions do not occur on the same interface.

You can specify the highest DLCI value permitted for a permanent connection by setting the *highestPermanentDlci* attribute under the *Signaling* component. All other values are implicitly reserved for switched connections. Therefore, the portion of available DLCIs for switched connections is in the range of $<\text{highest permanent DLCI} + 1>$ to 1007 inclusive.

See the figure [DLCI assignment for permanent and switched calls \(page 39\)](#).


With each call setup request for an SVC, a DLCI is assigned to the connection from the range permitted on the interface, both calling and called.

The network assigns DLCIs for S-SPVCs from the SVC DLCI range. You must configure both P-SPVC and PVC DLCIs with DLCI values in the permanent DLCI range.

If you define the highest permanent DLCI value as 0, no permanent calls are permitted on the interface. If you define it as 1007, no switched calls are permitted on the interface.



DLCI assignment for permanent and switched calls



FR NNI DLCI allocation policies

The calling FR NNI signals the value of the selected DLCI in its call setup message to the adjacent interface. The called FR NNI must select the same DLCI for the connection.

A discrepancy in DLCI allocation at adjacent interfaces can result in a portion of the DLCI range being unusable. If the DLCI range at the called FR NNI is different from the range configured at the calling FR NNI, the called FR NNI may not be able to select the same DLCI as indicated in the setup message. In this case, the call is cleared.

Configuring different DLCI allocation policies on adjacent interfaces helps to ensure that DLCI collisions do not occur between adjacent FR NNIs.

When a FR NNI receives a switched call request, it must assign the same DLCI as signaled in the setup message. A DLCI collision occurs if FR NNIs in two adjacent networks assign the same DLCI in separate call setup requests and forward the setup messages to each other. The called FR NNI clears the call because the signaled DLCI has already been assigned.

To avoid this type of DLCI collision, you can use different DLCI allocation strategies at each FR NNI. If the interface at one endpoint allocates switched DLCI values from the lower end of the permitted range (that is, lowest to highest), you can configure the interface at the other endpoint to allocate switched DLCI values from the high end of the permitted range (that is, highest to lowest).



Therefore, configure the highest PVC DLCI with the same value at both the calling and called NNIs, but configure the DLCI allocation strategies at each interface to be opposite to each other.

PVC status management

Frame relay uses LMI procedures for PVC and SPVC status management. The main attributes of the LMI procedures include

- bidirectional polling between frame relay networks for information on PVC and SPVC status
- verification of physical link integrity between frame relay networks
- A-bit signaling for notification of PVC and SPVC availability
- notification of PVC status change (inactive to active or active to inactive) on a local, remote, or adjacent network

For more information, see the following sections:

- [LMI protocols \(page 40\)](#)
- [FR UNI LMI application \(page 41\)](#)
- [FR NNI LMI application \(page 41\)](#)
- [Link integrity verification \(page 41\)](#)
- [Permanent connection status information \(page 41\)](#)
- [Asynchronous status reports \(page 42\)](#)
- [A-bit signaling \(page 42\)](#)
- [N-bit signaling \(page 43\)](#)
- [Multi-network PVC status \(page 44\)](#)
- [FR UNI automatic traffic parameter configuration \(page 44\)](#)
- [FR NNI LMI signaling exchange cycle \(page 44\)](#)

LMI protocols

Nortel Multiservice Switch frame relay service supports LMI procedures running the following protocols:

- ITU-T Q.933 Annex A
- ANSI T1.617 Annex D
- *Frame Relay Specification With Extensions*, Document Number 001-208966.

For LMI to be in normal condition, both physically linked services must be running the same LMI procedures.



See [FR UNI LMI application \(page 41\)](#) and [FR NNI LMI application \(page 41\)](#) for more information on how LMI protocols are applied to FR UNI and FR NNI.

FR UNI LMI application

You can configure the network side UNI to automatically determine the procedures implemented on the user side of the link. You can also disable all LMI activity for a particular frame relay service access point to allow interworking with user equipment that does not support the C-plane function. You can configure protocol options using the *procedures* attribute under the *FrUni Lmi* component.

The in-channel signaling link automatically sets to DLCI 0 when you select the ITU-T or ANSI protocols, and to DLCI 1023 when you select the Vendor Forum protocol.

FR NNI LMI application

You can configure the LMI protocol through the *procedures* attribute under the *FrNni Lmi* component by selecting either *itu*, *ansi*, or *vendorForum* option. You can also configure the NNI to automatically determine the procedures implemented at the remote end of the link. This type of configuration is performed by selecting the *autoConfigure* option. Disabling all LMI activity for a particular frame relay service access point in order to allow interworking with user equipment that does not support the C-plane function is performed by selecting the *none* option.

Link integrity verification

The LMI signaling protocol is based on a polling scheme. Periodically, the user and network exchange sequenced status enquiry and status messages to ensure that the in-channel signaling link and the physical connection are functioning properly. A polling cycle is one exchange of status enquiry and status messages across the interface.

A link integrity verification timer indicates how frequently a status enquiry message is generated. You can configure this value by changing the value for the *linkVerificationTimer* attribute under the *Lmi* component from the default of 10 seconds to another value. By having LMI running and configuring this attribute, you determine the duration of the time interval for reporting of PVC alarm warning messages. Full status reports contain information about PVCs that have changed status during the set time interval. However, the network can respond to any poll reporting a change in PVC status within the time frame denoted by the *linkVerificationTimer* attribute.

Permanent connection status information

In addition to status enquiry messages, each network sends full status enquiry messages to the other network to request information about all PVCs and end-to-end SPVCs configured at that interface. The network responds with a



full status report that provides information about permanent connection additions, deletions, presence, and bandwidth. It also provides information about the availability or unavailability of each configured permanent connection.

The PVC alarms feature provides four types of warning messages; inactive internal and external warning messages that display when a PVC becomes inactive and active internal and external warning messages that display when a PVC becomes active. Status change of a PVC can be due to either an internal or external event, this feature generates specific warning messages for status change caused by either internal or external events on a local, remote, or adjacent network.

If a PVC becomes active in a remote or adjacent network, its status is only reported if it is also active in a local network. A PVC must be active in a local network before its remote or adjacent network status is reported for that interface.

A full status polling cycle occurs less frequently than a status enquiry polling cycle. A full status polling cycle requests a full status message every *n391* polling cycles. You can configure this value by changing the value for the *fullStatusPollingCycles* attribute under the *Lmi* component from the default of 6 to another value. By having LMI running and configuring this attribute, you determine the duration of the time interval for the reporting of PVC alarm warning messages. Full status reports contain information about PVCs that have changed status during the set time interval. However, the network can respond to any poll reporting a change in PVC status within the time frame denoted by the *fullStatusPollingCycles* attribute.

Asynchronous status reports

When the status of a PVC segment or end-to-end SPVC connection changes, the network generates a single asynchronous status message to convey the change immediately. You can configure the asynchronous status message feature through the *asyncStatusReport* attribute under the *Lmi* component.

A-bit signaling

In addition to information about configured permanent connections, full status messages contain an active bit (A-bit) for each PVC and SPVC. The A-bit indicates the availability of a PVC segment or end-to-end SPVC connection.

For a PVC, when a failure occurs in any segment of the end-to-end connection, and A-bit warns end users not to send traffic on that PVC. A-bit = 0 means that the indicated PVC segment has failed or is inactive and cannot accept user frames from the other network. When the network sends an A-bit = 0 (inactive) message through a multi-network PVC, all data transmission in



the opposite direction should be terminated. A-bit = 1 means that the indicated PVC segment is operational and ready to accept user frames. When the network sends an A-bit = 1 (active) message, transmission can resume.

For an SPVC, the connection is cleared when a failure occurs in any segment of the end-to-end connection, resulting in an A-bit = 0. When the network reports an A-bit = 0 (inactive) message for an SPVC, all data transmission in both directions should be terminated. A-bit = 1 (active) means that the indicated SPVC end-to-end connection is operational and ready to accept user frames. When the network sends an A-bit = 1 (active) message, transmission can resume.

Because the A-bit warning feature is bidirectional, the A-bit can have different values in the two directions for a PVC. At the interface, the network can report that it can no longer accept data on a PVC, but can still offer traffic from the other network on that same PVC. For an SPVC, however, the A-bit represents the connectivity of the end-to-end connection, and unlike on a PVC, it must have the same value in both directions.

You can set the *ignoreActiveBit* attribute under the *Lmi* component so that A-bit status does not affect data transfer across the PVC for both directions of traffic flow.

D-bit signaling

In addition to information about configured PVCs and SPVCs, asynchronous status messages contain a deleted bit (D-bit) for each PVC and SPVC. The D-bit indicates whether the PVC segment or the end-to-end SPVC connection is currently being deleted.

A D-bit=1 message indicates that the PVC segment or SPVC connection is just deleted. The D-bit is 0 after the PVC segment or SPVC connection is newly created. The D-bit changes to 1 when the PVC segment or SPVC connection is currently being deleted.

N-bit signaling

In addition to information about configured PVCs and SPVCs, full status messages contain a new bit (N-bit) for each PVC and SPVC. The N-bit indicates whether the PVC segment or the end-to-end SPVC connection is newly configured.

An N-bit = 1 message indicates that the PVC segment or SPVC connection is newly configured. The N-bit changes to N = 0 status after the N-bit transmits successfully to the other network.



Multi-network PVC status

A multi-network PVC can consist of multiple PVC segments across different frame relay networks. This signaling information must reflect the current status of the whole multi-network PVC, not just of individual PVC segments or interfaces.

Before the network reports to the end user that a multi-network PVC is active, the PVC must meet all the following criteria:

- All UNIs and NNIs associated with the multi-network PVC are operational.
- All PVC segments are configured and operational.
- Link integrity verification is successful at all UNIs and NNIs that are associated with the multi-network PVC.

FR UNI automatic traffic parameter configuration

Smart Management Interface Local Exchange (SMILE) is an intelligent signaling interface based on Q.933 Information Elements. The interface exchanges service parameters and quality of service information between a Nortel Multiservice Switch node and attached user-side devices. This reduces the required bilateral configuration.

The LMI must be configured to run in this mode, where PVC configuration and QoS information is included in the full status message. The information is coded in Q.933 format information elements (IEs) and includes:

- committed information rate (CIR)
- committed burst size (Bc)
- excess burst size (Be)
- transfer priority (TP)
- discard priority (DP)

This allows the user-side device to automatically configure all relevant details of all PVCs to which it has access. Also, any updates to configuration for existing PVCs are propagated automatically to the user-side device.

In order to use this capability, the user-side device must be able to interpret the additional information. Nortel Networks routers running BayRS software at v13.30 or higher can be configured to process the additional information.

FR NNI LMI signaling exchange cycle

The figures [LMI signaling for PVC status \(page 46\)](#) and [LMI signaling for SPVC status \(page 47\)](#) show the signaling information exchange over the Nortel Multiservice Switch network's NNI.



A status enquiry message requests link integrity verification every T391 seconds. Both networks respond to each status enquiry message for link integrity verification with a status message that includes sequence numbers.


A status enquiry message for full status report of all configured PVCs and SPVCs on the interface transmits every N391 polling cycles. The network responds with a full status message that contains one PVC status information element with A-bit and N-bit indicators for each configured PVC or SPVC on the interface.

The absence of a previously reported PVC or P-SPVC from the full status report indicates that the connection is no longer present on the interface. An A-bit = 0 message transmits to the end users for that PVC or P-SPVC.

When the status of an existing PVC segment or SPVC connection changes, the network generates an asynchronous status message for that PVC segment or SPVC connection. Multiservice Switch networks can also handle asynchronous status reports from external frame relay networks.




LMI signaling for PVC status





LMI signaling for SPVC status



Frame relay call management

Nortel Multiservice Switch nodes' accounting system has a call management capability that prevents accounting record loss.

Call management prevents a logical processor (LP) from accepting new SVC calls when its resources are insufficient to process the associated accounting records.



Default call management values are set automatically for each logical processor that runs frame relay UNI or frame relay NNI. The default values are large numbers that prevent the call management system from limiting calls. You can configure call management attributes for frame relay connections on a logical processor.

For more information about call management, see NN10600-550 *Nortel Multiservice Switch 7400/15000/20000 Common Configuration Procedures*. For guidelines on setting call management parameters, see the Nortel Multiservice Switch Release Notes.

FR UNI switched virtual circuit (SVC) signaling

The following standards govern SVC signaling on the C-plane:

- ITU Q.933
- ITU Q.922

ITU Q.933 is a signaling standard based on Q.931.

ITU Q.922 is the layer two specification derived from Q.921 (LAPD). Q.922 is an acknowledged protocol with sequence checking, error detection, and recovery.

For more information, see the following sections:

- [Q.933 signaling \(page 48\)](#)
- [FR UNI Q.922 signaling \(page 48\)](#)
- [Call setup procedures for SVC calls \(page 49\)](#)

Q.933 signaling

ITU standard Q.933 specifies the signaling procedure to set up and tear down a frame relay SVC call. Nortel Multiservice Switch frame relay SVC uses Q.922 as the link layer protocol. The physical connection between the user and the network must be in place before frame mode connections can establish.

FR UNI Q.922 signaling

The Q.922 (LAPF) standard defines the data link layer for the frame relay bearer service. The LAPF protocol provides a reliable data link connection for the exchange of SVC and S-SPVC signaling messages across a FR UNI. LAPF uses the Q.922 Information (I) frame type to signal the layer 3 messages.




Call setup procedures for SVC calls

A calling user initiates the establishment of an SVC connection by issuing a call setup request. The FR UNI receives the request, assigns a DLCI for the call, and forwards the setup message to the destination.

At the destination side, a physical connection between the network and called user must be in place before the call can establish. To accept the call, the called user returns a connect message. See the figure [SVC call establishment \(page 49\)](#).

The call setup message carries the SVC's traffic parameters. If the minimum acceptable CIR cannot be met in either direction, the call clears. For more information, see [SVC call negotiation \(page 91\)](#).

SVC call establishment





FR NNI SVC and SPVC signaling

The following standards govern NNI switched signaling on the C-plane:

- ITU-T Q.922
- FRF.10

FR NNI switched connections use Q.922 as the data link layer protocol. All FRF.10 (X.76) signaling messages between the two networks are sent using DLCI 0.

FR NNI P-SPVCs use FRF.2 as the data link layer protocol. All LMI signaling messages between the two networks are sent using DLCI 0.

For more information, see the following sections:

- [FR UNI Q.922 signaling \(page 48\)](#)
- [FRF.10 signaling \(page 50\)](#)
- [Call setup procedures for SPVC calls \(page 50\)](#)
- [Call setup procedures for SVC calls \(page 49\)](#)

FR NNI Q.922 signaling

The Q.922 (LAPF) standard defines the data link layer for the frame relay bearer service. The LAPF protocol provides a reliable data link connection for the exchange of SVC and S-SPVC signaling messages across a FR NNI. LAPF uses the Q.922 Information (I) frame type to signal the layer 3 messages.

FRF.10 signaling

The ITU-T X.76 standard defines the signaling procedures to set up and tear down frame relay SVCs and S-SPVCs. The FRF.10 implementation agreement describes a subset of X.76 for SVC and S-SPVC signaling, and Annex A of the ITU-T X.76 standard defines the signaling procedures to set up and tear down SPVC connections.

For more information, see [FR UNI compliance with standards \(page 265\)](#) or [FR NNI compliance with standards \(page 273\)](#).

Call setup procedures for SPVC calls

Only an SPVC master endpoint can initiate an SPVC call request. The call request specifies the address of the called party and the destination DLCI, as well as the address of the master endpoint. The traffic parameters configured by the user at the master endpoint are encoded in the Link Layer Core Parameters information element (IE) and the Frame Transfer Priority IE.



When the signaling termination equipment (STE) receives an SPVC call request that is at a transit point, it creates a S-SPVC DLCI dynamically. The SPVC call request can originate from a P- SPVC, or from another S-SPVC. If the called STE can provide the service as requested, the call progresses with the parameters as signaled in the setup message. If the called STE is unable to provide the requested service, it releases the call.

Call setup procedures for SVC calls

The calling CPE initiates the establishment of an SVC to a CPE in another network by issuing a call setup request. The FR UNI receives the request, and forwards the setup information to the FR NNI or signaling termination equipment (STE). The STE in turn sends a setup message to the STE in an adjacent network. The calling STE selects the DLCI to be specified in the setup message and includes the destination address.

The call setup message carries the SVC's traffic parameters. If CAC is enabled, the traffic parameters can be negotiated at different points in the network. The traffic parameters may be reduced, but not increased from the values specified in the setup message. The setup message transmitted to the called STE in an adjacent network reflects any reduction performed by the intermediate interfaces (FR UNI or FR NNI) while processing the SVC setup request.

If the minimum acceptable CIR cannot be met in either direction, the call clears. There is no call negotiation if CAC is turned off, and the traffic parameters values as signaled in the setup message are used. For more information, see [SVC call negotiation \(page 91\)](#).

DLCI data loopback mode

You can use the data loopback tool to check connectivity for a DLCI or load test a connection before transmitting live traffic. Loopback is useful when you install new services or troubleshoot, and does not affect traffic for other DLCIs. Both permanent and switched connections support loopback.


A data loopback command returns all valid data frames to their source, for both the local and remote ends. A local loopback command (see the figure [Local DLCI loopback \(page 52\)](#)) returns valid frames received from the link. A remote loopback command (see the figure [Remote DLCI loopback \(page 52\)](#)) returns valid frames received from the subnet. Neither operation and maintenance (OAM) nor local management interface (LMI) frames support DLCI loopback.




In remote DLCI loopback mode, DLCI statistics counters increments in both directions of the datapath. Transmit statistic counters increment when a frame arrives from the subnet and receive statics counters increment as the frame is sent back to the subnet. However, DLCI statistics counters are not updated in local DLCI loopback mode.

Loopback commands do not reset the target DLCI. The PVC resiliency feature cannot force the loopback to move to another DLCI. For more information about how loopback works, see [Loopback and A-bit signaling \(page 54\)](#).

Local DLCI loopback



Remote DLCI loopback



You can run loopback locally, for frames received from the link, or remotely for frames received from the subnet.



You do not need to lock components to operate loopback mode. The table [Locked components in loopback mode \(page 53\)](#) summarizes how locked components behave in loopback mode.

When the VC is cleared (for example, if the remote DLCI is locked), the local loopback still functions. Remote loopback does not function until the VC is back up.

Locked components in loopback mode

Locked component	Loopback behavior
FR UNI or FR NNI	When you lock the interface, the effect on loopback is the same as when you lock the DLCI.
DLCI	When you lock a DLCI, you clear the VC and disable the DLCI datapath, so that it can no longer process data frames. Local loopback still functions on a locked DLCI, but you must unlock the DLCI to loop frames back remotely.
Port	When you lock the port, you disable local loopback mode, but remote loopback still functions. Frames from the subnet can still loop back to the remote end. No A-bit=0 message transmits through the subnet

Frame relay virtual circuit disconnection and recovery

Frame relay VC disconnection can be triggered by:

- the application.
- the detection of network failure: the loss of a path between two ends of a VC.

Disconnection triggered by application

A frame relay VC disconnects when it receives a disconnect message. When an application connected to a VC sends a disconnect request to it (for example lock FrUni/FrNni, lock FrUni/FrNni DLCI, or clear FrUni DLCI for SVC), the VC begins to exchange messages with the remote end to terminate the existing connection. The VC is disconnected a few seconds after the trigger event occurs.

Disconnection triggered by detection of network failure

A VC has an internal mechanism to detect a network environment. A VC will clear a connection if there is a persistent network failure.

The length of time required to detect network failure depends on an internal timer which runs every 60 seconds and on the status of the VC (that is, whether it is carrying data or is idle). Disconnection of a VC can take from 21 to 144 seconds after a network failure event occurs.



The above VC disconnect time considers only one end of the VC connection. Existing network delay is not considered.

Frame relay virtual circuit reconnection

Permanent VCs (PVCs) are designed to reconnect every 60 seconds. Reconnection times for switched VCs depend on the application making the SVC calls.

PVC reconnection can be divided into the following two situations:

- PVC automatic reconnect
- PVC reconnect triggered by an application

After a PVC is disconnected, it will automatically try to reconnect. The reconnect attempts continue indefinitely.

There are two timers that contribute to the recovery: `callSetupTimer` and `callRetryTimer`.

`CallSetupTimer` is set when a reconnect process is started. If a reconnect process cannot be finished during `callSetupTimer`, the VC will give up this try and set `callRetryTimer`. After `callRetryTimer` has expired, the VC begins a reconnect process again.

PVC reconnection triggered by application

Apart from Automatic Reconnection, a reconnection can be triggered by an application such as `FrUni/FrNni`. For example, the lock/unlock command under CAS can trigger a reconnection immediately. If this attempt is not successful, the VC will repeat the reconnect attempt every 60 seconds.

Loopback and A-bit signaling

For the DLCI data loopback feature to work, the application must be able to put traffic on the monitored DLCI, even when the LMI or the remote link fails.

When a DLCI receives a *start* loopback command, the LMI sends an A-bit = 1 (active) message to the local CPE and remote CPE. This indicates to the end devices that the DLCI is ready for data transfer. This change in the LMI interface allows the end devices to send data frames destined for that DLCI.

While the data loopback is in effect, any change in the A-bit status does not transmit to the CPEs. From the CPE's perspective, the DLCI is up for the duration of the loopback mode. In addition, a DLCI in loopback mode ignores A-bit signals from the subnet. A-bit signals that come from the subnet to a DLCI in loopback mode do not affect that DLCI's A-bit status.



When loopback stops, the PVC disconnects and reconnects, and restores the original A-bit status. For more information about A-bit signaling, see [PVC status management \(page 40\)](#) and [A-bit signaling \(page 42\)](#).

When data loopback is used with PVC backup and a failure occurs on the primary PVC, the DLCI on that PVC indicates that the connection is down via the A-bit status (A-bit=0) and all ingress traffic is discarded. In order to avoid traffic loss, you can stop loopback on the primary PVC before the connection is re-established. After the connection is re-established, loopback can resume on the primary PVC. If the primary PVC has already been re-established and the A-bit is active (A-bit=1), you can issue a Lock and then an Unlock on the DLCI of the primary PVC to prevent this type of condition from recurring.

Fast 1:1 FP equipment protection for Multiservice Switch 7400 FR UNI interfaces

This feature upgrades FR UNI services on the standby FP to enable a hot standby status. In a hot standby state, the FR UNI service has all its software loaded; all necessary provisioning data delivered and all network connections made. The standby FP is ready to rapidly go into service carrying user traffic in the event of failure to the currently active FP. This functionality is only available on a limited range of Nortel Multiservice Switch 7400 node FPs and only for PVC services.


The following Multiservice Switch 7400 nodes FPs running permanent virtual circuits (PVCs) and provisioned with frUniPvcDprs and frUniPvcSparingDprs support this feature:

- 32pDS1Msa, supports spareDNA
- 32pE1Msa, supports spareDNA
- 8pDS1
- 4pE1Ch, supports spareDNA
- 4pDS1Ch, supports spareDNA
- HSSI, supports FrUni Dlic Spare Dc component


You can only provision the frUniPvcDprs and frUniPvcSparingDprs features in the feature list. You cannot provision other features on an FP that is expected to perform Fast 1:1 FP equipment protection.



FP equipment protection before switchover



FP equipment protection after switchover





Frame relay UNI management

Nortel Multiservice Switch nodes' frame relay service supports functionality for redundancy and loadsharing among FR UNIs. In addition, you can partition your network logically by configuring closed user groups (CUGs).

Permanent connections can exist between Multiservice Switch node frame relay service and DPN frame relay service as long as the Nortel Multiservice Switch network includes a Multiservice Switch 7400 series node.

Navigation

- [Call redirection in a Multiservice Switch network \(page 57\)](#)
- [Hunt groups \(page 62\)](#)
- [Closed user groups \(CUGs\) \(page 65\)](#)

Call redirection in a Multiservice Switch network

Nortel Multiservice Switch nodes' frame relay UNI service supports redundancy for PVC, SPVC and SVC connections through call redirection. Call redirection shifts calls to alternate destinations if their primary destination is unavailable. This provides redundant paths for call attempts, so subscribers enjoy more highly reliable data transmission.

Call redirection requires the presence of a call redirection server in the Multiservice Switch network, or a DPN call redirection server in an interworked network. A destination is supported by a configured list of backup addresses. This redirection list includes a register of primary DNAs with their corresponding configurable backup DNAs. You can configure up to seven backup DNAs for each primary DNA.

For more information, see the following sections:

- [Call redirection for PVCs in a Multiservice Switch network \(page 58\)](#)
- [Call redirection for SPVCs in a Multiservice Switch network \(page 59\)](#)
- [Call redirection for SVCs in Multiservice Switch network \(page 60\)](#)
- [Situations for call redirection \(page 61\)](#)



- [Accounting records for call redirection \(page 62\)](#)

For information about configuring Multiservice Switch call redirection servers, see NN10600-410 *Nortel Multiservice Switch 7400/15000/20000 Operations: Call Redirection Server*. For information about the DPN call redirection server, see 241-1001-115 *DPN-100 Call Redirection System User Guide*.

Call redirection for PVCs in a Multiservice Switch network

For PVCs, call redirection occurs on each connection, but the DNA controls the actual redirection process. Call redirection capability requires the DLCIs for the backup DNAs to have the same configured attributes as the primary DNA.


The redirection process begins when a master endpoint attempts to establish a PVC connection with its slave endpoint. If the destination device cannot accept the PVC connection request, the request goes to the call redirection server (CRS). The CRS searches its call redirection list for an alternate address.

The CRS selects the designated backup DNA and forwards the connection request. If the backup DNA accepts the call, a connection occurs. If the DNA rejects the call, the CRS selects the next backup DNA in its redirection list and attempts to complete the call. The process continues until the call completes or until the CRS exhausts its redirection list. See the figure [Call redirection for PVCs \(page 59\)](#).

When the primary DNA comes up again, it returns to operation automatically. The primary DNA sends a call request to the master end, and clears the established backup PVC. The master end then sends a call request to the primary DNA to re-establish the call.



Call redirection for PVCs



Call redirection for SPVCs in a Multiservice Switch network

For SPVCs, call redirection is only supported by Nortel Multiservice Switch call redirection servers. DPN call redirection servers do not support SPVC calls. Call redirection capability also requires that the DLCIs for the backup SPVC slaves have the same configured attributes as the primary SPVC slaves.



The redirection process begins when a master endpoint attempts to establish an SPVC connection with its slave endpoint. If the destination device cannot accept the SPVC connection request, the request goes to the call redirection server (CRS).

The CRS searches its call redirection list for an alternate address, and forwards the connection request to the designated backup DNA. If the backup DNA accepts the call, a connection occurs. If the DNA rejects the call, the CRS selects the next backup DNA in its redirection list and attempts to complete the call. The process continues until the call completes or until the CRS exhausts its redirection list.

Unlike PVC redundancy, SPVC backups remain connected to the master endpoint, even if the primary DNA recovers and is back in service.

Call redirection for SVCs in Multiservice Switch network


For SVCs, call redirection occurs during the establishment phase of a call only. Call redirection does not occur once a call request goes to the called DTE or if the called DTE rejects the call request.

The redirection process begins when a calling DTE tries to establish a connection. If the called FR UNI cannot forward the request to the called DTE, the call request goes to the call redirection server (CRS). The CRS searches its redirection list and sends the request to one of the alternate DNAs.

If the alternate DNA accepts the call, a connection is established. If the call is rejected, the call request returns to the CRS. The CRS redirects the call request to the next alternate DNA on the call redirection list. The process continues until the call establishes or until the redirection list is exhausted. If the redirection list is exhausted, the call does not establish. See the figure [Call redirection for SVCs \(page 61\)](#).



Call redirection for SVCs



Situations for call redirection

Call redirection can occur under any of the following circumstances:

- node isolation or failure (Multiservice Switch trunk or CP failure)
- FP failure
- link failure
- LMI failure (for PVCs) or LAPF failure (for SVCs)
- locked interface (frame relay service not operational)



- memory congestion on the interface
- no SVC DLCI available on the interface (for SVCs only)
- failure of backup DNA
- primary DNA not configured
- frame relay signaling not supported on the interface (for SVCs only)

Accounting records for call redirection

Nortel Multiservice Switch frame relay PVC, SPVC, and SVC accounting capabilities are the same. The original DNA field can be part of the accounting records for redirected calls.

When the primary DNA fails and a call routes through a backup DNA, the network's accounting capability tracks this activity for both the primary and backup connections.

When a call redirects, the redirect flag is set to 1 in the accounting record; the accounting record also contains the called backup DNA. The primary DNA's account information registers in the set of *originalCalledAddressFax* values in the *accountingFax* (*afax*) attribute of the *Mod Vcs* component.

If you configure single-end accounting, an accounting record generates at the calling end only when a call redirects to a backup DNA. If you configure double-ended accounting, an accounting record generates at both the calling and primary called ends.

For more information about accounting records, see [Frame relay accounting \(page 177\)](#) or NN10600-560 *Nortel Multiservice Switch 7400/15000/20000 Accounting*.

Hunt groups

A hunt group is a single data network address (DNA) that represents a group of service DNAs. By using a hunt group, a group of FR UNI services can be accessed from a single DNA. Hunt groups also provide resiliency and allow the network to route around a failed or congested FR UNI.

A hunt group consists of a hunt group server and hunt group members. A hunt group supports a maximum of 63 hunt group members. A FR UNI can belong to either one or two hunt groups.

For more information, see the following sections:

- [Hunt group members \(page 63\)](#)
- [Multiservice Switch Hunt Group Server \(page 63\)](#)
- [Hunt group call setup \(page 63\)](#)



- [Hunt group member availability status \(page 64\)](#)
- [Backup hunt groups \(page 65\)](#)

If you set either the *ingressEirPoolAvailableBandwidth* or *egressEirPoolAvailableBandwidth* attribute to yes, the network establishes the EIR as the smaller value of the user-requested EIR or the line speed.

Hunt group members

To ensure that proper loadspreading occurs within a hunt group, you must configure similar service parameters on each hunt group member. Each member FR UNI must have

- the same address configured under the *FrUni Dna HuntGroupMember (Hgm)* component
- similar service parameters configured under the *FrUni Signaling (Sig)* component
- similar service parameters provisioned under the *FrUni Ca* component

Multiservice Switch Hunt Group Server

Nortel Multiservice Switch Hunt Group Server (HGS) resides on a Multiservice Switch node and has a unique DNA. The HGS maintains a list of hunt group members and their availability. The HGS receives call requests from the network and forwards them to the appropriate hunt group member. You can configure the HGS to select hunt group members based on different policies.

For information about installing and configuring the HGS, see NN10600-415 *Nortel Multiservice Switch 7400/15000/20000 Operations: Hunt Group Server*.

Hunt group call setup


When the HGS receives a call request from a DTE, the HGS searches its member list for an appropriate FR UNI and forwards the call to that hunt group member. If the FR UNI is available, it accepts the call and establishes a switched connection with the DTE. See the figure [Hunt group call set up process \(page 64\)](#).

If the first hunt group member cannot accept the call, it sends a call rejection message to the HGS. The HGS sends the call request to the next hunt group member in the list. You can configure the HGS to start at the top of the list, to start with the member coming immediately after the last selected member in the list, or to start with the member that reports the highest availability.



A call clears if hunt group members do not share the same service characteristics. In addition, if a hunt group member rejects a call because of invalid information elements, insufficient available CIR or incompatible call facilities, the call clears.

Hunt group call set up process



Hunt group member availability status

Each hunt group member sends an availability message packet (AMP) to the HGS when its availability status changes significantly. Availability status is defined in bits/s of available CIR.



The availability status is calculated from the available aggregate bandwidth, that is, the sum of the currently available ingress bandwidth pools. You can verify the available aggregate bandwidth by displaying the *availableAggregateCir (acir)* attribute under the *FrUni Dna HgM* component.

The availability value reported in the AMP ranges from 0 to 4095. The aggregate bandwidth must be converted from units of bits/s to units of 256 bytes (or 2048 bits) per second. Therefore, the *acir* attribute value is divided by 2048 before it is sent to the HGS:

$$\text{availabilityValue} = \text{availableAggregateCir} / 2048$$

A hunt group member reports an availability value of

- 0 if its available aggregate CIR is less than 1024 bits/s
- 1 if its available aggregate CIR is between 1024 and 2047 bits/s
- 4095 if its available aggregate CIR is greater than 8 386 560 bits/s

When a hunt group member becomes unavailable and the HGS is reset, the hunt group member will resend its availability status when the hunted incoming call is cleared.

Backup hunt groups

A call redirection server (CRS) can provide backup capability for a hunt group server. If a call request cannot access a hunt group server, the network transmits the call request to a CRS. The CRS uses the hunt group server's backup hunt group DNA.

Primary and backup DNAs must contain the same hunt group members and search mode. All members report availability status to both primary and backup hunt group servers.

To use the hunt groups backup capability, all Nortel Multiservice Switch 7400, Multiservice Switch 15000, and Multiservice Switch 20000 nodes must be able to route failed calls to a CRS.

See NN10600-410 *Nortel Multiservice Switch 7400/15000/20000 Operations: Call Redirection Server* for more information.

Closed user groups (CUGs)

A closed user group (CUG) is a group of DTE interfaces within a network that share incoming or outgoing access privileges. You can create CUGs to partition a network logically. CUG members can communicate only with members of the same CUG.



Nortel Multiservice Switch nodes support national and international CUGs. A national CUG is a group of DTEs in a single network. An international CUG defines a group of DTEs that span network boundaries.

For more information, see the following sections:

- [Simple CUGs \(page 66\)](#)
- [Selectable CUGs \(page 66\)](#)
- [CUG subscription at the FR UNI \(page 66\)](#)
- [Interlock codes \(page 70\)](#)
- [CUG signaling \(page 71\)](#)

Simple CUGs

In a simple CUG, the associated DTE can subscribe to only one CUG. If you add a second CUG where a simple CUG exists, the semantic check fails. Simple CUGs do not transmit CUG index numbers between the DTE and the DCE.

Simple CUGs allow incoming and outgoing calls. You must set the *incCalls* and *outCalls* attributes under the *FrUni Dna Cug* component to allowed. In addition, you must set the *preferential* attribute to yes. You cannot set the *preferential* attribute to yes unless the *incCalls* and *outCalls* attributes are set to allowed.

The concept of a simple CUG does not apply to PVCs.

Selectable CUGs

If you configure a selectable CUG, the associated DTE can subscribe to more than one CUG. A DTE can subscribe to as many as 256 CUGs. You assign different index numbers to each CUG to distinguish them. Interlock codes distinguish different CUGs on the same interface. The index numbers map to the DCE's table of interlock codes.

You can configure a CUG index for each SVC or PVC connection. For SVCs, you configure the CUG index under the *FrUni Dna* component. For PVCs, you configure the CUG index number under the *FrUni Dci Dc* component.

For SVCs, the DTE signals the CUG index number in the call setup request. For PVCs, the CUG index is not signaled because the PVC call setup request is provided by the DCE with no direct user signal.

CUG subscription at the FR UNI

For SVCs, the X.36 call setup message contains the CUG Information Element (IE). The IE is either CUG selection or CUG selection with outgoing access (CUG OA).



You can configure CUG signaling through the *cugType* attribute under the *FrUni Dci Dc* component. This attribute determines whether the *cugIndex* is signaled or not. If the *cugIndex* is signaled, you can set the CUG type to CUG selection or CUG selection with outgoing access (CUG OA).

The table [CUG subscription options \(page 67\)](#) summarizes the configuration options for each CUG subscription type. You configure the following attributes when choosing your subscription options:

- *incAccess* attribute (*FrUni Dna* component)
If this attribute set to allowed, the interface can accept incoming calls from the network to the DTE.
- *outAccess* attribute (*FrUni Dna* component)
If this attribute is set to allowed, the interface can accept outgoing calls from the DTE to the network.
- *incCalls* attribute (*FrUni Dna Cug* component)
If this attribute is set to disallowed, all incoming calls on the CUG are barred
- *outCalls* attribute (*FrUni Dna Cug* component)
If this attribute is set to disallowed, all outgoing calls on the CUG are barred.

Due to technical standards, the options outlined in the table, [CUG subscription options \(page 67\)](#) are the only supported combinations. For example, you could not configure all the attributes of a simple CUG (IA) to allow.

CUG subscription options

CUG type	FrUni Dna attributes		FrUni Dna Cug attributes		
	incAccess	outAccess	incCalls	outCalls	preferential
simple CUG (IA)	disallow	disallow	allow	allow	yes
simple CUG (OA)	disallow	disallow	allow	allow	yes
selectable CUG	disallow	disallow	allow	allow	no
selectable CUG (OCB)	disallow	disallow	allow	disallow	no
selectable CUG (ICB)	disallow	disallow	disallow	allow	no
selectable CUG (IA)	allow	disallow	allow	allow/ disallow	no
selectable CUG (OA)	disallow	allow	allow/ disallow	allow	no
selectable CUG (IA/OA)	allow	allow	disallow	disallow	no



The network processes a CUG call depending on the CUG subscription at both the calling and called ends. CUGs use the following ITU-T X.36 clear cause codes:

- (50): Requested facility not subscribed
- (87): User not member of CUG
- (90): Non-existent CUG

Call setup

The table, [Calling end \(page 68\)](#), shows the CUG signaling in the network, resulting from CUG signaling by the calling interface and the calling interface's CUG subscriptions. The table also describes the action selected by the DCE when call setup signals no CUG, CUG, or CUGOA. For example, if call setup is issued with CUG signaled and CUG subscription is "No CUG", then the interlock code does not match and the call is cleared.

Calling end

Type of call in the SETUP message [1]					
	Not a CUG call	CUG selection facility		CUG/OA selection facility	
CUG subscription of the calling interface		Match	No Match or Match/OCB [4]	Match	No match or Match/OCB [4]
No CUG	Normal call [5]	Clear call		Clear call	
CUG [3]	Clear call	Call with specified CUG	Clear call	Clear call	Clear call
CUG/OA [2]	Normal call [5]	Call with specified CUG	Clear call	Call with specified CUG + OA	Normal call [5]
[1] The inclusion of both the CUG and CUG/OA selection facility is not allowed in the call request. [2] OA: interface with outgoing access allowed (for example, DNA attribute "outAccess" provisioned to the value "allowed"). [3] Interface with outgoing access disallowed (for example, DNA attribute "outAccess" provisioned to the value "disallowed"). [4] OCB: CUG with outgoing calls disallowed (for example, the matching CUG found has its attribute "outCalls" provisioned to the value "disallowed"). [5] The outgoing call signaled into the network contains no CUG or CUG/OA signaling.					



Call completion

The table, [Called end \(page 69\)](#) shows the CUG signaling to the called interface resulting from CUG signaling in the network and the called interface's CUG subscription. The table also describes the action selected by a DCE when an incoming call contains no CUG, CUG, or CUGOA. For example, if an incoming call arrives with a CUG, and the subscription is "CUG", then the call is allowed because the interlock code matches. If the interlock code does not match, the call is cleared.

Called end

Type of call in the subnet message					
	Not a CUG call	CUG selection facility		CUG/OA selection facility	
CUG subscription of the called interface		Match	No match or Match/ICB [3]	Match	No match or Match/ICB [3]
No CUG	Normal call [4]	Clear call		Normal call [4]	
CUG [2]	Clear call	Call with specified CUG	Clear call	Call with specified CUG	Clear call
CUG/IA [1]	Normal call [4]	Call with specified CUG	Clear call	Call with specified CUG + OA	Normal call [4]
[1] IA: interface with incoming access allowed (for example, DNA attribute "incAccess" provisioned to the value "allowed"). [2] IA: interfaces with incoming access disallowed (for example, DNA attribute "incAccess" provisioned to the value "disallowed"). [3] ICB: CUG with incoming calls disallowed (for example, the matching CUG found has its attribute "incCalls" provisioned to the value "disallowed"). [4] The incoming call signaled to the interface contains no CUG or CUG/OA selection facility.					

Example of CUG subscription

In the figure [Example of CUG subscription \(page 70\)](#), calls can transmit between DTE1 and DTE2 and between DTE3 and DTE4. Their CUG identifiers match, indicating that they belong to the same CUG.

DTE1 has neither incoming nor outgoing access, so it communicates only with DTEs that have the same CUG identifier (DTE2).




DTE2 has incoming access and can receive calls from the open part of the network (DTE5) and from DTEs with outgoing access (DTE3 and DTE4).

DTE3 has outgoing access and can call the open part of the network (DTE5) and DTEs with incoming access (DTE2).

DTE4 has incoming and outgoing access, so it can call DTEs with incoming access (DTE2) and the open part of the network (DTE5). DTE4 can also accept calls from the open part of the network (DTE5) and from any DTEs with outgoing access.

Example of CUG subscription



Interlock codes

The subnet uses interlock codes to determine which calls have access permissions to a CUG. An interlock code can map to different index numbers on different nodes. Interlock codes remain constant throughout the network, but index numbers are specific to the node.

The FR UNI transmits interlock codes within the subnet; but transmits the associated index number to a calling or called DTE. Selectable CUGs transmit the CUG index between the DTE and the subnet. Simple CUGs do not transmit the CUG index.



All interlock codes of the same type (national or international) must be unique. If the interlock codes are not unique, the destination node cannot identify the CUG.


CUG signaling

The DTE transmits the CUG index number to the FR UNI. The FR UNI searches its CUG index table for the transmitted index number, and maps the index number to its interlock code. The FR UNI then transmits the interlock code across the subnet to the appropriate interface.

The receiving interface searches its CUG index table for the interlock code. The interface maps the interlock code to the CUG index, and transmits the CUG index to the called DTE. See the figure [Selectable CUG signaling \(page 71\)](#).

The network verifies the destination interface's access privileges. First, the network verifies incoming and outgoing access, then it verifies barred incoming and outgoing calls. The network accepts calls only if the subscription options are set to allow the call.

Selectable CUG signaling





Frame relay NNI management

Nortel Multiservice Switch nodes' frame relay service supports functionality for redundancy and loadsharing among FR NNIs.

Navigation

- [Call redirection in a Multiservice Switch network \(page 72\)](#)
- [Gateway call routing to external networks \(page 76\)](#)

Call redirection in a Multiservice Switch network

Nortel Multiservice Switch nodes' frame relay service NNI supports redundancy for PVC and SPVC connections through call redirection. Call redirection shifts calls to alternate destinations if their primary destination is unavailable. This provides redundant paths for call attempts, so subscribers enjoy more highly reliable data transmission.

Call redirection requires the presence of a call redirection server in the Multiservice Switch network, or a DPN call redirection server in an interworked network. A destination is supported by a configured list of backup addresses. This redirection list includes a register of primary DNAs with their corresponding configurable backup DNAs. You can configure up to seven backup DNAs for each primary DNA.

Permanent connections can exist between Multiservice Switch frame relay service and DPN frame relay service as long as the network includes a Multiservice Switch 7400 series node.

For more information, see the following sections:

- [Call redirection for PVCs in a Multiservice Switch network \(page 73\)](#)
- [Call redirection for SPVCs in a Multiservice Switch network \(page 74\)](#)
- [Manual call redirection \(page 75\)](#)
- [Situations for call redirection \(page 75\)](#)
- [Accounting records for call redirection \(page 76\)](#)



For information about configuring the call redirection server, see NN10600-410 *Nortel Multiservice Switch 7400/15000/20000 Operations: Call Redirection Server*. For information about the DPN call redirection server, see 241-1001-115 *DPN-100 Call Redirection System User Guide*.

Call redirection for PVCs in a Multiservice Switch network

For PVCs, call redirection occurs on each connection, but the DNA controls the actual redirection process. Call redirection capability requires that the DLCIs for the backup DNAs have the same configured attributes as the primary DNA.


The redirection process begins when a master endpoint attempts to establish a PVC connection with its slave endpoint. If the destination device cannot accept the PVC connection request, the request goes to the call redirection server (CRS).

The CRS searches its call redirection list for an alternate address, and forwards the connection request to the designated backup DNA. If the backup DNA accepts the call, a connection occurs. If the DNA rejects the call, the CRS selects the next backup DNA in its redirection list and attempts to complete the call. The process continues until the call completes or until the CRS exhausts its redirection list. See the figure [Call redirection for FR NNI PVCs \(page 74\)](#).

When the primary DNA comes up again, it returns to operation automatically. The primary DNA sends a call request to the master end, and clears the established backup PVC. The master end then sends a call request to the primary DNA to re-establish the call.



Call redirection for FR NNI PVCs



Call redirection for SPVCs in a Multiservice Switch network

For SPVCs, call redirection is only supported by Nortel Multiservice Switch call redirection servers. DPN call redirection servers do not support SPVC calls. Call redirection capability also requires that the DLCIs for the backup SPVC slaves have the same configured attributes as the primary SPVC slaves.

The redirection process begins when a master endpoint attempts to establish an SPVC connection with its slave endpoint. If the destination device cannot accept the SPVC connection request, the request goes to the call redirection server (CRS).

The CRS searches its call redirection list for an alternate address, and forwards the connection request to the designated backup DNA. If the backup DNA accepts the call, a connection occurs. If the DNA rejects the call, the CRS selects the next backup DNA in its redirection list and attempts to complete the call. The process continues until the call completes or until the CRS exhausts its redirection list.

Unlike PVC redundancy, SPVC backups remain connected to the master endpoint, even if the primary DNA recovers and is back in service.




Manual call redirection

When two Nortel Multiservice Switch networks are connected, calls redirect automatically. You need the manual locking capability only when you connect to a non-Multiservice Switch network. Different vendors' networks do not always react in the same manner when they detect NNI failure, so you can redirect calls manually.

When you configure the network for manual redirection, redirection occurs only after you lock the primary slave. If you do not lock the primary slave, redirection does not take place, and the end-to-end connection remains down until the NNI recovers.

The figure [Manual call redirection \(page 75\)](#) shows a locked primary slave. A failure can occur at any point in the network. The Multiservice Switch network cannot detect the type of failure in networks B or C. In such a case, you can use manual intervention.

Manual call redirection



Situations for call redirection

Call redirection can occur under any of the following circumstances:

- node isolation or failure (Nortel Multiservice Switch trunk or CP failure)
- FP failure
- link failure
- LMI failure
- locked interface (frame relay service not operational)



- memory congestion on the interface
- failure of backup DNA
- primary DNA not configured

Accounting records for call redirection

Nortel Multiservice Switch nodes' frame relay PVC, SVC, and SPVC accounting capabilities are the same. The original DNA field can be part of the accounting records for redirected calls.

When the primary DNA fails and a call routes through a backup DNA, the network's accounting capability tracks this activity for both the primary and backup connections.

When a call redirects, the redirect flag is set to 1 in the accounting record; the accounting record also contains the called backup DNA. The primary DNA's account information registers in the set of *originalCalledAddressFax* values in the *accountingFax* (*afax*) attribute of the *Mod Vcs* component.

If you configure single-end accounting, an accounting record generates at the calling end only when a call redirects to a backup DNA. If you configure double-ended accounting, an accounting record generates at both the calling and primary called ends.

For more information about accounting records, see NN10600-560 *Nortel Multiservice Switch 7400/15000/20000 Accounting*.

Gateway call routing to external networks

To establish SVCs and S-SPVCs to external networks, the call must be routed to an appropriate NNI through a gateway call router. You can use the Hunt Group Server (HGS) to provide the gateway call routing functionality necessary for end-to-end connectivity.

The DPN Hunt Group Server (for interworking between Nortel Multiservice Switch 7440 node and DPN) supports SVC calls only. The Hunt Group Server must reside on a Multiservice Switch node for SPVC call support.

For more information, see the following sections:

- [Hunt groups \(page 77\)](#)
- [Multiservice Switch Hunt Group Server \(page 77\)](#)
- [Hunt group call setup \(page 78\)](#)
- [Hunt group member availability status \(page 79\)](#)
- [Backup hunt groups \(page 80\)](#)



Hunt groups

A hunt group is a single data network address (DNA) that represents a group of service DNAs. By using a hunt group, a group of FR NNI services can be accessed from a single DNA. Hunt groups also provide resiliency and allow the network to route around a failed or congested FR NNI.

A hunt group consists of a hunt group server and one or more hunt group members. A hunt group supports a maximum of 64 hunt group members. A FR NNI can belong to one or two hunt groups.

To ensure that proper loadspreading occurs within a hunt group, you must configure similar service parameters on each hunt group member. Each member FR NNI must have

- the same address configured under the *FrNni Dna HuntGroupMember (Hgm)* component
- similar service parameters configured under the *FrNni Signaling (Sig)* component
- similar service parameters provisioned under the *FrNni Ca* component

If you set either the *ingressEirPoolAvailableBandwidth* or *egressEirPoolAvailableBandwidth* attribute to yes, the network establishes the EIR as the smaller value of the user-requested EIR or the line speed.

Multiservice Switch Hunt Group Server

A Nortel Multiservice Switch Hunt Group Server (HGS) resides on a Multiservice Switch node and has a unique DNA. For gateway call routing to external networks, you must configure its DNA with a unique prefix that identifies the external network.

The HGS maintains a list of hunt group members and their availability. The HGS receives call requests from the network and forwards them to the appropriate hunt group member. You can configure the HGS to select hunt group members based on different policies.

You must configure a separate instance of the HGS for each external network that you want to reach through SVCs or SPVCs. The DNA of each HGS instance must correspond to a unique prefix that identifies the associated external network. For example, to connect to another Multiservice Switch network, you must configure the hunt group DNA with the same value as the network identifier (DNIC) of the external network (that is, the value of the *networkIdCode* attribute under the *Mod Vcs* component). You must also configure the FR NNIs that connect to the external network as members of the hunt group.



For information about installing and configuring the HGS, see NN10600-415 *Nortel Multiservice Switch 7400/15000/20000 Operations: Hunt Group Server*.

Hunt group call setup


When the network receives a call destined for an external network, it forwards the request to the HGS whose DNA corresponds to the unique prefix of the external network. The HGS searches its member list for an appropriate FR NNI and forwards the call to that hunt group member. If the FR NNI is available, it accepts the call and establishes a switched connection with the FR NNI in the adjacent network. See the figure [Hunt group call set up process \(page 79\)](#).

If the first hunt group member cannot accept the call, it sends a call rejection message to the HGS. The HGS sends the call request to the next hunt group member in the list. You can configure the HGS to start at the top of the list, to start with the member coming immediately after the last selected member in the list, or to start with the member that reports the highest availability.

A call clears if hunt group members do not share the same service characteristics. In addition, if a hunt group member rejects a call because of invalid information elements, insufficient available CIR or incompatible call facilities, the call clears.



Hunt group call set up process



Hunt group member availability status

Each hunt group member sends an availability message packet (AMP) to the HGS when its availability status changes significantly. Availability status is defined in bits/s of available CIR.

The availability status is calculated from the available aggregate bandwidth, that is, the sum of the currently available ingress bandwidth pools. You can verify the available bandwidth in the ingress bandwidth pool by displaying the *ingressPoolAvailableBandwidth (ingAvailBw)* attribute under the *FrNni Ca* component. You can verify the available aggregate bandwidth by displaying the *availableAggregateCir (acir)* attribute under the *FrNni Dna HgM* component.



The availability value reported in the AMP ranges from 0 to 4095. The aggregate bandwidth must be converted from units of bits/s to units of 256 bytes (or 2048 bits) per second. Therefore, the *acir* attribute value is divided by 2048 before it is sent to the HGS:

$$availabilityValue = availableAggregateCirr/2048$$

A hunt group member reports an availability value of

- 0 if its available aggregate CIR is less than 1024 bits/s
- 1 if its available aggregate CIR is between 1024 and 2047 bits/s
- 4095 if its available aggregate CIR is greater than 8 386 560 bits/s

Backup hunt groups

A call redirection server (CRS) can provide backup capability for a hunt group server. If a call request cannot access a hunt group server, the network transmits the call request to a CRS. The CRS uses the hunt group servers's backup hunt group DNA.

Primary and backup DNAs must contain the same hunt group members and search mode. All members report availability status to both primary and backup hunt group servers.

To use the hunt groups backup capability, all Nortel Multiservice Switch nodes must be able to route failed calls to a CRS.

See NN10600-410 *Nortel Multiservice Switch 7400/15000/20000 Operations: Call Redirection Server* for more information.



Bandwidth management

Nortel Multiservice Switch nodes' frame relay service supports traffic parameter negotiation for switched calls to protect network resources from becoming oversubscribed.

In addition, you can restrict the amount of bandwidth allocated to a given type of traffic in both the ingress and egress directions.

If ingress or egress Call Admission Control (CAC) is enabled, the interface only accepts a call request if there is enough bandwidth available in the assigned pool to accommodate the requested service. If there are insufficient resources available, the interface negotiates traffic parameters for switched connections.

Navigation

- [Bandwidth pool management \(page 81\)](#)
- [Connection admission control \(page 86\)](#)
- [SVC call negotiation \(page 91\)](#)
- [FR UNI CAC migration \(page 95\)](#)

Bandwidth pool management

You can share port capacity among 16 different pools for both ingress and egress committed information rate (CIR) and excess information rate (EIR) traffic. Each pool services one or more transfer priorities.

The capacity allocated to each transfer priority accommodates switched connections only. Permanent connections are not accounted for in the bandwidth pools.

For more information, see the following sections:

- [Bandwidth pools \(page 82\)](#)
- [Over- and under-subscribing port capacity \(page 83\)](#)
- [Full sharing of port capacity \(page 84\)](#)



- [Transfer priority and bandwidth pool assignment \(page 85\)](#)

For information about frame relay transfer priorities, see [Transfer priority \(page 151\)](#).

Bandwidth pools

You can partition the port capacity into a maximum of 16 different CIR and EIR bandwidth pools for traffic in the ingress and egress directions. You can assign each bandwidth pool a percentage of the link capacity through the following attributes of the *FrUni Ca* or *FrNni Ca* component:


- *ingressCirBandwidthPool*
- *ingressEirBandwidthPool*
- *egressCirBandwidthPool*
- *egressEirBandwidthPool*

FR UNI or FR NNI CAC, if enabled, uses the configured values to determine whether resources are available for a new connection.

The figure [Port capacity to bandwidth pool mapping \(page 83\)](#) shows bandwidth pools and their associated port capacity.



Port capacity to bandwidth pool mapping



Over- and under-subscribing port capacity


The percentages assigned to the 16 CIR and EIR bandwidth pools do not have to add up to 100 percent. The percentage of link capacity for a bandwidth pool can vary between 0 and 1000 percent. This flexibility allows for over- and under-subscribing the port as shown in the figure [Over- and under-subscribing \(page 84\)](#). You can use this flexibility to increase or decrease port usage.

Over-subscription is a set of percentages totalling more than 100 percent. It assumes that only a fraction of the subscriber connections are active simultaneously, and provides a statistical gain through variable user traffic flow.

Under-subscription is a set of percentages totalling less than 100 percent. It implicitly reserves bandwidth for connections with a specific transfer priority. If you know the amount of bandwidth required by a traffic type, set aside the necessary bandwidth capacity and protect it from being used by traffic with other transfer priorities.



Over- and under-subscribing




Full sharing of port capacity

You can also configure the port capacity for full sharing among bandwidth pools without pre-defined partitions between transfer priorities. You can assign 100 percent (or more) of the capacity to one common bandwidth pool (for example, pool 0) from which connections reserve their required CIR or EIR, or some of both, regardless of their transfer priority. See the figure [Full sharing of bandwidth \(page 85\)](#). In this configuration, the entire port capacity is available on a first-come, first-serve basis to any transfer priority.



Full sharing of bandwidth



Transfer priority and bandwidth pool assignment

You can map each transfer priority to a given bandwidth pool, for ingress and egress traffic, either globally or locally for each interface. If you have assigned a fixed amount of port capacity to the bandwidth pool, this implicitly reserves bandwidth for traffic assigned to that transfer priority. For information about transfer priorities, see [Transfer priority \(page 151\)](#).

To configure the relationship between transfer priorities and bandwidth pools globally, set the *assignedIngressBandwidthPool* and *assignedEgressBandwidthPool* attributes under the *Mod Frs DprsNet Tpm* component.

You can override the global parameter by assigning the transfer priority to a bandwidth pool for an individual FR UNI interface, under the *FrUni Ca Tpm* component.

You can override the global parameter by assigning the transfer priority to a different bandwidth pool for an individual interface. To specify bandwidth pools for ingress and egress traffic locally, set the *assignedIngressBandwidthPool* and *assignedEgressBandwidthPool* attributes under the *FrUni Ca Tpm* or *FrNni Ca Tpm* component.



Connection admission control

Connection admission control (CAC) is a mechanism used to determine whether or not an interface can accommodate a connection request. CAC provides quality of service (QoS) guarantees to traffic with various transfer priorities (TPs). CAC operates on switched connections only.


Connection admission or rejection is based on the attributes of both the requested connection and the existing connections. Ingress and egress CAC use several factors to determine resource availability, such as the link rate, the frame relay traffic descriptors (CIR, Bc, Be, and Tc), and the current available link capacity. See the figure [FR UNI call admission control \(CAC\) algorithm \(page 87\)](#) or [FR NNI call admission control \(CAC\) algorithm \(page 88\)](#).

For more information about ingress and egress connection admission control, see the following sections:

- [Assigned bandwidth pools \(page 88\)](#)
- [Available link capacity \(page 89\)](#)
- [Bandwidth requirements for new connections \(page 89\)](#)
- [Ingress CAC behavior \(outgoing calls\) \(page 89\)](#)
- [Egress CAC behavior \(incoming calls\) \(page 90\)](#)




FR UNI call admission control (CAC) algorithm





FR NNI call admission control (CAC) algorithm



Assigned bandwidth pools

You can assign a bandwidth pool for ingress traffic and egress traffic. If CAC is enabled for the ingress and egress directions, it checks the requested bandwidth for a connection against the bandwidth available in the assigned pool. If there is sufficient bandwidth available, CAC reserves it for the connection. The amount of bandwidth that CAC reserves for a connection is the CIR and EIR.

You can display information about available bandwidth under the *FrUni Ca* or *FrNni Ca* component. The following attributes indicate the total admitted bandwidth from each pool in the ingress and egress directions:

- *ingressCirPoolAdmittedBandwidth*
- *ingressEirPoolAdmittedBandwidth*



- *ingressEirPoolAdmittedBandwidth*
- *egressEirPoolAdmittedBandwidth*

It is the sum of the CIR and EIR values for all active connections, both permanent and switched. The attributes listed above indicate the bandwidth still available for new connections.

Available link capacity

The assigned bandwidth pool is defined as a percentage of the port capacity, or link rate. The operational *linkRate* attribute under the *FrUni Ca* or *FrNni Ca* component reflects the actual link rate of the interface. You can specify a different link rate for ingress and egress CAC by configuring the *overrideLinkRate* attribute. CAC uses this value to calculate the bandwidth assigned to each pool, and from there the bandwidth available in each pool for new connections.

If you are using the *FrUni Framer* or *FrNni Framer* component, and the *overrideLinkRate* attribute is 0, the value of the operational *linkRate* attribute reflects the speed of the port of the connection. If you are using the *FrUni VirtualFramer* or *FrNni VirtualFramer* component, you must configure the *overrideLinkRate* attribute with a link rate.

Bandwidth requirements for new connections

The ingress and egress CAC algorithms use several parameters including CIR, Bc, Be, and EIR to determine how much bandwidth to reserve for a connection. You can configure the interface to apply CAC policies to switched connections only.

In addition, you can limit the number of admitted calls requesting EIR-only bandwidth through the *maximumIngressEirOnlyCalls* and *maximumEgressEirOnlyCalls* attributes.

For switched calls, ingress and egress CAC use the traffic parameters signaled in the call request to calculate the required bandwidth for a connection. For SVCs that request bandwidth of CIR>0 (that is, not EIR-only bandwidth), these parameters can be negotiated to a minimum level of acceptable service in either direction. For more information, see [SVC call negotiation \(page 91\)](#).


Ingress CAC behavior (outgoing calls)

Ingress call admission control occurs toward the network, on calls received from an adjacent FR UNI or FR NNI. See the figure [Ingress CAC on outgoing calls \(page 90\)](#).

When ingress CAC is enabled, the interface applies CAC policies to outgoing SVC connections.



Ingress CAC on outgoing calls



Outgoing SVCs

For switched calls, CAC calculates the requested bandwidth from the traffic parameters signaled in the setup message.

If the setup message does not specify a transfer priority, CAC uses the transfer priority configured at the interface to determine the assigned ingress bandwidth pool. CAC then determines if the assigned pool has sufficient bandwidth available to accommodate the connection's requested service parameters.

If there is sufficient bandwidth, the call proceeds; if there is insufficient bandwidth, CAC determines if the assigned pool has sufficient bandwidth to accommodate the minimum requested service. If the minimum requested bandwidth is available, the call proceeds; if there is insufficient bandwidth available, the call clears.


Egress CAC behavior (incoming calls)

Egress call admission control occurs toward the link, on calls received from within the subnet. See the figure [Egress CAC on incoming calls \(page 91\)](#).

When egress CAC is enabled, the interface applies CAC policies to incoming SVC connections.



Egress CAC on incoming calls



Incoming SVCs

For switched calls, CAC calculates the requested bandwidth from the traffic parameters signaled in the call request.

The forward TP determines the assigned egress bandwidth pool for the call. CAC then determines if the assigned pool has sufficient bandwidth available to accommodate the connection's requested service parameters. If there is sufficient bandwidth, the call proceeds; if there is insufficient bandwidth, the call clears.

SVC call negotiation

Nortel Multiservice Switch nodes' frame relay SVC service enforces rate control on traffic that proceeds from the user to the network. Therefore, specific parameters at the network interface require traffic parameter negotiation. The network and the called user can reduce them if necessary, or reject the call if the minimum requirement is not met.

The FR NNI Link Layer Core Parameter information element (IE) in the setup message carries traffic parameters for the requested call. The network and the called user can reduce them if necessary, or reject the call if the minimum requirement is not met. If the network accepts the call, final traffic parameter values transmit to the calling user in the connect message. See the figure [Call negotiation points for a FR NNI SVC connection \(page 92\)](#).


FR NNI SPVC connections do not support call negotiation. All traffic parameters configured at the P-SPVC DLCIs remain unchanged during call establishment procedures. If the configured traffic parameters cannot be accommodated at any interface along the path of the end-to-end SPVC connection, the call is cleared.



For more information, see the following sections:

- [Negotiated traffic parameters \(page 92\)](#)
- [FR UNI CIR negotiation \(page 93\)](#)
- [FR UNI EIR negotiation \(page 94\)](#)
- [FR UNI maximum frame size negotiation \(page 94\)](#)
- [FR UNI default traffic parameters for SVCs \(page 94\)](#)
- [FR NNI ingress and egress call negotiation \(page 95\)](#)

Call negotiation points for a FR NNI SVC connection



Negotiated traffic parameters

The Link Layer Core Parameter IE in the setup message carries the following traffic parameters, as requested by the SVC:

- maximum outgoing and incoming frame information field size
- outgoing and incoming CIR
- minimum acceptable outgoing and incoming CIR
- outgoing and incoming committed burst size (Bc)
- outgoing and incoming excess burst size (Be)

Other than the transfer priority and the minimum acceptable CIR, the network may reduce the traffic parameters at the negotiation points (FR UNI and FR NNI) if necessary. If the parameters are negotiated, the interface updates the setup request accordingly before sending to the next interface. If the call is accepted by the end user, the final values of the traffic parameters are transmitted back to the source. If the minimum acceptable incoming or outgoing CIR cannot be met, the call clears.



The Frame Transfer Priority IE contains the outgoing and incoming transfer priority for the call:

The calling user does not include the Link Layer Core Parameter IE in its setup message; the FR UNI includes the LLCPE IE in its call setup message, with the default traffic parameters configured at its interface.

You can define default traffic parameters at the interface level. The network uses these values if the calling user does not specify them in the setup message:

- default maximum frame information field size
- default committed information rate (CIR)
- default committed burst size (Bc)
- default excess burst size (Be)

See the following sections for more information:

- [FR UNI CIR negotiation \(page 93\)](#)
- [FR UNI EIR negotiation \(page 94\)](#)
- [FR UNI maximum frame size negotiation \(page 94\)](#)
- [FR NNI ingress and egress call negotiation \(page 95\)](#)

FR UNI CIR negotiation

When the Nortel Multiservice Switch frame relay network receives an incoming or outgoing call, it determines if it can accommodate the requested bandwidth specified in the setup message. The network establishes the EIR as the smaller value of the user-requested EIR or the line speed.

CIR negotiation occurs in either an ingress or egress direction if you have set the *ingressApplyToCos* or *egressApplyToCos* attribute to a value other than none.

The FR UNI checks the bandwidth allocated to existing SVCs against the maximum ingress or egress CIR permitted for SVCs, as configured by the network administrator. If there is sufficient bandwidth available, the call proceeds with the requested CIR value. If there is insufficient bandwidth, the FR UNI determines if it can accommodate the minimum requested service. If the minimum requested CIR is available, the call proceeds; if the network cannot support the minimum acceptable CIR, the call clears.



If the DTE does not specify a value in the setup message, the minimum acceptable CIR sets to the default of zero. If the user-requested CIR is greater than the line speed, the Multiservice Switch frame relay network reduces the CIR to line speed.

FR UNI EIR negotiation

You can enforce EIR negotiation for incoming and outgoing call requests by setting the *ingressApplyToCos* or *egressApplyToCos* attributes under the *FrUni Ca* component.

The network derives the requested EIR according to the equation $EIR = Be / Tc$. If you set either the *ingressApplyToCos* or *egressApplyToCos* attributes to none, the FR UNI checks the EIR bandwidth allocated to existing SVCs against the maximum ingress or egress EIR permitted for SVCs (as configured by the network administrator). The FR UNI can adjust the requested excess burst size (Be) to ensure that the sum of the requested EIR and the current ingress or egress EIR does not exceed the maximum allowed.

If you set either the *ingressApplyToCos* or *egressApplyToCos* attribute to yes, the network establishes the EIR as the smaller value of the user-requested EIR or the line speed.

FR UNI maximum frame size negotiation

You can configure the maximum supported frame size of the information field. The network enforces the frame size only in the user to network direction.

The network reduces the user-requested frame size if it exceeds the maximum frame size permitted for SVCs on the interface. For traffic that uses normal and high priority Nortel Multiservice Switch trunks, the network reduces the frame size to the maximum frame size permitted. For traffic that uses interrupting Multiservice Switch trunks, the network reduces the frame size to the minimum permitted frame size.

If the transfer priority is mapped to a multimedia RCOS, the frame size cannot exceed 128 bytes.

FR UNI default traffic parameters for SVCs

You can define default traffic parameters at the interface level. Defined parameters specify traffic characteristics in the user to network direction. The network uses these values if the calling user does not specify them in the setup message:

- default maximum frame information field size
- default committed information rate (CIR)



- default committed burst size (Bc)
- default excess burst size (Be)

FR NNI ingress and egress call negotiation

Ingress and egress call negotiation for FR NNI SVCs is based on the traffic's transfer priority (forward and backward directions) and the assigned bandwidth pool. Ingress call negotiation always occurs towards the network; egress call negotiation always occurs towards the link.

You can define transfer priority and bandwidth pools for ingress and egress traffic at the interface level. Call negotiation occurs only if ingress and egress CAC are enabled on the FR NNI. Otherwise, no negotiation of call parameters occurs, and the call is accepted with the parameters signaled in the LLCPE.

When the FR NNI receives the call request, it assigns bandwidth to the connection from the assigned bandwidth pool, as configured at the interface. If sufficient bandwidth for that transfer priority is not available, the call clears.

You can specify the bandwidth pool allocated to each transfer priority for both the ingress and egress directions. For more information, see [Bandwidth pool management \(page 81\)](#).

FR UNI CAC migration

Frame relay call administration control (CAC) migration occurs when upgrading Nortel Multiservice Switch node releases. Multiservice Switch frame relay software for each release provides a migration code that makes migration automatic. There are substantial differences between the functionality of CAC for different Multiservice Switch releases. Old CAC is any Multiservice Switch frame relay software released prior to Multiservice Switch release 1.2 while new CAC is Multiservice Switch release 1.2 or any subsequent Multiservice Switch frame relay software. For a comparison of the two CAC types, see the table [Comparison of old and new CAC \(page 95\)](#).

Comparison of old and new CAC

Old CAC	New CAC
CAC for SVCs only.	CAC for SVCs only.
CAC applies in ingress direction only.	CAC applies in ingress and egress direction.
Possesses three bandwidth pools for transfer mapping to interrupting, normal, and high trunk queues with values ranging from 0 to 1000% of link speed.	Possesses sixteen bandwidth pools ranging from 0 to 1000% of link speed. The bandwidth pool assigned to the connection is determined by the provisioned transfer priority.
(1 of 2)	



Comparison of old and new CAC (continued)

Old CAC	New CAC
The CIR CAC default disable setting can not be changed. EIR CAC can be enabled or disabled.	Ingress and egress CAC for CIR and EIR can be enabled or disabled independently of each other.
Bandwidth reservation for both CIR and EIR traffic.	Bandwidth reservation for both CIR and EIR traffic.
For connections with CIR and EIR traffic, there is bandwidth reservation for only the CIR portion of the traffic.	For connections with CIR and EIR traffic, provision whether CAC is applied to only the CIR portion of the traffic or to both CIR and EIR traffic.
(2 of 2)	

Assigned bandwidth pool mapping

Three possible bandwidth pool mapping assignment scenarios exist that determine the status of your transfer priority following migration to Nortel Multiservice Switch release 1.2.

- If you are migrating from Multiservice Switch 7400 Release 5.0 or earlier, the module-wide transfer priorities are defined for the first time.
- If you are migrating from Multiservice Switch 7400 Release 5.1 or a later Multiservice Switch release that does not support the FR NNI SVC feature, the module-wide transfer priorities exist but the *assignedIngressBandwidthPool* and *assignedEgressBandwidthPool* attributes are defined for the first time.
- If you are migrating from Multiservice Switch 7400 Release 5.1 or a later Multiservice Switch release that supports the FR NNI SVC feature, the module-wide transfer priorities exist and are already provisioned. In order to preserve the behavior of both FR NNI and FR UNI SVCs, the sixteen transfer priorities under the *frUni Ca* component are added and override module-wide provisioning.

Regardless of which Multiservice Switch release you migrate from, the *assignedIngressBandwidthPool* and *assignedEgressBandwidthPool* attributes are mapped to one of the following transfer priorities:

- 0 for transfer priorities 0 to 5
- 1 for transfer priorities 6 to 10
- 2 for transfer priorities 11 to 15

CAC bandwidth pool mapping

CAC bandwidth pool mapping determines how the 16 ingress and egress committed information rate (CIR) and excess information rate (EIR) bandwidth pools are assigned following Nortel Multiservice Switch release migration.



This mapping only effects the behavior of ingress CAC for SVC connections. Ingress CAC for PVC and SPVC connections and egress CAC for all connections types is always off by default.

Old CIR CAC is enabled by default while old EIR CAC is disabled when the *unlimitedAggregateEir* attribute has a value of yes. Old CIR or EIR CAC means any CIR or EIR CAC existing on a Multiservice Switch node pre-dating Multiservice Switch release 1.2. For more information on the *unlimitedAggregateEir* attribute, see NN10600-060 *Nortel Multiservice Switch 7400/15000/20000 Component Reference*.

When the old CIR CAC is disabled, the resulting migrations occur:

- *ingressCirBandwidthPool 0* is assigned the value of *maximumAggregateSvcCirNormalQ*
- *ingressCirBandwidthPool 1* is assigned the value of *maximumAggregateSvcCirHighQ*
- *ingressCirBandwidthPool 2* is assigned the value of *maximumAggregateSvcCirInterruptQ*

When the old EIR CAC is enabled, the resulting migrations occur:

- *ingressCirBandwidthPool 0* is assigned the value of *maximumAggregateSvcCirNormalQ*
- *ingressCirBandwidthPool 1* is assigned the value of *maximumAggregateSvcCirHighQ*
- *ingressCirBandwidthPool 2* is assigned the value of *maximumAggregateSvcCirInterruptQ*
- *ingressEirBandwidthPool 0* is assigned the value of *maximumAggregateSvcEirNormalQ*
- *ingressEirBandwidthPool 1* is assigned the value of *maximumAggregateSvcEirHighQ*
- *ingressEirBandwidthPool 2* is assigned the value of *maximumAggregateSvcEirInterruptQ*

The value of the *applyIngressCacToCos* attribute is updated to CIR and EIR CAC.

The 13 other ingress CIR and EIR bandwidth pools are set to the default of 0%.



BNX interworking overview

This section provides an overview of the interworking functionality between a network element running hyperstream software (BNX) and a Nortel Multiservice Switch 15000 or Multiservice Switch 20000 node, which supports frame relay over IP.

Interworking between a BNX and a Multiservice Switch 15000 or Multiservice Switch 20000 node supporting frame relay over IP offers the following benefits:

- provide interworking of frame relay applications between a Multiservice Switch node and a BNX
- facilitates replacement of BNXs with Multiservice Switch nodes for increased performance
- provide frame relay service access over IP

Navigation

- [BNX-lwf overview \(page 98\)](#)
- [BNX-lwf DLCI overview \(page 100\)](#)
- [BNX-lwf features \(page 101\)](#)
- [Virtual router \(page 103\)](#)
- [Frame relay IP server \(page 103\)](#)
- [Cross-net signaling and LMI \(page 104\)](#)
- [PVC redirect \(page 111\)](#)

BNX-lwf overview

BNX-lwf allows the network to set up permanent virtual circuits (PVC) between Nortel Multiservice Switch 15000 or Multiservice Switch 20000 nodes and BNX frame relay interfaces across an IP backbone. At this time, BNX-lwf is supported on channelized 4-port DS3 function processors installed on Multiservice Switch 15000 or Multiservice Switch 20000 nodes.




BNX-lwf is an enhancement of the standards-based FrUni/FrNni for Multiservice Switch 15000 and Multiservice Switch 20000 nodes. BNX-lwf is added as a type of DLCI known as a BNX-lwf DLCI.

The service provider can combine BNX-lwf DLCIs with existing DLCI types—PVCs, switched virtual circuits (SVC), and switched permanent virtual circuits (SPVC) across the dynamic packet routing system (DPRS). That is, a single FrUni/FrNni can have a combination of circuits across DPRS to another Multiservice Switch FrUni/FrNni and PVCs across IP to a BNX. The Multiservice Switch FrUni/FrNni accesses the IP subnet through a protocol port of a Multiservice Switch virtual router.

Multiservice Switch 15000 or Multiservice Switch 20000 node frame relay interface requires a valid IP address and a valid data network address (DNA), either X.121 or E.164. The node uses the DNA to setup PVCs across DPRS and uses the IP address to setup PVCs across IP. To support a combination of PVCs across DPRS and IP and SVCs across DPRS, all PVCs (across DPRS or IP) must be within the PVC range. This range is configured through the *FrUni Signaling highestPvcDlci* attribute, through which the range of DLCIs reserved for PVCs is defined as 16 to the value of *highestPvcDlci*. The figure [PVC and SVC DLCI range \(page 99\)](#) shows how various types of DLCI can be combined under the same FrUni/FrNni.

PVC and SVC DLCI range



Furthermore, the service provider can overlay portions of the DPRS and IP subnets. Both the DPRS and IP subnets can be implemented on an ATM infrastructure since Multiservice Switch 15000 or Multiservice Switch 20000 virtual routers can be connected to other virtual routers (either Nortel Multiservice Switch or third-party) using ATM virtual channel connections (VCC). DPRS trunks can also be connected together through ATM VCCs.




The figure [PVCs across DPRS and IP subnet \(page 100\)](#) shows a sample network configuration.

The BNX-lwf DLCI does not support recombination. The interworking function operates under the assumption that neither the BNX nor the IP network segment incoming frame relay frames into multiple datagrams. The BNX-lwf DLCIs drop segmented frames.

Lastly, BNX-lwf DLCIs can generate accounting records.

PVCs across DPRS and IP subnet



BNX-lwf DLCI overview

BNX-lwf DLCIs encapsulate the incoming frame relay frames in two headers:

- a proprietary header called the BNX header
- an IP version 4 header

The BNX header carries end-to-end quality of service (QOS) information such as traffic priority and discard level. The IP header routes the frames from the source DLCI to the destination DLCI. The remote end points of the PVC are uniquely identified in the network with the IP address of the frame relay port and the DLCI number.

BNX-lwf DLCIs support these key capabilities:

- BNX frame relay cross-net polling mechanism to report end-to-end status of PVCs.
- standards-based Nortel Multiservice Switch 15000 and Multiservice Switch 20000 FrUni/FrNni congestion management mechanisms like dual-leaky bucket, setting FECN, BECN and DE bits



BNX-lwf DLCI does not support recombination and expects that neither the BNX or the IP network segments the incoming frame relay frames into multiple datagrams. BNX-lwf DLCIs drop segmented frames.

BNX-lwf features

The functionality for interworking between the network element running hyperstream software (BNX) and a Nortel Multiservice Switch 15000 or Multiservice Switch 20000 node supporting frame relay over IP supports the following capabilities and characteristics:

- frame relay UNI/NNI PVCs
- physical interface: channelized DS3 4-port function processor
 - channelized 4-port DS3 3/1/0
 - $n \times \text{DS0}$, where n is between 1 and 1024 inclusive
 - 112 DS1 terminations
 - 1024 frame relay UNI services
 - maximum of 5000 DLCIs per function processor
 - maximum frame size of 8 Kbytes
 - 1:N sparing, where $N \leq 6$
- backup and redirection for each DLCI
- FrUni/FrNni port numbering to accommodate administration of large numbers of logical interfaces; the FrUni/FrNni instance number has 10 digits
- 2-byte DLCIs within the range 16 to 1007 inclusive
- multi-link frame relay
- traffic parameter ranges:
 - CIR 0 to 155 000 000 bits/s, with 800-bit/s resolution
 - Bc 0 to 155 000 000 bits, with 8-bit/s resolution
 - EIR 0 to 155 000 000 bits/s, with 800-bit/s resolution
 - Tc 0.001 to 25.5 sec, 0.001 sec resolution
 - Be 0 to 155 000 000 bits, 8-bit resolution
 - Dn 1 to 8189 sec, with 1-byte resolution.
- service classes:
 - BNX 4 traffic priority
 - BNX 4 discard level
- address plans:



- internet protocol (IP)
- E.164
- X.121
- LMI control plane (FRF.1):
 - user side, network side, bidirectional side
 - asynchronous status report
 - ITU-T Annex A
 - ANSI Annex D
 - Vendor Forum Specification 001-208966
 - a-bit ignore
 - auto-detection of LMI type
- support for value-added functionality and network management
 - alarm conditions
 - PVC alarms
- implementation and configuration on the Multiservice Switch 15000 or Multiservice Switch 20000 node only. No modification is required on BNX hardware or software to allow interworking between Multiservice Switch and BNX frame relay applications
- BNX-lwf DLCI and DPRS DLCI can be configured concurrently under the same FrUni/FrNni
- establishes PVCs only across the IP subnet; however, a mix of PVCs, SVCs and SPVCs across DPRS can be configured concurrently under the same FrUni/FrNni
- frame relay carrier-grade warm standby functionality
- BNX header format version 5
- BNX frame relay cross-net signaling version 1
- accounting
- minimal impact on customers with existing FrUni/FrNni over DPRS implementations. BNX-lwf software is packaged as an optional feature that can be added to frame relay UNI features. There is no data model change on existing FrUni/FrNni configuration.
- end-to-end FRF.12 DTE at the local Multiservice Switch node FrUni interface



Virtual router

Nortel Multiservice Switch 15000 or Multiservice Switch 20000 node's virtual routing system enables the node to function as a router. Routing IP traffic is accomplished through either static routing or the following routing protocols:

- routing information protocol (RIP)
- open shortest path first (OSPF)
- border gateway protocol (BGP)


Multiservice Switch 15000 or Multiservice Switch 20000 node's virtual router supports frame relay, ATM, PPP, and gigabit Ethernet.

Frame relay IP server

From the perspective of the virtual router, the frame relay IP server (FRIP) is implemented as a media application. The FRIP is bound to a protocol port on a virtual router, and is considered by the virtual router as one subnet. The virtual router has no knowledge of what application is connected to a FRIP and relays to the FRIP all incoming traffic destined for any address in the associated subnet. The FRIP looks up the destination address in its own look up table to determine the end destination. and relays the traffic to the appropriate FrUni/FrNni.

The figure [Example of the relationship between the frame relay IP server and the virtual router \(page 103\)](#) shows an example of address resolution through the FRIP.

Example of the relationship between the frame relay IP server and the virtual router





Routing protocol

OSPF is the only IP routing protocol that the BNX interworking function (BNX-lwf) supports. When establishing the protocol port associated with the FRIP, configure the protocol port for passive OSPF interface. Passive OSPF interfaces ensure the following:

- the node does not send OSPF messaging in the direction of the link
- the area border router (ABR) can aggregate the routes from the Nortel Multiservice Switch 15000 or Multiservice Switch 20000 node

Routing export policies advertise the subnet associated with the protocol port to the network.

Choosing IP addresses

Nortel Multiservice Switch 15000 or Multiservice Switch 20000 nodes' FrUni/FrNni accesses the IP network by binding itself to a FRIP. For this reason, The IP address of the FrUni/FrNni must be in the same subnet as its FRIP. During configuration, this rule is enforced by a semantic check.

IP has special meaning for the host portion of any address with all ones (in this example, 138.10.5.255) and all zeros (in this example 138.10.5.0). See the figure [Example of the relationship between the frame relay IP server and the virtual router \(page 103\)](#). The "all ones" host portion of an address is reserved for broadcasting the incoming datagram to all the hosts in the subnet. The "all zeros" host portion is reserved for anycast (the incoming packet would be sent to the first available host in the subnet). Therefore, these addresses cannot be assigned to protocol ports, FrUnis, or FrNnis.

Each FRIP requires a protocol port and therefore virtually owns the IP subnet configured on the Multiservice Switch 15000 or Multiservice Switch 20000 virtual router. The FRIP ensures that any two FrUni/FrNnis on one FRIP have different addresses.

Cross-net signaling and LMI


This section describes PVC status reporting. Information is organized into the following sections:

- [LMI signaling \(page 105\)](#)
- [Cross-net signaling \(page 105\)](#)
- [State change of a DLCI \(page 106\)](#)
- [PVC tear down \(page 108\)](#)
- [Cross-net message format \(page 109\)](#)



The figure [End-to-end PVC status reporting \(page 105\)](#) shows a PVC across an IP subnet between a Nortel Multiservice Switch 15000 or Multiservice Switch 20000 node and a BNX, and identifies the segments where LMI and cross-net signaling events occur.

End-to-end PVC status reporting



LMI signaling

A FrUni/FrNni with BNX-lwf DLCI supports LMI with the *side* attribute set to user, network, or both.

All the existing LMI options can be used when configuring a FrUni/FrNni with BNX-lwf DLCI with the exception of the Smart Management Interface Local Exchange (SMILE) capability. SMILE is turned off by setting the *pvcConfigParmsInFsr* attribute under the LMI component to the value no.

For more information on SMILE, see NN10600-901 *Nortel Multiservice Switch 7400/15000/20000 Frame Relay Configuration Management*.

Cross-net signaling

If the *crossNetEnabled* attribute is set on the frame relay interface where the DLCI is added, the system periodically sends cross-net messages to the remote DLCI. Cross-net messages tell the remote end that the sender of the cross-net is ready or not ready to accept traffic in its direction. The period between two cross-net messages is configurable through the *crossNetPollingInterval* attribute under the *FrUni BnxProtocol* and *FrNni BnxProtocol* components. If the *crossNetListenEnabled* attribute is set on the frame relay interface, the DLCI expects a cross-net message from the subnet. Again, the period between two cross-net messages is configurable through the *crossNetPollingInterval* attribute.



The *crossNetErrorTreshold* attribute represents the number of polling intervals, without a cross-net signal from the subnet, tolerated before declaring the PVC down.

The format of the cross-net message is described in [Cross-net message format \(page 109\)](#).


State change of a DLCI

The *ABitStatus* operational attributes group under the *FrUni DlcI* and *FrNni DlcI* components keeps track of DLCI status. This group represents:

- the administrative state of a DLCI (lock/unlock)
- the state of the PVC
- the status signaled through the LMI

The figure [Points of failure of a BNX-lwf DLCI \(page 106\)](#) shows the points of failure (labelled 1 through 7). These labels appear in the tables [Reasons for failure reported by the adjacent LMI \(page 107\)](#) and [Reasons for dropping ingress traffic \(from the local DTE\) \(page 107\)](#) to indicate the reason for failure. These reasons are provided through the *aBitReasonFromIf* and *aBitReasonToIf* operational attributes.

Points of failure of a BNX-lwf DLCI



The *aBitReasonFromIf* attribute represents the status reported by the adjacent LMI. This attribute applies only when the FrUni/FrNni LMI is configured as user or both (user and network) side procedure. The table [Reasons for failure reported by the adjacent LMI \(page 107\)](#) shows possible values for the *aBitReasonFromIf* attribute and the point of failure.



Reasons for failure reported by the adjacent LMI

Value for aBitReasonFromIf	Description	Point of failure
notApplicable	status of the DLCI is active	
remoteUserSignaled	remote user is down	1
localLmiError	LMI is down	3
localLinkDown	link to the FrUni or FrNni is down	2
missingFromLmiReport	LMI is in service but the DLCI was not reported by the adjacent LMI	1
The point of failure is identified in the figure Points of failure of a BNX-lwf DLCI (page 106) .		

The *aBitReasonTolf* attribute indicates the reason why ingress traffic from the local data terminal equipment (DTE) is dropped. The cross-net messages indicate only the status of the active or inactive PVC without providing reasons. The *remoteLmiError* and *remoteLinkDown* events are reported as *remoteUserSignaled*. The table [Reasons for dropping ingress traffic \(from the local DTE\) \(page 107\)](#) shows possible values for the *aBitReasonTolf* attribute and the point of failure.

Reasons for dropping ingress traffic (from the local DTE)

Value for aBitReasonTolf	Description	Point of failure
notApplicable	aBitStatusTolf is active	
remoteUserSignaled	last cross-net message received by this DLCI had the status field set to inactive	5,6,7
localLmiError	LMI is down	3
remoteLmiError	n/a for BNX-lwf DLCI	
localLinkDown	local link is down	2
remoteLinkDown	n/a for BNX-lwf DLCI	
pvcSpvcDown	crossNetListen attribute is enabled and cross-net polling messages are not received as expected	4
userNotAuthorized	n/a for BNX-lwf DLCI	
resourceNotAvailable	n/a for BNX-lwf DLCI	
dlciCollisionAtNni	n/a for BNX-lwf DLCI	
The point of failure is identified in the figure Points of failure of a BNX-lwf DLCI (page 106) .		



PVC tear down

When a BNX-lwf DLCI is deleted, it sends an inactive cross-net message to the remote end. The remote responds by stopping traffic to the deleted BNX-lwf DLCI.

If this cross-net message is lost, the remote end continues to send data for several polling cycles before declaring the PVC out of service. At this time, the remote end stops sending data to the deleted BNX-lwf DLCI. The number of polling cycles is determined by the *crossNetPollingInterval* attribute (indicating the duration of a cross-net polling cycle) and the *crossNetErrorThreshold* (indicating the number of failures before declaring the PVC out of service). These attributes are configurable under the *FrUni BnxProtocol* and *FrNni BnxProtocol* components.

Locking a BNX-lwf DLCI or locking the FrUni/FrNni also causes an inactive cross-net message sent to the remote end. The near-end does not send cross-net messages until an unlock command is issued.

The table [Status of a DLCI advertised to the link and in the subnet \(page 109\)](#) summarizes status for these network elements.



Status of a DLCI advertised to the link and in the subnet

Input				Output	
Status of port, FrUni/FrNni, LMI	DLCI state	Status received from the last LMI status report (see [1])	Status received in last cross-net message	Status sent in LMI reports	Cross-net message sent to remote end
All are in service	enabled	up	active	active	active
			inactive	inactive	active
		down	active	active	inactive
			inactive	inactive	inactive
	disabled	up	active	inactive	inactive
			inactive		
		down	active		
			inactive		
one or more is out of service	enabled	up	active	does not report to link	
			inactive		
		down	active		
			inactive		
	disabled	up	active		
			inactive		
		down	active		
			inactive		
[1] Only when user side procedure is running.					

Cross-net message format

The cross-net message is encapsulated inside an IP packet. The characteristics for the IP packet header are described in the following points:

- version of the IP packet is always set to 4
- IP header length represents, in 32 bits words, the length of the IP header, not counting the data
- for the cross-net message, the IP header has a fix length of 5 words since no options are used
- type of service is set to 0
- total length is expressed in bytes and it is the total length of the datagram, including payload



- for the cross-net message, the total length is 36 bytes
- given the small size of the datagram, there is no requirement for fragmentation or recombination of cross-net messages, so the identification field is 0
- the D-bit means “don’t fragment” and is set to 1
- the M-bit means “more fragments” and is set to 0
- fragmentation offset is 0 since fragmentation is not used
- time-to-live field is a decreasing hop count; the field is initially set to 30 and decremented at each hop
- protocol field is decimal 202, which indicates to the receiving end that this datagram is a cross-net message and not user data (for which the protocol field would be set to 200)
- header checksum is calculated following RFC1071 and no options or padding are used
- source IP address indicates if A-bit is from the primary or backup


The cross-net message has the following characteristics:

- version field is set to 1
- message type for cross-net messages is 1
- number records field is 1
- DLCI data link connection identifier represents the destination for which the cross-net message intended
- status field represents the state of the sending end and is either 0 (inactive) or 1 (active)

The figure [Cross-net datagram format \(page 111\)](#) shows the layout of the cross-net message encapsulated inside an IP packet.



Cross-net datagram format



PVC redirect


Nortel Multiservice Switch BNX-lwf PVC can be configured with a backup destination. This backup requires a valid backup called IP address and a valid backup called DLCI. You configure the backup called IP address and backup called DLCI when you configure the PVC or at any time thereafter. This capability is known as PVC redirect, and has the following options:

- redirect to primary: this option allows an operator to redirect a PVC source from its backup destination to its primary destination using a command line
- redirect to backup: this option allows an operator to redirect a PVC source from its primary destination to its backup destination using a command line
- switch-on-demand: this option allows automatic redirection a PVC source to either its backup destination or its primary destination based on the cross-net status received from the primary destination. Operator intervention is an option.
 - When the current destination is primary, the PVC source is automatically redirected to the backup destination when an inactive cross-net status message is received from the primary or a cross-net status message has not been received for a configurable number of polling cycles.
 - When the current destination is the backup, the PVC is automatically redirected back to the primary destination when an active cross-net status message is received from the primary.



The figure [PVC redirect \(page 112\)](#) illustrates a PVC redirect configuration.

PVC redirect



The remainder of this section provides information on PVC redirect characteristics.

The BNX supports two other options: switch-on-demand to primary and switch-on-demand to backup. Nortel Multiservice Switch 15000 and Multiservice Switch 20000 nodes do not support these options at this time.

Manual and switch-on-demand redirect mode

The options for redirect to primary and redirect to backup are known as manual redirect modes. In the manual modes, only an operator command can initiate PVC redirection.

In switch-on-demand mode, both operator commands and the changes in the cross-net status message from the primary can cause a redirection.

Cross-net signaling for PVC redirect

When PVC redirection is initiated, the PVC source in the Nortel Multiservice Switch 15000 or Multiservice Switch 20000 node stops sending cross-net status messages to its current destination and starts sending messages to the new destination. In either manual mode or switch-on-demand mode, when the *crossNetListenEnabled* attribute is set to yes and a backup is configured, the PVC source listens to the cross-net status messages from both the primary and the backup.



Switch-on-demand prerequisites and enhancement

The *crossNetListenEnabled* attribute under the *FrUni BnxProtocol* and *FrNni BnxProtocol* components must be set to yes to retrieve the change in cross-net status of the primary, which will cause a redirection.

In some situations, the PVC can be redirected from the backup to the unstable primary destination. To avoid this unstable condition, which results in frequent redirections, BNX-lwf can be configured for a minimum number of consecutive active time intervals. In each time interval, an active cross-net message is received and the node waits for the specified number of intervals before redirecting the PVC back to the primary.

This feature is available only on Nortel Multiservice Switch 15000 and Multiservice Switch 20000 nodes and is not available on the BNX. By default, this feature is turned off.

Redirection DTE notification

Nortel Multiservice Switch 15000 or Multiservice Switch 20000 BNX-lwf PVC uses LMI asynchronous update to notify its DTE when redirection occurs. This use of LMI status reports is similar to their use on the BNX. Multiservice Switch nodes support three options for LMI status reports:

- transparent: sends no notification to the DTE
- intrusive N: sends asynchronous update to the DTE with the N-bit set
- intrusive A: sends asynchronous update to the DTE with A-bit not set

When the *procedure* attribute under the *FrUni Lmi procedure* or *FrNni Lmi procedure* components is set to none, only the transparent option is available.



Frame relay ISDN switched access overview

Use the following sections to learn more about frame relay ISDN switched access service.

Navigation


- [Frame relay ISDN switched access overview \(page 114\)](#)
- [Frame relay ISDN switched access protocols \(page 117\)](#)
- [Frame relay ISDN switched access application examples \(page 118\)](#)
- [Switched access to frame relay \(page 119\)](#)
- [ISDN protocol stack \(page 121\)](#)
- [Detailed description \(page 123\)](#)
- [Software modules \(page 124\)](#)
- [Frame relay virtual circuit call handling \(page 133\)](#)
- [Interaction between ISDN and virtual circuit connections \(page 136\)](#)
- [End-to-end call establishment for VCs \(page 137\)](#)
- [Compatibility issues \(page 138\)](#)

Frame relay ISDN switched access overview

Integrated services digital network (ISDN) technology provides a digital switched access path through an ISDN network. See the figure [ISDN switched access to a frame relay network \(page 115\)](#). The frame relay ISDN switched access feature supports the on-demand establishment of B-channel circuit connections over primary rate interface (PRI) links. The B-channel supports a data service: the frame relay user network interface (FRUNI).



ISDN switched access to a frame relay network



Summary of capabilities

Frame relay ISDN switched access supports

- up to 4 PRI links on a four-port DS1C or E1C function processor (FP)
- up to 4 PRI links (configured on ports 0 to 3) on a 32-port MSA FP
- up to 92 ISDN frame relay UNI services on a four-port DS1C FP and 32-port DS1MSA FP in channelized mode
- up to 120 ISDN frame relay UNI services on a four-port E1C FP and 32-port E1MSA FP in channelized mode
- up to 1000 different calling party numbers on each FP
- up to 9 ISDN call setups and clears per second at 100% CPU
- up to 23 (DS1 and 32-port DS1MSA FPs) or 30 (E1 and 32-port E1MSA FPs) ISDN calls up at the same time on a PRI
- frame relay UNI service, including such features as LMI, rate control, egress accounting and others, by way of the B-channel (this service is not supported on the D-channel)
- switched virtual circuits (SVCs) by way of Q.922 and Q.933 signaling protocols running on the B-channel
- permanent virtual circuits (PVCs) by way of the B-channel
- clear channel 64 kbit/s access
- 64 kbit/s bearer capabilities



- one D-channel for each PRI to control access to the B-channels
- layer 2 and 3 protocols of Q.921 and Q.931 (ITU-T), respectively, for D-channels
- termination of Q.931/Q.921 protocols at the PRI (the frame relay ISDN switched access service does not support end-to-end signaling)
- ISDN incoming calls (the calls originate in the ISDN network)
- ISDN call validation on calling line party number (optional)
- en-bloc signaling (the frame relay ISDN switched access service does not support overlap signaling)
- provisioning of acceptable calling line party numbers for each PRI (optional)
- channel assignment by the Nortel Multiservice Switch node or ISDN switch
- Australian Technical Standard-014 (TS014) and European ISDN (EURO-ISDN) protocols on E1
- Japanese ISDN (INS-NET) and North American ISDN (NIS and Ni2) protocols on T1

Benefits of frame relay ISDN switched access

A frame relay service connected to ISDN technology by way of switched access is an economical way for remote offices with branch devices to connect to corporate internetworks.

Below a certain level of daily connection time, the switched access solution is less expensive than dedicated access. This economic benefit point varies with the amount of frame relay traffic and local tariffs.

Frame relay ISDN switched access provides the following benefits:

- end-to-end digital transmission
ISDN provides end-to-end digital transmission, unlike analog modems that convert analog-to-digital and digital-to-analog signals. Substantial throughput and reliability improvements result.
- outband signaling capabilities
ISDN uses common channel signaling (CCS), allowing value-added services, such as call security and end-to-end calling line identification.
- lower access cost
Compared to dedicated ISDN lines, switched access provides a quick, economical connection to ISDN lines that are required only occasionally.



- simple architecture multimedia service
Switched access uses separate B-channels to transmit, video, voice, and other data across an ISDN link with low discard rates, delay, and delay variance. Although B-ISDN provides more bandwidth management facilities, ISDN is fast and economical for multimedia services.
- bandwidth sharing
Any B-channel can serve many calling parties that have frequent, but short duration access to frame relay service. Substantial bandwidth savings result.
- bandwidth-on-demand
You can add B-channels as required. Switched access does not use bonding; each B-channel is a different data stream.
- provisioning simplification
Switched ISDN connections need less provisioning on the ISDN network than dedicated connections do.
- resiliency
Switched connections provide an alternate path if the primary path fails.

Frame relay ISDN switched access protocols

Frame relay ISDN switched access supports protocols for

- Australia
- Europe
- Japan
- North America

See the following sections for information about each variant's compliance to the standards.

- [Australian TS014 compliance statement \(page 288\)](#)
- [European ISDN compliance statement \(page 305\)](#)
- [Japanese ISDN compliance statement \(page 338\)](#)
- [North American NI-2 compliance statement \(page 353\)](#)
- North American NIS A211-1

North American NIS A211-1 is the ISDN PRI specification between the Nortel Networks ISDN DMS-100 switch and the user equipment.



Frame relay ISDN switched access application examples

This section provides examples of how you can use frame relay ISDN switched access. Typical applications include

- bandwidth on demand
- disaster recovery
- occasional access

Bandwidth on demand


As bandwidth requirements increase, switched access can establish ISDN calls to off-load the dedicated line. It is more efficient to establish switched access connections than to dedicate lines. The CPE can terminate the ISDN calls when the extra bandwidth is unnecessary.

Disaster recovery

Frame relay ISDN switched access can provide an emergency backup ISDN connection when a dedicated line fails. See the figure [Disaster recovery scenario \(page 118\)](#).

The configuration at the remote customer premises equipment (CPE) is consistent regardless of the B-channel to which the equipment attaches. So, multiple remote devices can contend for the backup ports. Normally, calls transmit on the dedicated lines and switch to the backup ports and channels only when a dedicated line fails.

Disaster recovery scenario





CPE requirements for resiliency

The remote CPE redirects the call from the dedicated line to the ISDN backup line. The remote CPE (Branch Router/FRAD or Terminal Adapter function) must be able to manage two interfaces into the network (dedicated line and ISDN line). During periods of dedicated line downtime (which interface signals or LMI failure indicate) switched access reroutes frame relay traffic to the established B-channel. When the dedicated line restores, the remote CPE drops the B-channel connection. The central CPE must be able to accommodate multiple DLCIs from the same remote site.

Occasional access

A branch office may not need full-time access to the corporate backbone. For example, a branch office update process that runs for a short time each day requires only a brief connection to the corporate backbone. Frame relay ISDN switched access provides the dynamic allocation of B-channels that makes occasional ISDN access possible.

Switched access to frame relay

Frame relay ISDN switched access uses the signaling channel (D-channel) to provide an interface to the frame relay network over a switched access path. This service uses a PRI link with the D-channel configured on one timeslot and the B-channels on the other timeslots. This allows the network administrator to offer B-channel circuit service dynamically, based on a primary rate interface (PRI) for calls that originate on a basic rate interface (BRI) circuit.

The figure, [ISDN switched access to a frame relay network \(page 115\)](#) shows an end-to-end frame relay connection that accepts ISDN calls. When an ISDN call sets up, it creates a complete ISDN VC between the router and the Nortel Multiservice Switch subnet, on the B-channel. During call setup, a VC connection establishes between the routers. When the called end accepts a call, the router on the left side of the diagram sends ISDN calls to the Multiservice Switch node. The VC completes the connection on the B-channel.

In the subnet, frame relay calls transmit on the B-channel. ISDN calls transmit on the D-channel.


The Multiservice Switch edge node establishes VCs to the other nodes in the subnet. To complete the call, the nodes in the subnet can establish VCs outside the subnet.

ISDN connection

The ISDN network side of the network initiates the switched connection. See the figure [End-to-end ISDN connection \(page 120\)](#).



End-to-end ISDN connection



The D-channel provides ISDN signaling. The Q.931 protocol on the D-channel controls the B-channel connections. A B-channel must establish before the frame relay interface can activate. Once the B-channel connects, the frame relay CPE (router) on the ISDN BRI and the frame relay user network interface (UNI) at the Nortel Multiservice Switch PRI can send and receive HDLC frames between them.

Virtual circuit (VC) connection

Frame relay UNI is the only feature that this service supports.


When you connect the B-channel, you activate the frame relay UNI service, enable the local management interface (LMI), and can establish virtual circuits (VCs). VCs connections are internal to the frame relay network. The Nortel Multiservice Switch frame relay network supports two types of virtual circuit: the permanent virtual circuit (PVC), and the switched virtual circuit (SVC).

Permanent virtual circuits (PVCs) establish when the B-channel is connected; this process requires no additional signaling.

Switched virtual circuits (SVCs) establish by way of Q.922 and Q.933 signaling inband on the B-channel. The remote CPE on the BRI sends the Q.933 SETUP message. The message transmits transparently through the ISDN network. Then, the B-channel frame relay interface interprets the message. Once the VCs connect, data frames can transmit over the connection.



Virtual circuit connection



ISDN protocol stack

The figure [ISDN protocol stack \(page 122\)](#) shows the ISDN protocol stack tailored for switched access to the frame relay service. For simplicity, the figure shows only two B-channels: a PVC and SVC.

In ISDN, the D-channel controls access to the B-channels. The D-channel's layer 2 and 3 protocols are ITU-Q.921 (LAPD) and ITU-Q.931 respectively. The layer 2 protocol transmits frames reliably while the layer 3 protocol connects and disconnects the B-channels. The Q.931 protocol has several variants, including the Australian Technical Standard-014 (TS014) protocol.



- 1 establish a D-channel connection. The D-channel corresponds to a physical timeslot in the PRI.
- 2 request a B-channel connection. The B-channel request includes two items: calling party number and the channel number of the B-channel that the calling party wants to use.

Nortel Multiservice Switch nodes allocate the B-channel and determines whether to validate the call. If the B-channel request succeeds, the network

- enables the physical timeslot for data transmission
- activates frame relay
- can establish VCs

The B-channel supports the standard frame relay UNI, and frame relay capabilities (for example, LMI, rate enforcement, rate adaptation, and egress accounting).



The kind of virtual connection you provision determines which protocol stack operates:

- PVC connections use the Q.922 (layer 2) protocol.
- SVC connections use the Q.922 (layer 2) and the Q.933 (layer 3) protocols.

When you enable the B-channel, you establish the PVC connection.

The SVC connection establishes differently than the PVC connection establishes. The Q.933 and Q.922 protocols run on the B-channel. To make an SVC connection, allocate the B-channel first. The calling party then sends the destination frame relay UNI's data network address (DNA) to the B-channel. If the destination frame relay UNI accepts the call, the SVC connection completes and data transfer follows.

A frame relay UNI can support both SVC and PVC connections at the same time.

Detailed description


Frame Relay ISDN switched access consists of the following modules:

- call control
- Q.921/Q.931 protocol
- access service (frame relay UNI)

All modules reside on the same DS1C, E1C, or 32-port MSA functional processor. The figure [Software and hardware architecture of frame relay ISDN switched access \(page 124\)](#) shows the software and hardware components of the frame relay ISDN switched access service.



Software and hardware architecture of frame relay ISDN switched access



ISDN call handling


The figure [Software modules for frame relay ISDN switched access \(page 125\)](#) shows the hardware and software components. The software components establish and release the ISDN call between the ISDN BRI and Nortel Multiservice Switch frame relay B-channel on the DS1C, E1C, or 32-port MSA card.

Software modules

The figure [Software modules for frame relay ISDN switched access \(page 125\)](#) shows the software modules that make up frame relay ISDN switched access. This section describes each of these modules.



Software modules for frame relay ISDN switched access



The data ISDN module

Each D-channel has one data ISDN module. This module interfaces with the Q.931 protocol module and controls all the associated B-channels on the same PRI as the D-channel.

The data ISDN module supports up to 30 simultaneous calls on E1C and 32-port E1MSA cards and up to 23 simultaneous calls on DS1C and 32-port DS1MSA cards. This ISDN module can establish only one 64 k/bit channel per call.

The Q.931 protocol module

Each Q.931 protocol module controls one ISDN PRI interface. This module implements the call control signaling protocol on the D-channel of the ISDN PRI interface.

The PRI link is divided (by time division multiplexing) into a number of timeslots. The Nortel Multiservice Switch node side and the ISDN side communicate over the timeslot that the D-channel establishes B-channel connections on other timeslots.

The Q.931 protocol module is the universal ISDN protocol engine (UIPE).

The universal ISDN protocol engine

The universal ISDN protocol engine provides a single software state machine for different ISDN protocols. The UIPE receives the ISDN message and translates it into a proprietary format, using a lookup table that contains all supported protocols.



ISDN message handling in the receive direction

The UIPE includes both layer 2 and layer 3 functions according to the OSI model. The layer 2 processes the layer 2 aspects of the message (such as N(R) and N(S)), and strips off the layer 2 envelope. It passes the message to layer 3.

Layer 3 parses the message. Initial header checks validate the message versus state and message versus call reference. If valid, the UIPE ensures the mandatory information elements (IE) are present and uncorrupted. Missing and corrupted mandatory IEs are fatal errors that trigger corrective action as defined in the Q.931 protocol. If no fatal errors occur, the message passes to the call control module.

ISDN message handling in the send direction

Although the Nortel Multiservice Switch node does not initiate ISDN calls, it does send other messages relating to the call establishment or calling sequence. UIPE builds the specified message and sends it to layer 2. Layer 2 encapsulates it into a layer 2 frame, and sends it to the framer.

The Call control module

The call control module validates the Q.931 SETUP according to the screening indicator and the calling party number. It also allocates and de-allocates B-channels.

The channel group

The channel group associates a group of B-channels with a group of users. Different groups of users can require different behaviors, so you can group channels to customize parameters.

Each call control needs at least one channel group.

To define a list of channels that are available without validation of the calling party number, provision a channel group with no calling party number. When the call control receives an unprovisioned calling party number (in a Q.931 SETUP message), the call uses this group of channels and parameters.


ISDN call request validation

When the SETUP message is valid and the basic parameters are acceptable, the call control module compares the Q.931 SETUP calling party number, screening indicator, and channel identification parameters with the provisioned values. See the figure [ISDN call request validation \(page 127\)](#).

If the calling party number is not provisioned, and the call control module can not locate an empty calling group, the call clears immediately.



ISDN call request validation



Calling Party Number provisioning

One B-channel can be associated with one or many calling party numbers. So, many users can share one channel.

Alternately, a calling party number can be associated with one or many B-channels. So, the same group of users can use different channels.



The call control module builds a local database of calling party numbers that apply to every channel. You provision the list of possible calling party numbers and their associate B-channels for each D-channel.


The network validates the calling party number against the list of provisioned CGPNs. If the CGPN is in the provisioned CGPN list, then the network allocates a B-channel provisioned for that CGPN. If validation fails or no B-channels are available, the incoming call request clears.

In the figure [A Calling Party Number provisioning example \(page 129\)](#), three groups of users (see below) produce requests. Consider the figure together with the table [ISDN call request \(page 129\)](#), which summarizes the request and results. The following three groups make call requests:

- Group #1:
 - calling party numbers: 1231230, 1231231 and 1231232
 - Fruni:904 (channel 4)
- Group #2:
 - calling party number: 2222220
 - Fruni:905 (channel 5)
- Group #3:
 - calling party numbers: 1111111, 2222222, 3333333, 4444444 and 5555555
 - Fruni:901 (channel 1), 902 (channel 2) and 903 (channel 3)



A Calling Party Number provisioning example



ISDN call request

Calling number of the request	Condition	Result	Channel number assigned	Comment
1231230	channel 4 free	call up	4	
1231232	channel 4 free	call up	4	
1111111	channel 1 free	call up	1	
1111111	channel 1 busy channel 2 free	call up	2	
2222222	channel 1 busy channel 2 busy channel 3 free	call up	3	
2222222	channel 1 busy channel 2 busy channel 3 busy	call refused		all channels are busy

(1 of 2)



ISDN call request (continued) (continued)

Calling number of the request	Condition	Result	Channel number assigned	Comment
1231230	channel 4 busy	call refused		the only channel acceptable is busy
1234567		call refused		non-acceptable number
(2 of 2)				

The above examples assume that

- the channel identification information element (IE) specifies each channel
- the screening indicator is acceptable
- B-channel allocation is provisioned for ascending order
- the channel selection procedure is absent (see [Channel selection \(page 131\)](#))

Screening Indicator validation

The screening indicator is one of the fields of the calling party number IE.

The network uses provisioned values to determine the acceptability of the screening indicator in the Q.931 SETUP message. See NN10600-060 *Nortel Multiservice Switch 7400/15000/20000 Component Reference* for information about the ChannelGroup (Cg) component.

A B-channel call request clears immediately when it has no provisioned screening mode in the channel group.

Screening indicator values

A provisioning component specifies the calling party number's acceptable screening values. The possible screening values are

- 1 user-provided: not screened
- 2 user-provided: verified and passed
- 3 user-provided: verified and failed
- 4 network provided

The TS014 protocol supports only (2) and (4).

You can provision any combination of these values. For example, if you provision (2) and (4) and the screening indicator in the incoming ISDN call request is either (2) or (4), then this validation is successful and the call request proceeds to calling party number validation. On the other hand, if the screening indicator is (1) or (3), the call clears immediately.



The access service module

The access service module is a standard frame relay UNI that you can provision with PVCs or SVCs. The VCs activate only when the B-channel is up. The frame relay service characteristics are exactly the same as the non-ISDN frame relay service on an E1C, DS1C or 32-port MSA function processor.

The ISDN component

You can add the ISDN component to the frame relay UNI. This component specifies the value of the idle timer (*T320* attribute).

Security

The calling party number and screening indicator provide remote access security for frame relay networks.

Origin of the ISDN call

This service supports only dial-in ISDN calls. The ISDN network always generates the B-channel request.

Channel selection

To connect and disconnect calls, Nortel Multiservice Switch node and ISDN network switches at both ends of the PRI must refer to timeslots in the same way (in their layer 3 messages). Channel numbers perform this role. Channel numbers are not necessarily the same as timeslot numbers. For example: channels 1..15 can identify timeslots 1..15, timeslot 16 can be the D-channel, and timeslots 17..31 can apply to channels 16..30.

Ensure that both the Multiservice Switch and ISDN sides of the interface observe the proper convention as defined in the ISDN standard.

In the SETUP message, the channel selection procedure can signal one of three options: absent, exclusive, or preferred. When the channel selection message signals

- **Absent:** The message can use any channel. In this case, the message uses the first free channel with a calling party number equal to the calling party number of the SETUP message. If no channel has a matching calling party number or all matching channels are busy, a RELEASE COMPLETE message transmits over the D-channel with a specified clear cause.
- **Preferred:** The message can use any channel, but uses the specified channel if it is available. In this case, if the specified channel is free and if it gets a matching calling party number, the message uses this channel. Otherwise, the call uses the first free channel with a calling party number equal to the that of the SETUP message. If no channel has a matching



calling party number or all matching channels are busy, a RELEASE COMPLETE message transmits over the D-channel.

- Exclusive: The call cannot use any channel but the specified one. In this case, if the specified channel is free and if it gets a matching calling party number, the call uses this channel. Otherwise, a RELEASE COMPLETE message transmits over the D-channel with a specified clear cause.

Release of an idle B-channel

If you lock the Fruni or the idle timer expires, an internal message transmits to the call control module to request the B-channel's release.

An appropriate message on the D-channel can also release the active B-channel.

When the B-channel returns to idle, VCs clear and the LMI protocol is disabled.



Frame relay virtual circuit call handling

Frame relay ISDN switched access supports both types of frame relay virtual circuits:

- permanent virtual circuit (PVC)
- switched virtual circuit (SVC)

ISDN requires that if the physical layer is down, all the active virtual circuits should also go down. The frame relay interface can now accommodate this requirement. Now, all active virtual circuits (that is, permanent and switched virtual circuits) disconnect if the physical layer fails. Similarly, when the B-channel becomes idle, the virtual circuits disconnect.

Permanent virtual circuit (PVC)

In the current PVC implementation, a permanent virtual circuit establishes once you provision the direct call component (DC) that corresponds to the PVCs two ends. The PVC cannot be up if the physical layer is down, so if you provision a PVC alone, you do not bring the call up.

To optimize Nortel Multiservice Switch network resources, Nortel Networks recommends that you provision the Multiservice Switch ISDN end of the PVC as master.

Switched virtual circuits (SVC)

Frame relay switched virtual circuits (SVC) are established over the B-channel using the Q.933 and Q.922 protocols.

If the ISDN side initiates the SVC, the B-channel's physical layer should already be running (to receive the setup frame). The ISDN calling party number adds to the call request packet and transmits to the remote VC for signaling and accounting purposes.

Data network address


Add this provisionable field to specify which DNA is the calling DNA in the Q.933 SETUP message. Some applications require that the BRI calling party number is the calling DNA. On the other hand, other applications require just the Nortel Multiservice Switch calling DNA. See the figure [Address handling in the SVC case \(ISDN option\) \(page 134\)](#). In the ISDN case, the B-channel's data network address (DNA) can be meaningless.

ISDN DNA

If you provision the ISDN DNA option at the FR UNI, and if the screening indicator is user provided verified and passed or network provided, and you set the presentation indicator to non-restricted, the ISDN calling party number and ISDN screening indicator are the calling party number IE in the Q.933 setup frame.



Address handling in the SVC case (ISDN option)



Presentation indicator

The ISDN network can use this option to permit or prevent the call control module using the ISDN calling party number.

If you set the presentation indicator to restricted, the call control module signals only the Nortel Multiservice Switch DNA. If you set the presentation indicator to non-restricted, the call control module can signal either the ISDN DNA or the Multiservice Switch DNA.



DNA and screening mode signaled for SVC

Address signaling: provisioning	Presentation Indicator IE in Q.931 SETUP	Screening mode IE in Q.931 SETUP	Accounting DNA	Calling DNA signaled in Q.931 SETUP (SVC)	Screening mode signaled in Q.931 SETUP (SVC)
normalBehavior	non-restricted	User provided, not screened (upns)	Multiservice Switch	Multiservice Switch	no change
normalBehavior	non-restricted	User provided, verified and passed (upvp)	Multiservice Switch	Multiservice Switch	no change
normalBehavior	non-restricted	User provided, verified and failed (upvf)	Multiservice Switch	Multiservice Switch	no change
normalBehavior	non-restricted	Network provided (np)	Multiservice Switch	Multiservice Switch	no change
normalBehavior	restricted	upns	Multiservice Switch	Multiservice Switch	no change
normalBehavior	restricted	upvp	Multiservice Switch	Multiservice Switch	no change
normalBehavior	restricted	upvf	Multiservice Switch	Multiservice Switch	no change
normalBehavior	restricted	np	Multiservice Switch	Multiservice Switch	no change
IsdnDna	non-restricted	upns	Multiservice Switch	Multiservice Switch	no change
IsdnDna	non-restricted	upvp	ISDN	ISDN	upvp
IsdnDna	non-restricted	upvf	Multiservice Switch	Multiservice Switch	no change
IsdnDna	non-restricted	np	ISDN	ISDN	np
IsdnDna	restricted	upns	Multiservice Switch	Multiservice Switch	no change
IsdnDna	restricted	upvp	Multiservice Switch	Multiservice Switch	no change
IsdnDna	restricted	upvf	Multiservice Switch	Multiservice Switch	no change
IsdnDna	restricted	np	Multiservice Switch	Multiservice Switch	no change



Accounting

In the case of ISDN, Nortel Multiservice Switch DNAs are not sufficient to identify the end-user, so the calling party number in the Q.931 SETUP message replaces the channel DNA in the accounting record (see [ISDN DNA \(page 133\)](#)). If both sides of the call are on an ISDN channel, both calling and called DNA change.

Interaction between ISDN and virtual circuit connections

No VCs can establish if the ISDN connection is not up; VCs clear if the ISDN call clears.

ISDN connection versus VCs

The network monitors the FRUNI to see if any VCs are in data transfer. If a FRUNI has no active VCs for longer than the provisioned time period, the network deems the B-channel idle and disconnects. This action saves bandwidth.

The optional *T320* attribute terminates idle ISDN calls. If you use this attribute, provision it for every B-channel and the timer will

- start when the B-channel comes up
- stop when the first VC comes up
- restart when the last VC terminates
- clear the ISDN call when the timer expires

The *T320* attribute is an optional way to provide at least one valid VC on an ISDN channel. When the ISDN call establishes, at least one VC should connect within a set time or the ISDN connection ends. When the last VC of the ISDN connection drops, a new one should establish before a set time, or the ISDN connection releases.

If you do not use the *T320* attribute, the ISDN call can stay up forever, even if the ISDN call does not have running VCs.

For frame relay, accounting occurs on a per VC basis, so an idle ISDN channel generates no accounting records.

Clearing of virtual connections

In addition to clearing in response to an explicit clear request message, PVCs and SVCs can clear when

- the B-channel connection releases in response to a request from a ISDN network on the D-channel
- the B-channel connection releases due to a sustained failure in the LAPD layer or the D-channel's physical layer



- the VC itself has communication problems because an intermediate Nortel Multiservice Switch trunk fails

End-to-end call establishment for VCs

The figure [DNA and screening mode signaled for SVC \(page 138\)](#) shows how an ISDN and SVC call establishes as follows:

- a request to establish an ISDN call on a B-channel arrives through the D-channel as a Q.931 SETUP message
- the network extracts the calling party number and the screening indicator from the SETUP message
- the network validates the ISDN call request
- after successful call validation and B-channel identification, the channel state changes to busy, the channel comes up, and the Q.931 CONNECT message transmits on the D-channel with the selected channel number

On the B-channel


- the physical layer is comes up
- the LMI protocol starts (if you have provisioned it)
- the Q.922 layer 2 establishes on the B-channel
- a Q.933 SETUP message comes in and establishes a SVC call on the B-channel
- the network extracts the called party number from the SETUP message to find out the remote interface
- the Q.933 SETUP message transmits on the remote interface with the appropriate calling party number
- when the remote CPE sends the Q.933 CONNECT, the Q.933 CONNECT transmits on the B-channel

Data can now transmit on the frame relay DLCI.

As with the SVC, a PVC call establishes after the channel comes up.



DNA and screening mode signaled for SVC



Compatibility issues

The frame relay ISDN switched access feature can work with earlier frame relay versions. Previous versions, though, do not support the change in the accounting fields.

The ISDN component is optional; you do not need it for non-ISDN frame relay service.

Initial frame relay switched access supports Nortel Multiservice Switch frame relay UNI only. You can establish VCs between frame relay UNI services and

- frame relay UNI (non-ISDN variant)



- Both ends support accounting changes.
- ISDN frame relay UNI
 - Both ends support accounting changes.
- frame relay UNI (not supporting ISDN)
 - The non-upgraded end does not support accounting modifications.
- ISDN frame relay NNI (non-upgraded)
 - The non-upgraded end does not support accounting modifications.
- ISDN frame relay NNI (upgraded)
 - Both ends support accounting changes.
- DPN-100 frame relay UNI
 - The DPN-100 end does not support accounting changes.

Frame relay ISDN switched access has some limitations in its interworking with DPN-100 frame relay service. Frame relay DPN-100 does not generate the new accounting fields. Nor does DPN-100 support ISDN, but PVCs can connect ISDN Multiservice Switch frame relay UNI and DPN-100 frame relay services (DPN-100 frame relay does not support SVCs). To avoid this limitation, use single end accounting and define the charged end on the Multiservice Switch node side.



Traffic management

The frame relay service uses various techniques to avoid and control congestion in the network. In addition, network engineering provides a coordinated provisioning strategy across UNIs and NNIs. For maximum throughput with minimum frame loss, you need to configure the frame relay service with the appropriate subscription options for traffic priority and class of service. Make sure that only networks that connect directly to users control rate enforcement across a multi-network.

Navigation

- [Congestion control \(page 140\)](#)
- [Traffic subscription options \(page 151\)](#)
- [Multimedia traffic \(page 161\)](#)
- [Class of service types \(page 163\)](#)
- [Traffic management for BNX interworking \(page 164\)](#)

Congestion control

Frame relay uses several mechanisms to control congestion in the network. It notifies other devices of congestion, and also modifies traffic flow to enforce service level agreements.

The discard eligibility (DE) bit in a frame relay frame helps to manage congestion. The DE bit determines whether the network should discard one frame in preference to others. The discard of DE frames during congestion ensures that frames within the CIR limit get transmission priority.

For more information, see the following sections:

- [Discard priority \(page 141\)](#)
- [Congestion notification \(page 142\)](#)
- [Congestion avoidance policy \(page 142\)](#)
- [Rate enforcement \(page 143\)](#)
- [Rate adaptation \(page 144\)](#)



Discard priority


Discard priority (DP) specifies traffic discard preference. A connection can have a DP of high or normal. Normal discard priority (NDP) frames discard before high discard priority (HDP) frames. Network discard priority (NDP or HDP) operates within the subnet and has no bearing on the CIR/EIR rate enforcement at the access.

If a network resource experiences persistent mild congestion, frame relay discards NDP frames marked discard eligible (DE bit=1). If moderate congestion becomes heavy, the network discards NDP frames with the DE bit=0 and HDP frames marked discard eligible (DE bit=1). The network discards HDP frames with the DE bit=0 only during severe network congestion.

For switched calls, the discard priority is configured at the interface level (under the *Dna* component), and applies to all switched connections on that interface. Since the DP can only be signaled within the subnet and not across networks, it is defined for the SVC segment in each network, not for the entire connection. See the figure [Discard priority across networks \(page 141\)](#).

For permanent calls, you can override the DP configured for the interface on an individual DLCI basis. You can configure a different DP through the *discardPriority* attribute under the *Dlci Dc* component. Use the *deDiscardPriority* attribute under the *Dlci Dc* component to specify whether the configured discard priority applies to DE bit set frames for the DLCI. If you set the *deDiscardPriority* attribute to lowest, the network discards DE=1 frames first when the network is congested, regardless of the connection's discard priority.

Discard priority across networks





Congestion notification

The network uses proprietary internal congestion bits, the forward congestion indication (FCI) bit and the backward congestion indication (BCI) bit, to indicate congestion in the subnet and at the access.

The network uses the forward explicit congestion notification (FECN) and backward explicit notification (BECN) bits to indicate that frames are encountering congestion. The FECN and BECN bits are set during mild congestion, when resource levels are low but sufficient to maintain operation. This mechanism notifies devices to reduce user demand on resources, so that the network can return to normal operation mode.

The frame relay service at the Nortel Multiservice Switch node evaluates frames from the subnet for FCI and BCI bits, then sets the correct FECN or BECN bit on the frame it sends to the user.

Congestion notification in the forward direction proceeds as follows:

- 1 The subnet sets the FCI bit when congestion prevents frame forwarding. The Multiservice Switch node recognizes the set FCI bit and sets the FECN bit on the frame that transmits to the user.
- 2 The FECN bit warns the end device when the incoming frames encounter congestion. Some application protocols use the destination end to control the amount of traffic the sender can transmit. In this case, the network sends a set FECN indicator to the destination device, and the destination device instructs the sender to slow or stop transmission. If congestion arises in either the subnet or the access, the FCI bit can set the FECN bit.

Congestion notification in the backward direction proceeds as follows:

- 1 The subnet sets the BCI bit to warn the sending device of congestion. The Multiservice Switch node recognizes the set BCI bit and sets the BECN bit on the frame that transmits to the user.
- 2 The BECN bit warns the end device that the frames it transmits will encounter congestion. Then, the sending end can adjust its sending rate. The BCI bit can set the BECN bit when congestion conditions arise in either the subnet or at the access.

Congestion avoidance policy

Depending on the severity of the congestion, the network chooses one of four congestion control methods. See the table [Congestion levels \(page 143\)](#).

For more information about normal and high discard priority (NDP and HDP), see [Discard priority \(page 141\)](#).



Only congestion triggers FECNs and BECNs. The network can discard frames without prior FECN and BECN notice if traffic exceeds the subscribed CIR/EIR limits and you have enabled rate enforcement. For information about rate enforcement, see [Rate enforcement \(page 143\)](#).

Congestion levels

Congestion level	Congestion control method
Mild	The network sends FECN or BECN, or both.
Moderate	The network discards all new normal discard priority (NDP) DE = 1 (discard eligible) frames arriving at the congestion point to protect new DE=0 frames.
Heavy	The network discards high discard priority (HDP) DE=1, NDP DE=0.
Severe	The network discards all new frames that arrive at the congestion point to protect the network.

Rate enforcement


Rate enforcement is an optional feature that prevents a user from using an inequitable share of network resources. With rate enforcement on, the network confines the data rate to a predefined throughput. This throughput is based on several parameters, including the committed information rate (CIR), committed burst size (Bc), excess information rate (EIR), and committed rate measurement interval (Tc).

Rate enforcement operates on each connection. Each DLCI on a frame relay port can have a CIR value, which reflects the sustainable throughput that the network supports under normal conditions. Throughput in excess of CIR is not allowed if the excess burst size (Be) is 0. In this way, users or applications on different PVCs on a single frame relay port have fair access to network resources, according to their subscribed CIR.

Appropriate CIR allocation can manage the allowable amount of sustainable throughput on a frame relay port. As a result, a single port cannot monopolize the switch's resources at the expense of other users. In the figure [Rate enforcement \(page 144\)](#), the line speed of port 1 is 256 kbit/s, but the sum of the CIRs of Users A, B and C is only 192 kbit/s. Rate enforcement prevents users from consuming more than their subscribed CIR, even though demand is less than the available resources.



Rate enforcement



You do not need to implement rate enforcement at the NNI if the UNI has applied it. If you configure rate enforcement at the NNI, an external network can alter the original traffic characteristics. In this case the CIR/Be/Bc mechanism reports a violation when no violation exists. When you configure the frame relay NNI service, turn rate enforcement off to disable the CIR (and EIR) mechanisms.

Rate adaptation

Rate adaptation is an optional feature that controls congestion during traffic overload. Because many devices do not respond to FECNs and BECNs, the network cannot always rely on the user to perform congestion control. To reduce congestion, rate adaptation limits the number of frames that enter the network at the access point. This process prevents the resource waste that occurs when frames enter the network, only to be discarded when congestion occurs. Rate adaptation operates on both permanent and switched connections.

Rate adaptation reduces the available information rate (AIR), which is the sum of the EIR and the CIR, at the onset of congestion. Frames are discarded at the access point if the ingress rate exceeds the AIR.

You must turn rate enforcement on when you enable rate adaptation.

Rate adaptation operates in a normal, waiting, or monitoring state. In each state, rate adaptation uses a timer and counts the number of backward frames (frames that pass through the DLCI travelling in the opposite direction to the



congested traffic). A state reaches its congestion threshold when the timer elapses, and the specified number of backward frames for that state has passed through the DLCI. Rate adaptation then moves into the next state.

For more information on rate adaptation, see the following sections:

- [Rate adaptation options \(page 145\)](#)
- [Rate adaptation sensitivity \(page 148\)](#)
- [Using rate adaptation across networks \(page 149\)](#)

Rate adaptation options

Three rate adaptation options accommodate the differences in service definitions among network providers:


- **on**: rate adaptation reduces EIR traffic, and CIR traffic if necessary
- **eirOnly**: rate adaptation reduces EIR traffic only
- **cirProportionate**: rate adaptation reduces EIR traffic calculated to a rate that is proportional to the configured CIR

For all options, set the *raSensitivity* attribute to define the time interval between rate adjustments.

If you set the *ra* attribute to **on**, the EIR is reduced to its minimum value of 2400 bits/s at the onset of congestion. If congestion persists, the CIR is reduced by increments of 25%, to a minimum of 50% of its configured value. As congestion subsides, the CIR is increased by increments of 6.25% until it reaches its configured value. Once the CIR is restored, the EIR increases by increments of 6.25% of its configured value. See the figure [Rate adaptation option - on \(page 146\)](#).




Rate adaptation option - on



If you set the *ra* attribute to *eirOnly*, rate adaptation does not affect CIR traffic. At the onset of congestion, EIR traffic is reduced by increments of 25% of its current value, to a minimum of 2400 bit/s. As congestion subsides, the EIR is increased by increments of 6.25% until it reaches its configured value. See the figure [Rate adaptation for the eirOnly option \(page 147\)](#).



Rate adaptation for the eirOnly option



If you set the *ra* attribute to *cirProportionate*, rate adaptation applies to EIR traffic only, but the available EIR is calculated from the configured CIR. The configured EIR does not change and rate adaptation occurs in proportion to the amount of service provided to each user. At the onset of congestion, the service reduces available EIR by increments of 25%, to a minimum of 2400 bit/s. When congestion subsides, the service increases the available EIR by 6.25% of the configured EIR, until it reaches the value of the configured EIR.

The value of the *raSensitivity* attribute has a corresponding *cirProportionate* multiple value. The configured CIR is multiplied by this value to calculate the available EIR. See the table [cirProportionate multiples \(page 148\)](#) for the *raSensitivity* attribute values and corresponding *cirProportionate* multiples.



cirProportionate multiples

raSensitivity value	cirProportionate multiple
1-6	4
7-14	2
15-20	1

Rate adaptation sensitivity

Rate adaptation sensitivity allows you to control and adjust the speed at which rate adaptation reduces the AIR for each connection. To change the reaction speed of rate adaptation for a DLCI, set the value of the *raSensitivity* attribute under the *Dlci Dc* component.

A higher number indicates a higher sensitivity, and therefore a quicker reaction time to network conditions. A higher setting allows the network to discard frames at the access and tends to favor protection of the network over the user's allowed throughput. Traffic on a DLCI set to a higher sensitivity has a lower priority and discards EIR before DLCIs set to a lower sensitivity.

When you change the value of the *raSensitivity* attribute, you change the amount of time that elapses before the AIR adjusts. The table [Rate adaptation sensitivity adjustment intervals \(page 148\)](#) shows the time intervals between AIR adjustments for each *raSensitivity* attribute value.

Rate adaptation sensitivity adjustment intervals

Rate adaptation sensitivity	Rate adaptation adjustment interval (milliseconds)
20	40
19	80
18	120
17	160
16	200
15	300
14	440
13	600
12	900
11	1 200
10	1 800
(1 of 2)	



Rate adaptation sensitivity adjustment intervals (continued)

Rate adaptation sensitivity	Rate adaptation adjustment interval (milliseconds)
9	2 500
8	3 500
7	5 000
6	7 000
5	10 000
4	14 000
3	20 000
2	28 000
1	40 000
(2 of 2)	

Using rate adaptation across networks

You can use rate adaptation to manage congestion across connected networks by setting the *updateBCI* attribute under the *Dlci* component. Congestion information between networks transmits by way of the BECN bit. When the *updateBCI* attribute is turned on, the FR NNI interface converts all incoming BECN bits to BCI bits. Nortel Multiservice Switch nodes use BCI bits to transmit congestion notification within the subnet.

If the *updateBCI* attribute is turned on and there is congestion in a destination network, frames can discard at the ingress, even though the originating network is not congested. This process prevents transmission of traffic through an entire source network. Otherwise, traffic is discarded in transit or at a destination network.

If the *updateBCI* attribute is turned off, the BCI bit for the DLCI is set to zero and BECN bits do not map to BCI bits.

Overflow routing

The standard congestion management mechanism discards frames when congestion occurs. Frame relay overflow routing is an optional feature that responds to congestion by rerouting frames instead of discarding them.

You can specify either high or normal reliability for services through the *outdefaultPathReliability* attribute under the *Dna* component. At the high-reliability level, the network tries to overflow traffic onto other routes when congestion occurs. At the normal-reliability level, the network does not reroute packets.




Overflow can occur only when the primary and alternate paths have the same path characteristics for the related data's class of service. For traffic that uses delay class of service, the paths must have the same total path delay. For throughput-based traffic, the paths must provide the same total bandwidth. Typically, high reliability causes packets to duplicate and lose their order in the network when errors or link congestion occur. The VC corrects the duplication and disordering to ensure that packets arrive intact at the DTE.

Multimedia applications do not support overflow. Frames that arrive from the subnet go directly to the application. Frames do not re-order, and the network discards out-of-order frames.

In the figure [Overflow routing for the frame relay service \(page 150\)](#), if traffic is going to path A-C-B and path A-C is congested, the high reliability traffic overflows onto path A-D-B; normal reliability traffic discards. Overflow can occur only when both paths, A-C-B and A-D-B, have the same path characteristics for the related data's class of service. The path characteristics include the same total path delay for traffic that uses the delay class of service, or the same total bandwidth for throughput-based traffic.

Overflow routing for the frame relay service





Traffic subscription options

You can configure different subscription options for traffic to match application needs. In addition, Nortel Multiservice Switch FR UNI service uses the weighted fair queue scheme (WFQ) to ensure fair service to all queues, but without depriving low priority queues of service.

For more information, see the following sections:

- [Transfer priority \(page 151\)](#)
- [FR UNI routing class of service \(page 152\)](#)
- [FR NNI routing class of service \(page 153\)](#)
- [Emission priority \(page 153\)](#)
- [Flexible QoS mapping \(page 154\)](#)
- [IP class of service \(CoS\) mapping \(page 155\)](#)
- [IP packet fragmentation \(page 156\)](#)
- [Large frame transmission \(page 156\)](#)
- [End-to-end FRF.12 frame fragmentation \(page 157\)](#)
- [FR UNI weighted fair queue scheme \(page 159\)](#)

Transfer priority

Transfer priority (TP) specifies transmission path and order. High transfer priority frames transmit within the network along the path with the least delay, and take precedence over normal transfer priority traffic.

Transfer priority values map to emission priority queues that define the traffic's urgency. The network tags urgent traffic with high emission priority and transmits it before lower emission priority traffic. A combination of three network priority queues and four egress emission queues service the transfer priorities.

You can choose between 16 transfer priority indices (0 to 15). Transfer priority on a connection can be different for both the ingress and egress directions. In addition, you can map each TP index to a routing class of service (RCOS) and emission priority (EP) that are appropriate for your service. For more information, see [Flexible QoS mapping \(page 154\)](#).

You can set the default transfer priority for the interface and all its connections through the *defaultTransferPriority* attribute under the *Dna* component. You can override this value on an individual DLCI basis by setting the *transferPriority* attribute under the *Dlci Dc* component.



You can configure a different transfer priority for each end of the PVC connection through the *transferPriorityOverride* attribute under the *Dna* component. By default, the transfer priority configured at the master end (as signalled in the call request) overrides the transfer priority configured at the slave end. If you set the *transferPriorityOverride* attribute to no, the slave end retains its configured transfer priority, whether at the interface or DLCI level.

For information about assigning a transfer priority to a particular bandwidth pool, see [Transfer priority and bandwidth pool assignment \(page 85\)](#).

FR UNI routing class of service

Routing class of service (RCOS) determines how the network treats a frame relay packet as it traverses the network to its destination. Different priorities of traffic can transmit on different network paths, depending on their RCOS value.

Nortel Multiservice Switch nodes support three RCOS values for frame relay traffic:

- delay
- throughput
- multimedia

For traffic with a transfer priority mapped to the delay RCOS (typically interactive applications), the network uses the delay metrics to route packets. The delay metric reflects the routing cost on the basis of delay. The least delay has the best metric.

For traffic with a transfer priority mapped to the throughput RCOS (typically interactive and file transfer applications), the network uses the throughput metrics to route packets. The throughput metric reflects the routing cost in terms of bandwidth. The highest bandwidth has the best metric.

Transfer priorities mapped to the multimedia RCOS transmit with the highest quality of service. Because multimedia traffic is delay-sensitive, the delay metric helps route it along the least delay path. Each node automatically determines the best delay and throughput path to every other node in the network. These routing tables are constantly updated to accommodate topology changes.

As traffic transmits along its network path, its transfer priority and RCOS values affect its transmission order onto each Multiservice Switch trunk and access link at the destination node. Higher priority traffic has greater access to transmission resources than lower priority traffic.



The network always services the higher emission priority queue before the lower priority queue. For a Multiservice Switch trunk, traffic going out on the interrupting priority queue can interrupt traffic on other lower priority queues. However, a frame at normal priority runs to completion before a frame at high priority gets service.

You can map each transfer priority (TP) index to RCOS and emission priority values that are appropriate for your frame relay service. The table [Default frame relay transfer priority definitions \(page 153\)](#) summarizes the default transfer priority definitions.

Default frame relay transfer priority definitions

Typical application	Transfer priority indices	Internal transfer priority in MPS	Routing class of service	If 2 egress emission priority queues are used	If 4 egress emission priority queues are used
multimedia (voice, FAX)	15, 14, 13, 12, 11	interrupting	multimedia	high	3
	10	high	delay	high	3
interactive2	9	high	delay	high	2
interactive1	8, 7, 6	high	delay	normal	1
	5, 4, 3	normal	throughput	normal	1
file transfer	2, 1, 0	normal	throughput	normal	0

FR NNI routing class of service

Routing classes of service determine the emission priority queue at the egress trunk over which traffic is routed. The routing class of service can be either multimedia, delay, or throughput. Traffic with multimedia RCOS has highest priority and is forwarded on the interrupting trunk queue. Traffic with delay RCOS is forwarded on the high trunk queue. Traffic with throughput RCOS has lowest priority and is forwarded on the normal trunk queue.

You can define a routing class of service and emission priority for each transfer priority index. For more information, see [Flexible QoS mapping \(page 154\)](#).

Emission priority

Emission priority maps different transfer priorities to egress emission priority queues. The emission priority queues help avoid bottlenecks at the access, by placing different priorities of traffic in separate queues at the egress.



Traffic connections on a high emission priority queue transmit with less delay than traffic connections on lower emission priority queues. [Transfer priority \(page 151\)](#)

You can configure two or four emission priority queues for each frame relay interface. Two default queues (normal and high) are available. You can use four egress emission priority queues (where zero signifies the lowest priority and three the highest) to further differentiate traffic at the network's egress.

If you use two emission priority queues (high and normal), control-plane traffic on DLCI 0 or DLCI 1023 transmits to the link on a high egress priority queue and with a high discard priority. This process applies to Q.922, Vendor Forum, Q.933 Annex A and ANSI Annex D LMI frames.

If you use four emission priority queues, control-plane traffic on DLCI 0 and DLCI 1023 transmits to the link on the second highest egress emission priority queue (queue 2) and with a high discard priority.

For information on the hardware involved in emission priority queues, see the Nortel Multiservice Switch Release Notes.

Flexible QoS mapping

For flexible networking over DPRS connections, you can map individual transfer priorities to DPRS routing class of service (RCOS) and emission priority (EP) values. This practice allows you to

- tailor a connection to a particular QoS
- introduce a transfer priority with unique RCOS and EP mappings for a new application

In a networking scenario, the TP-to-RCOS mapping occurs at the ingress port and the TP-to-EP mapping occurs at the egress port. The TP-to-RCOS and EP mappings must be consistent across the network so that the QoS provided by the network is the same for traffic flowing in both directions.

The table [Default QoS mapping \(page 154\)](#) shows the default mapping of transfer priorities to RCOS and EP values.

Default QoS mapping

Transfer priority	RCOS	Egress EP (4 queues)	Egress EP (2 queues)
11 to 15	multimedia	3	1
10	delay	3	1
(1 of 2)			




Default QoS mapping (continued)

Transfer priority	RCOS	Egress EP (4 queues)	Egress EP (2 queues)
9	delay	2	1
6 to 8	delay	1	0
3 to 5	throughput	1	0
0 to 2	throughput	0	0
(2 of 2)			

To change these mappings, you can configure your own TP definitions for the Nortel Multiservice Switch node. See the figure [QoS mapping \(page 155\)](#). Nortel Multiservice Switch node E receives multimedia traffic from Multiservice Switch nodes A, B, C, and D. Although all the traffic uses the multimedia RCOS and travels on the frame-cell trunk interrupting mode queue, the network operator has configured a different egress emission priority queue for each transfer priority.

QoS mapping



IP class of service (CoS) mapping

Frame relay services can provide variable QoS within a single PVC connection for IP-initiated calls. By interworking IP CoS and frame relay QoS, you can ensure consistent, end-to-end IP class of service between user devices.

You can configure the interface to assign different discard and emission priorities to individual frames at the egress, based on the type of service (TOS) field in the IP header. The TOS octet is an 8-bit field that specifies IP class of service parameters based on delay, throughput and reliability.



You can map any TOS value (0 to 255) to a default discard priority (DP) and egress emission priority (EP) at the module level. You enable or disable the IP-aware egress policy on individual connections. With this option turned on, the interface prioritizes egress traffic based on the TOS field (if present), instead of the discard and emission priority values configured for the connection.

The IP CoS mapping feature applies to egress traffic management on PVC connections only. In addition, the TOS mappings you configure override any configured flexible QoS mappings at the egress.

For detailed information about IP classes of service, see NN10600-801 *Nortel Multiservice Switch 7400/15000/20000 IP Configuration Management*.

IP packet fragmentation

If an IP packet is fragmented at layer 3, each fragment carries an IP header with the same TOS value. If the TOS mapping feature is enabled on the PVC, the interface transmits the packets with the corresponding discard and emission priorities (as defined for the module).

If an IP packet is FRF.12-fragmented, the TOS value is not present in each fragment. Packets fragmented at layer 2 transmit with the configured discard priority (DP) and emission priority (EP) for the connection.

Large frame transmission

A user can transmit frames with payloads as large as 8187 bytes. To allow large frame transmission on a connection, configure the *maximumFrameSize* attribute under the *Dlci Sp* component. When you configure this attribute, you reduce the number of frame discards that occur when frames from the link exceed the configured maximum.

Large frame transmission causes long serialization delays. These delays affect the quality of service for all connections that share the same egress link or channel. CRC-16 effectively detects errors in frames of 4096 bytes or less, but is less effective for frames greater than 4096 bytes. Therefore, larger frames are more prone to undetected corruption.

Follow these guidelines for large frame transmission:

- Set the maximum subnet packet size to 512 bytes or higher to reduce message block exhaustion or frame discard. For high speed Nortel Multiservice Switch trunks, use 4096 as the value for the *Mod Vcs maxSubnetPacketSize* attribute.
- Set the transfer priority to normal for connections transferring data greater than 4096 bytes.



- Minimize the number of connections that transfer large frames, to reduce their impact on each other and on the other traffic that shares the same egress link or channel.
- Verify that the function processor can be suitably provisioned for large frame transmission. Refer to the NN10600-551 *Nortel Multiservice Switch 7400/15000/20000 FP Configuration Reference* for information on individual FPs.

To reduce delay for other traffic that uses the same egress link

- avoid transmitting delay-sensitive traffic
- do not transmit multimedia traffic
- set other traffic to high emission priority to reduce delay. Traffic can still be affected, however.

Where possible, keep egress link use low. When other traffic must share the same physical interface at the egress link, increase the egress link capacity.

End-to-end FRF.12 frame fragmentation

End-to-end fragmentation, according to the Frame Relay Forum FRF.12 Implementation Agreement, is a frame fragmentation feature that can be used between two fragmentation peers: an external DTE and a Nortel Multiservice Switch DTE. The transmitting frame relay DTE fragments long frames into a sequence of shorter frames. These shorter frames make up frame fragments. The frame fragments are reassembled into their original frame by the receiving DTE. The fragmentation process occurs transparently to the frame relay network between the peers.

End-to-end fragmentation is necessary to control delay and delay variation when voice traffic is carried across the same interface as data. End-to-end fragmentation is configured on a per DLCI basis. It is recommended that DLCIs dedicated to transporting data traffic on a FRUNI carrying voice traffic have *frf12EndToEnd* turned on. DLCIs dedicated to carrying high priority voice traffic should not have *frfEndToEnd* turned on. Voice frames should be transported in complete frames rather than in fragments.

The following list outlines the benefits of configuring end-to-end FRF.12:

- voice frames can be interleaved with FRF.12 data frames during transmission
- voice jitter is reduced when data and voice traffic is carried across the same slow link
- FRF.12 data frames are discarded before voice frames if congestion begins to occur
- external FRF.12 DTE can interwork with IP, BNX, and DPRS networks



- interworks with IP-optimized DLCI to provide an efficient access mechanism from FRF.12 interfaces to the Multiservice Switch virtual router

If *frf12EndToEnd* is turned on, the EP, DP, and RCOS of the DLCI dedicated to voice traffic should be set higher than the EP, DP, and RCOS of a DLCI dedicated to other types of traffic. This allows voice frames to be enqueued ahead of data frames in the transmission queue.


The BNXiwf service does not support the DP being set on a per DLCI basis, and therefore, uses CIR and EIR to distinguish between voice traffic and data traffic. DLCIs dedicated to carrying voice traffic should have a setting of DE=0 and be placed in the CIR bucket, while DLCIs dedicated to carrying data traffic should have a setting of DE=1 and be placed in the EIR bucket.

For more information on available frame relay DPRS EP, DP, and RCOS values, see [Discard priority \(page 141\)](#), [Emission priority \(page 153\)](#) and [FR UNI routing class of service \(page 152\)](#), respectively. For information on DP and how it relates to IP-optimized BNX-interworking or DLCI, see [BNX interworking overview \(page 98\)](#) and NN10600-801 *Nortel Multiservice Switch 7400/15000/20000 IP Configuration Management*.

See figure [End-to-end FRF.12 reference diagram \(page 159\)](#) for an illustration of how fragmentation and reassembly occurs with end-to-end FRF.12.



End-to-end FRF.12 reference diagram



End-to-end FRF.12 can only be configured for FR UNI PVCs and SPVCs on PQC cards.

For information on configuring end-to-end FRF.12, see NN10600-901 *Nortel Multiservice Switch 7400/15000/20000 Frame Relay Configuration Management*.

FR UNI weighted fair queue scheme

The weighted fair queue (WFQ) scheme ensures fair share of service, relative to traffic priority, by using a fixed algorithmic technique to schedule connections. For each DLCI, WFQ uses the traffic's transfer priority (TP) to provide fair service (transmission priority) across all emission queues.

Transfer priority values map to emission priority (EP) queues that define the traffic's urgency. The network tags urgent traffic with high EP and transmits it before lower EP traffic.

Flexible QoS provides the ability to map the DLCI's transfer priority (TP) to the EP. For more information about QoS, see [Flexible QoS mapping \(page 154\)](#).



An EP of 0 indicates the lowest emission priority and an EP of 3 indicates the highest emission priority. The EP is mapped to the emission queue (EQ) of the cards. See tables [Emission priority to queue number mappings for SBIC and PQC cards \(page 160\)](#) for mappings between EP and EQ.

MSA32 cards use a strict priority WFQ algorithm.

Emission priority to queue number mappings for SBIC and PQC cards

Emission Priority	Emission Queue of SBIC cards	Emission Queue of PQC cards
0	0	3
1	1	2
2	2	1
3	3	0
EP of 0 = lowest priority		

Traffic assigned an EP of 3 will be serviced before traffic assigned an EP of 2, followed by an EP of 1, and an EP of 0. If there is no traffic with an EP of 3, then the order in which the queues are serviced is EP 2, followed by EP 1, then EP 0.

Traffic serviced on the highest priority queue will experience the least amount of delay in transmission.

Each EQ is automatically assigned a weight. The weight indicates at what percentage that traffic will be transmitted in relation to other queues.

Table [Weight assigned to MSA32, 4pDS3Ch and 1pSTM1Ch cards \(page 160\)](#) lists the weight assigned to MSA32, 4pDS3Ch and 1pSTM1Ch cards.

Weight assigned to MSA32, 4pDS3Ch and 1pSTM1Ch cards

Emission Priority	Emission Queue	Weight
3	0	91
2	1	3
1	2	3
0	3	3
An EP of 0 = lowest priority		



For more information on transfer and emission priorities, see [Traffic subscription options \(page 151\)](#).

Multimedia traffic

For information about how the frame relay service handles multimedia traffic, see the following sections:

- [Multimedia traffic classes \(page 161\)](#)
- [Packet forwarding of multimedia traffic \(page 161\)](#)
- [Frame-cell trunk interrupting mode engineering guidelines \(page 162\)](#)

Multimedia traffic classes

Multimedia traffic classes (MMTC) are subscriber options that let you select traffic handling preferences from the network to meet your applications' requirements. MMTC handle delay-sensitive multimedia traffic and minimize the delay transfer time and variance through the Nortel Multiservice Switch network.

The following parameters define traffic class and provide flexibility to discriminate frame relay traffic delay requirements:

- transfer priority (TP)
- discard priority
- discard eligibility (DE) bit

You can map transfer priorities to emission priorities and routing class of service (RCOS) values that are appropriate for your frame relay service. You can set a transfer priority for an entire interface, or for each DLCI on the interface. The transfer priority for the DLCI takes precedence over the TP for the interface. For more information, see [Flexible QoS mapping \(page 154\)](#).

Configure discard priority as normal or high for each DLCI and interface. Either your CPE or the network tags the DE bit.

Packet forwarding of multimedia traffic

The dynamic packet routing system (DPRS) forwards Nortel Multiservice Switch frame relay service traffic across a Multiservice Switch network. The *forwardingPolicy* attribute, which you define for each module, determines the forwarding policy for traffic with transfer priorities mapped to the multimedia RCOS. Loadspreading and loadsharing are the forwarding policy options. For details about the *forwardingPolicy* attribute, see NN10600-450 *Nortel Multiservice Switch 7400: Operations: DPN-100 Interworking*.

Multiservice Switch nodes handle multimedia traffic on a separate forwarding algorithm. This algorithm is independent of the forwarding policy that the other traffic classes use.



The DPRS examines the multimedia packet header to check the emission priority (EP) bit and to see if it requires RID, MID, or local forwarding. The EP bit ensures that the quality of service for delay-sensitive multimedia traffic is guaranteed, even when crossing Multiservice Switch trunks over ATM. If the packet requires RID or MID forwarding, an algorithm selects the link for the multimedia packets.

Packet forwarding of multimedia traffic involves three steps:

- selecting a forwarding path
- selecting a forwarding link
- testing the selected link for congestion

The RID or MID selects a delay link group to forward the multimedia packet. A delay link group is the route to the path that provides the lowest delay characteristics. This is the path on which the packet forwards from the source node to the destination node. At each node in a network, DPRS maintains up to two equal-cost delay link groups. Half of the DLCIs use one link group, and the remaining DLCIs use the other link group to achieve traffic balancing.

The packet header generates a link selector. To balance traffic statistically within a link group, the link selector directs DLCIs to different preferred links, based on their destination and their unique process identity.

If the selected link is not congested, the packet forwards to the link and the forwarding algorithm completes. If the link is congested, the network compares the packet discard priority with the link congestion level. If the network allows it to continue, the packet transmits on the congested link. Otherwise, the network discards the packet.

DPRS does not support multimedia traffic overflow either across links within a link group, or across link groups. As much as possible, this ensures order-preserving while multimedia traffic traverses the Multiservice Switch network.

Frame-cell trunk interrupting mode engineering guidelines

Persistent multimedia oversubscription degrades quality of service (QoS) in applications. Follow these guidelines when you engineer the Nortel Multiservice Switch trunking bandwidth and network topology for the MMTC feature:

- For normal no-failure cases, restrict multimedia traffic to 30% of individual Multiservice Switch trunking bandwidth within the network.
- For multiple Multiservice Switch trunks or node failures, restrict multimedia traffic to 60% of individual Multiservice Switch trunking bandwidth within the network.



- When interrupting PORS and multimedia traffic use the same Multiservice Switch trunking facility, reduce the PORS traffic to satisfy the above restrictions.

Class of service types

You can configure frame relay for one of four access classes of service. The different discard policies result in different qualities of service. Frame relay traffic parameters such as CIR, Bc, Be, and Tc determine the possible access classes of service.

The table [Access class of service types \(page 163\)](#) summarizes the access class of service types.

For more information on class of service types, see the following sections:

- [Unrestricted CIR \(page 163\)](#)
- [Class of service A: Discard eligibility only \(best effort\) \(page 163\)](#)
- [Class of service B: CIR only \(guaranteed\) \(page 164\)](#)
- [Class of service C: CIR plus DE \(premium\) \(page 164\)](#)

Access class of service types

Class	Type of service	CIR	Bc	Be	Tc	EIR
A	Discard eligibility only	= 0	= 0	> 0	> 0	Be/Tc
B	CIR only	> 0	> 0	= 0	= Bc/CIR	0
C	CIR plus discard eligibility	> 0	> 0	> 0	= Bc/CIR	Be/Tc

Unrestricted CIR

Nortel Networks recommends this class of service for the Nortel Multiservice Switch frame relay NNI service. In this case, turn rate enforcement off to disable the CIR (and EIR) mechanisms. You can use this setting for low line speeds (for example, 64 kbit/s or less with T1 trunking), where the expected sustainable throughput is the line speed. However, maximum benefits of frame relay occur where line speeds are much greater than CIR (for example, 64 kbit/s CIR with a 256 kbit/s access line).

Class of service A: Discard eligibility only (best effort)

With this class of service, you can send data at a rate that, under normal conditions, is equal to the EIR. The network tags all frames as discard eligible. In the event of congestion, all new frames that arrive at the congestion point in the network discard. You can engineer the network to accommodate some percentage of all the subscribed throughput (for example 3:1) in this class. In this way, you can provide a best effort throughput class of service.



Class of service B: CIR only (guaranteed)

With this class of service, you can send data at a rate that, under normal conditions, is equal to the CIR. The network does not tag any frames for discard eligibility. New frames discard only at the onset of congestion. This service is a higher priority frame relay service than Class of service A. If you engineer the network to accommodate all of the subscribed CIR throughput, you can provide a guaranteed throughput class of service.

Class of service C: CIR plus DE (premium)

With this class of service, you can send data at a rate that, under normal conditions, exceeds CIR. The network transmits the allowed traffic in excess of CIR (that is, EIR) and tags those frames as discard eligible. Upon congestion, these discard-eligible frames discard immediately to protect network resources. This class of service offers the greatest flexibility, because you can use it to provide guaranteed throughput, and additional throughput if the network has spare bandwidth capacity.

Traffic management for BNX interworking

This section provides information on the traffic management characteristics on interworking between the network element running hyperstream software (BNX) and a Nortel Multiservice Switch 15000 or Multiservice Switch 20000 node, in which interworking supports frame relay over IP.

Information is organized into the following sections:

- [Data flow \(page 165\)](#)
- [IP datagram in FrUni/FrNni interworking with BNX switches \(page 166\)](#)
- [BNX header field settings for incoming frame relay traffic \(page 168\)](#)
- [Processing of BNX header fields in outgoing traffic \(page 170\)](#)
- [Quality of service \(page 171\)](#)
- [Congestion management \(page 173\)](#)

There is also relevant traffic management-related information in the following documents.

- NN10600-800 *Nortel Multiservice Switch 7400/15000/20000 IP Technology Fundamentals*
- NN10600-801 *Nortel Multiservice Switch 7400/15000/20000 IP Configuration Management*
- NN10600-802 *Nortel Multiservice Switch 7400/15000/20000 IP VPN Technology Fundamentals*



Data flow

The BNX interworking functionality (BNX-lwf) involves the Nortel Multiservice Switch 15000 or Multiservice Switch 20000 node's frame relay service and the BNX at the remote access end. The discussion of data flow focuses on the Multiservice Switch network edges (ingress and egress Multiservice Switch 15000 or Multiservice Switch 20000 FrUni/FrNni) of the IP network since frame relay frames and the PVC functionality exist only at the access ends.


The figure [Data flow in an IP network involving Multiservice Switch 15000 or Multiservice Switch 20000 FrUni Service \(page 166\)](#) illustrates how frame relay frames are encapsulated in IP datagrams before transmission across an IP subnet. This scenario also includes the possibility of using ATM in the backbone, which is transparent to the BNX-lwf and not discussed in this document.

The BNX-lwf is based on the premise that one end of the permanent virtual circuit (PVC) is always a Multiservice Switch FrUni/FrNni while the other end is always handled by a BNX supporting frame relay as an access protocol. In the figure 7, [Data flow in an IP network involving Multiservice Switch 15000 or Multiservice Switch 20000 FrUni Service \(page 166\)](#), the IP source address of the Multiservice Switch FrUni/FrNni is 192.4.5.6. The IP destination address is defined as a DLCI. In the figure, DLCIs are specified with the same destination address (192.1.2.3).

The BNX-lwf supports only FrUni and FrNni PVCs.



Data flow in an IP network involving Multiservice Switch 15000 or Multiservice Switch 20000 FrUni Service



IP datagram in FrUni/FrNni interworking with BNX switches

When the Nortel Multiservice Switch 15000 or Multiservice Switch 20000 FrUni/FrNni service interworks with BNX switches, incoming frames are formatted with a BNX header before encapsulation into an IP datagram (with an IP header). That is, a BNX header is added to frames received at the ingress node. The BNX header is inserted into user traffic only.


In this release, the BNX header has these characteristics:

- the BNX-lwf supports only the BNX Version 5 header, which is the most recent version
- the three octets following the version field octet were changed in Version 5



The figure [Format of an IP datagram with a BNX header \(page 167\)](#) illustrates the format of the BNX header. The table [Fields in the BNX header \(page 167\)](#) describes the fields in the BNX header.

Format of an IP datagram with a BNX header



Fields in the BNX header

Bits in the BNX header (mapping to IP datagram)	NAME	Value	Meaning
24-31 (Bit 0 to 7 in IP datagram)	version	5	BNX version header
16-23 (Bit 8 to 15 in IP datagram)	reserved	0	reserved
12-15 (Bit 16 to 19 in IP datagram)	traffic protocol type	3	The value 3 identifies frame relay traffic.
11 (Bit 20 in IP datagram)	discard eligible (DE)	0 or ignored	obsolete

(1 of 2)



Fields in the BNX header (continued)

Bits in the BNX header (mapping to IP datagram)	NAME	Value	Meaning
8-10 (Bit 21 to 23 in IP datagram)	discard level (DL)	1/2/3/4	N-th Discard level once N-1 congestion level threshold is reached for frame relay traffic marked as Discard Eligible (DE=1).
0-7 (Bit 24 to 31 in IP datagram based on a 32 bits representation)	priority level	0/1/2/3	0 - Critical (cross-net signaling datagrams) 1 - High 2 - Normal 3 - low

(2 of 2)

BNX header field settings for incoming frame relay traffic

The fields of the BNX header (as shown in table [Fields in the BNX header \(page 167\)](#)) are set according to the following description.

- 1 The value of the version field is always set to 5.
- 2 The value of the reserved field is always set to 0.
- 3 The value of the traffic protocol type field is always set to 3 for frame relay user traffic.
- 4 The value of the DE field is always set to 0.
- 5 The *discardLevelEnabled* attribute under the *FrUni BnxProtocol* and *FrNni BnxProtocol* components controls the scheme selection. When the *discardLevelEnabled* attribute is set to false, the discard level field is set to 2 for DE=1 frames and to 4 for DE=0 frames. When the *discardLevelEnabled* attribute is set to true, the discard level field is tagged with the default value provided by the *trafficDiscardLevel* attribute, also under the *FrUni BnxProtocol* and *FrNni BnxProtocol* components. The *trafficDiscardLevel* attribute is configured on an interface basis (one value for each FrUni/FrNni). and is found under the *FrUni BnxProtocol* and *FrNni BnxProtocol* components. The table [Setting discard level at the ingress end of the IP network \(page 169\)](#) summarizes this relationship.
- 6 When the *trafficPrioritizationEnable* attribute under the *FrUni BnxProtocol* or *FrNni BnxProtocol* component is set to false, the priority level field of the BNX header is set to 2 for user traffic. The value 2 is the default. When the *trafficPrioritizationEnable* attribute is set to true, the priority level field is set based on an attribute defined at the FrUni/FrNni level or at the DLCI level.



The source is either the value of the *defaultTrafficPriority* attribute under the *FrUni BnxProtocol* and *FrNni BnxProtocol* components or the *trafficPriority* attribute under the *FrUni DlcI BnxInterworkingFunction* and *FrNni DlcI BnxInterworkingFunction* components. The *trafficPriority* attribute, when provisioned under the DLCI, overrides the *defaultTrafficPriority* attribute at the FrUni/FrNni level. The table [Setting traffic priority at the ingress end of the network \(page 169\)](#) summarizes this relationship.

Setting discard level at the ingress end of the IP network

discardLevelEnabled attribute	Discard Level field of the BNX header for
false	2 for DE = 1 frames 4 for DE=0 frames
true	value provisioned in the <i>trafficDiscardLevel</i> attribute (a value between 1 and 4)

Setting traffic priority at the ingress end of the network

trafficPrioritization-Enable attribute	trafficPriority attribute	Priority Level field of the BNX header
false	Not applicable	2
true	useDefaultTrafficPriority	defaultTrafficPriority value (1 through 3)
	1,2 or 3	trafficPriority value (1 through 3)

When the BNX header is formatted, the IP header is added to the packet to complete the datagram. The figure [Format of an IP datagram with a BNX header \(page 167\)](#) shows the layout of the IP header. Field values are derived as follows:

- 1 The version field is set to 4. No other IP header versions are supported by for the BNX-lwf.
- 2 The IHL field is set to 5, independently of the protocol field value.
- 3 The total length field is calculated.
- 4 The identification field is a different value for each packet.
- 5 The D-bit is set to 0 to allow segmentation in the IP network.
- 6 The fragmentation offset field is set to 0.



- 7 The protocol field is set to 200 for user traffic and to 202 for cross-net datagrams that the FrUni/FrNni generates.
- 8 The TOS field is set to 0.
- 9 The TTL field is set to 30.
- 10 The header checksum is computed according to RFC1071 and inserted into the header checksum field.

Processing of BNX header fields in outgoing traffic

In this description, the term “outgoing traffic” refers to frame relay frames that are encapsulated in IP datagrams received from the IP network.

When the Nortel Multiservice Switch 15000 or Multiservice Switch 20000 FrUni/FrNni service receives IP datagrams from a BNX router supporting frame relay, the switch analyzes the IP and BNX headers to determine validity. The analysis is first performed on the IP header and then, in the case of user traffic, on the BNX header.

Processing for the IP header

Processing of the IP header proceeds through these stages:

- 1 The version field is checked to ensure that it is set to 4.
- 2 The IHL field is checked independently of the protocol field value to ensure that it is set to 5.
- 3 The total length field is verified against their minimum and maximum values allowed at the interface.
- 4 The identification field is a different value for each packet.
- 5 The segmented frames will be recombined before being sent to the link.
- 6 The protocol field is verified. Only datagrams containing a value of either 200 or 202 are accepted.
- 7 The TOS field header is ignored.
- 8 The source IP address is verified by the destination DLCI.

If the IP header is invalid, the frame is dropped.

If the IP header is valid, and the datagram is associated with user traffic (that is, the IP protocol id is 200), the system processes the BNX header.

Processing for the BNX header

Processing of the BNX header proceeds through these stages:

- 1 The value of the version field is checked to ensure that it is set to 5.
- 2 The reserved field is ignored.



- 3 The traffic protocol type is checked to ensure that it is set to 3 for user traffic
- 4 The DE field (which is expected to be 0) is ignored by the Multiservice Switch FrUni/FrNni service.
- 5 The discard level field is checked to ensure that the value is between 1 and 4.

Discard information is used in the IP network and at the egress to determine if the datagram or frame can be discarded or enqueued.
- 6 The priority level field is checked independently of the setting of the *trafficPrioritizationEnable* attribute at the interface to ensure that the value is between 0 and 3.

Quality of service

Quality of service (QOS) is limited at ingress to the layer 2 DLCI classification. The QOS associated with the frame is translated into a priority level and a discard level. Markers for these levels are inserted into the BNX header attached to the payload. At egress, QOS is translated to an emission priority and a discard priority. These topics are described in the following sections:

- [Quality of service at the ingress end \(page 171\)](#)
- [Quality of service at the egress end \(page 172\)](#)

Quality of service at the ingress end

The QOS associated with the Nortel Multiservice Switch 15000 or Multiservice Switch 20000 FrUni/FrNni in a DPRS networking environment is organized around the routing class of service (RCOS) functionality that is implemented in the DPRS network. A functionality similar to RCOS is available in an IP network that is based on BNX switches. The BNX header that is carried in each datagram makes this equivalent functionality possible.

In the BNX-based IP network, the QOS functionality is distributed at each BNX switch by performing packet classification and prioritization on a the basis of the datagram. The packet classification and prioritization mechanisms are available through the prioritized queue selection mechanism and the congestion management techniques defined through configuration.

Decisions to discard or forward datagrams are made at the egress port where the transmit queues (both software and hardware) are located. This decision is based on the discard and priority levels associated with each datagram compared to the congestion level experienced in a particular queue. (These levels are defined in the discard level and priority level fields of the BNX header.



Quality of service at the egress end

At the egress node, a class-based packet forwarding mechanism is available as part of the BNX-lwf in the Multiservice Switch 15000 or Multiservice Switch 20000 FrUni/FrNni service. This mechanism converts the priority level that was part of the BNX header received into an emission priority. Emission priority determines the emission priority queue associated with the level of QOS desired.

The priority level and the discard level both determine the level of QOS attributed to a frame.

The conversion of priority level to emission priority queue is shown in the table [Priority level field of the BNX header versus emission priority \(page 172\)](#).

Priority level field of the BNX header versus emission priority

Priority level defined in the BNX header	Emission priority with 4 EP queues	Emission priority with 2 EP queues
0	-	-
1	2 This queue is also used to carry LMI traffic.	1 This queue is also used to carry LMI traffic.
2	1	0
3	0	0

In BNX-lwf, four values are available to specify traffic priority in which numerically low values are associated with high priority traffic. The operator can configure data traffic for the following priorities:

- 1 = high
- 2 = normal
- 3 = low

On Nortel Multiservice Switch nodes, Priority 0 is not used for data traffic.

The discard level field in the BNX header received determines the ability of the switch to discard the frame when congestion is experienced in a transmit queue. A numerically low value for a discard level index in a BNX environment means that the frame is more easily discarded. Hence, a discard level of 1 means that as soon as the switch experiences congestion (reaching congestion level 1), the frame is discarded.




Congestion management

Throughout this description of congestion management, reference is made to numbered nodes and queues. These numbers refer to the labels in the figure [Data flow in the context of frame relay over IP \(page 173\)](#).

In BNX-lwf, congestion management includes congestion avoidance and congestion handling schemes. Frame relay frames and IP datagrams can cause congestion at the following points in the network:

- at the ingress end of the Nortel Multiservice Switch 15000 or Multiservice Switch 20000 FrUni/FrNni service, in the receive queue for incoming traffic also known as traffic from link (Point 11)
- within the network (Point 2)
- at the egress end of the network (transmit queue) for outgoing or egress traffic, which is also known as traffic to the link (Point 5)

Data flow in the context of frame relay over IP



Congestion management at the ingress

This description on congestion management refers to numbered nodes and queues in the figure [Data flow in the context of frame relay over IP \(page 173\)](#). This example is based on the FrUni.



When ingress traffic is received on the Multiservice Switch 15000 or Multiservice Switch 20000 FrUni (Point 1), processing goes through the following stages.

- 1 If the rate enforcement mechanism is configured, traffic compliance is verified through rate enforcement. If verification is not successful, the frames are discarded, otherwise the process continues. By default, DE=0 frames are put in the committed information rate (CIR) bucket. When the CIR bucket is full, the discard eligibility bit (DE-bit) is set to 1 and the frame is put in the excess information rate (EIR) bucket. DE=1 traffic is sent directly in the EIR bucket and is discarded if EIR bucket is full. As a result, DE=1 traffic is always discarded before DE=0 traffic.
- 2 If the rate enforcement mechanism is not configured, the system assumed that traffic compliance is met (in terms of CIR and EIR) and no DE-bit is set. However, the DE-bit setting in the frame received is carried and also considered as described in the following stages.
- 3 The system checks the congestion level in the receive queue at node 1. If the receive queue experiences severe congestion and the DE-bit is set in the frame just received then the frame is discarded. Otherwise, the FECN bit is set in the header.
- 4 The system checks the congestion level in the transmit queue. If there is any congestion, then the BECN bit is set in the header.
- 5 If a proportion of FECN frames has been received from the subnet, the system can set the BECN bit. BECN frames can be sent in response to FECN frames.

These concepts are summarized in the following tables:

- [Frame relay traffic policing at the ingress of the IP network \(page 174\)](#)
- [Receive queue congestion at the ingress \(page 175\)](#)
- [Transmit queue congestion at the ingress \(page 175\)](#)

Frame relay traffic policing at the ingress of the IP network

Rate enforcement	Traffic throughput	Outcome
On	$TT \leq CIR$	DE-bit not set
	$CIR < TT \leq EIR$	DE-bit set
	$TT \geq EIR$	Frame may be discarded
Off	ignored	DE-bit not set



Receive queue congestion at the ingress

Receive queue congestion	DE-bit in received frame	Outcome
< severe	DE = 0 or 1	no impact
>= severe	DE = 0	FECN = 1
	DE = 1	Frame is discarded

Transmit queue congestion at the ingress

Transmit queue congestion	Outcome
no congestion	no impact
>= mild	BECN = 1

Congestion management at the egress

The application of traffic management controls is triggered after the IP header and BNX header have been processed and removed. The priority level and discard level contained in the BNX header at that time are internally translated and carried all along the process phase. The priority level and discard level are now being considered at this stage.

In the context of the BNX-lwf, the egress link is connected to a device that handles frame relay traffic.

The congestion management is applied at the queue level, where a queue corresponds to a traffic priority. The system applied management strategies through the following stages.

- 1 The congestion level of the receive queue is verified. The BECN bit is set in the frame relay header when the receive queue experiences severe congestion. See the table [Receive queue congestion at the egress \(page 176\)](#).
- 2 Based on the congestion level of the transmit queue and the discard level (as recorded in the datagram BNX header), the system can discard the frame at this stage. See [Transmit queue congestion at the egress \(page 176\)](#). This behavior matches the approach implemented by the BNX congestion control mechanism and is different from the Nortel Multiservice Switch 15000 or Multiservice Switch 20000 FrUni congestion management approach.
- 3 The discard levels used at stage 2 are in reference with the discard levels as supported by the BNX software (four levels).



- 4 If the frame can be placed in the transmit queue, that queue corresponds to the priority level associated with the frame. The FECN field in the frame header is set if any congestion is experienced at the transmit queue. Possible settings are mild, low, moderate, and severe.
- 5 The frame is dequeued and sent to the link.

Receive queue congestion at the egress

Receive queue congestion	DE-bit in frame received	Outcome
no	DE = 0 or 1	no impact
mild	DE = 0 or 1	no impact
moderate	DE = 0 or 1	no impact
severe	DE = 0 or 1	BECN = 1
total	DE = 0 or 1	BECN = 1

Transmit queue congestion at the egress

Congestion level in the Multiservice Switch 15000 or Multiservice Switch 20000 Tx queue	Discard Level in the BNX header	Action
0 (No Congestion)	1 through 4	transmit
1 (Mild)	1	discard
	2 through 4	transmit with FECN = 1
2 (Moderate)	1 and 2	discard
	3 and 4	transmit with FECN = 1
3 (Severe)	1 through 3	discard
	4	transmit with FECN = 1
4 (Total)	1 through 4	discard



Frame relay accounting

Since the frame relay service is a best-effort delivery service (that is, it does not guarantee data delivery through the network), traditional billing methods require the network operator to correlate separate statistics gathered from each end point to determine the actual data transfer. The frame relay accounting system provides an option for the user to eliminate this manual correlation and in either one or two accounting records provides an accurate measure of data delivery across the entire end-to-end virtual connection.

Frame relay accounting allows a service provider to bill end users based on the amount of network resources they use. It also allows the generation of end user reports that can be used to validate quality of service parameters. These reports contain relevant usage data such as the committed information rate (CIR), the excess information rate (EIR), and byte and frame/segment counts. Billing can be based on end user data traffic that is collected at the exit (egress) points of a frame relay network. The end user data traffic is put into one or two accounting records per virtual circuit. The accounting records are then spooled to the Nortel Multiservice Switch node's disk.

Accounting records are transferred from the disk to the MDP. These accounting files are either converted to bulk format or published format and stored on the MDP. After the files are collected and stored on the MDP, they can be transferred to either a billing host or a network engineering host. For more information, refer to the 241-6001-309 *Nortel Multiservice Data Manager Management Data Provider*.

Navigation

- [Benefits of frame relay accounting \(page 178\)](#)
- [Frame relay accounting meter \(page 179\)](#)
- [Frame relay accounting record \(page 181\)](#)
- [CircuitId attribute and the accounting record \(page 185\)](#)
- [Egress protocol \(page 186\)](#)
- [Call clear \(page 187\)](#)
- [Peak water mark \(page 187\)](#)



- [Call management for switched virtual circuits \(page 193\)](#)
- [Generating reports \(page 194\)](#)
- [Frame relay billing policies \(page 200\)](#)
- [Accounting for BNX-lwf \(page 204\)](#)

Benefits of frame relay accounting

There are many benefits to frame relay accounting.

The benefits of this feature are as follows:

- The set of billing parameters for the frame relay service is concise and supports a variety of billing policies.
- The generation of end user reports can be generated to validate quality of service parameters, peak to average traffic ratio, and service contract delivery. Frames from both ends are synchronized, providing an accurate tabulation of frames/bytes that are successfully delivered.
- The service provider can gather detailed information about end user traffic and can develop billing policies based on this detailed information.
- Single-ended frame relay accounting (when provisioned) reduces the number of accounting records generated and spooled to the downstream processing system.
- It relates data from the end-points of a virtual circuit to each other.
- Frame relay accounting takes a snapshot of the billing data on both ends of the virtual circuit every 15 minutes and exchanges this information between the endpoints of a virtual circuit. In the event of a path or node failure, the network operator has a record of network activity which is at most 15 minutes old.
- When there is no data exchanged in the accounting interval, snapshots are not exchanged. This conserves network bandwidth.
- By provisioning a correlation tag identifier with a value (the default is null), the identifier is assigned to a frame relay PVC and becomes part of the accounting record and is listed as the circuitId attribute. The value provisioned for the correlation tag is the same value that is reported for the circuitId attribute in the accounting report. Only when the correlation tag identifier is provisioned with a value other than the default does the circuitId attribute appear in the accounting record. The use of a circuit identifier facilitates off-switch correlation of data for the interfaces involved in the same permanent connection.

For more information about Multiservice Switch accounting, see *NN10600-560 Nortel Multiservice Switch 7400/15000/20000 Accounting*.



Frame relay accounting meter

The accounting meter is a software entity associated with a frame relay DLCI which collects end user data traffic at the DLCI and generates accounting records. There are two accounting meters in each frame relay virtual circuit (VC), one at each DLCI. For more information on DLCIs, see [DLCI assignment \(page 37\)](#).

The generation of an accounting record is controlled by the meter's internal timer, which expires after 12 hours or by a provisionable time-of-day accounting (TODA) schedule. For more information on provisioning the TODA schedule, refer to NN10600-560 *Nortel Multiservice Switch 7400/15000/20000 Accounting*. If no TODA is provisioned, the meter's internal 12-hour timer is used.

The operation of the accounting meter is determined by accounting-specific provisionable attributes in the module wide *VirtualCircuitSystem (Vcs)*, the *FrUni* (or *FrNni*), *DataNetworkAddress (Dna)*, and the frame relay *DLCI* components. After the meter is created, provisioning data is sent to the meter. Changes to provisioned data, while the virtual circuit (VC) is up and the meter is in operation, are not automatically sent to the meter. Calls established before the new provisioning change continue to use old provisioned information. The current call must be disconnected and a new call must be established in order for the new provisioned changes to take effect on the accounting meter.

Accounting meter synchronization

To generate an accounting record (which includes relevant data from opposite ends of the virtual circuit), the required information from each end must be retrieved at the same time. In other words, the accounting system must take a simultaneous snapshot of accounting data at each end of the virtual circuit and use this information to generate an accounting record. The egress protocol utilizes the in-band synchronization concept to ensure the information collected by both accounting meters are taken at the same time.

In-band synchronization

In-band synchronization can best be illustrated using the figure, [In-band synchronization—the bookmark concept \(page 181\)](#). A virtual circuit joins two physical network access points. Numbers R1 through R4 denotes reference points for discussion purposes. The master end of the protocol is located at the charged end of the virtual circuit.

To best illustrate how the in-band synchronization concept works, it is necessary to examine the exchange of accounting information from one meter to the next. This exchange occurs at the end of a meter interval. The meter



interval is the time period after which accounting information is exchanged between accounting meters located at each end of the virtual circuit. The master end of the virtual circuit tracks these meter intervals.

The meter interval is hardcoded to 15 minutes.

At the end of the meter interval, the master end will send an accounting inquiry signal. This in-band signal will act as a place holder (hence the term bookmark) in the data stream. This signal is used to capture a snapshot in time of the traffic.

The sequence of events shown in the figure, [In-band synchronization—the bookmark concept \(page 181\)](#) occur as follows:

- When issuing the inquiry, the master end of the virtual circuit will capture its ingress usage data (I_m) counts and send the ingress usage data (I_m) counts (reference point R1) to the slave end in-band with other end user data, right after the last counted end user frame.
- After receiving the accounting inquiry, the slave end will store the ingress usage data from the master end (I_m), capture its own egress and ingress counts (E_s, I_s) (reference point R2), and send the captured egress and ingress counts (E_s, I_s) to the master end in-band, right after the last counted end user frame (reference point R3).
- After receiving a response to the inquiry, the master end will store its egress count (E_s) received from the slave end and capture its own egress counts (E_m).
 - At this point, the master end has a complete view of all counts on both ends of the virtual circuit. This view consists of an in-band snapshot of the egress and ingress data traffic in both directions. This information is used to generate a single-ended usage-based accounting record.
- Both ends of the virtual circuit have the same captured snapshot which further increases the reliability of the accounting data. Therefore the master end sends a final in-band message which contains the master end's egress counts (E_m) to the slave end (reference point R4).


There are two types of egress protocol sequences: the update sequence and the reset sequence. The two types of sequence differ only in the meaning of the final step listed above. In the update sequence, the master end sends a message to the slave end to update the master end's egress counts. The update sequence is used to periodically exchange usage counts between accounting meters. In the reset sequence, the master end notifies the slave end to reset its usage counters. The reset sequence is used at the time of record generation under the following conditions:

- When usage counts have to be exchanged between meters
- When one or both meters generate an accounting record



- Both accounting meters' usage counters have to be reset

In-band synchronization—the bookmark concept



Frame relay accounting record

An accounting record is a data record containing information on end user traffic over a certain period of time. This time period is known as the accounting interval. The accounting record is created by an accounting meter at the end of an accounting interval. The accounting record can be either singleEnd (only one accounting record is created for a given accounting interval from a VC) or bothEnds (when two accounting records are created for



a given accounting interval for a VC—one at each end of the VC). A singleEnd accounting record contains metered data from the DLCIs at both ends of the VC. BothEnds accounting records contain metered data only from the local end of the VC (DLCI).

The figure, [SingleEnd accounting record generation \(page 183\)](#) shows singleEnd accounting record generation. Accounting data is collected from the non-generating accounting meter using the egress protocol. The accounting meter then generates an accounting record with collected data from the generating and non-generating end-points.


The figure, [BothEnds accounting record generation \(page 183\)](#) shows how accounting data is collected from both accounting meters. Each accounting meter then generates an accounting record.

Nortel Multiservice Switch node's accounting record is formatted in an internal type-length-value format. Each record contains information such as start time, end time, call reference number and other information details that identify the record. For the frame relay service, the accounting records also contain provisioned data such as CIR and EIR, specific counts, such as byte counts or frame and segment counts. The counts shown in the figure, [Possible byte or frame/segment counts in the frame relay accounting record \(page 184\)](#) are contained in the frame relay accounting record.


Since the counts from the opposite ends of the virtual circuit are synchronized, both service providers and service users can use the generated accounting record as a performance or quality of service (QoS) report.



SingleEnd accounting record generation



BothEnds accounting record generation





Possible byte or frame/segment counts in the frame relay accounting record

Ingress	Egress
Total number of bytes received from the link	Total number of bytes sent to the link
Total number of frames/segments received from the link	Total number of frames/segments sent to the link
Number of EIR bytes sent to the subnet	Number of EIR bytes sent to the link
Number of EIR frames/segments sent to the subnet	Number of EIR frames/segments sent to the link
Number of bytes in discarded invalid frames	
Number of invalid discarded frames/segments (discarded in the following order: frames/segments discarded when A-bit is on, frame errors (too long or too short), rate enforcement, or local processor and Multiservice Switch trunk congestion)	

The Multiservice Switch system counts all bytes and invalid discarded frames (or segments). During times of congestion, Multiservice Switch nodes can temporarily lower its effective CIR/EIR rate (achieved through provisioning specific attributes in the *DataLinkConnectionIdentifier* component). In other words, it “adapts” its effective CIR/EIR rate to the level of congestion. Invalid discarded bytes and frames/segments are counted. This value is included in the accounting record to assist network engineering. The incidence of rate adaptation at the frame relay DLCI is recorded in the accounting record.

To allow carriers and end users to differentiate between invalid frames/segments discarded and frames/segments discarded due to network faults (such as congestion or frame corruption), the number of invalid frames/segments is also collected and included in the accounting record.

All frame relay usage data is contained in two optional fields: the calling snapshot field and the called snapshot field. In charged-end mode, the Calling Snapshot contains the last snapshot of frame relay data from the calling end taken by the last egress protocol, and the called snapshot contains the last snapshot of frame relay data from the called end taken by the egress protocol. In bothEnds mode, the calling and called snapshots are not synchronized. Snapshots are exchanged under the following conditions:

- when a TODA change-over is detected
- when the charged-end 12-hour timer expires
- when a call is cleared

There are three important parameters collected in the accounting record that pertain to congestion. The first parameter is the total number of frames/segments or bytes dropped due to congestion (total in - invalid dropped - total



out). The second parameter measured is the EIR dropped due to congestion (EIR to subnet - EIR out). The third parameter measured is the CIR discarded due to congestion, which is derived from the following equation: (total in - EIR to subnet) - (total out - EIR out).

These three parameters are measured because they allow the user to differentiate between frames/segments or bytes dropped by the network due to congestion and ones that were invalid at the interface.

To minimize the performance and memory effects, either segment counts or frame counts are collected in the accounting record (not both at the same time).

CircuitId attribute and the accounting record

When frame relay accounting is enabled and the correlation tag identifier is provisioned on a connection by a user, the identifier is added to the accounting record for frame relay UNI and NNI PVCs as the circuitId attribute. However, the PVC circuit id feature is not essential to frame relay accounting. A correlation tag can be provisioned against a frame relay DLCI without enabling accounting and accounting can be enabled without having to provision a correlation tag. A user can not use a correlation tag that is provisioned against a connection to control or access the endpoint of that connection. A correlation tag value can be changed on an active permanent connection without bringing down the connection. The new correlation tag value takes effect immediately but is only visible if a display of the specific DLCI is performed. Changing a correlation tag value does not result in the immediate generation of an accounting record; instead, the value is included in the next accounting record that is generated.

The circuitId attribute is part of the accounting record generated at the endpoint where accounting is enabled. The circuitId is local to the frame relay DLCI where it is provisioned by the user. The value for the circuitId attribute is the same value that was provisioned as the value for the correlation tag. Provisioning bothEnds accounting with egress protocol disabled and accounting enabled, results in the generation of an accounting record at both the master and slave ends of the connection. Provisioning singleEnd accounting with both egress protocol and accounting enabled, results in the generation of an accounting record only at the master end of the connection. In this case, the accounting record contains both local and remote accounting information. When singleEnd accounting is provisioned with egress protocol disabled, only local accounting information is included in the accounting record.

The circuitId attribute is part of the accounting record generated at the endpoint where accounting is enabled. The circuitId is local to the frame relay DLCI where it is provisioned by the user. The value for the circuitId attribute is



the same value that was provisioned as the value for the correlation tag. Provisioning bothEnds accounting with egress protocol disabled and accounting enabled, results in the generation of an accounting record at both the master and slave ends of the connection. Provisioning singleEnd accounting with both egress protocol and accounting enabled, results in the generation of an accounting record only at the master end of the connection. In this case, the accounting record contains both local and remote accounting information. When singleEnd accounting is provisioned with egress protocol disabled, only local accounting information is included in the accounting record.

The circuitId is a part of the bulk data format (BDF) information of an accounting record sent to the data collection system (DCS) only when the correlation tag is provisioned. The use of a circuit identifier facilitates off-switch correlation of data for the interfaces involved in the same connection.

There is no system enforcement of the correlation tag value. The user can provision different values for the correlation tag at the frame relay DLCI for a connection. Nortel Multiservice Switch nodes do not require that the value of the string be unique. It is the user's responsibility to provision a unique correlation tag value, if required, or provision the same value for two endpoints for a connection, if correlation of accounting records for that connection is required. Once provisioned, the correlation tag value is static and can only be changed by a further provisioning change.

Since the DLCI for switched endpoints are created dynamically, the circuit identifiers for switched connections are not provisionable.

Egress protocol

The egress protocol is used to synchronize the accounting meters at the opposite ends of the connection and to facilitate the exchange of data collected by the accounting meters.

The egress protocol offers the following advantages:

- It minimizes the loss of accounting information if a node becomes inoperational. The egress protocol does not modify the collected accounting information.
- It synchronizes the accounting meters at the opposite ends of the virtual circuit.
- For ease of introduction into existing networks, the egress protocol can be switched on or off by provisionable attributes at each interface.



Call clear

When a virtual circuit is clearing, the in-band messages cannot be used to carry egress protocol messages since the data path is down. Therefore, the virtual circuit call take-down sequence (disconnect and disconnect confirm messages) is used to exchange accounting data between the accounting meters at the ends of the virtual circuit.

A single-ended accounting record contains usage data from both the local and remote ends of a virtual circuit.

Call termination occurs under the following conditions:

- call clear issued by either end of the connection
- call clear issued when the other end is not responding

When a call is terminated from one end of the connection, the final usage counts have to be delivered from the non-generating meter to the generating meter. If both ends accounting is provisioned, the final usage counts are exchanged between the two generating meters. These usage counts are sent with the disconnect and disconnect confirm messages.

Since the virtual circuit is not able to determine which meter generates an accounting record, it will always receive usage counts from the meter at the clearing end and put them in the disconnect message. Similarly, usage counts from the meter at the non-clearing end are put in the disconnect confirm message. The usage counts are then delivered to the meter at the opposite end of the virtual circuit.

When a call clear message is issued while the other end is not responding, clear request and clear confirm will not reach their destinations. The remote final usage counts will not be updated. The accounting record will be generated with the last egress protocol snapshot. This accounting record will contain both the latest local usage counts and remote usage counts which are at most 15 minutes old. The snapshot valid flag in the accounting record will be set to 0. (See "Troubleshooting" for more information on clear cause codes.)

Peak water mark


The peak water mark (PWM) is defined as the highest or peak traffic level (either CIR traffic, EIR traffic, or total traffic) measured over a defined time interval within the accounting interval. This time interval defined as the peak water mark interval represents an incremental time period when accounting data is collected by the accounting meter. The peak water mark represents the highest measured CIR traffic, EIR traffic, or total traffic that entered the DLCI over the peak water mark interval. The PWM interval is provisionable in the



range of 1 to 15 minutes. The *PWMinterval* attribute located under the *Collector/Account* component is used to provision the time interval. This time interval is set for all DLCIs on the node.

The actual peak traffic can be higher than the recorded peak water mark since the recorded peak water mark is dependent on the traffic profile. See the figure, [Peak water marks \(page 188\)](#) for a graphical interpretation of the peak water mark. For the sample intervals in the figure, [Peak water marks \(page 188\)](#), see [Example of peak water mark \(page 192\)](#).

Peak water marks



The PWM allows the network provider to understand the profile of the end user (in particular the peak traffic to average traffic ratio) and to verify that the provisioned CIR and EIR values are within the end user's requirements for available throughput at the DLCI. The PWM benefits the end user since it provides a profile of their DLCI usage. This allows the end users to better assess their current and future network bandwidth requirements on every DLCI.

As mentioned previously, there are three types of high water marks:

- CIR ingress byte PWM
- EIR ingress byte PWM
- total ingress byte PWM




The CIR ingress byte PWM measures the highest CIR byte traffic from the end user to the DLCI. The EIR ingress PWM measures the highest EIR byte traffic from the end user to the DLCI. The total ingress PWM measures the total byte traffic from the end user. All of these marks allow the network provider and the end user to compare the provisioned CIR and EIR values with the actual CIR and EIR values used by the end user.

The three peak water marks measure their respective traffic types at the reference points shown in the figure, [Accounting meter reference points \(page 190\)](#). Although traffic is shown in the figure as proceeding in only one direction, measurement of the three peak water marks also occurs in the other traffic direction at the opposite DLCI. The figure, [Accounting meter reference points \(page 190\)](#) only depicts traffic metering reference points where accounting data is collected.

A traffic peak may not be captured by using just one PWM per accounting interval. A PWM represents the average traffic through a DLCI over a duration of the PWM interval. If a short burst spans two PWM intervals (as shown in the figure [Peak water marks \(page 188\)](#)), the average traffic as indicated by the PWM value is smaller than the actual traffic peak. In the worst case scenario, the PWM value would be half the size of the peak traffic value.



Accounting meter reference points



Peak water mark interval

As mentioned previously, the peak water mark (PWM) interval is the time period over which the PWM is measured. If the TODA schedule is provisioned, the peak water mark interval can be provisioned in the range from 1 to 15 minutes. The *PWMinterval* attribute located under the *Collector/Account* component is used to provision the time interval. One PWM interval can be specified for one accounting interval. If TODA is not provisioned, then the PWM functionality is not supported. The PWM data can only be accessed by the accounting Application Programming Interface (API).



PWM interval provisioning is module-wide. In other words, all DLCIs on the Nortel Multiservice Switch node have the same PWM intervals, which are synchronized with the TODA schedule. PWM intervals are resynchronized with the next TODA changeover each time a TODA changeover occurs (see the figure, [TODA and PWM interval synchronization \(page 191\)](#)).


For example, if the current TODA changeover is at 9:57, the next TODA changeover is at 11:00 and the PWM interval is 5-minutes, the PWM changeovers are scheduled as depicted in the figure, [TODA and PWM interval synchronization \(page 191\)](#). The time interval between TODA 1 and TODA 2 is divided into 13 intervals (12 intervals of 5-minute duration and one interval of 3-minute duration). The short PWM interval is scheduled at the beginning so that PWM metering (for a Multiservice Switch node that is just starting up) is synchronized with other switches in the network. The PWM intervals are scheduled at the following times:

- TODA 2 - 12 * 5 (=10:00 am)
- TODA 2 - 11 * 5 min (=10:05 am)
- TODA 2 - 10 * 5 min (=10:10 am)
- and so on

The PWM interval starts 5 minutes before the TODA 2 changeover and it ends when TODA 2 changeover occurs.

Each time TODA is provisioned or the switch time is changed such that a TODA changeover occurs, the PWM intervals are resynchronized with the next TODA changeover.

TODA and PWM interval synchronization





Example of peak water mark


The use of the peak water mark is demonstrated using the following example. Consider the four peak water mark intervals as shown in the figure, [Sample traffic profile \(page 192\)](#). Assume that there was no traffic on the DLCI for all other PWM intervals except the four intervals shown. The provisioned CIR and EIR values are 4 kbit/s and 6 kbit/s respectively.

The actual measured PWM are as follows:

- total ingress PWM = 10 kbit/s
- CIR ingress PWM = 6 kbit/s
- EIR ingress PWM = 6 kbit/s

The figure, [Sample traffic profile \(page 192\)](#) shows the breakdown of total ingress PWM, CIR PWM, and EIR PWM. A summary of the PWM values (in kbits per second) appears in the table, [Values for total, CIR and EIR traffic \(page 193\)](#). The measured total, CIR, and EIR ingress PWMs appear in bold type.

Sample traffic profile





Values for total, CIR and EIR traffic

Traffic type	Interval 1	Interval 2	Interval 3	Interval 4
Total kbit/second	3	8	Total PWM =10	7
CIR kbit/second	1	CIR PWM = 6	5	1
EIR kbit/second	2	2	5	EIR PWM = 6

The EIR PWM (6 kbit/s) is within the provisioned limit, which poses no problem. The CIR PWM (6 kbit/s) is over the provisioned limit, which requires closer examination. At first glance, these PWM values suggest that an increase for the provisioned CIR is required:

$$\text{total} = \text{CIR} + \text{EIR} = 6 + 6 = 12 \text{ kbit/s}$$

However from the total PWM, we can see that the total traffic never exceeds the upper provisioned limit for the DLCI (10 kbit/s). Therefore, the end user or the network provider has two options:

- increase the provisioned CIR to 6 kbit/second to make sure that no CIR traffic overflows into the EIR bucket
- leave the CIR/EIR provisioning as is, because frames are not being discarded due to rate enforcement

Call management for switched virtual circuits

With frame relay switched virtual circuits (SVC), you can engineer call management attributes to limit the number of SVC calls allowed on the logical processor. When properly engineered, call management only accepts new SVC calls if the logical processor has enough resources to process the SVC call and its accounting records.

When you set up call management for switched virtual circuits you must consider the importance of accounting records and the impact of preventing new calls. Preserving all accounting records can mean rejecting new calls. Accepting all new calls can mean losing accounting records.

You can configure the SVC call management for the frame relay service running on a particular logical processor. The following attributes control the call management:

- *currentCalls*—the number of frame relay calls (PVC and SVC) currently capable of generating accounting records. Calls do not have to transfer data to generate accounting records.
- *maxCalls*—the maximum number of permanent virtual circuit (PVC) and SVC frame relay calls on the logical processor



- *headroomCalls*—a value that limits the new SVC call setup rate when there is a backlog in the accounting queue.
- *callsRefused*—the number of frame relay SVC calls that call management has refused

Call management does not affect PVC calls. PVC calls are set up regardless of the values of the call management attributes.

The defaults for *maxCalls* and *headroomCalls* are set to high values which prevent the call management system from limiting calls.

Generating reports

The meter generates accounting records from the collected data. The mode of generation can be provisioned to either *singleEnd* or *bothEnds*. In *singleEnd* generation, only the meter at the master end of the DLCI generates an accounting record which contains data from both accounting meters at a given VC. In *bothEnds* generation, both the master and the slave meters generate accounting records which contain only data from the respective local meters. The generation mode is controlled by the *generationMode* attribute in the *mod Vcs* component.

During *singleEnd* generation, the collection of data by the accounting meters are synchronized using the egress protocol. *SingleEnd* generation is triggered by the following conditions:

- TODA change-over (all meters react to the same change-over)
- the expiry of the accounting meter timer (each meter contains its own accounting meter timer)
- call clear
- change of the provisioned CIR or EIR values, at either end of the DLCI
- turning rate adaptation or rate enforcement on/off at either end of the DLCI

The egress protocol is automatically turned off when *bothEnds* generation is selected.

When *bothEnds* generation is turned on, record generation is triggered by the charged end by the following conditions:

- TODA change over (all meters react to the same change over)
- The expiry of the accounting meter timer (each meter contains its own accounting meter timer)

Accounting intervals at a given VC do not overlap. The end of one accounting interval is also the start of the subsequent accounting interval. Accounting intervals are defined for the whole node by the TODA provisioning schedule.



If TODA is not provisioned individually for each VC (DLCI), the accounting interval is set to 12 hours, starting either from the end of a previous accounting interval or from the beginning of a call.

An accounting interval can be shorter than 12 hours if the call is cleared before.

Accounting records generated in bothEnds accounting generation mode contain either calling end counts only (when the node that generated the record is the calling end) or called end counts only (when the node that generated the record is the called end). Accounting records generated in singleEnd generation mode contain counts from both ends (calling and called counts).

A singleEnd accounting record consists of a synchronized snapshot of the local and remote data taken when the record was generated at a TODA change over, or at the expiry of the 12-hour timer, or the normal termination of the connection. If a connection clears because of a Nortel Multiservice Switch trunk or remote node failure, the meter will generate an accounting record with local and remote snapshots which are not synchronized. The local snapshot will contain the latest counts taken when the connection went down and the remote snapshot will contain data taken by the last egress protocol sequence (at most 15 minutes). The snapshot valid flag in the accounting record will be set to 0.

The accounting record is generated by the meter at the non-charged end (even in singleEnd mode) to prevent loss of data during a Multiservice Switch trunk or remote node failure. During a Multiservice Switch trunk failure, an accounting record is generated at each end of the virtual circuit. This allows the accounting records to be checked against one another for accuracy. During a node failure, no accounting record exists so the local and remote snapshots taken by the egress protocol can be used to bill the end user.

If the frame relay backup PVC capability has been provisioned, and a call is redirected to the backup slave, an accounting record is generated on the master end only (if singleEnd accounting is provisioned) and on the master and primary slave end (if bothEnds accounting is provisioned).

When call redirection occurs, the redirect flag is set to 1 in the new accounting record which also contains the called backup DNA. The optional will display the primary DNA.

This appendix provides a summary of the accounting data generated by the Management Data Provider (MDP) accounting system based on a number of provisioning scenarios.



By using the tables in this section, you can determine both the location of the accounting record (whether it is generated at the charged end or the non-charged end) and the content of the accounting record (either the fixed portion, or the local or remote snapshot) depending on the provisioning data of the accounting system.

The location and content of the accounting record depends on the following:

- The *module VirtualCircuitSystem (mod vcs)* component which specifies segment size, the types of units counted, either frames or segments and the generation mode, either chargedEnd/bothEnds generation
- The collection reason bit selected in the *Dna* subcomponent (one of billing, test, audit, study, or force bit). The type of collection bit set indicates the intended purpose of the accounting record.
- The *DataNetworkAddress (Dna)* subcomponent of the *FrUni* and *FrNni* component which turns the frame relay accounting feature On or Off and indicates the intended purpose of the collected accounting record
- The *AccountingOn* attribute defined on each frame relay DLCI. This attribute is used to define whether accounting (and record generation) on a particular DLCI is turned on. The default value is yes (on).

For more information about frame relay accounting record generation modes, see the following sections:

- [Accounting record generation mode - singleEnd \(page 196\)](#)
- [Accounting record generation mode - bothEnds \(page 198\)](#)

Accounting record generation mode - singleEnd

The table, [SingleEnd mode with normal call termination \(page 197\)](#) summarizes the behavior of the accounting system in singleEnd mode.



SingleEnd mode with normal call termination

Provisioning singleEnd mode				As a result of generation mode and accounting records		
Charged end DNA (C)		Non-charged end DNA (D)		Egress protocol	Accounting record generation	
Collection reason	Frame Relay Accounting	Collection reason	Frame Relay Accounting		Location (reason)	Contents
b	off	n or b	off or on	off	C(b)	FP
b	on	n or b	off	off	C(b)	FP + LS
b	on	n or b	on	on	C(b)	FP + LS + RS
b	off	X or bX	off	off	C(b),D(X,bX)	FP
b	off	X or bX	on	off	C(b),D(X,bX)	FP at C, FP + LS at D
b	on	X or bX	off	off	C(b),D(X,bX)	FP + LS at C, FP at D
b	on	X or bX	on	off	C(b),D(X,bX)	FP + LS
X	off	n or b	off or on	off	C(X)	FP
X	on	n or b	off	off	C(X)	FP + LS
X	on	n or b	on	on	C(X)	FP + LS + RS
X	off	X or bX	off	off	C(X),D(X,bX)	FP
X	off	X or bX	on	off	C(X),D(X,bX)	FP at C, FP + LS at D
X	on	X or bX	off	off	C(X),D(X,bX)	FP + LS at C, FP at D
X	on	X or bX	on	off	C(X),D(X,bX)	FP + LS
bX	off	n or b	off or on	off	C(bX)	FP
bX	on	n or b	off	off	C(bX)	FP + LS
a = audit, b = billing, s = study, t = test, f = forced generation, n = no collection reasons set X means any combination of collection reasons: a, s, t, f C(b) means a Multiservice Switch node's accounting record is generated at the location of DNA-C with billing (b) indicated as a collection reason FP = Fixed Portion, LS = Local Snapshot, RS = Remote Snapshot In charged end mode, the Multiservice Switch node's accounting records produced are independent of whether C calls D or D calls C.						
(1 of 2)						



SingleEnd mode with normal call termination (continued)

Provisioning singleEnd mode				As a result of generation mode and accounting records		
Charged end DNA (C)		Non-charged end DNA (D)		Egress protocol	Accounting record generation	
Collection reason	Frame Relay Accounting	Collection reason	Frame Relay Accounting		Location (reason)	Contents
bX	on	n or b	on	on	C(bX)	FP + LS + RS
bX	off	X or bX	off	off	C(bX),D(X,bX)	FP
bX	off	X or bX	on	off	C(bX),D(X,bX)	FP at C, FP + LS at D
bX	on	X or bX	off	off	C(bX),D(X,bX)	FP + LS at C, FP at D
bX	on	X or bX	on	off	C(bX),D(X,bX)	FP + LS
n	off or on	n or b	off or on	off	no record generated	
n	off or on	X or bX	off	off	D(X, bX)	FP
n	off	X or bX	on	off	D(X, bX)	FP + LS
n	on	X or bX	on	on	D(X, bX)	FP + LS + RS
a = audit, b = billing, s = study, t = test, f = forced generation, n = no collection reasons set X means any combination of collection reasons: a, s, t, f C(b) means a Multiservice Switch node's accounting record is generated at the location of DNA-C with billing (b) indicated as a collection reason FP = Fixed Portion, LS = Local Snapshot, RS = Remote Snapshot In charged end mode, the Multiservice Switch node's accounting records produced are independent of whether C calls D or D calls C.						
(2 of 2)						

Accounting record generation mode - bothEnds

The table, [BothEnds mode with normal call termination \(page 199\)](#) summarizes the behavior of the accounting system in bothEnds mode.



BothEnds mode with normal call termination

Provisioning bothEnds mode		As a result of generation mode and accounting records			
Charged end DNA (C)		Non-charged end DNA (D)		Accounting record generation	
Collection reason	Frame Relay Accounting	Collection reason	Egress protocol	Location (reason)	Contents
b	off	n	off	C(b)	FP
b	on	n	off	C(b)	FP + LS
b	on	n	off	C(b)	FP + LS
b	off	b or X or bX	off	C(b), D	FP
b	off	b or X or bX	off	C(b), D	FP at C, FP + LS at D
b	on	b or X or bX	off	C(b), D	FP + LS at C, FP at D
b	on	b or X or bX	off	C(b), D	FP + LS
X	off	n	off	C(X)	FP
X	on	n	off	C(X)	FP + LS
X	on	n	on	C(X)	FP + LS + RS
X	off	b or X or bX	off	C(X), D	FP
X	off	b or X or bX	off	C(X), D	FP at C, FP + LS at D
X	on	b or X or bX	off	C(X), D	FP + LS at C, FP at D
X	on	b or X or bX	off	C(X), D	FP + LS
bX	off	n	off	C(bX)	FP
bX	on	n	off	C(bX)	FP + LS
bX	on	n	off	C(bX)	FP + LS

a = audit, b = billing, s = study, t = test, f = forced generation, n = no collection reasons set

X means any combination of collection reasons: a, s, t, f

C(b) means a Multiservice Switch node's accounting record is generated at the location of DNA-C with billing (b) indicated as a collection reason

FP = Fixed Portion, LS = Local Snapshot, RS = Remote Snapshot

In charged end mode, the Multiservice Switch node's accounting records produced are independent of whether C calls D or D calls C.

(1 of 2)



BothEnds mode with normal call termination (continued)

Provisioning bothEnds mode		As a result of generation mode and accounting records			
Charged end DNA (C)		Non-charged end DNA (D)		Accounting record generation	
Collection reason	Frame Relay Accounting	Collection reason	Egress protocol	Location (reason)	Contents
bX	off	b or X or bX	off	C(bX), D	FP
bX	off	b or X or bX	off	C(bX), D	FP at C, FP + LS at D
bX	on	b or X or bX	off	C(bX), D	FP + LS at C, FP at D
bX	on	b or X or bX	off	C(bX), D	FP + LS
n	off or on	n	off	no record generated	
n	off or on	b or X or bX	off	D	FP
n	off	b or X or bX	off	D	FP + LS
n	on	b or X or bX	on	D	FP + LS + RS
a = audit, b = billing, s = study, t = test, f = forced generation, n = no collection reasons set X means any combination of collection reasons: a, s, t, f C(b) means a Multiservice Switch node's accounting record is generated at the location of DNA-C with billing (b) indicated as a collection reason FP = Fixed Portion, LS = Local Snapshot, RS = Remote Snapshot In charged end mode, the Multiservice Switch node's accounting records produced are independent of whether C calls D or D calls C.					
(2 of 2)					

Frame relay billing policies

This section describes how data in the accounting records can be used to develop flat-rate and usage-sensitive billing policies for the frame relay service. This section also describes how data can be used to monitor service at a data link connection identifier (DLCI).

- [Flat-rate billing \(page 201\)](#)
- [Usage-sensitive billing \(page 202\)](#)
- [Quality of service parameters \(page 202\)](#)



Flat-rate billing

Flat-rate billing policies are developed by the service provider based on specific parameters and values provisioned by the end user and not on the amount of data delivered. This flexibility allows the service provider to develop a variety of flat-rate tariff bands based on the following elements:

- CIR, EIR values
- class-of-service parameters
- rate enforcement and rate adaptation parameters

The table, [Parameter and value effects on cost of service \(page 201\)](#) shows the effects on the cost of service for varying parameters and values.

Each time the parameters and values for the frame relay service are changed, a new accounting record is generated. The exception is changing the value of a correlation tag on an active frame relay permanent connection which does not result in the generating of a new accounting record.

Parameter and value effects on cost of service

Parameters and values	Cost of service
CIR, EIR values	
- higher	increase cost of service
- lower	decrease cost of service
class of service	
- high emission priority	increase cost of service
- low emission priority	decrease cost of service
- high discard priority	increase cost of service
- normal discard priority	decrease cost of service
rate enforcement	
- yes	decrease cost of service
- no	increase cost of service
rate adaptation	
- yes	decrease cost of service
- no	increase cost of service



Usage-sensitive billing

Usage-sensitive billing policies are developed by the service provider based on the amount of data delivered to the end user. The amount of data delivered to the end user can be based on the number of segments or frames delivered or on the number of kilobytes or megabytes of data.

For billing policies based on the number of bytes or segments/frames delivered, the following counts can be used:

- total egress bytes (or frames or segments)
- EIR egress bytes (or frames or segments)
- CIR egress bytes (or frames or segments) + EIR egress bytes (or frames or segments)

The table, [Resultant billing policies based on count types \(page 202\)](#) shows the billing policy that results from various count types.

Resultant billing policies based on count types

Count types	Resultant billing policy
Total egress bytes	Bill per kilobyte of data delivered to the end user
EIR bytes	Bill flat rate for CIR data and bill per kilobyte of EIR data delivered to the end user
CIR egress bytes + EIR egress bytes	Bill a specific rate for CIR data and bill a different rate for EIR traffic
Total egress frames or segments	Bill per number of frames or segments delivered to the end user
EIR egress frames or segments	Bill per number of frames or segments delivered to the end user based on EIR traffic
CIR egress frames or segments + EIR egress frames or segments	Bill a specific rate for CIR traffic and bill a different rate for EIR traffic

Quality of service parameters

Data in the accounting record can be used to validate quality-of-service (QOS) parameters. The QOS parameters are valid for the billing period captured in the accounting record. These parameters can serve as proof to the end user that the negotiated service (CIR, EIR, bandwidth, and so on) has been delivered. A possible list of quality of service parameters are as follows:

- ratio of calling total egress bytes to called total ingress bytes
- ratio of called total egress bytes to calling total ingress bytes



- ratio of calling total egress segments or frames to called total ingress segments or frames
- ratio of called total egress segments or frames to calling total ingress segments or frames
- calling total ingress bytes, calling discarded bytes, called egress bytes
- called total ingress bytes, called discarded bytes, calling egress bytes

The ratio of calling total egress bytes to called total ingress bytes determines the percentage of bytes that have been successfully delivered from the called end of the connection to the calling end.

The ratio of called total egress bytes to calling total ingress bytes determines the percentage of bytes that have been successfully delivered from the calling end of the connection to the called end.

The ratio of calling total egress segments or frames to called total ingress segments/frames determines the percentage of segments or frames that have been successfully delivered from the called end of the connection to the calling end.

The ratio of called total egress segments or frames to calling total ingress segments/frames determines the percentage of segments or frames that have been successfully delivered from the calling end of the connection to the called end.

The calling total ingress bytes, the calling discard bytes, and the called egress bytes can be used to compute the number of bytes lost in the network due to congestion (as opposed to bytes in invalid frames discarded at the service) in the calling-to-called direction.

The called total ingress bytes, the called discarded bytes, and the calling egress bytes can be used to compute the number of bytes lost in the network due to congestion (as opposed to bytes in invalid frames discarded at the service) in the called-to-calling direction.

Other possible quality of service parameters are as follows:

- Ratio of (calling or called) EIR egress bytes to (called or calling) EIR ingress bytes: determines the QOS parameter for the EIR traffic
- The CIR QOS parameters can be computed as follows:

$$QOS_{CIR} = (Egress\ Count_{TOTAL} - Egress\ Count_{EIR}) / (Ingress\ Count_{TOTAL} - Ingress\ Count_{EIR})$$



Accounting for BNX-lwf

For interworking between a Nortel Multiservice Switch 15000 or Multiservice Switch 20000 node and a network element running hyperstream software (BNX), existing DPRS accounting is provided for each data link connection identifier (DLCI). Accounting has the following characteristics:

- can be turned on or off on for each interface or for each DLCI
- is single-ended or double-ended

If the egress protocol for the connection is enabled, the accounting record generated at the master end (that is, the charged end) contains the local and remote accounting information, which is signaled using egress protocol. If the accounting is double ended (record generated at master & slave end), only the local information is contained in the accounting record.

The FrUni/FrNni with the BNX interworking function (BNX-lwf) DLCI provides the existing accounting functionality for the non-BNX-lwf DLCIs. If accounting is turned off at the interface level, it is turned off for both BNX-lwf and non-BNX-lwf DLCIs. Accounting can be independently turned on or off for each BNX-lwf DLCI. At the BNX-lwf DLCI, only the local accounting data is included in the accounting record for both single-ended or double-ended accounting configurations. This characteristic arises because the egress protocol is not supported.

An accounting record is generated for each connection at the following times, provided there has been traffic on that connection:

- expiry of the 12-hour timer
- on time-of-day accounting (TODA) intervals
- call clearing
- DLCI deletion
- change of the configured values for any critical attribute under the DLCI or the associated interface

For additional information about BNX-lwf accounting, see the following sections:

- [Mapping of attributes in the PVC accounting record \(page 204\)](#)
- [Attribute mapping for BNX-lwf DLCI accounting records \(page 209\)](#)

Mapping of attributes in the PVC accounting record

The table [Mapping Multiservice Switch 15000 or Multiservice Switch 20000 nodes to BNX accounting attributes for frame relay PVC DLCI \(page 205\)](#) lists the Nortel Multiservice Switch 15000 and Multiservice Switch 20000 node frame relay PVC accounting attributes and the corresponding BNX frame



relay PVC accounting attributes. Most attributes are available for both BNX-lwf and non-BNX-lwf accounting requirements; shaded table cells indicate accounting attributes that are specific to BNX-lwf only.

Mapping Multiservice Switch 15000 or Multiservice Switch 20000 nodes to BNX accounting attributes for frame relay PVC DLCI

Multiservice Switch 15000 or Multiservice Switch 20000 node attribute	BNX attribute	BNX size	Notes
Not applicable. Multiservice Switch nodes do not provide a header. These attributes can be generated off-switch.	Network element identifier	4 bytes	A 32-bit backplane hardware serial identifier. Provides a unique ID for this Baystream platform within the switched network.
	Version	1 byte	Billing entity: 7 for PVC.
	Time zone	1 byte	Time zone is the offset from Greenwich mean time.
	Usage record length	2 bytes	Number of bytes in appended billing records.
	Start timestamp	6 bytes	Time at which the billing information was gathered.
	Number of records	2 bytes	Number of appended billing records.
	End timestamp	6 bytes	Time at which the billing file was written to the PCMCIA SRAM media.
	Polling interval	2 bytes	Value assigned to the Site Manager PVC Update Interval parameter <i>wfFrSw- UsageUpdateInterval</i> .
startTime	Start timestamp	6 bytes	BNX: Time at which billing was started.
Multiservice Switch nodes always uses a 2 octet DLCI.	Address type	1 byte	BNX: DLCI address format (2, 3 or 4 byte DLCI).
	Circuit Number	1 byte	BNX: Not used. Definition: Baystream platform assigned number that identifies this VC.
Not present	Source IP Address	4 bytes	Multiservice Switch node: not present. BNX: local IP address associated with the VC.
(1 of 2)			



Mapping Multiservice Switch 15000 or Multiservice Switch 20000 nodes to BNX accounting attributes for frame relay PVC DLCI (continued)

Multiservice Switch 15000 or Multiservice Switch 20000 node attribute	BNX attribute	BNX size	Notes
callingDlci	Source DLCI	4 bytes	Multiservice Switch node: software supports a 2-byte DLCI, not 4-byte. For PQC-12, the hardware supports the 4-byte DLCI. BNX: local DLCI
Not present	Destination IP Address	4 bytes	Multiservice Switch node: not present. BNX: local IP address associated with the VC.
calledDlci	Destination DLCI	4 bytes	Multiservice Switch node: software supports a 2-byte DLCI, not 4-byte. For PQC-12, the hardware supports the 4-byte DLCI. BNX: local DLCI
callingIngressFrames calledIngressFrames callingEgressFrames calledEgressFrames	Sent PDUs	4 bytes	BNX: total number of FR frames transmitted over the VC since billing data was last reported,
callingIngressBytes calledIngressBytes callingEgressBytes calledEgressBytes	Sent octets	6 bytes	
(callingIngressFrames - callingEirIngressFrames - callingDiscardedFrames) (calledIngressFrames - calledEirIngressFrames - calledDiscardedFrames)	Sent priority PDUs	4 bytes	BNX: number of non-DE FR frames transmitted over this VC since billing was last reported. Multiservice Switch node provides the number of ingress discarded frames. Can subtract it from total to get the number of priority frames transmitted.
(callingIngressBytes - callingEirIngressBytes - callingDiscardedBytes) (calledIngressBytes - calledEirIngressBytes - calledDiscardedBytes)	Sent priority octets	6 bytes	
(2 of 2)			



Multiservice Switch 15000 and Multiservice Switch 20000 node frame relay PVC DLCI and BNX accounting fields

	Multiservice Switch node Field	BNX Field	BNX Size	Notes
H	N/A Multiservice Switch nodes do not provide a header. These fields can be generated off-switch.	Network element identifier	4 bytes	32 bit backplane hw serial identifier. Provides a unique ID for this Baystream platform within the switched network
H		Version	1 byte	Billing entity - 7 for PVC
H		Time zone	1 byte	Time zone is the offset from Greenwich mean time
H		Usage record length	2 bytes	# bytes in appended billing records
H		Start timestamp	6 bytes	Time at which the billing information was gathered
H		Number of records	2 bytes	Number of appended billing records
H		End timestamp	6 bytes	Time at which the billing file was written to the PCMCIA SRAM media
H		Polling interval	2 bytes	Value assigned to the Site Mgr PVC Update Interval parameter wfFrSwUsageUpdateInterval
P	startTime	Start timestamp	6 bytes	BNX: Time at which billing was started
P	Multiservice Switch node always uses a 2 octet DLCI for now	Address type	1 byte	BNX : DLCI address format (2, 3 or 4 byte DLCI)
P		Circuit Number	1 byte	BNX : Not being used. Defn - Baystream platform assigned number that IDs this vc
P	Not present	Source IP Address	4 bytes	Multiservice Switch node: not present BNX : local IP address associated with the vc
P	callingDlci	Source DLCI	4 bytes	Multiservice Switch node: Software only supports a 2 byte DLCI, not 4. In PQC-12 the hardware supports the 4 byte DLCI. BNX : local DLCI
(1 of 3)				



Multiservice Switch 15000 and Multiservice Switch 20000 node frame relay PVC DLCI and BNX accounting fields (continued)

	Multiservice Switch node Field	BNX Field	BNX Size	Notes
P	Not present	Destination IP Address	4 bytes	Multiservice Switch node: not present BNX : remote IP address associated with the vc
P	calledDlci	Destination DLCI	4 bytes	Multiservice Switch node: Software only supports a 2 byte DLCI, not 4. In PQC-12 the hardware supports the 4 byte DLCI. BNX : remote DLCI
P	callingIngressFrames calledIngressFrames callingEgressFrames calledEgressFrames	Sent PDUs	4 bytes	BNX : total number of FR frames transmitted over the vc since billing data was last reported
P	callingIngressBytes calledIngressBytes callingEgressBytes calledEgressBytes	Sent octets	6 bytes	
(2 of 3)				



Multiservice Switch 15000 and Multiservice Switch 20000 node frame relay PVC DLCI and BNX accounting fields (continued)

	Multiservice Switch node Field	BNX Field	BNX Size	Notes
P	(callingIngressFrames - callingEirIngressFrames - callingDiscardedFrames) (calledIngressFrames - calledEirIngressFrames - calledDiscardedFrames)	Sent priority PDUs	4 bytes	BNX: number of non-DE FR frames transmitted over this vc since billing was last reported Multiservice Switch node provides the ingress discarded frames. Can subtract it from total to get the # of priority frames transmitted.
P	(callingIngressBytes - callingEirIngressBytes - callingDiscardedBytes) (calledIngressBytes - calledEirIngressBytes - calledDiscardedBytes)	Sent priority octets	6 bytes	
(3 of 3)				

Attribute mapping for BNX-lwf DLCI accounting records

The table [Multiservice Switch 15000 and Multiservice Switch 20000 node frame relay accounting attributes for BNX-lwf DLCI accounting \(page 210\)](#) lists attributes for the Nortel Multiservice Switch 15000 and Multiservice Switch 20000 node frame relay accounting record and identifies availability for the BNX-lwf DLCI. Note the following:

- There is no concept of calling end versus called end for the BNX-lwf DLCI, and the BNX-lwf DLCI does not have knowledge of the other end
- All segment counts are the same as the frame counts since frame:segment = 1:1 (no segmentation)

The main difference between Nortel Multiservice Switch node accounting and BNX accounting is the time interval granularity:

- the accounting time granularity is 1 sec for Multiservice Switch nodes and 0.01 sec for BNX
- the accounting time granularity is 1 sec for the Multiservice Switch node's frame relay BNX-lwf DLCI



Multiservice Switch 15000 and Multiservice Switch 20000 node frame relay accounting attributes for BNX-lwf DLCI accounting

Multiservice Switch 15000 and Multiservice Switch 20000 node attributes	Available/not available for the BNX-lwf DLCI	Notes
recordType	available	This attribute identifies the record: first (record for the call), intermediate, and last.
recordIdentifier	available	This attribute is set when the record is formatted. Its value depends on the type of record.
callType	not available not needed - DPRS specific	This attribute defines the type of the call for which this accounting record is issued: non gateway call or gateway call.
callReferenceNumber	not needed - DPRS specific	This attribute contains the value of the call reference number assigned to each new call. Accounting records are uniquely defined by a call reference number on each function processor. In double ended accounting the <i>callReferenceNumber</i> attribute is the same for both accounting records. Call reference number is unique only on function processors.
callingNodeIdentifier	available	This attribute identifies the calling (originating) Multiservice Switch node by a unique number assigned to each node through the <i>ModuleData</i> component.
callingCustomerIdentifier	available	This attribute contains the number assigned in each service to a customer and identifies the customer at the calling end of the call in the customer network management.
callingPortIdentifier	available	This attribute contains the function processor number used by the service issuing the accounting record for the calling (originating) part of the call. The pe, port, and pi numbers are encoded as 5, 5, and 6 bits. The pi number, if from a Multiservice Switch node, is always 0.
callingNpi	available	This attribute contains the calling numbering plan indicator (NPI) of the calling address.
(1 of 11)		



Multiservice Switch 15000 and Multiservice Switch 20000 node frame relay accounting attributes for BNX-lwf DLCI accounting (continued)

Multiservice Switch 15000 and Multiservice Switch 20000 node attributes	Available/not available for the BNX-lwf DLCI	Notes
callingAddress	available	This attribute contains the calling data network address (DNA) of the interface. The <i>callingAddress</i> attribute includes DNIC (first 4 digits) for X.121, and CC (Country Code) for the E.164 address of the calling end. The maximum length is 15 digits.
callingDlci	available	This attribute contains the calling DLCI number of the interface. The DLCI is sometimes called logical channel number (LCN). This number is assigned to the call request at the calling end.
calledNodeIdentifier	not available at called data end	Same as the calling end.
calledCustomerIdentifier	not available at called data end	Same as the calling end.
calledPortIdentifier	not available at called data end	Same as the calling end.
calledNpi	not available at called data end	Same as the calling end.
calledAddress	not available at called data end	Same as the calling end.
calledDlci	available	Same as for calling end.
segmentSizeIndex	not available - DPRS specific	This attribute is the segment size index configured on the module. This attribute is used to calculate the number of segments sent and received in the accounting record. Two to the power of the value for the <i>segmentSizeIndex</i> attribute is equal to the segment size (in bytes) configured under the <i>VirtualCircuitSystem</i> component. For example, if the segment size is 128, then segment size index is 7.
clearCause	not available - DPRS specific	This attribute is the clear cause signaled at the end of a connection. The <i>clearCause</i> attribute is set to H.b0 for intermediate records.
(2 of 11)		



Multiservice Switch 15000 and Multiservice Switch 20000 node frame relay accounting attributes for BNX-lwf DLCI accounting (continued)

Multiservice Switch 15000 and Multiservice Switch 20000 node attributes	Available/not available for the BNX-lwf DLCI	Notes
diagnosticCode	not available - DPRS specific	This attribute is the diagnostic code signalled at the end of a connection. The <i>diagnosticCode</i> attribute is set to 0 for intermediate records.
callingCir	available	This attribute contains the calling end configured CIR.
callingEir	available	This attribute contains the calling end configured EIR.
callingServiceType	available	This attribute indicates the service supported by the port used by the calling end.
callingCollectReason	available	This attribute indicates the reason for the accounting collection of the calling end (can be any combination of Audit, Study, Test, Billing, Force).
callingDse	available	This attribute defines the calling end data service exchange configured in Dna as an arbitrary number in the <i>serviceExchange</i> attribute.
callingRateAdaptation	available	provisioned: rate adaptation configured on the calling end of the call. occurred: rate adaptation occurred on the calling end of the call. enforced: rate enforcement configured on the calling end of the call
callingClassOfService	not available	emissionPriority - emission priority configured. For BNX, priority is per packet and does not apply to the BNX-lwf DLCI.
calledCir	not available at called data end	Same as for calling end.
calledEir	not available at called data end	Same as for calling end.
(3 of 11)		



Multiservice Switch 15000 and Multiservice Switch 20000 node frame relay accounting attributes for BNX-lwf DLCI accounting (continued)

Multiservice Switch 15000 and Multiservice Switch 20000 node attributes	Available/not available for the BNX-lwf DLCI	Notes
calledServiceType	not available at called data end	Same as for calling end.
calledCollectReason	not available at called data end	Same as for calling end.
calledDse	not available at called data end	Same as for calling end.
calledRateAdaptation	not available at called data end	Same as for calling end.
calledClassOfService	not available at called data end	Same as for calling end.
flags	individual flags are not available since they are DPRS specific or do not apply to the BNX-lwf DLCI	This attribute includes values for the call characteristics. • calledEndCharging: called end is charging • ~calledEndCharging: calling end is charging • highPriority: call is high priority • pvcType: call is pvc • ~pvcType: call is svc • hunted: call was hunted • redirected: call was redirected • calledGenerated: called end generated accounting record • ~calledGenerated: calling end generated accounting record • calledCleared: call was cleared by called end • ~calledCleared: call was cleared by calling end
(4 of 11)		



Multiservice Switch 15000 and Multiservice Switch 20000 node frame relay accounting attributes for BNX-lwf DLCI accounting (continued)

Multiservice Switch 15000 and Multiservice Switch 20000 node attributes	Available/not available for the BNX-lwf DLCI	Notes
flags (continued)		• unique: accounting record is unique • reachedDataTransfer: call reached data transfer • collReasonValid: collection reasons are valid • callingUni: interface is callingUni • ~callingUni: interface is callingNni • calledUni: interface is calledUni • ~calledUni: interface is calledNni • spareOne: not used • spareTwo: not used • callingAbit: called Abit set • calledAbit: calling Abit set • userDataInCallReq: user data present in call request
startTime	available	This attribute represents the start time in the accounting record. For SVCs, the attribute records the time at which the call request packet is created. For PVC, the attribute is set by the Master VC when the master call is formatted. On the subsequent accounting record, the start time corresponds to the end time of the previous record. If TODA is used, start time (and end time) correspond to times set in the schedule (set in <i>Col/accounting collectionTimes</i> attribute). Set by the BNX-lwf DLCI.
(5 of 11)		



Multiservice Switch 15000 and Multiservice Switch 20000 node frame relay accounting attributes for BNX-lwf DLCI accounting (continued)

Multiservice Switch 15000 and Multiservice Switch 20000 node attributes	Available/not available for the BNX-lwf DLCI	Notes
endTime	available	This attribute represents the end time in the accounting record. If TODA is used, the end time in non-final records is the time of the TODA period. For 12-hour accounting, the time on non-final records is the time when the 12-hour timer expires at the charged end. On final accounting records, the end time is the time when the disconnect packet is first received. If TODA is used, end time (and start time) correspond to times set in the schedule (set in <i>Collector collectionTimes</i> attribute). Set by the BNX-lwf DLCI.
elapsedTime	available	This attribute represents the duration of the call in 0.1 second increments. This time interval is started when call request is processed. If call does not reach data transfer, then elapsed time is set to 0. Elapsed time represents the true measure of the accounting period. It is measured independently of start time and end time. Since start time and end time can be changed externally by adjusting the clock, elapsed time can be shorter or longer than the difference between start time and end time. This difference can show as much as 60 seconds on accounting records for the virtual circuits which are idle (carry no traffic) and when TODA is used (in which case it is caused by the internal TODA detection mechanism). When a 12 hour accounting period is used, elapsed time is exact or differs by a few seconds from a 12 hour duration.
callingEgressBytes	available	This attribute contains the total number of bytes received at the calling end.
callingIngressBytes	available	This attribute contains the total number of bytes sent at the calling end.
(6 of 11)		



Multiservice Switch 15000 and Multiservice Switch 20000 node frame relay accounting attributes for BNX-lwf DLCI accounting (continued)

Multiservice Switch 15000 and Multiservice Switch 20000 node attributes	Available/not available for the BNX-lwf DLCI	Notes
callingEgressFrames	available	This attribute contains the total number of frames received at the calling end.
callingIngressFrames	available	This attribute contains the total number of frames sent at the calling end.
callingEgressSegments	not available - DPRS specific	This attribute contains the total number of segments received at the calling end.
callingIngressSegments	not available - DPRS specific	This attribute contains the total number of segments sent at the calling end.
callingEirEgressBytes	available	This attribute contains the number of EIR bytes received from the subnet at the calling end
callingEirIngressBytes	available	This attribute contains the number of EIR bytes sent to the subnet at the calling end.
callingEirEgressFrames	available	This attribute contains the number of EIR Frames received from the subnet at the calling end.
callingEirIngressFrames	available	This attribute contains the number of EIR Frames sent to the subnet at the calling end.
callingEirEgressSegments	not available - DPRS specific	This attribute contains the number of EIR segments received from the subnet at the calling end.
callingEirIngressSegments	not available - DPRS specific	This attribute contains the number of EIR segments sent to the subnet at the calling end
callingDiscardedBytes	available	This attribute contains the number of bytes discarded at the calling end due to rate enforcement when no rate adaptation occurred.
callingDiscardedFrames	available	This attribute contains the number of frames discarded at the calling end due to rate enforcement when no rate adaptation occurred.
(7 of 11)		



Multiservice Switch 15000 and Multiservice Switch 20000 node frame relay accounting attributes for BNX-lwf DLCI accounting (continued)

Multiservice Switch 15000 and Multiservice Switch 20000 node attributes	Available/not available for the BNX-lwf DLCI	Notes
callingDiscardedSegments	not available - DPRS specific	This attribute contains the number of segments discarded at the calling end due to rate enforcement when no rate adaptation occurred.
calledEgressBytes	not available at called data end	This attribute contains the total number of bytes received at the called end
calledIngressBytes	not available at called data end	This attribute contains the total number of bytes sent at the called end.
calledEgressFrames	not available at called data end	This attribute contains the total number of frames received at the called end.
calledIngressFrames	not available at called data end	This attribute contains the total number of frames sent at the called end.
calledEgressSegments	not available at called data end	This attribute contains the total number of segments received at the called end.
calledIngressSegments	not available at called data end	This attribute contains the total number of segments sent at the called end.
calledEirEgressBytes	not available at called data end	This attribute contains the number of EIR bytes received from the subnet at the called end
calledEirIngressBytes	not available at called data end	This attribute contains the number of EIR bytes sent to the subnet at the called end
calledEirEgressFrames	not available at called data end	This attribute contains the number of EIR Frames received from the subnet at the called end
calledEirIngressFrames	not available at called data end	This attribute contains the number of EIR Frames sent to the subnet at the called end
calledEirEgressSegments	not available at called data end	This attribute contains the number of EIR segments received from the subnet at the called end
calledEirIngressSegments	not available at called data end	This attribute contains the number of EIR segments sent to the subnet at the called end
(8 of 11)		



Multiservice Switch 15000 and Multiservice Switch 20000 node frame relay accounting attributes for BNX-lwf DLCI accounting (continued)

Multiservice Switch 15000 and Multiservice Switch 20000 node attributes	Available/not available for the BNX-lwf DLCI	Notes
calledDiscardedBytes	not available at called data end	This attribute contains the number of bytes discarded at the called end due to rate enforcement when no rate adaptation occurred.
calledDiscardedFrames	not available at called data end	This attribute contains the number of frames discarded at the called end due to rate enforcement when no rate adaptation occurred.
calledDiscardedSegments	not available at called data end	This attribute contains the number of segments discarded at the called end due to rate enforcement when no rate adaptation occurred.
callingTotalPwmBytes		This attribute contains the calling end peak water mark (PWM) for the total number of ingress bytes.
callingEirPwmBytes		This attribute contains the calling end peak water mark for the number of ingress bytes carried on frames with DE-bit = 1.
callingCirPwmBytes		This attribute contains the calling end peak water mark for the number of ingress bytes carried on frames with DE-bit = 0.
calledTotalPwmBytes	available	This attribute contains the called end peak water mark for the total number of ingress bytes.
calledEirPwmBytes	available	This attribute contains the called end peak water mark for the number of ingress bytes carried on frames with DE-bit = 1.
calledCirPwmBytes	available	This attribute contains the called end peak water mark for the number of ingress bytes carried on frames with DE-bit equals 0.
(9 of 11)		



Multiservice Switch 15000 and Multiservice Switch 20000 node frame relay accounting attributes for BNX-lwf DLCI accounting (continued)

Multiservice Switch 15000 and Multiservice Switch 20000 node attributes	Available/not available for the BNX-lwf DLCI	Notes
callingTotalPwmInterval-Number	available	This attribute contains the number of the PWM interval in which the PWM occurred and is not part of accounting record if TODA is not running (that is, the <i>Collector collectionTimes</i> attribute for accounting is not set.) The value is relative to <i>Collector collectionTimes</i> attribute together with <i>Collector peakWaterMarkInterval</i> attribute. For example if the <i>collectionTimes</i> attribute is set to 16:00 and the value of the <i>callingTotalPwmInterval-Number</i> attribute is reported as 3 and the <i>Collector peakWaterMarkInterval</i> attribute is set to 5 minutes, then peak traffic occurred between 16:15 and 16:20.
callingEirPwmIntervalNumber	available	This attribute contains the number of the PWM interval in which the PWM occurred and is not part of accounting record if Time of Day Accounting (TODA) is not running. (The <i>collectionTimes</i> attribute under the <i>Collector</i> component for accounting is not set.)
callingCirPwmIntervalNumber	available	This attribute contains the number of the PWM interval in which the PWM occurred and is not part of accounting record if TODA is not running. (The <i>collectionTimes</i> attribute under the <i>Collector</i> component for accounting is not set.)
calledTotalPwmInterval-Number	not available at called data end	This attribute contains the number of the PWM interval in which the PWM occurred and is not part of accounting record if TODA is not running. (The <i>collectionTimes</i> attribute under the <i>Collector</i> component for accounting is not set.)
(10 of 11)		



Multiservice Switch 15000 and Multiservice Switch 20000 node frame relay accounting attributes for BNX-lwf DLCI accounting (continued)

Multiservice Switch 15000 and Multiservice Switch 20000 node attributes	Available/not available for the BNX-lwf DLCI	Notes
calledEirPwmIntervalNumber	not available at called data end	This attribute contains the number of the PWM interval in which the PWM occurred and is not part of accounting record if TODA is not running. (The <i>collectionTimes</i> attribute under the <i>Collector</i> component for accounting is not set.)
calledCirPwmIntervalNumber	not available at called data end	This attribute contains the number of the PWM interval in which the PWM occurred and is not part of accounting record if TODA is not running. (The <i>collectionTimes</i> attribute under the <i>Collector</i> component for accounting is not set.)
originalNpi	not needed	This attribute contains numbering plan indicator for originally called address.
originalAddress	not needed	This attribute contains originally called Data Network Address of the interface if call was redirected.
originalServiceType	not needed	This attribute indicates the service supported by the port used by the calling end. This number is configured on each interface Dna as accountClass.
callingTransferPriority	available	This attribute indicates the transfer priority used in the direction towards the network. For BNX-lwf DLCI this is the priority level.
calledTransferPriority	not available at called data end	This attribute indicates the transfer priority used in the direction from the network.
circuitId	available	Configured circuit identifier to uniquely identify the connection. Configured as the attribute correlationTag
callingIpAddress	available	This attribute contains the calling IP address for the BNX-lwf DLCI.
calledIpAddress	available	This attribute contains the called IP address for the BNX-lwf DLCI.
(11 of 11)		



Troubleshooting

The following sections provide guidelines for resolving some of the problems you may encounter when using frame relay. Your network problems may also be associated with non-frame relay component.

In addition, Preside Multiservice Data Manager offers a comprehensive set of fault management tools to help you detect and isolate problems. When you use Preside Multiservice Data Manager to troubleshoot, use the Integrated Alarms Display to collect alarms and the Integrated Command Console to access frame relay NNI operational attributes.

For information about Preside Multiservice Data Manager tools, see 241-6001-023 *Nortel Multiservice Data Manager Configuration Tools*.

Navigation

- [OSI state information for FR UNI components \(page 222\)](#)
- [OSI state information for FR NNI components \(page 223\)](#)
- [Alarms \(page 225\)](#)
- [Pinging a virtual circuit \(page 225\)](#)
- [Tracing data frames \(page 226\)](#)
- [Troubleshooting connection failure \(page 226\)](#)
- [Troubleshooting traffic problems \(page 230\)](#)
- [Troubleshooting accounting problems \(page 231\)](#)
- [Troubleshooting hardware resources \(page 232\)](#)
- [Clear cause codes \(page 233\)](#)
- [Diagnostic codes \(page 237\)](#)
- [FR UNI frame link monitor for E1 frame relay \(page 241\)](#)
- [BNX-lwf performance monitoring and fault management \(page 242\)](#)
- [Frame relay ISDN switched access monitoring and troubleshooting \(page 254\)](#)



OSI state information for FR UNI components

Nortel Multiservice Switch nodes use component state definitions according to the Open Systems Interconnection (OSI) standards. For a description of OSI state and status attributes, see NN10600-520 *Nortel Multiservice Switch 7400/15000/20000 Fault and Performance Management: Troubleshooting*.

The following tables summarize the OSI state attributes specific to the Multiservice Switch nodes' FR UNI service:

- [OSI state information for the FR UNI \(page 222\)](#)
- [OSI state information for the LMI \(page 222\)](#)
- [OSI state information for a DLCI \(page 223\)](#)

OSI state information for the FR UNI

State	Meaning
adminState: unlocked operationalState: disabled usageState: idle	The FR UNI is waiting for the access facility to become available.
adminState: unlocked operationalState: enabled usageState: idle	The FR UNI is ready to provide service, but no user DLCIs have been defined on the interface.
adminState: unlocked operationalState: enabled usageState: active	The FR UNI is in service and supports at least one user DLCI.
adminState: locked operationalState: disabled usageState: idle	An operator lock command is in effect. There are no operational user DLCIs on the interface.
On a 4-port DS3 channelized frame relay FP, provisioning the timeslot of the associated frame interface to the value of none and not locking the FrUni and channel result in the OSI state of unlocked, enabled, idle.	

OSI state information for the LMI

State	Meaning
adminState: unlocked operationalState: disabled usageState: idle	The <i>Lmi</i> component is waiting for the FR UNI to become available. An access facility may have failed, or the <i>FrUni</i> component may be locked.
adminState: unlocked operationalState: enabled usageState: busy	The <i>Lmi</i> component is in service.



OSI state information for a DLCI

State	Meaning
adminState: locked operationalState: disabled usageState: idle	An operator lock command is in effect. The DLCI cannot provide service.
adminState: unlocked operationalState: disabled usageState: idle	Any one of the following circumstances can cause the <i>Dlci</i> component to be out of operation: • an access facility failure • a lock command of the <i>FrUni</i> component • a failure at the LMI (for PVCs) or LAPF (for SVCs) • a broken connection across the network The <i>Dlci</i> component is not operational until the following conditions are satisfied: • the interface becomes available • the LMI (for PVCs) or LAPF (for SVCs) becomes operational • the connection is established
adminState: unlocked operationalState: enabled usageState: busy	The DLCI is operational and in service.
adminState: locked operationalState: disabled usageState: idle availabilityStatus: failed	Underlying hardware resources are unavailable. See Troubleshooting hardware resources (page 232) .

OSI state information for FR NNI components

The following tables summarize the Nortel Multiservice Switch nodes' OSI state attributes specific to the FR NNI service:

- [OSI state information for the FR NNI \(page 224\)](#)
- [OSI state information for the LMI \(page 222\)](#)
- [OSI state information for a DLCI \(page 223\)](#)



OSI state information for the FR NNI

State	Meaning
adminState: unlocked operationalState: disabled usageState: idle	The FR NNI is waiting for the access facility to become available.
adminState: unlocked operationalState: enabled usageState: idle	The FR NNI is ready to provide service, but no user DLCIs have been defined on the interface.
adminState: unlocked operationalState: enabled usageState: active	The FR NNI is in service and supports at least one user DLCI.
adminState: locked operationalState: disabled usageState: idle	An operator lock command is in effect. There are no operational user DLCIs on the FR NNI.
On a 4-port DS3 channelized frame relay FP, provisioning the timeslot of the associated frame interface to the value of none and not locking the FrNni and channel result in the OSI state of unlocked, enabled, and idle.	

OSI state information for the LMI

State	Meaning
adminState: unlocked operationalState: disabled usageState: idle	The <i>Lmi</i> component is waiting for the FR NNI to become available. An access facility may have failed, or the <i>FrNni</i> component may be locked.
adminState: unlocked operationalState: enabled usageState: busy	The <i>Lmi</i> component is in service.



OSI state information for a DLCI

State	Meaning
adminState: locked operationalState: disabled usageState: idle	An operator lock command is in effect. The DLCI cannot provide service.
adminState: unlocked operationalState: disabled usageState: idle	Any one of the following circumstances can cause the <i>Dlci</i> component to be out of operation: • an access facility failure • a lock command of the <i>FrNni</i> component • a failure at the LMI • a broken connection across the network The <i>Dlci</i> component is not operational until the following conditions are satisfied: • the interface becomes available • the LMI becomes operational • the connection is established
adminState: unlocked operationalState: enabled usageState: busy	The DLCI is operational and in service.
adminState: locked operationalState: disabled usageState: idle availabilityStatus: failed	Underlying hardware resources are unavailable. See Troubleshooting hardware resources (page 232) .

Alarms

Once your frame relay application is operational, alarms can appear at the user interface to indicate faults or failure conditions on the node.

Nortel Multiservice Switch node components generate alarms asynchronously. The alarms contain detailed information to help you monitor your network. For detailed descriptions of individual alarms, see NN10600-500 *Nortel Multiservice Switch 6400/7400/15000/20000 Alarms Reference*.

Pinging a virtual circuit

You can also use the VC ping command to determine the path of a particular frame relay VC and its round trip delay. See NN10600-520 *Nortel Multiservice Switch 7400/15000/20000 Fault and Performance Management: Troubleshooting* for more information about this diagnostic tool.



Tracing data frames

PTS is a user-controlled, diagnostic tool that monitors incoming and outgoing data frames. This tool traces data from a centralized location, in real time, without disconnecting the customer. The PTS can conduct multiple trace sessions simultaneously.

The PTS copies data from frames defined for trace, appends a trace header to the copied data, and sends the traced data to a receiver over a virtual circuit.

For more information, refer to NN10600-510 *Nortel Multiservice Switch 7400/15000/20000 Operations: Trace System*.

Troubleshooting connection failure

The following tables provides guidelines for resolving some of the problems you may encounter with frame relay connections.

- [Troubleshooting FR UNI PVC connection failure \(page 226\)](#)
- [Troubleshooting FR UNI SVC connection failure \(page 227\)](#)
- [Troubleshooting FR NNI connection failure \(page 229\)](#)

Troubleshooting FR UNI PVC connection failure

Problem	Possible causes	Corrective measures
Link does not come up.	Operator error	Check the physical layer attributes. Make sure essential components are not locked.
	Bad cable	Use port and line tests to test cables. Replace bad cables.
	Incorrect configuration	Ensure that both ends of the connection have compatible configuration data.
LMI is not operating properly.	LMI configuration at both ends of the connection do not match.	Ensure that the user and network are running the same LMI protocol.
	LMI <i>side</i> attributes for the network and access sides are mismatched.	If you are using a user/network designation, ensure that one side is configured as user and the other side is configured as network. Otherwise, ensure that the LMI <i>side</i> attribute on both the network and access devices is set to both.
Number of LMI status and full status messages does not match the number of status and full status enquiries.	Low resource availability.	Re-engineer the function processor to use fewer resources (reduce the number of configured DLCIs).
(1 of 2)		



Troubleshooting FR UNI PVC connection failure (continued)

Problem	Possible causes	Corrective measures
The PVC connection does not come up.	No connectivity between nodes. Incorrect service data for the PVC. The Call Server Resource Module (CSRM) is not operational or is unreachable. The call router (CR) is not configured with the slave's prefix (Multiservice Switch node-only network).	Ensure that the necessary Multiservice Switch trunks are up. Verify that the following parameters are compatible at both ends of the connection: - throughput and packet size (DPN) - transfer priority - DNA - DLCI - master/slave assignments Verify connectivity with the CSRM. Ensure that all other nodes can reach all the Multiservice Switch nodes directly connected to the CRSM. Configure the slave's prefix under the CR. (For details, see NTP NN10600-405 <i>Nortel Multiservice Switch 7400/15000/20000 Operations: Call Server.</i>)
The PVC is up, but not active.	The DLCI configured on the remote network is not present on the local network.	Configure a PVC on the local network with the missing DLCI.
(2 of 2)		

Troubleshooting FR UNI SVC connection failure

Problem	Possible causes	Corrective measures
Link does not come up.	Operator error Bad cable Incorrect configuration	Check the physical layer attributes. Make sure essential components are not locked. Use port and line tests to test cables. Replace bad cables. Ensure that both ends of the connection have compatible configuration data.
The <i>FrUni Signaling</i> component is not visible in the operational view.	The <i>Signaling</i> component is not configured.	Configure the <i>FrUni Signaling</i> component.
(1 of 2)		



Troubleshooting FR UNI SVC connection failure (continued)

Problem	Possible causes	Corrective measures
The SVC data link does not enter informationTransfer state, or an alarm issues to indicate that the SVC data link did not establish.	The user side does not support SVC calls.	Ensure that the attached device supports SVC calls and that signaling is enabled.
	SVC data link not properly configured.	Ensure that both ends of the link have compatible SVC data link configuration data and timer values.
The SVC call cannot establish.	The <i>FrUni Signaling</i> component is not configured properly.	Ensure that the <i>highestPvcDlci</i> attribute value is sufficient to allow for SVC DLCI assignment. If the interface uses default values for call setup, ensure that the default CIR, Bc, and Be are set to non-zero values.
	Maximum data link frame size at the user side is too small.	Increase the maximum data link frame size at the user side.
	Quality of service is not available.	Ensure that the values of the <i>maximumAggregateSvcCir</i> and <i>maximumAggregateSvcEir</i> attributes at the network and user side are sufficient to support the requested quality of service.
	No connectivity between nodes.	Ensure that the necessary Multiservice Switch trunks are up.
	The Call Server Resource Module (CSRM) is not operational or is unreachable.	Verify connectivity with the CSRM. Ensure that all other nodes can reach all the Multiservice Switch nodes directly connected to the CRSM.
	The call router (CR) is not configured with the slave's prefix (Multiservice Switch node-only network).	Configure the slave's prefix under the CR. (For details, see NTP NN10600-405 <i>Nortel Multiservice Switch 7400/15000/20000 Operations: Call Server</i> .)
The PVC is up, but not active.	The DLCI configured on the remote network is not present on the local network.	Configure a PVC on the local network with the missing DLCI.
(2 of 2)		



Troubleshooting FR NNI connection failure

Problem	Possible causes	Corrective measures
Link does not come up.	Operator error Bad cable Incorrect configuration	Check the physical layer attributes. Make sure essential components are not locked. Use port and line tests to test cables. Replace bad cables. Ensure that both ends of the connection have compatible configuration data.
LMI is not operating properly. LMI does not run bidirectionally when connected to a FR-ATM interface.	LMI configuration at both ends of the connection do not match. LMI procedures on the FR-ATM interface are not configured to run in both directions.	Ensure that both networks are running the same LMI protocol. Ensure that the <i>side</i> attribute under the <i>FrAtmLmi</i> component is set to both.
Number of LMI status and full status messages does not match the number of status and full status enquiries.	Low resource availability.	Re-engineer the function processor to use fewer resources (reduce the number of configured DLCIs).
(1 of 2)		



Troubleshooting FR NNI connection failure (continued)

Problem	Possible causes	Corrective measures
The PVC connection does not come up.	No connectivity between nodes. Incorrect service data for the PVC. The Call Server Resource Module (CSRM) is not operational or is unreachable. The call router (CR) is not configured with the slave's prefix (Multiservice Switch node-only network).	Ensure that the necessary Multiservice Switch trunks are up. Verify that the following parameters are compatible at both ends of the connection: - throughput and packet size (DPN) - transfer priority - DNA - DLCI - master/slave assignments Verify connectivity with the CSRM. Ensure that all other nodes can reach all the Multiservice Switch nodes directly connected to the CRSM. Configure the slave's prefix under the CR. (For details, see NTP NN10600-405 <i>Nortel Multiservice Switch 7400/15000/20000 Operations: Call Server.</i>)
The PVC is up, but not active.	The DLCI configured on the remote network is not present on the local network.	Configure a PVC on the local network with the missing DLCI.
(2 of 2)		

Troubleshooting traffic problems

The table below provides guidelines for resolving some of the problems you may encounter with frame relay traffic management.

Troubleshooting traffic problems

Problem	Possible causes	Corrective measures
The <i>FrUni</i> or <i>FrNni</i> component generates an engineering alarm.	Configured DLCIs are consuming too much memory.	Reduce the number of configured DLCIs.
The PVC is up, but all frames discard.	PVC status is inactive.	Check the PVC A-bit=0 causes to and from the interface.
The PVC is up and active, but only frames with the DE =0 transmit. Frames with DE = 1 discard.	Excess burst size (Be) for the connection is set to 0.	Configure the excess burst size (Be) to a non-zero value.
(1 of 2)		



Troubleshooting traffic problems (continued)

Problem	Possible causes	Corrective measures
Only partial data gets through (frame discard occurs).	Connection traffic parameters are not configured to accommodate your network traffic requirements. Check the • committed information rate (CIR) • committed burst size (Bc) • excess burst size (Be)	The preceding network is altering traffic characteristics. Turn rate enforcement off or adjust burst size (Bc) to account for traffic characteristics.
The number of calling EIR egress bytes for a connection is not 0.	Either the user or the network is tagging frames from the calling end as discard eligible (DE).	Determine why the excess burst size (Be) is configured at the calling end.
The number of calling EIR egress bytes for a connection is always 0.	The excess burst size (Be) at the calling end is 0.	Configure the excess burst size (Be) at the calling end to a non-zero value.
The number of called EIR egress bytes for a connection is not 0.	Either the user or the network is tagging frames from the called end as discard eligible (DE).	Determine why the excess burst size (Be) is configured at the called end.
The number of called EIR egress bytes for a connection is always 0.	The excess burst size (Be) at the called end is 0.	Configure the excess burst size (Be) at the called end to a non-zero value.
(2 of 2)		

Troubleshooting accounting problems

The table below provides guidelines for resolving some of the problems you may encounter with frame relay accounting. For detailed information on frame relay accounting, see [Frame relay accounting \(page 177\)](#) and NN10600-560 *Nortel Multiservice Switch 7400/15000/20000 Accounting*.



Troubleshooting accounting problems

Problem	Possible causes	Corrective measures
The API or published record format does not contain the frameRelay service fields.	Egress accounting is not enabled.	Turn egress accounting on under the <i>Dna</i> component (set the <i>egressAccounting</i> attribute to on).
When double-ended accounting is enabled, the number of called egress frames is less than the number of calling ingress frames.	The configured CIR/EIR values are too low, or the network is congested.	Ensure that the remote link is up. If the link is up, increase the Multiservice Switch trunk bandwidth or increase the user rate.
The PVC is up, but the number of called egress frames is 0 when double-ended accounting is enabled.	The PVC is inactive.	Check the link and LMI status of the remote PVC.
When the average traffic measured at the frame relay DTE is less than the peak water mark (PWM) CIR, frames discard during rate enforcement.	The DTE measures the PWM CIR over a 5 minute interval. The DTE measures its average traffic using a criteria different than that used for the PWM CIR.	Review how the DTE measures traffic. Increase the configured CIR value.
The accounting record contains no PWM values.	You must configure time-of-day accounting (TODA) is not configured.	You must configure time-of-day accounting (TODA) to generate PWM values. See NN10600-560 <i>Nortel Multiservice Switch 7400/15000/20000 Accounting</i> .
You cannot use the published format in the accounting data server (ADS) to access the PWM.	The ADS does not support PWM functionality.	Only the accounting application programming interface (API) can access the PWM.

Troubleshooting hardware resources

When the hardware resources for a DLCI are unavailable, the system will lock the DLCI. Attempts to unlock the DLCI will result in an alarm.

Hardware resources can become unavailable due to either of the following two reasons:

- More than the maximum supported DLCIs are provisioned on the Lp.
- All sub-connection resources are used on the Lp.

To determine whether these states exist, see the “Monitoring the availability of hardware resources” procedure in NN10600-901 *Nortel Multiservice Switch 7400/15000/20000 Frame Relay Configuration Management*.



Clear cause codes

This section lists clear cause codes that are proprietary to the Nortel Multiservice Switch node. The *lastClearCause* attribute under the *Sig* component converts SVC clear cause codes to the proprietary format. It displays values in decimal format. The PTS uses the diagnostic codes to display responses in hexadecimal format. Accounting records use the following values.

For more information, see the following tables:

- [Restart cause codes \(page 233\)](#)
- [Reset cause codes \(page 234\)](#)
- [Clear cause codes \(page 234\)](#)
- [Subnet clear cause codes \(page 235\)](#)
- [Multiservice Switch node internal clear cause codes \(page 235\)](#)

The *lastClearCause* attribute is significant only after a call clears and the accounting system generates the final accounting record. The default value of the *lastClearCause* attribute is 176, and remains 176 for the first and any intermediate accounting record. The attribute value does not change until the final accounting record has been generated for a call clear (if the call clear reason is not local user clear). The *recordType* field in the accounting record specifies whether the record is the first, intermediate, final, or one (first and last).

Restart cause codes

Cause code (hex)	Cause code (decimal)	Restart cause code description
00	0	restart DTE originated
01	1	restart local procedure error
03	3	restart network congestion
07	7	restart network operational
8F	143	restart reg or cancellation conf



Reset cause codes

Cause code (hex)	Cause code (decimal)	Reset cause code description
00	0	DTE originated
01	1	out of order
03	3	remote procedure error
05	5	local procedure error
07	7	network congestion
09	9	remote DTE operational
0F	15	network operational
11	17	incompatible destination
1D	29	network out of order

Clear cause codes

Cause code (hex)	Cause code (decimal)	Clear cause code description
00	0	DTE originated
01	1	number busy
03	3	invalid fax request
05	5	network congestion
09	9	out of order
0B	11	access barred
0D	13	not obtainable
11	17	remote procedure error
13	19	local procedure error
15	21	rpoa out of order
19	25	revr chrg invalid
21	33	incompatible destination
29	41	fast select invalid
39	57	ship absent
FE	254	node Id Pid
FF	255	undefined



Subnet clear cause codes

Cause code (hex)	Cause code (decimal)	Subnet clear cause code description
A0	160	network congested
A1	161	* RID invalid
A2	162	RID unreachable
A3	163	* Tm not responding
A4	164	* fast select Acp not subscribed
A5	165	* sn RID invalid
A6	166	address not in tree
A7	167	service down
AD	173	X.36 clear cause 50 (requested facility not subscribed)
AE	174	X.36 clear cause 90 (non-existent CUG)
The Multiservice Switch node may receive the codes marked with an * in this table from a DPN-100 node; however, it does not generate them. For detailed descriptions of these DPN-100 codes, see 241-1001-509 <i>DPN-100 Call Clear, Reset and Diagnostic Codes</i> .		

Multiservice Switch node internal clear cause codes

Cause code (hex)	Cause code (decimal)	Internal clear cause code description
B0	176	local user clear
B1	177	remote user clear
B2	178	* cm requested
B3	179	local process error
B4	180	remote process error
B5	181	* message not wanted
B6	182	packet not wanted
B8	184	call collision
B9	185	network congestion
BA	186	* common block failure
BB	187	* local block failure
BD	189	invalid call
The Multiservice Switch node may receive the codes marked with an * in this table from a DPN-100 node; however, it does not generate them. For detailed descriptions of these DPN-100 codes, see 241-1001-509 <i>DPN-100 Call Clear, Reset and Diagnostic Codes</i>		
(1 of 3)		



Multiservice Switch node internal clear cause codes (continued)

Cause code (hex)	Cause code (decimal)	Internal clear cause code description
BE	190	access barred
BF	191	not obtainable
C0	192	out of order
C1	193	* DTE clear bfr acp
C4	196	* nui barred
C5	197	incomplete destination
C6	198	* rpoa required
C7	199	* rpoa invalid
D0	208	* packet network address error
D1	209	service not up
D2	210	service to go down
D3	211	* no links up
D4	212	* rpoa prohibit
D5	213	link out of service
D6	214	* no more calls
D8	216	no free subchannels
D9	217	* no existing cup
DA	218	* call counted create fail
DB	219	* application process busy
DE	222	collect refused
DF	223	cug violation
E0	224	illegal fax
E7	231	device disabled
E8	232	* service option not subscribed
EA	234	* prohibited util
EB	235	abort with ID
EC	236	utility missing
The Multiservice Switch node may receive the codes marked with an * in this table from a DPN-100 node; however, it does not generate them. For detailed descriptions of these DPN-100 codes, see 241-1001-509 <i>DPN-100 Call Clear, Reset and Diagnostic Codes</i>		
(2 of 3)		



Multiservice Switch node internal clear cause codes (continued)

Cause code (hex)	Cause code (decimal)	Internal clear cause code description
ED	237	bad utility value
EE	238	call not allowed
The Multiservice Switch node may receive the codes marked with an * in this table from a DPN-100 node; however, it does not generate them. For detailed descriptions of these DPN-100 codes, see 241-1001-509 <i>DPN-100 Call Clear, Reset and Diagnostic Codes</i>		
(3 of 3)		

For longer descriptions of each of the restart, reset, and clear cause codes, see the ITU-T standards Q.931 and X.76. For longer descriptions of each of the subnet, and Multiservice Switch nodes' internal clear cause codes that originate from DPN-100 switches, see 241-1001-509 *DPN-100 Call Clear, Reset and Diagnostic Codes*.

Diagnostic codes

The *lastDiagnosticCode* attribute under the *FrUni Sig* component converts SVC diagnostic codes to the X.25 format and displays values in decimal format. The PTS uses the diagnostic codes to display responses in hexadecimal format. Accounting records use the following values.

See the table [X.25 format diagnostic codes \(page 237\)](#) lists X.25 diagnostic codes.

X.25 format diagnostic codes

Diagnostic code (hex)	Diagnostic code (decimal)	Diagnostics description
00	0	NO ADDITIONAL INFORMATION
01	1	invalid Ps
02	2	invalid Pr
1X	16 - 19	PACKET TYPE INVALID
20	32	PACKET NOT ALLOWED
21	33	unidentifiable packet
22	34	call on one-way logical channel
23	35	invalid packet type on a permanent virtual circuit
24	36	packet on unassigned logical channel
25	37	reject not subscribed to
26	38	packet too short
(1 of 5)		



X.25 format diagnostic codes (continued)

Diagnostic code (hex)	Diagnostic code (decimal)	Diagnostics description
27	39	packet too long
28	40	invalid general format identifier
29	41	restart or registration packet with nonzero in bits 1 to 4 of octet 1, or bits 1 to 8 of octet 2
2A	42	packet type not compatible with facility
2B	43	unauthorized interrupt confirmation
2C	44	unauthorized interrupt
2D	45	unauthorized reject
30	48	TIME EXPIRED
31	49	for incoming call
32	50	for clear indication
33	51	for reset indication
34	52	for restart indication
40	64	CALL SET UP, CALL CLEARING OR REGISTRATION PROBLEM
41	65	facility registration code not allowed
42	66	facility parameter not allowed
43	67	invalid called address
44	68	invalid calling address
45	69	invalid facility/registration length
46	70	incoming call barred
47	71	no logical channel available
48	72	call collision
49	73	duplicate facility requested
4A	74	bad address length
4B	75	bad facility length
4C	76	facility not provided when expected
4D	77	invalid ITU-T-specified DTE facility
50	80	MISCELLANEOUS
51	81	improper cause code from DTE
52	82	not aligned octet
(2 of 5)		



X.25 format diagnostic codes (continued)

Diagnostic code (hex)	Diagnostic code (decimal)	Diagnostics description
53	83	inconsistent Q bit setting
54	84	invalid Nui
61	97	Dnic unsupported
62	98	Tnic mismatch
64	100	bad utility parameter
65	101	bad utility length
67	103	M bit error
70	112	INTERNATIONAL PROBLEM
71	113	remote network problem
72	114	international protocol problem
73	115	international link out of order
74	116	international link busy
75	117	transit network facility problem
76	118	remote network facility problem
77	119	international routing problem
78	120	temporary routing problem
79	121	unknown called Dnic
7A	122	international link refused
81	129	frame congestion
82	130	invalid clear cause
83	131	incorrect packet size
87	135	operator terminated
8B	139	no source address
8C	140	bad source address
8D	141	link disconnect
8F	143	level 3 idle probe timeout
91	145	destination address supplied for direct call
96	150	subnet Interrupt Request error
97	151	subnet Interrupt Confirmation error
98	152	restricted fast select calls only allowed
(3 of 5)		



X.25 format diagnostic codes (continued)

Diagnostic code (hex)	Diagnostic code (decimal)	Diagnostics description
99	153	incompatible PVC
9A	154	local window negotiation error
9B	155	mandatory fields in Call Request absent
9E	158	incomplete field in Clear packet
9F	159	illegal throughput class
A0	160	hunt group not updated
A1	161	hunt group unavailable
A2	162	hunt group disallowed
A3	163	hunt group Dna insertion error
A4	164	hunt group Dna insertion error
A8	168	utility marker missing
A9	169	block same service
AA	170	routing tables unavailable
AB	171	Nui required
AC	172	Nui required for fast select calls
AF	175	call rejected
B0	176	destination out of order
B1	177	channel unavailable
B2	178	QoS not available
B4	180	call threshold count exceeded
B5	181	unsuccessful call threshold count exceeded
B8	184	database unavailable
B9	185	input / output collision
BD	189	underlying hardware resources are unavailable
C0	192	invalid backup Dna
C1	193	invalid protocol id
C2	194	invalid user data
C3	195	no RFS signal on dial-out modem
C4	196	bad dial-out modem
C5	197	bad CFI frame on dial-out modem
(4 of 5)		



X.25 format diagnostic codes (continued)

Diagnostic code (hex)	Diagnostic code (decimal)	Diagnostics description
C8	200	dial-out connection established
C9	201	dial-out port connecting
F0	240	higher layer initialized
F9	249	invalid protocol id in call user data
(5 of 5)		

For longer descriptions of each of the X.25 format diagnostic codes that originate from DPN-100 nodes, see 241-1001-509 *DPN-100 Call Clear, Reset and Diagnostic Codes*.

FR UNI frame link monitor for E1 frame relay

A frame link monitor (FLM) helps determine the state of an HDLC connection to remote equipment. This is done by monitoring HDLC flags and E1 CAS A-bit information. If no HDLC frames are received for the connection, the FLM helps the operator determine whether this is caused by a line break, or by a fault or shutdown of the remote equipment.

An FLM also lets an operator determine the state of the HDLC connection through an SNMP MIB attribute.

An FLM can be used on the following FPs with any service that uses an HDLC connection:

- 32-port DS1/E1 MSA FP on Nortel Multiservice Switch 7400 node
- 1-port STM-1 FP on Multiservice Switch 15000 node



BNX-lwf performance monitoring and fault management

This chapter provides information on performance monitoring and fault management for interworking between the network element running hyperstream software (BNX) and a Nortel Multiservice Switch 15000 or Multiservice Switch 20000 node, in which interworking supports frame relay over IP.

See the following sections for additional information:

- [Displaying BNX protocol operational attributes \(page 243\)](#)
- [Displaying BNX-lwf operational attributes \(page 244\)](#)
- [Use of the REDIRECT operational verb \(page 245\)](#)
- [Statistics \(page 252\)](#)
- [Frame relay statistics for Multiservice Switch-BNX \(page 247\)](#)
- [Multiservice Switch-BNX frame relay alarms \(page 246\)](#)
- [Behavior on control processor switchover \(page 254\)](#)



Displaying BNX protocol operational attributes

Display BNX protocol operational attributes to view the following attributes

- *invalidBnxHeader*
- *lastInvalidBnxHeader*
- *invalidIpfProtocolId*
- *discardedIpfFragment*
- *sourceIpfAddressLastDropped*
- *droppedUnknownDlci*
- *lastUnknownDlci*
- *droppedUnknownDlciXnet*
- *lastUnknownDlciXnet*

Prerequisites

- For clarity, the following steps show configuration of BNX-lwf for a FrUni. The steps are identical for configuring BNX-lwf for a FrNni.
- See NN10600-060 *Nortel Multiservice Switch 7400/15000/20000 Component Reference* under *FrUni Dlci BnxProtocol* for a description of each of these attributes.

Procedure steps

Step	Action
1	Use the display command to lists the values for the BNX protocol operational attributes. display FrUni/<fr_uni> Dlci/<dlci> BnxProtocol
--End--	

Variable definitions

Variable	Value
<dlci>	is the instance number of the DLCl and is a decimal between 16 and 1007.
<fr_uni>	is the instance number of the FrUni and is a decimal between 1 and 4 294 967 295.



Displaying BNX-lwf operational attributes

Display BNX-lwf operational attributes to view the following attributes:

- *timeOfLastStateChange*
- *crossNetSent*
- *crossNetStatus*
- *crossNetReceived*
- *crossNetErrors*
- *backupCrossNetSent*
- *backupCrossNetStatus*
- *backupCrossNetReceived*
- *backupCrossNetErrors*
- *currentDestination*
- *redirectMode*
- *remoteIpAddress*
- *remoteDlci*
- *lastMismatchedIpAddress*
- *ipAddressMismatches*

Prerequisites

- For clarity, the following steps show configuration of BNX-lwf for a FrUni. The steps are identical for configuring BNX-lwf for a FrNni.
- See NN10600-060 *Nortel Multiservice Switch 7400/15000/20000 Component Reference* under *FrUni Dlci BnxInterworkingFunction* for a description of each of these attributes.

Procedure steps

Step	Action
1	Use the display command to lists the values for the BNX-lwf operational attributes. display FrUni/<fr_uni> Dlci/<dlci> BnxInterworkingFunction
--End--	



Variable definitions

Variable	Value
<dlci>	is the instance number of the DLCI and is a decimal between 16 and 1007.
<fr_uni>	is the instance number of the FrUni and is a decimal between 1 and 4 294 967 295.

Use of the REDIRECT operational verb

The redirect operation verb redirects the PVC source from the primary destination to the backup destination, or vice versa. This verb is also used to change the PVC redirection mode between the manual mode and switchOnDemand mode.

The -force option redirects the PVC to a either a primary or backup remote destination. The redirection proceeds even if the cross-net status of the remote is inactive or cannot be received. When the value of the -destination option is onDemand, the -force option is ignored.

The -destination option values have the following affects:

- primary: changes the redirectMode attribute to manual; the PVC is redirected from the backup destination to the primary destination
- backup: changes the redirectMode attribute to manual; the PVC is redirected from the primary destination to the backup destination
- onDemand: changes the redirectMode attribute to switchOnDemand; the PVC is redirected to the primary destination or the backup destination based on the cross-net status of the primary destination

The format of the verb is

```
redirect -force -destination(<dest>) FrUni/<fruni> Dlci/  
<dlci> Bnx
```

```
redirect [-force] [-destination(<dest>)] FrNni/<frnni>  
Dlci/<dlci> Bnx
```

Variable	Value
<dest>	is one of primary, backup, or onDemand. Note that there are parentheses around the value for the -destination option.
(1 of 2)	



Variable	Value
<dlci>	is the instance number of the DLCI and is a decimal between 16 and 1007.
<fr_uni>	is the instance number of the FrUni and is a decimal between 1 and 4 294 967 295.
(2 of 2)	

Multiservice Switch-BNX frame relay alarms

The alarms reported on a BNX are known as event messages. Each event message is uniquely identified by an entity name and an event code in the format <entity name>/<event code> (for example, FRSW/1). The entity name identifies which software entity issues the event message; the event code identifies what particular event reported by the entity. Each event message contains the following information:

- event message identifier (<entity name>/<event code>)
- decimal identifier (equivalent to the 32-bit SNMP trap number that the BNX sends to the SNMP manager)
- severity (one of fault, warning, info, trace, or debug)
- message (description of the event message)
- meaning (explanation on why the BNX generated the event message)
- action required (recommended remedial action in response to the event message)

The display of event messages can be filtered by entity name, severity, slot, or IP address. For the details of event messages supported on BNX and their usages, see the document Part No. 115412-C Rev. A, "Event Messages", (February 1998, Bay Networks).

In network configurations that involve BNX-lwf, the most concerned event messages are FRSW events. The FRSW events related to configuration are not required for mapping between Nortel Multiservice Switch 15000 and Multiservice Switch 20000 nodes and the BNX, since they have local significance on only BNX. In general, events that are locally significant to BNX only do not need to be mapped. For BNX-lwf, alarms related to PVC connections and cross-net signaling must be mapped to equivalent BNX events since they are part of interworking between Multiservice Switch nodes and BNX.

The PVC alarms enhancement provides a comprehensive set of alarms for reporting the status of PVC connections that are configured in both local and external networks. The external networks can be either remote or adjacent. The PVCs are either inactive or active. BNX-lwf supports these alarm enhancements.



The cross-net signaling protocol is specific to the frame relay service on the Multiservice Switch 15000 and Multiservice Switch 20000 nodes for interworking with BNX. The alarms are required on the Multiservice Switch node for the following:

- for cross-net signaling failure in which the Multiservice Switch node cannot receive the cross-net status message from BNX within a pre-defined cross-net polling interval
- recovery of the cross-net signaling failure, in which the Multiservice Switch node is again able to receive the cross-net status message from BNX

The state change notification (SCN) of the frame relay Uni/Nni BnxProtocol (active to inactive, or inactive to active) can reflect the status of cross-net signaling procedure.

To communicate with a BNX through frame relay across a IP network, the Multiservice Switch 15000 and Multiservice Switch 20000 node requires the following for the frame relay UNI/NNI interface

- an IP address (the source IP address)
- a DLCI associated with the called IP address (the destination IP address)

The source IP address must be registered with the frame relay IP service (FRIP). The 7007 9001 alarm addresses the event of a registration failure. See NN10600-500 *Nortel Multiservice Switch 6400/7400/15000/20000 Alarms Reference* for information on this alarm.

Frame relay statistics for Multiservice Switch-BNX

Nortel Multiservice Switch 15000 or Multiservice Switch 20000 nodes' frame relay virtual circuit statistical counts do not apply to BNX-lwf. The reason is the BNX-lwf DLCI virtual circuit is not established through the Multiservice Switch network using DPRS. Instead, the virtual circuit is established through the IP network based on the frame relay IP server using the virtual router. The virtual circuit statistics are set to zero and display zeros.

BNX-lwf supports almost all other frame relay statistical counts on the Multiservice Switch node. Unsupported statistical counts are set to zero and display zeros. New statistical counts are needed in some cases for correlation with the BNX.

The table [Mapping for frame relay statistics between BNX and Multiservice Switch 15000 or Multiservice Switch 20000 \(page 248\)](#) maps frame relay statistics between the BNX and Multiservice Switch nodes. Multiservice Switch node statistical counts or levels that are specific to BNX-lwf are indicated through highlighted table cells.



This table addresses statistics-related data only. The LMI-related statistical counts are not required for mapping since LMI is locally significant to the interface and its remote end. Cross-net signaling-related statistical counts must be synchronized through configuration. See NN10600-901 *Nortel Multiservice Switch 7400/15000/20000 Frame Relay Configuration Management*.

Mapping for frame relay statistics between BNX and Multiservice Switch 15000 or Multiservice Switch 20000

BNX		Multiservice Switch 15000 or Multiservice Switch 20000	
Statistical data	Count (MIB object)	Count	Level
Number of VCs currently configured on this interface	wfFrSwDlcmiVCsInUse	Display all VCs on the interface (use the command <code>d FrUni/<n> Dlci/* Vc</code>)	DLCI
Number of packets dropped due to unknown IP address in the header	wfFrSwSwitchHdrErrors	unknownIpAddressErr	FRIP
Number of frames dropped due to insufficient number of bytes	wfFrSwDlcmiFrameToo-ShortErrors	errorShortFrmFromIf	DLCI
Number of frames dropped due to too long	wfFrSwDlcmiFrameToo-LongErrors	errorLongFrmFromIf	DLCI
Number of frames dropped due to illegal DLCI (out of range DLCI)	wfFrSwDlcmiIllegalDlciErrors	unknownDlciFramesFromIf	Inter- face
Number of frames dropped due to unknown DLCI	wfFrSwDlcmiUnknownDlci-Errors	unknownDlciFramesFromIf	Inter- face
Number of frames dropped due to Q922 control byte error	wfFrSwDlcmiControlByte-Errors	invalidHeaderFramesFromIf	Inter- face
Number of frames dropped due to format error	wfFrSwDlcmiFormatErrors	invalidHeaderFramesFromIf	Inter- face
Number of frames dropped due to other errors	wfFrSwDlcmiOtherErrors	invalidHeaderFramesFromIf	Inter- face
(1 of 4)			



Mapping for frame relay statistics between BNX and Multiservice Switch 15000 or Multiservice Switch 20000 (continued)

BNX		Multiservice Switch 15000 or Multiservice Switch 20000	
Statistical data	Count (MIB object)	Count	Level
Number of DE=0 frames sent to interface	wfFrSwVcTxNoneDeFrames	frmTolf - deFrmTolf	DLCI
Number of DE=0 bytes sent to interface	wfFrSwVcTxNoneDeOctets	bytesTolf - deBytesTolf	DLCI
Number of DE=1 frames sent to interface	wfFrSwVcTxDeFrames	deFrmTolf	DLCI
Number of DE=1 bytes sent to interface	wfFrSwVcTxDeOctets	deBytesTolf	DLCI
Number of frames sent to interface with FECN set by the frame relay service due to local congestion detected	wfFrSwVcSetFecnFrames	No Multiservice Switch equivalent.	DLCI
Number of bytes sent to interface with FECN set by the frame relay service due to local congestion detected	wfFrSwVcSetFecnOctets	No Multiservice Switch equivalent.	DLCI
Number of frames sent to interface with BECN set by the frame relay service due to local congestion detected	wfFrSwVcSetBecnFrames	becnFrmSetByService	DLCI
Number of bytes sent to interface with BECN set by the frame relay service due to local congestion detected	wfFrSwVcSetBecnOctets	No Multiservice Switch equivalent.	DLCI
Number of frames sent to interface with DE bit set by the frame relay service	wfFrSwVcSetDeFrames	excessFrmFromIf	DLCI
(2 of 4)			



Mapping for frame relay statistics between BNX and Multiservice Switch 15000 or Multiservice Switch 20000 (continued)

BNX		Multiservice Switch 15000 or Multiservice Switch 20000	
Statistical data	Count (MIB object)	Count	Level
Number of bytes sent to interface with DE bit set by the frame relay service	wfFrSwVcSetDeOctets	excessBytesFromIf	DLCI
Number of DE=0 frames received from interface, which were discarded	wfFrSwVcDropNonDeFrames	discCongestedFromIf - discDeCongestedFromIf	DLCI
Number of DE=0 bytes received from interface, which were discarded	wfFrSwVcDropNonDeOctets	discCongestedFromIfBytes - discDeCongestedFromIfBytes	DLCI
Number of DE=1 frames received from interface, which were discarded	wfFrSwVcDropDeFrames	discDeCongestedFromIf	DLCI
Number of DE=1 bytes received from interface, which were discarded	wfFrSwVcDropDeOctets	discDeCongestedFromIfBytes	DLCI
Number of frames dropped due to PVC inactive	wfFrSwVcInactiveVcDrop-Frames	discFrameAbit	DLCI
Number of bytes dropped due to PVC inactive	wfFrSwVcInactiveVcDrop-Octets	discByteAbit	DLCI
Number of DE=0 frames received from the interface	wfFrSwVcRecvNonDeFrames	frmFromIf - deFrmFromIf	DLCI
Number of DE=0 bytes received from the interface	wfFrSwVcRecvNonDeOctets	bytesFromIf - deBytesFromIf	DLCI
Number of DE=1 frames received from the interface	wfFrSwVcRecvDeFrames	deFrmFromIf	DLCI
(3 of 4)			



Mapping for frame relay statistics between BNX and Multiservice Switch 15000 or Multiservice Switch 20000 (continued)

BNX		Multiservice Switch 15000 or Multiservice Switch 20000	
Statistical data	Count (MIB object)	Count	Level
Number of DE=1 bytes received from the interface	wfFrSwVcRecvDeOctets	deBytesFromIf	DLCI
Number of FECN=1 frames received from the interface	wfFrSwVcRecvFecnFrames	fecnFrmFromIf	DLCI
Number of FECN=1 bytes received from the interface	wfFrSwVcRecvFecnOctets	No Multiservice Switch equivalent.	DLCI
Number of BECN=1 frames received from the interface	wfFrSwVcRecvBecnFrames	becnFrmFromIf	DLCI
Number of BECN=1 bytes received from the interface	wfFrSwVcRecvBecnOctets	No Multiservice Switch equivalent.	DLCI
Number of consecutive errors in cross-net signaling	wfFrSwVcXnetErrors	crossNetErrors	BNX-lwf
Number of excess burst frames dropped	wfFrSwVcDropExcessBurst-Frames	discExcessFromIf	DLCI
Number of excess burst bytes dropped	wfFrSwVcDropExcessBurst-Octets	discExcessFromIfBytes	DLCI
Multiservice Switch statistical counts or levels that are specific to BNX-lwf are indicated through highlighted table cells.			
(4 of 4)			

The table [Statistical counts for Multiservice Switch-BNX interworking \(page 252\)](#) summarizes components on the Multiservice Switch 15000 or Multiservice Switch 20000 node that are specific to BNX-lwf.



Statistical counts for Multiservice Switch-BNX interworking

Component (or Level)	Statistical Count	Description
FrIpServer	datagramsFromPp	Counts the number of datagrams received from the protocol port.
	invalidIpHeaderVersion	Counts the number of datagrams dropped because of invalid IP header version.
	UnknownIpAddressErr	Counts the number of datagrams dropped because of unknown IP address.
FrUni BnxProtocol	invalidBnxHeader	Counts the number of datagrams dropped due to invalid BNX header version.
FrNni BnxProtocol	invalidProtocolId	Counts the number of datagrams dropped due to invalid protocol ID in the IP header.
	discardedIpFragments	Counts the number of datagrams dropped.
	droppedUnknownDlci	Counts the number of datagrams dropped due to unknown DLCI.
	droppedUnknownDlci-Xnet	Counts the number of cross-net signals dropped due to unknown DLCI.
FrUni Dlci BnxInterworking-Function	crossNetErrors	Counts the number of consecutive errors in cross-net signaling.
FrNni Dlci BnxInterworking-Function	failMatchedSourceIp-Address	Counts the number of packets received with unknown source IP address.

Statistics

Statistics are collected on-switch to provide information for monitoring the behavior and performance of the node or network. Operators retrieve statistics counts through the operational commands on-switch or through the Preside Multiservice Data Manager (MDM).

Operators and engineers use statistics can be used for traffic analysis. This analysis is the basic for optimizing node and network resources that best serve customer needs.



Nortel Multiservice Switch 15000 or Multiservice Switch 20000 nodes spool statistics to the Management Data Provider (MDP) through the Data Collection System (DCS). Spooling intervals are defined through configuration. The MDP transfers the statistics to Preside MDM for generating the reports and graphs required for traffic analysis and monitoring.

Multiservice Switch 15000 and Multiservice Switch 20000 node frame relay statistics and BNX-lwf

Only the frame relay framer statistics are spooled to the MDP and then to the Preside MDM for further processing. The other statistics can be obtained by the operational commands on-switch or off-switch. The frame relay services (FrUni and FrNni) provide the statistical counts at the components or levels described in the following sections. Additional requirements or exceptions for the BNX interworking function (BNX-lwf) over and above standards-based frame relay service are noted.

Framer statistics

Framer statistics are generated for each logical link on a port. There are no additional requirements for BNX-lwf.

Interface statistics

Interface statistics are generated for each service. The *frmTolIfByQueue* and *octetTolIfByQueue* counts are supported only when four software emission queues are configured. The channelized 4-port DS3 function processor supported by BNX-lwf uses four hardware (not software) emission queues and therefore does not support the *frmTolIfByQueue* and *octetTolIfByQueue* counts.

LMI statistics

LMI statistics are generated for each interface. There are no additional requirements for BNX-lwf.

DLCI statistics

The node generates DLCI statistics for each connection, relative to link side. BNX-lwf does not support the following counts for the channelized 4-port DS3 function processor since this function processor never encounters congestion:

- *discardCongestedFromIf*
- *discardCongestedFromIfBytes*
- *discardDeCongestedFromIf*
- *discardDeCongestedFromIfBytes*

The following counts are used for accounting. After the accounting report is generated, these counts are reset to zero.

- *totalIngressBytes*



- *totalEgressBytes*
- *eirIngressBytes*
- *eirEgressBytes*
- *discardedBytes*
- *totalIngressSegFrm*
- *totalEgressSegFrm*
- *eirIngressSegFrm*
- *eirEgressSegFrm*
- *discardedSegFrm*

Virtual circuit statistics

Virtual circuit statistics are generated for each connection, relative to the network side for dynamic packet routing system (DPRS).

BNX-lwf does not support the *cannotForwardToSubnet* count for the channelized 4-port DS3 function processor.

Behavior on control processor switchover

Following a control processor switchover, the frame relay function processor remains in service but the IP routing cache is flushed. The system removes from service the BNX-lwf DLCIs until the associated FRIP can successfully re-register with the newly active control processor.

Frame relay ISDN switched access monitoring and troubleshooting

This section helps you monitor, maintain, and troubleshoot frame relay ISDN switched access

- [Monitoring the service \(page 254\)](#)
- [Maintaining the service \(page 256\)](#)
- [Detecting frame relay ISDN switched access problems \(page 257\)](#)

Monitoring the service

To determine how well your frame relay ISDN switched access service is operating, examine its operational statistics.

Operational commands

Use the display command to monitor the frame relay ISDN switched access components. For information on operator and provisioning commands, see NN10600-050 *Nortel Multiservice Switch 7400/15000/20000 Command Reference*.



Examples

To display incoming call statistics, use any of the following commands:

```
display DataSigChan/n callControl
display DataSigChan/n callControl bch/*
display DataSigChan/n callControl channelGroup/*
display DataSigChan/n callControl channelGroup/m
display DataSigChan/n callControl channelGroup/* bch/*
display DataSigChan/n callControl channelGroup/m bch/*
```

To display statistics for an ISDN FrUni, use the following command:

```
display fruni/<x> isdn
```

Operational attributes

For information about operational attributes, see NN10600-060 *Nortel Multiservice Switch 7400/15000/20000 Component Reference*.

OSI state

The tables in this section describe the OSI states that the network supports for the *DataSigChan* and *<Q.931 variant> Framer* components, their meaning, and the events that cause their state change.

DataSigChan component state combinations

Combination (Administrative, Operational, Usage)	Details
Unlocked, Enabled, Busy	The signaling channel is enabled (Layer 2 and Layer 3 are enabled).
Unlocked, Disabled, Idle	The signaling channel is disabled.
Locked, Disabled, Idle	The signaling channel is locked and disabled.

<Q.931 variant> Framer component state combinations

Combination (Administrative, Operational, Usage)	Details
Unlocked, Enabled, Busy	The physical layer is enabled.
Unlocked, Disabled, Idle	The physical layer is disabled



SNMP management

Frame relay ISDN switched access supports Enterprise MIBs. For more information see NN10600-300 *Nortel Multiservice Switch 7400/15000/20000 Operations: SNMP*.

Maintaining the service

This section describes

- [Locking and unlocking DataSigChan components \(page 256\)](#)
- [Managing interfaces, components and connections \(page 256\)](#)

Locking and unlocking DataSigChan components

To lock and unlock *DataSigChan* components, use the following commands:

```
lock DataSigChan/<n>
```

The lock command disconnects any BCH(s) that setup, then it disconnects the DCH. In the locked mode, both the layer 2 and layer 3 DCH protocols are disabled.

```
unlock DataSigChan/<n>
```

The unlock command transmits a RESTART message to the link. A RESTART then transmits to the link. When the RESTART ACK reply arrives, new ISDN calls can occur. Previously connected BCH calls do not restore.

When you lock the port associated with the D-channel, you disable the framer. Unlocking enables it. When the framer is disabled, the network ignores all data bits that arrive from the DCH physical layer.

Managing interfaces, components and connections

This section describes how to manage frame relay ISDN switched access interfaces, components, and connections.

Primary Rate Interface

To disable (enable) administratively the DS1 or E1 port issue the lock (unlock) command to the *Lp DS1* or *Lp E1* component. When you lock the DS1 or E1 interface, you completely disable the ISDN link. This forces the DCH and all the existing BCH connections down. VCs that terminate on the BCH links disconnect.

ISDN D-channel

To disable (enable) the ISDN DCH administratively, issue the lock (unlock) operator command to the *Lp DS1 Chan* or *Lp E1 Chan* component. This action disconnects the BCHs.

ISDN B-channel



To disable (enable) the ISDN BCHs administratively, issue the lock (unlock) command to the *Lp DS1 Chan* or *Lp E1 Chan* component or use the DCH control mechanism to lock (unlock) it automatically.

Frame relay UNI

To disable (enable) the FR UNI administratively, lock (unlock) the *FrUni* component.

Frame relay DLCI

To disable (enable) the DLCI administratively, issue the lock (unlock) command to the *FrUni DlcI* component.

Detecting frame relay ISDN switched access problems

This section lists problems that can occur when you provision and use frame relay ISDN switched access service. The table [Detecting DataSigChan component problems \(page 257\)](#) lists problems, probable causes, and corrective actions to help you trouble-shoot.

Detecting DataSigChan component problems

Problem	Probable Cause	Corrective Action
The <i>DataSigChan</i> component is in disabled state (layer 1 of DCH is down).	faulty or improper cabling	Check the cable and physical connection.
	faulty configuration	Check the configuration of the remote equipment.
	The <i>DataSigChan</i> component is locked.	Lock and unlock the <i>DataSigChan</i> component, as follows: lock dataSig/x unlock dataSig/x
	The LP, LP port, or channel associated with the <i>DataSigChan</i> component is locked.	Lock and unlock the LP, LP port and channel, as follows: lock lp/x unlock lp/x
	DataSigChan/x <Q.931 variant> is not properly provisioned.	Check the parameters, as follows: display Data/x <Q.931 variant> side
(1 of 2)		



Detecting DataSigChan component problems (continued)

Problem	Probable Cause	Corrective Action
The <i>DataSigChan</i> component was enabled and then changed to disabled permanently.	The <i>DataSigChan</i> component, LP, LP port or the channel is locked.	Lock and unlock the <i>DataSigChan</i> component, LP, LP port and channel, as follows: lock dataSig/x unlock dataSig/x lock lp/x unlock lp/x lock lp/x e1/y unlock lp/x e1/y lock lp/x e1/y chan/0 unlock lp/x e1/y chan/0
Alarm 7094 0000 related to <i>DataSigChan</i> component issues.	The <i>DataSigChan</i> component changed to disabled state. The LP, LP port, or channel associated with the <i>DataSigChan</i> component is locked.	Lock and unlock the LP, LP port and channel, as follows: lock lp/x unlock lp/x lock lp/x e1/y unlock lp/x e1/y lock lp/x e1/y chan/0 unlock lp/x e1/y chan/0 A clear message is issued when the <i>DataSigChan</i> component returns to an enabled state.
(2 of 2)		



Detecting B-channel-related problems

Problem	Probable Cause	Corrective Action
An ISDN call can not establish. (You receive a Q931 Release Complete message in response to a Q931 Setup message.)	One: operationalState is disabled and the keep-alive RRs frame is no longer transmitting and receiving on the link.	Check that the <i>frmTolIf</i> and <i>frmFromIf</i> attributes of the <i>Framer</i> component are being incremented, as follows: display Data/x <Q.931 variant> framer If these statistics are being incremented, then proceed to Probable Cause Two. If these statistics are not being incremented, then proceed to Probable Cause Four.
	Two: DataSigChan is "disabled".	Check that the <i>totalValidInCalls</i> attribute of the <i>CallControl</i> component is being incremented, as follows: display Data/x cc If this statistic is being incremented, then proceed to Probable Cause Three. If this statistic is not being incremented, then proceed to Probable Cause Four.
	Three: The call IE field is not acceptable.	Check that <i>successfulInCalls</i> attribute is being incremented, as follows: display Data/x cc If this statistic is being incremented, then proceed to Probable Cause Six. If this statistic is not being incremented, then proceed to Probable Cause Seven.
(continued)	Four: The operationalState attribute is disabled. Layer 2 or Layer 3 of the signaling channel may be out of service.	Check the state of the <i>DataSigChan</i> component, as follows: display dataSigChan/x If disabled, lock and unlock the channel, as follows: lock lp/x e1/y chan/0 unlock lp/x e1/y chan/0 The <i>dataSigChan</i> component links to the channel. If the problem persists, check your cabling and pinouts. If the <i>DataSigChan</i> component is enabled, proceed to Probable Cause Five.
(1 of 4)		



Detecting B-channel-related problems (continued)

Problem	Probable Cause	Corrective Action
	Five: The SETUP message did not comply to the ISDN standard is use.	Check that the SETUP message is valid. Turn on the trace buffer to see the internal processing of the ISDN messages, as follows: set Data/1 <Q.931 variant> tracing q931Summary (Different levels of trace detail are available. Use the help command on the attribute for more information.) Use the trace information and the clearCause command to perform a detailed analysis of the incompatible items and correct them in your SETUP message. If the SETUP message did not have the mandatory IEs (for example, bearer capability, calling party number, or channel information) but the basic value complied to the ISDN standard (for example, TS014), then the message rejects without CallControl software processing.
(continued)	Six: Message block congestion exists on the node.	Check that none of the invalid counters are being incremented, as follows: display Data/x cc Look for a <i>clear cause 42</i>. Reduce the traffic load on the processor card.
	Seven: SETUP is compliant to the ISDN standard, but contains an bearer service that this feature does not support.	Check whether other counters are being incremented. Diagnostic rejections include the following: • Counter, lastClearCause • inChannelBusy, 34 • inInvalidScreen, 21 • inInvalidCgpn, 28 • inInvalidCapability, 57
	The channel is busy.	Make the call later, when the channel is freed up. Alternately, request another channel if one is available or provision more channels.

(2 of 4)



Detecting B-channel-related problems (continued)

Problem	Probable Cause	Corrective Action
	The Q931 SETUP message has a non-acceptable screening indicator. (The <code>inInvalidScreen</code> is being incremented and the <code>lastClearCause</code> is 21.)	Check for a screening indicator that is not acceptable as provisioned in the <i>ChannelGroup</i> component.
	A <i>cgpn</i> component that is not allowed in any of the channel groups is received.	Check each channel group for acceptable calling group party numbers, as follows: display -p DataSigChan/x cc ChannelGroup/x cgpn/x Or, if appropriate, add the incoming calling party number to a <i>ChannelGroup</i> component.
(continued)	(continued) An invalid capability was requested.	Check the bearer capability and lower layer compatibility IE to verify that the service is unrestricted digital and the transfer rate is 64 kbit/s.
B-channel is in disabled state.	The <i>DataSigChan</i> component is disabled.	See the table DataSigChan component state combinations (page 255) .
	The <i>FrUni</i> component is locked.	Unlock the <i>FrUni</i> component.
B-channel stays in idle state after a call is made for this B-channel.	The call was invalid.	Refer to Probable Cause Three.
	A call was made, but another B-channel was allocated for it.	Check that the requested <i>FrUni</i> component is provisioned and that the FR UNI ISDN state is idle.
	A call was made and accepted, but there are no VCs in data transfer state on this B-channel.	Check that the frame relay PVC is provisioned.
The provisioned <i>FrUni</i> is not displayed when using: display data/1 cc bch/*	The <i>FrUni</i> component does not have the <i>Isdn</i> component.	Add the <i>Isdn</i> component under the <i>FrUni</i> component.
The provisioned channel is not displayed when using: display data/1 cc bch/*	The channel list of this <i>CallControl</i> component does not include this channel.	Add the channel to the <i>channelList</i> attribute under the <i>CallControl</i> component.
(3 of 4)		



Detecting B-channel-related problems (continued)

Problem	Probable Cause	Corrective Action
The VC does not display the PSTN number.	The <i>FrUni Isdn</i> component <i>addressSignaling</i> attribute is provisioned to normalBehaviour	Set this attribute to IsdnDna.
	The presentation indicator in the Q.931 SETUP message is set to restricted.	Change the presentation indicator in the Q.931 SETUP message to displayed.
The B-channel is disconnected by the Multiservice Switch node.	The T320 timer expired when the B-channel stayed in idle state for T320 seconds after the call was up.	Locate and resolve problems that prevent the VCs from coming up.
(4 of 4)		

Detecting physical-port-related problems

Problem	Probable Cause	Corrective Action
DS1 or E1 port does not come up.	Impedance level error Bit coding algorithm error	Refer to the following documents: NN10600-520 <i>Nortel Multiservice Switch 7400/15000/20000 Fault and Performance Management: Troubleshooting</i>
	The port is locked.	Unlock the port.
(1 of 2)		



Detecting physical-port-related problems (continued)

Problem	Probable Cause	Corrective Action
The framer <i>operationalState</i> attribute is in the disabled state.	Bad cables	Use port and line tests to locate bad cables.
	Incorrect provisioning data	Check that both ends of the link have compatible provisioning data.
	Essential components are locked.	Unlock the components.
Timeslot or B-channel errors.	The requested channel number is assigned to an unexpected timeslot.	Check B-channel and timeslot provisioning, as follows: d -p lp/<n> ds1/<m> ch/<p> (for a DS1 function processor) or d -p lp/<n> e1/<m> ch/<p> (for an E1 function processor) Check that each timeslot number and channel number are matched correctly. Refer to the appropriate clause of the ISDN standard in use.
(2 of 2)		

Detecting ChannelGroup- and CallingPartyNumber- related problems

Problem	Probable Cause	Corrective Action
The <i>DataSigChan</i> component accepts calls with a non-provisioned calling party number (cgpn).	Unused channel groups exist.	Check for empty channel groups (that is, they have no subcomponents). For each channel group type: list DataSigChan/<n> Cc Cg/<m> Cgpn/* Delete all channel groups with no subcomponents.
Alarm 7048 2001 related to cgpn issues.	Duplicate <i>cgpn</i> attributes exist.	Remove the duplicate cgpn(s). A MSG alarm issues whenever a cgpn exists under more than one <i>CallingGroup</i> component within the same <i>CallControl</i> component
Calls (with a provisioned cgpn) are accepted under the wrong channel group.	Duplicate <i>cgpn</i> attributes exist.	Remove the duplicate cgpn(s). Use only the first <i>cgpn</i> attribute provisioned. If you provision the same cgpn under different channel groups, calls entering the Multiservice Switch node will route to the channel group with the lowest instance number.
(1 of 2)		



Detecting ChannelGroup- and CallingPartyNumber- related problems (continued)

Problem	Probable Cause	Corrective Action
Calls with a non-provisioned cgpn are rejected even if an empty channel group exists and all B-channels listed in the <i>ChannelList</i> attribute are available.	The SETUP message is invalid, or the corresponding B-channel is not provisioned properly.	Refer to <Q.931 variant> Framer component state combinations (page 255) .
	Duplicate empty channel groups exist.	Check for duplicate empty channel groups. Use only the empty channel group with the lowest number. Remove duplicate empty channel groups.
(2 of 2)		

Detecting memory-related problems

Problem	Probable Cause	Corrective Action
An engineering alarm generates for one of the components.	The FP memory that is available for code and components is at or near its limit.	Choose from these three options: • Reduce the feature package so that only essential features are loaded. • Reduce the number of provisioned components. (Revisit your network engineering decisions.) • Upgrade to a processor with more memory, if applicable. The feature package you use and the number of provisioned components affect available memory. On the application side, the following components require memory: • FrameRelayUni, some Sub-components (like isdn) and related hardware (Fruni/1209 and Ip/12/E1/0 ch/9) • all DLCIs On the DataSignalling side (the D-channel), the following components require memory: • all DataSignalling and related hardware • all Calling Party Numbers.



FR UNI compliance with standards

Use the following sections to learn more about FR UNI compliance with standards:

Navigation

- [Compliance with Vendor Forum specification \(page 265\)](#)
- [Compliance with joint specification \(page 266\)](#)
- [Compliance with Annex D of T1.617 \(page 266\)](#)
- [Compliance with Annex A of ITU-T Q.933 \(page 267\)](#)
- [Compliance with frame relay UNI Interface Implementation Agreement \(page 268\)](#)
- [Compliance with FRF.4 \(page 268\)](#)
- [Compliance with ITU-T Q.922 \(page 269\)](#)
- [Compliance with ITU-T X36 \(page 271\)](#)
- [Compliance with FRF.12 \(page 272\)](#)
- [Compliance with FRF.16 \(page 272\)](#)

Compliance with Vendor Forum specification

This feature complies with *Frame Relay Specification with Extensions*, Issue I of Joint Specification, Revision 1.0, September 1990, Document # 001-208966.

The joint efforts of Nortel Networks, StrataCom Inc., Digital Equipment Corporation, and Cisco produced this specification.

Detailed compliance statement

The tables in this section use the following terms to indicate compliance with the joint specification:

- **Noted:** The specification provides clarification, non-specific information, or details that do not relate directly to the frame relay service.



- Fully complies: Nortel Multiservice Switch frame functionality fully complies with the text for this section.
- Compliant with these exceptions: Multiservice Switch frame relay service does not comply completely with the text.
- Not supported: Multiservice Switch FRS does not support the functionality described in the text.

Compliance with joint specification

Nortel Multiservice Switch frame relay complies with Issue I of Joint Specification.

Frame relay compliance with joint specification

Section	Section title	Compliance
Section 1	Introduction	Noted
Section 2	References	Noted
Section 3	Overview	Noted
Section 4	Physical Interfaces	Fully complies.
Section 5	Frame Relay Data Link Interface	Noted
Section 5.1	Specifications	Fully complies. The value of dN1 is 2100.
Section 5.2	Procedures	Fully complies.
Section 5.3.1	Standard Addressing Convention	Fully complies.
Section 5.3.2	Configurable Parameters	The value of dN1 is configurable from 1 to 2100 bytes.
Section 6	Local Management Interface - Common Extensions	Fully complies.
Section 7	Local Management Interface - Optional Extensions	Not supported

Compliance with Annex D of T1.617

Nortel Multiservice Switch frame relay complies with American National Standard for Telecommunications (ANSI) Annex D of T1.617.



Frame relay compliance to Annex D of T1.617 ANSI (June 1991)

Section	Section title	Compliance
Section D.1	Messages Used for PVC Status	Fully complies.
Section D.1.1	Status	Fully complies. Asynchronous PVC status report also supported.
Section D.1.2	Status Enquiry	Fully complies.
Section D.2.1	Protocol Discriminator	Fully complies. The protocol discriminator is H.08.
Section D.2.2	Call Reference	Fully complies.
Section D.2.3	Message Type	Fully complies. Only Status and Status Enquiry messages are supported.
Section D.3.1	- Report Type	Fully complies.
Section D.3.2	- Link Integrity Verification	Fully complies.
Section D.3.3	- PVC Status	Compliant with exceptions Only two-octet DLCIs supported.
Section D.4:	Procedures	Noted
Section D.4.1	Periodic Polling	Fully complies.
Section D.4.2	Link Integrity Verification	Fully complies.
Section D.4.3	Reporting New PVCs	Fully complies.
Section D.4.4	Reporting the Availability of a PVC	Fully complies.
Section D.5	Error Conditions	Noted
Section D.5.1	Network Operation Errors	Fully complies.
Section D.5.2	User Equipment Operation Errors	Noted
Section D.6	Optional Bidirectional Network Procedures	Fully complies.
Section D.7	System Parameters	Default for N392 is 3. Possible values of T392 range from 5 to 30 seconds in increments of 5 seconds. Default for T392 is 15.

Compliance with Annex A of ITU-T Q.933

Nortel Multiservice Switch frame relay complies with ITU-T (formerly CCITT) standards Annex A of Q.933.



Frame relay compliance to Annex A of ITU-T Q.933 (1992)

Section	Section title	Compliance
Section A.1	Messages Used for PVC Status	Fully complies.
Section A.1.1	Status	Fully complies. Asynchronous PVC status report also supported.
Section A.1.2	Status Enquiry	Fully complies.
Section A.2.1	Protocol Discriminator	Fully complies.
Section A.2.2	Call Reference	Fully complies.
Section A.2.3	Message Type	Compliant with exceptions Only Status and Status Enquiry messages are supported.
Section A.3.1	Report Type	Fully complies.
Section A.3.2	Link Integrity Verification	Fully complies.
Section A.3.3	PVC Status	Fully complies. Only two-octet DLCIs supported.
Section A.4	Procedures	Noted
Section A.4.1	Periodic Polling	Fully complies.
Section A.4.2	Link Integrity Verification	Fully complies.
Section A.4.3	Reporting New PVCs	Fully complies.
Section A.4.4	Reporting the Availability of a PVC	Fully complies.
Section A.5	Error Conditions	Noted
Section A.5.1	Network Operation Errors	Fully complies.
Section A.5.2	User Equipment Operation Errors	Noted
Section A.6	Optional bi-directional Network Procedures	Fully complies.
Section A.7	System Parameters	Fully complies

Compliance with frame relay UNI Interface Implementation Agreement

This feature complies with “Frame Relay Forum User to Network Interface Implementation Agreement, Document No FRF.1”.

Compliance with FRF.4

This feature complies with “Frame Relay Forum User-to-Network SVC Implementation Agreement, Document FRF TC 93.043R1 June, 1993.



This implementation agreement is a subset of ITU-T Recommendation Q.9333.

Frame relay compliance with FRF.4

Section	Section title	Compliance
Section 1	Introduction	Noted.
Section 2	Relevant Standards	Noted.
Section 3	Implementation Agreements	Sections listed below.
Section 3.1	Physical Layer Interface Guidelines	Compliant with exceptions. Refer to Port Management System Functional Specification.
Section 3.2	Data Transfer	Fully complies.
Section 3.3	Permanent Virtual Connection (PVC) Management Procedures	Fully complies. Q.933 Annex A and ANSI T1.617 Annex D are the supported options.
Section 3.4	Switched Virtual Connection (SVC) Procedures	Sections listed below.
Section 3.4.1	Purpose, scope and structure	Fully complies. The optional switched ISDN access is not supported.
Section 3.4.2	Overview of frame mode connection control	Fully complies.
Section 3.4.3	Messages functional definitions	Fully complies. The network does not initiate STATUS ENQUIRY message.
Section 3.4.4	General message format and information element coding	Fully complies.
Section 3.4.5	Frame mode call control procedures	Fully complies.
FRFTC/96-025	Attachment A to FRF. 4	Fully complies.
FRFTC/96-025	Attachment B to FRF. 4	Fully complies.

Compliance with ITU-T Q.922

Nortel Multiservice Switch frame relay complies with ITU-T (formerly CCITT) standards Q.922.



Frame relay compliance to ITU-T Q.922

Section	Section title	Compliance
Section 1	General	Noted.
Section 2	Frame-structure for peer-to-peer communication	Sections listed below.
Section 2.1	General	Fully complies.
Section 2.2	Flag sequence	Fully complies.
Section 2.3	Address Field	Fully complies. Two-octet DLCIs supported only.
Section 2.4	Control Field	Fully complies.
Section 2.5	Information Field	Fully complies.
Section 2.6	Transparency	Fully complies.
Section 2.7	Frame check sequence	Fully complies.
Section 2.8	Format convention	Fully complies.
Section 2.9	Invalid frames	Fully complies.
Section 3	Elements of Procedures and Formats of Fields for Data Link Layer Peer-to-Peer Communication	Sections listed below.
Section 3.1	General	Noted.
Section 3.2	Address field format	Fully complies. Two-octet DLCIs supported only.
Section 3.3	Address field variables	Fully complies. Two-octet DLCIs supported only.
Section 3.4	Control field formats	Fully complies.
Section 3.5	Control field parameters and associated state variables	Fully complies.
Section 3.6	Frame types	Fully complies. XID frames not supported.
Section 4	Elements for Layer-to-Layer Communication	Noted.
Section 5	Definition of the Peer-to-Peer Procedures of the Data Link Layer	Sections listed below.
Section 5.1	Procedure for the P/F bit	Fully complies.
Section 5.2	Procedure for unacknowledged information transfer	Fully complies.
Section 5.3	DLCI management	Fully complies.
(1 of 2)		



Frame relay compliance to ITU-T Q.922 (continued)

Section	Section title	Compliance
Section 5.4	Automatic negotiation of data link layer parameters	Fully complies. Default values are used and parameters are not negotiated. XID not supported.
Section 5.5	Procedures for establishment and release of multiple frame operation	Fully complies.
Section 5.6	Procedures for information transfer in multiple frame operation	Fully complies. CLLM not supported.
Section 5.7	Re-establishment of multiple frame operation	Fully complies.
Section 5.8	Exception condition reporting and recovery	Fully complies.
Section 5.9	List of system parameters	Fully complies. Parameters are not negotiated.
(2 of 2)		

Compliance with ITU-T X36

The basis for multiple CUGs is the ITU-T X36 amendment 1, section 10.5.11, "Closed user group".

The basis for the CUG functionality is the ITU-T X36 section 10.6.8.

Frame relay compliance to ITU-T X36

Section	Section title	Compliance
Section 10.6.8.1	general	complies
Section 10.6.8.2	subscription options	complies
Section 10.6.8.3	per call options	complies
Section 10.6.8.4	simple CUG	complies
Section 10.6.8.5	CUG selection	complies
		The signalPreferentialCugToLink attribute under the Dna component controls CUG signaling to the called DTE.
Section 10.6.8.6	no CUG	complies



Compliance with FRF.12

Nortel Multiservice Switch 15000 and Multiservice Switch 20000 node's implementation of multilink frame relay (MLFR) as described in this document is compliant with FRF.12, "Frame Relay PVC Fragmentation Implementation Agreement.", December 1997, Frame Relay Forum Technical Committee.

Nortel Multiservice Switch 7400, Multiservice Switch 15000, and Nortel Multiservice Switch 20000 node's implementation of end-to-end FRF.12 DTE as described in this document is compliant with FRF.12, "Frame Relay PVC Fragmentation Implementation Agreement", December 1997, Frame Relay Forum Technical Committee.

Compliance with FRF.16

Nortel Multiservice Switch 15000 and Multiservice Switch 20000 implementation of multilink frame relay (MLFR) as described in this document is compliant with FRF.16, "Multilink Frame Relay UNI/NNI Implementation Agreement.", August 1999, Frame Relay Forum Technical Committee.



FR NNI compliance with standards

Use the following sections to learn more about FR NNI compliance with standards.

Navigation

- [Detailed compliance statement \(page 273\)](#)
- [Standards guiding frame relay development \(page 274\)](#)
- [Compliance with Annex D of T1.617 \(page 274\)](#)
- [Compliance with Annex A of ITU-T Q.933 \(page 275\)](#)
- [Compliance with Frame Relay NNI Phase 1 Implementation Agreement \(page 276\)](#)
- [Compliance with FRF.10 \(subset of ITU-T X.76\) \(page 277\)](#)
- [Compliance with ITU-T X.76 \(year 2000 revision\) \(page 278\)](#)
- [Compliance with FRF.12 \(page 286\)](#)
- [Compliance with FRF.16 \(page 287\)](#)

Detailed compliance statement

The following terms indicate compliance with the joint specification:

- **Noted:** The specification provides clarification, non-specific information, or details that do not relate directly to the frame relay service.
- **Fully complies:** Nortel Multiservice Switch frame relay service functionality fully complies with the text for this section.
- **Compliant with these exceptions:** Multiservice Switch frame relay service does not comply completely with the text.
- **Not supported:** Nortel Multiservice Switch frame relay service does not support the functionality described in the text.



Standards guiding frame relay development

Nortel Multiservice Switch frame relay service conforms to the following standards and implementation agreements:

- American National Standard Institute for Telecommunications (ANSI) T1.617 Annex D (as recited in the Frame Relay Forum Implementation agreement)
- International Telecommunications Union (ITU) Q.933 Annex A ((as recited in the Frame Relay Forum Implementation agreement)
- Frame Relay Forum Technical Committee Phase 1 Implementation Agreement FRF.2 (for NNI only)
- Frame Relay Forum FRF.10 (subset of ITU-T X.76)
- [Compliance with ITU-T X.76 \(year 2000 revision\) \(page 278\)](#)

Compliance with Annex D of T1.617

Nortel Multiservice Switch frame relay complies with American National Standard for Telecommunications (ANSI) Annex D of T1.617.

Frame relay compliance with Annex D of T1.617 ANSI (June 1991)

Section	Section title	Compliance
Section D.1	Messages Used for PVC Status	Fully complies
Section D.1.1	Status	Fully complies. Asynchronous PVC status report also supported.
Section D.1.2	Status Enquiry	Fully complies
Section D.2.1	Protocol Discriminator	Fully complies
Section D.2.2	Call Reference	Fully complies
Section D.2.3	Message Type	Fully complies
Section D.3.1	Report Type	Fully complies
Section D.3.2	Link Integrity Verification	Fully complies
Section D.3.3	PVC Status	Fully complies. Only two-octet DLCIs supported.
Section D.4	Procedures	Noted
Section D.4.1	Periodic Polling	Fully complies
Section D.4.2	Link Integrity Verification	Fully complies
Section D.4.3	Reporting New PVCs	Fully complies
Section D.4.4	Reporting the Availability of a PVC	Fully complies
Section D.5	Error Conditions	Noted
(1 of 2)		



Frame relay compliance with Annex D of T1.617 ANSI (June 1991) (continued)

Section	Section title	Compliance
Section D.5.1	Network Operation Errors	Fully complies
Section D.5.2	User Equipment Operation Errors	Noted
Section D.6	Optional Bi-directional Network Procedures	Fully complies
Section D.7	System Parameters	Default for N392 is 3. Possible values of T392 range from 5 to 30 seconds in increments of 5 seconds. Default for T392 is 15.
(2 of 2)		

Compliance with Annex A of ITU-T Q.933

Nortel Multiservice Switch frame relay complies with ITU-T (formerly CCITT) standards Annex A of Q.933.

Frame relay compliance with Annex A of ITU-TQ.933 (1992)

Section	Section title	Compliance
Section A.1	Messages Used for PVC Status	Fully complies
Section A.1.1	Status	Fully complies. Asynchronous PVC status report also supported.
Section A.1.2	Status Enquiry	Fully complies
Section A.2.1	Protocol Discriminator	Fully complies
Section A.2.2	Call Reference	Fully complies
Section A.2.3	Message Type	Fully complies
Section A.3.1	Report Type	Fully complies
Section A.3.2	Link Integrity Verification	Fully complies
Section A.3.3	PVC Status	Fully complies. Only two-octet DLCIs supported.
Section A.4	Procedures	Noted
Section A.4.1	Periodic Polling	Fully complies
Section A.4.2	Link Integrity Verification	Fully complies
Section A.4.3	Reporting New PVCs	Fully complies
Section A.4.4	Reporting the Availability of a PVC	Fully complies
Section A.5	Error Conditions	Noted
Section A.5.1	Network Operation Errors	Fully complies
(1 of 2)		



Frame relay compliance with Annex A of ITU-TQ.933 (1992) (continued)

Section	Section title	Compliance
Section A.5.2	User Equipment Operation Errors	Noted
Section A.6	Optional Bi-directional Network Procedures	Fully complies
Section A.7	System Parameters	Default for N392 is 3. Possible values of T392 range from 5 to 30 seconds in increments of 5 seconds. Default for T392 is 15.
(2 of 2)		

Compliance with Frame Relay NNI Phase 1 Implementation Agreement

The frame relay service complies with “Frame Relay Network to Network Interface Phase 1 Implementation Agreement, Document No FRF.2” by Frame Relay Forum Technical Committee, August 21, 1992.

Frame relay compliance with FRF.2

Section	Section title	Compliance
Section 1.0	Introduction	Noted
Section 2.0	Relevant Standards	Noted
Section 3.0	Implementation Agreements	Noted
Section 3.1	Physical Layer Interface Guidelines	Noted
Section 3.2	Data Transfer	Noted
Section 3.2.1	Interface Time Fill	Fully complies
Section 3.2.2	Frame Relay Information Field	Fully complies
Section 3.2.3	Congestion Management	Fully complies
Section 3.2.4	Congestion Management	Fully complies
Section 3.2.5	Consolidated Link Layer Management (CLLM) message	Noted
Section 3.3	Control (Signalling) Procedure	Noted
Section 3.3.1	Permanent Virtual Connection (PVC) Procedures	Fully complies
Section 3.3.2	Switch Virtual Connection (SVC) Procedures	Noted
Section 3.4	Network Performance parameters	Noted
Section 3.5	PVC Parameter Coordination	Fully complies
(1 of 2)		



Frame relay compliance with FRF.2 (continued)

Section	Section title	Compliance
Section 4	Application of Bidirectional Procedures for Use at the Network-to-Network Interface	Noted
Section 4.1	Introduction	Noted
Section 4.2	Bidirectional Network Procedures and Multi-network PVCs	Fully complies
Section 4.3	Polling Requirements of Network-to-Network Interface	Compliant with Asynchronous PVC status report also supported.
Section 4.4	Initial NNI Status	Compliant with the first criterion is supported.
Section 4.5	Multi-Network PVC Active Status Criteria	Fully complies
Section 4.6	Adding a Multi-Network PVC	Fully complies
Section 4.7	Deleting a Multi-Network PVC	Fully complies
Section 4.8	Response to UNI Failure	Fully complies
Section 4.9	Response to PVC Segment Failure	Fully complies
Section 4.10	Response to NNI Failure	Fully complies
Section 4.11	Examples of Multi-Network PVC Status Signalling	Noted and complied with
(2 of 2)		

Compliance with FRF.10 (subset of ITU-T X.76)

Nortel Multiservice Switch frame relay NNI SVC feature complies with FRF.10, "Frame Relay Network-to-Network SVC Implementation Agreement", September 1996.

Frame relay compliance with FRF.10 (1996)

Section	Section title	Compliance
Section 1	Introduction	Noted.
Section 2	Relevant Standards	Noted.
Section 3	Implementation Agreements	
Section 3.1	Scope of Agreement	Noted.
Section 3.2	Physical Layer	Compliant with exceptions.
Section 3.3	Data Transfer	Compliant with exceptions. DLCI values of 2 octets only are supported.
Section 3.4	PVC Management Procedures	Fully complies.
(1 of 2)		



Frame relay compliance with FRF.10 (1996) (continued)

Section	Section title	Compliance
Section 3.5	SVC Signaling Procedures	Fully complies.
Section 3.5.1	Numbering Plans	Fully complies.
Section 3.5.2	General	Fully complies.
Section 3.5.3	State Definitions	Fully complies.
Section 3.5.4	Message Definitions	Fully complies.
Section 3.5.5	General Message Format & IE Coding	Fully complies.
Section 3.5.6	Procedures	Fully complies.
Section 3.6	Switched PVCs	Compliant with exceptions. Single-ended configuration is not supported. Only a UNI or NNI P-SPVC master endpoint on a Multiservice Switch node can initiate an SPVC call. If a SPVC master endpoint receives a call request from an external network that supports call initiation by a called endpoint, the master endpoint will clear the call. Multiservice Switch node only supports a selection type of "specific DLCI". Network call requests that specify "any DLCI" or "Specific SPVC Correlator" are cleared at the SPVC slave endpoint.
(2 of 2)		

Compliance with ITU-T X.76 (year 2000 revision)

Nortel Multiservice Switch nodes' implementation of FRF.10 ("*Frame Relay Network-to-Network SVC Implementation Agreement*", September 1996) is compliant with ITU-T X.76.

Frame relay compliance with ITU-T X.76 (Year 2000 revision)

Section	Section title	Compliance
Section 10	Frame relay SVC signaling	
Section 10.1	General	Noted.
(1 of 9)		



Frame relay compliance with ITU-T X.76 (Year 2000 revision) (continued)

Section	Section title	Compliance
Section 10.2	Signaling channel	Compliant. The following signaling channel link layer parameters can be configured on each Multiservice Switch node's NNI: • T200 • T203 • N200 • k • N201
Section 10.3.1	Frame relay call states at the NNI	Compliant with exceptions. The following ISDN states defined in Q.933 (and not used in Recommendation X.36 for frame relay UNI) are not supported: • Call delivered (NN4) • Call received (NN7)
Section 10.3.2	States associated with restart	Fully complies.
Section 10.4	Message definitions	Noted.
Section 10.4.1	Alerting	Not compliant. ISDN Q.933 message not used in Recommendation X.36.
Section 10.4.2	Call proceeding	Fully complies.
Section 10.4.3	Connect	Compliant with exceptions. 1. The following ISDN information elements defined in Q.933 (and not used in Recommendation X.36 for frame relay UNI) are not supported: • progress indicator • end-to-end transit delay • packet layer binary parameter • link layer protocol parameter • X.213 priority 2. Transit network identification 3. Generic application transport (new to X.76 year 2000 revision)
Section 10.4.4	Progress	Not compliant. ISDN Q.933 message not used in Recommendation X.36.
(2 of 9)		



Frame relay compliance with ITU-T X.76 (Year 2000 revision) (continued)

Section	Section title	Compliance
Section 10.4.5	Release	Compliant with exceptions: • transit network identification information element • clearing network identification • generic application transport (new to X.76 year 2000 revision)
Section 10.4.6	Release complete	Compliant with exceptions: • transit network identification information element • clearing network identification • generic application transport (new to X.76 year 2000 revision)
Section 10.4.7	Restart	Compliant with exceptions: • Multiservice Switch nodes' support the Restart indicator according to the current version of X.76, but does not generate it.
Section 10.4.8	Restart acknowledge	Compliant with exceptions: • Multiservice Switch nodes' support the Restart indicator according to the current version of X.76, but does not generate it.
(3 of 9)		



Frame relay compliance with ITU-T X.76 (Year 2000 revision) (continued)

Section	Section title	Compliance
Section 10.4.9	Setup	Compliant with exceptions. 1. The following ISDN information elements defined in Q.933 are not used in Recommendation X.36: • progress indicator • end-to-end transit delay • packet layer binary parameter • link layer protocol parameter • X.213 priority 2. Reverse charging indication 3. Transit network identification 4. Call identification 5. Low layer compatibility (only one instance according to the existing version of X.76) 6. Generic application transport (new to X.76 year 2000 revision)
Section 10.4.10	Status	Fully complies.
Section 10.4.11	Status enquiry	Fully complies.
Section 10.5	General message format and IE coding	Noted.
Section 10.5.1	Protocol discriminator	Fully complies.
Section 10.5.2	Call reference	Fully complies.
Section 10.5.3	Message type	Compliant with exceptions: • alerting • progress
Section 10.5.4	Bearer capability	Fully complies.
Section 10.5.5	Call identification	Not compliant.
Section 10.5.6	Call state	Compliant with exceptions: • Call delivered (NN4) • Call received (NN7)
(4 of 9)		



Frame relay compliance with ITU-T X.76 (Year 2000 revision) (continued)

Section	Section title	Compliance
Section 10.5.7	Called party number	Compliant with exceptions: - type of number = alternative address (101) - numbering plan identification = ISO NSAP (0001)
Section 10.5.8	Called party sub-address	Fully complies.
Section 10.5.9	Calling party number	Compliant with exceptions: - type of number = alternative address (101) - numbering plan identification = ISO NSAP (0001)
Section 10.5.10	Calling party sub-address	Fully complies.
Section 10.5.11	Cause	Fully complies.
Section 10.5.8	Clearing network identification	Not supported.
Section 10.5.13	Closed user group interlock code	Compliant with exception: interlock code using E.164 country code
Section 10.5.14	Connected number	Fully complies.
Section 10.5.15	Connected sub-address	Fully complies.
Section 10.5.16	Data link connection identifier	Compliant with exception: DLCI values of 2 octets only are supported
Section 10.5.17	End-to-end transit delay	Not supported. ISDN Q.933 information element not used in Recommendation X.36.
Section 10.5.17 (bis)	Generic application transport	Not supported. New to X.76 year 2000 revision.
Section 10.5.18	Higher layer compatibility	Fully complies.
Section 10.5.19	Link layer core parameters	Fully complies.
Section 10.5.20	Link layer protocol parameters	Not supported. ISDN Q.933 information element not used in Recommendation X.36.
Section 10.5.21	Low layer compatibility	Fully complies.
(5 of 9)		



Frame relay compliance with ITU-T X.76 (Year 2000 revision) (continued)

Section	Section title	Compliance
Section 10.5.22	Packet layer binary parameters	Not supported. ISDN Q.933 information element not used in Recommendation X.36.
Section 10.5.23	Priority and service class parameters	Compliant with exception: service class
Section 10.5.24	Progress indicator	Not supported. ISDN Q.933 information element not used in Recommendation X.36.
Section 10.5.2x	Restart indicator	Not supported. Information element removed from X.76 (year 2000 revision) but is present in the existing X.76 Recommendation (1997 version).
Section 10.5.25	Reverse charging indicator	Not supported.
Section 10.5.26	Transit network identification	Not supported.
Section 10.5.27	Transit network selection	Compliant with exception: carrier/network identification using E.164 country code
Section 10.5.28	User-user	Not supported. ISDN Q.933 information element not used in Recommendation X.36.
Section 10.6.1	Call establishment at the calling STE	Compliant as per the following subsections.
Section 10.6.1.1	Initiating a call setup request	Fully complies.
Section 10.6.1.2	Call proceeding	Fully complies.
Section 10.6.1.3	Alerting and call progressing	Not supported as per the following subsections.
Section 10.6.1.3.1	Handling of call alerting	Not supported. ISDN Q.933 information element not used in Recommendation X.36.
Section 10.6.1.3.2	Handling of call progressing	Not supported. ISDN Q.933 information element not used in Recommendation X.36.
Section 10.6.1.4	Call established	Fully complies.
Section 10.6.2	Call establishment at the called STE	Compliant as per the following subsections.
Section 10.6.2.1	Receiving a call setup request	Fully complies.
(6 of 9)		



Frame relay compliance with ITU-T X.76 (Year 2000 revision) (continued)

Section	Section title	Compliance
Section 10.6.2.2	Alerting and call progressing	Not supported as per the following subsections.
Section 10.6.2.2.1	Handling of call alerting	Not supported. ISDN Q.933 information element not used in Recommendation X.36.
Section 10.6.2.2.2	Handling of call progressing	Not supported. ISDN Q.933 information element not used in Recommendation X.36.
Section 10.6.2.3	Call established	Fully complies.
Section 10.6.3	Normal call clearing	Compliant as per the following subsections.
Section 10.6.3.1	Initiation of the clearing of a call	Fully complies.
Section 10.6.3.2	Receipt of a release message	Fully complies.
Section 10.6.3.3	Clearing in the null state	Fully complies.
Section 10.6.3.4	Clearing collision	Fully complies.
Section 10.6.4	Restart procedures	Compliant as per the following subsections.
Section 10.6.4.1	Sending a restart message	Fully complies.
Section 10.6.4.2	Receipt of a restart message	Fully complies.
Section 10.6.5	Status enquiry and status procedures	Compliant as per the following subsections.
Section 10.6.5.1	Status enquiry procedure	Fully complies.
Section 10.6.5.2	Receiving a status message	Fully complies.
Section 10.6.5.3	Receipt of the status message with the global call reference	Fully complies.
Section 10.6.6	Handling of error conditions	Fully complies.
Section 10.6.7	DLCI management	Compliant as per the following subsections.
Section 10.6.7.1	DLCI allocation between SVC and PVC	Fully complies.
Section 10.6.7.2	DLCI collision at the NNI	Fully complies.
Section 10.6.8	List of timers at the NNI	Fully complies.
Section 10.6.9	Frame relay NNI facilities	Compliant as per the following subsections.
Section 10.6.9.1	Transit network identification	Not supported.
(7 of 9)		



Frame relay compliance with ITU-T X.76 (Year 2000 revision) (continued)

Section	Section title	Compliance
Section 10.6.9.2	Call identification	Not supported.
Section 10.6.9.3	Closed user group interlock code	Fully complies.
Section 10.6.9.4	Reverse charging indication	Not supported.
Section 10.6.9.5	Clearing network indication	Not supported.
Section 10.6.9.6	Transit network selection	Fully complies.
Section 10.6.9.7.1	Frame transfer priority	Fully complies.
Section 10.6.9.7.2	Frame discard priority	Fully complies.
Section 10.6.9.7.3	Procedures for priorities	Fully complies.
Section 10.6.9.7.4	Frame relay service class	Not supported.
Section 10.6.9.7.4.4	Interaction between service class and priorities	Not supported.
Section 10.6.9.8	Generic application transport	Not supported (new to X.76 year 2000 revision).
Annex A	Signaling for switched PVC (SPVC)	Introductory text: Noted.
A.1	Messages needed for SPVC establishment	Noted.
A.2	Called party SPVC information element	Compliant with exceptions: • specific SPVC correlator, Any DLCI, Assigned DLCI, ATM endpoint of the called endpoint selection • ATM called endpoint selection type • VPI and VCI identifiers • DLCI lengths other than 10 bits
A.2	Calling party SPVC information element	Compliant with exceptions: • VPI and VCI identifiers • DLCI lengths other than 10 bits
A.3.1	Initiating SPVC establishment	Compliant with exception: • single-endpoint SPVC establishment is not supported
(8 of 9)		



Frame relay compliance with ITU-T X.76 (Year 2000 revision) (continued)

Section	Section title	Compliance
A.3.2	Receiving a setup message at the called endpoints	Fully complies.
A.3.2.1	Call collision	Compliant with exception. Only a master endpoint on a Multiservice Switch node can initiate SPVC setup. If a master endpoint receives a SPVC setup request initiated by another node, it rejects the setup request.
A.3.2.2	Confirmation of SPVC configuration	Compliant with exceptions: specific SPVC correlator, any DLCI, assigned DLCI.
A.3.2.3	Allocation of DLCI at called PVC UNI	Compliant with exceptions: specific SPVC correlator, any DLCI, assigned DLCI.
A.3.2.4	Called endpoint availability	Fully complies.
A.3.2.4.1	Incompatible destination error case	Not supported. ATM endpoint is not supported.
A.3.3	Receiving a connect message	Fully complies.
A.3.4	Receiving a release or release complete message	Fully complies.
A.3.5	Coordination with PVC signaling procedures	Fully complies.
A.3.5.1	PVC addition	Compliant with exception: any DLCI and specific DLCI.
A.3.5.2	PVC deletion	Fully complies.
A.3.5.3.1	PVC availability -- poll response	Compliant with exception: any DLCI and specific DLCI.
A.3.5.3.2	PVC availability -- poll initiation	Fully complies.
Annex B	Usage of cause and location	Compliant as per the following subsections.
B.1	Location field generation	Fully complies.
B.2	Cause values	Fully complies.
B.3	Coding of diagnostic field	Fully complies.
(9 of 9)		

Compliance with FRF.12

Nortel Multiservice Switch 15000 and Multiservice Switch 20000 nodes' implementation of multilink frame relay (MLFR) as described in this document is compliant with FRF.12, "Frame Relay PVC Fragmentation Implementation Agreement.", December 1997, Frame Relay Forum Technical Committee.



Compliance with FRF.16

Nortel Multiservice Switch 15000 and Multiservice Switch 20000 nodes' implementation of multilink frame relay (MLFR) as described in this document is compliant with FRF.16, "Multilink Frame Relay UNI/NNI Implementation Agreement.", August 1999, Frame Relay Forum Technical Committee.



Australian TS014 compliance statement

Use the following sections to learn more about Nortel Multiservice Switch frame relay ISDN switched access Compliance with (TS014) *Technical Standard 014: General Requirements for Customer Equipment Connected to an ISDN Primary Rate Interface*, issued by AUSTEL - Australian Telecommunication Authority, 1995.

The Technical Standard 014 (TS014) is the Australian specification that is based on ISDN ITU-T Recommendations and that includes Recommendation Q.921 (for layer 2), Recommendation Q.931 Red Book (for layer 3), and all related layer 1 ITU-T recommendations.

TS014 specifies the minimum requirements for customer equipment (CE) that connects to the public network by way of an integrated service digital network (ISDN) primary rate access interface, at the T reference point. TS014 also specifies conformance testing procedures for testing the compliance of the CE to this standard.

Navigation

- [Terminology \(page 288\)](#)
- [General information and definitions \(page 289\)](#)
- [General requirements \(page 289\)](#)
- [Layer 1 requirements \(page 290\)](#)
- [Layer 2 requirements \(page 290\)](#)
- [Layer 3 requirements \(page 296\)](#)
- [Compatibility requirements \(page 303\)](#)
- [Appendices requirements \(page 304\)](#)

Terminology

In this appendix, the following terms define Compliance with TS014:

- Noted. The standard provides background information or clarification.



- Fully complies. Nortel Multiservice Switch frame relay ISDN switched access fully complies with this section.
- Complies, with exception. Multiservice Switch frame relay ISDN switched access does not completely comply with the text. Exceptions are listed.
- Not supported. Nortel Multiservice Switch frame relay ISDN switched access does not support this section of the standard.
- N/A. Not applicable to Nortel Multiservice Switch frame relay ISDN switched access.
- Title. This is a clause number and title of the standard.

General information and definitions

The clauses related to this group are in TS014 chapters 2 to 4.

Compliance with TS014 general Information and definitions

Clause	Title	Compliance
2	Interpretative Guidelines	Title
2.1	Categories of Requirements	Noted
2.2	Compliance Statements	Noted
2.3	Definitions, Expressions and Terms	Noted
2.4	Notations	Noted
2.5	Notes	Noted
2.6	References	Noted
2.7	Units and Symbols	Noted
3	Scope	Noted
4	Abbreviations and Technical Definitions	Title
4.1	Abbreviations	Noted

General requirements

The clauses related to this group are in TS014 chapter 5, paragraph 1.

Compliance with TS014 general requirements

Clause	Title	Compliance
5.1	General	Noted
5.1.1	Safety	Noted
(1 of 2)		



Compliance with TS014 general requirements (continued)

Clause	Title	Compliance
5.1.2	Identification Markings	Noted
5.1.3	Fail-safe Operation	Noted
5.1.4	Operating Voltage and Environmental Conditions	Noted
5.1.5	Line Polarity and Line Conductor Polarization	Noted
5.1.6	Emission of Electromagnetic Interference	Noted
5.1.7	Voice Frequency Performance	N/A for data
(2 of 2)		

Layer 1 requirements

See NN10600-170 *Nortel Multiservice Switch 7400 Hardware Description* for details on how the DS1C, E1C and 32-port MSA function processors comply to the applicable sections of the ISO physical layer (layer 1) specification.

Layer 2 requirements

The clauses related to the ISO network layer (Layer 2) specification are in TS014 chapter 5, paragraph 3.

Compliance with TS014 Layer 2 - data link layer requirements

Clause	Title	Compliance
5.3	Data Link Layer - Layer 2	Title
5.3.1	General	Title
5.3.1.1	Link Access Procedure	Noted
5.3.1.2	Concepts and Terminology	Noted
5.3.1.3	Overview of LAPD Functions and Procedures	Title
5.3.1.3.1	Messages	Noted
5.3.1.3.2	Acknowledged Operation	Noted
5.3.1.3.3	Data Link Connection Identification	Noted
5.3.1.3.4	Data Link States	Noted
5.3.1.3.5	Establishment of Multiple Frame Operation	Noted
(1 of 7)		



Compliance with TS014 Layer 2 - data link layer requirements (continued)

Clause	Title	Compliance
5.3.1.4	Service Characteristics	Title
5.3.1.4.1	Services Provided to Layer 3	Noted
5.3.1.4.2	Acknowledged Information Transfer Service	Noted
5.3.1.4.3	Administrative Services	Noted
5.3.1.4.4	Model of the Data Link Service	Title
5.3.1.4.4.1	General	Noted
5.3.1.4.5	Data Link Layer Representation as Seen by Layer 3	Title
5.3.1.4.5.1	Data Link Connection Endpoint States	Noted
5.3.1.4.5.2	Point-to-point Data Link Connection Endpoint Services	Noted
5.3.1.4.5.3	Sequences of Primitives at Point-to-point Data Link Connection Endpoint	Noted
5.3.1.4.5.4	Services Required from the Physical Layer	Noted
5.3.1.5	Data Link - Management Layer Structure	Title
5.3.1.5.1	General	Noted
5.3.1.5.2	Data Link Procedure	Noted
5.3.1.5.3	Multiplex Procedure	Noted
5.3.1.5.4	Structure of Data Link Procedure	Noted
5.3.1.6	Testing	Noted
5.3.2	Data Link Layer Specification	Title
5.3.2.1	General	Noted
5.3.2.2	Frame Structure for Peer-to-peer Communication	Title
5.3.2.2.1	Formats	Fully complies
5.3.2.2.2	Flag Sequence	Fully complies
5.3.2.2.3	Address Field	Fully complies
5.3.2.2.4	Control Field	Fully complies
5.3.2.2.5	Information Field	Fully complies
(2 of 7)		



Compliance with TS014 Layer 2 - data link layer requirements (continued)

Clause	Title	Compliance
5.3.2.2.6	Transparency	Fully complies
5.3.2.2.7	Frame Checking Sequence (FCS) Field	Fully complies
5.3.2.2.8	Format Convention	Title
5.3.2.2.8.1	Numbering Convention	Noted
5.3.2.2.8.2	Order of Bit Transmission	Fully complies
5.3.2.2.8.3	Field Mapping Convention	Fully complies
5.3.2.2.9	Invalid Frames	Fully complies
5.3.2.2.10	Frame Abort	Fully complies
5.3.2.2.11	InterFrame Timefill	Fully complies
5.3.2.3	Elements of Procedures and Formats of Field for Data Link Layer Peer-to-peer Communication	Title
5.3.2.3.1	General	Noted
5.3.2.3.2	Address Field Format	Fully complies
5.3.2.3.3	Address Field Variables	Title
5.3.2.3.3.1	Address Field Extension Bit (EA)	Fully complies
5.3.2.3.3.2	Command/Response (C/R) Field Bit	Fully complies
5.3.2.3.3.3	Service Access Point Identifier (SAPI)	Fully complies
5.3.2.3.3.4	Terminal Endpoint Identifier	Fully complies
5.3.2.3.3.4.1	TEI for Point-to-point Data Link Connection	Fully complies
5.3.2.3.4	Control Field Format	Fully complies
5.3.2.3.4.1	Information Transfer Format - I	Fully complies
5.3.2.3.4.2	Supervisory Format - S	Fully complies
5.3.2.3.4.3	Unnumbered Format - U	Fully complies
5.3.2.3.5	Control Field Parameters and Associated State Variables	Noted
5.3.2.3.5.1	Pool/Final (P/F) Bit	Fully complies
5.3.2.3.5.2	Multiple Frame Operation - Variable and Sequence Numbers	Title
(3 of 7)		



Compliance with TS014 Layer 2 - data link layer requirements (continued)

Clause	Title	Compliance
5.3.2.3.5.2.1	Modulus	Fully complies
5.3.2.3.5.2.2	Send State Variable V(S)	Fully complies
5.3.2.3.5.2.3	Acknowledge State Variable V(A)	Fully complies
5.3.2.3.5.2.4	Send Sequence Number N(S)	Fully complies
5.3.2.3.5.2.5	Receive State Variable V(R)	Fully complies
5.3.2.3.5.2.6	Receive Sequence Number N(R)	Fully complies
5.3.2.3.6	Frame Types	Title
5.3.2.3.6.1	Commands and Responses	Fully complies
5.3.2.3.6.2	Information (I) Command	Fully complies
5.3.2.3.6.3	Set Asynchronous Balanced Mode Extended (SABME) Command	Fully complies
5.3.2.3.6.4	Disconnect (DISC) Command	Fully complies
5.3.2.3.6.5	Receive Ready (RR) Command/Response	Fully complies
5.3.2.3.6.6	Reject (REJ) Command/Response	Fully complies
5.3.2.3.6.7	Receive Not Ready (RNR) Command/Response	Fully complies
5.3.2.3.6.8	Unnumbered Acknowledgment (UA) Response	Fully complies
5.3.2.3.6.9	Disconnected Mode (DM) Response	Fully complies
5.3.2.3.6.10	Frame Reject (FRMR) Response	Fully complies
5.3.2.4	Elements of Layer-to-layer Communication	Title
5.3.2.4.1	General	Title
5.3.2.4.1.1	Primitives	Noted
5.3.2.4.1.2	Generic Names	Noted
5.3.2.4.1.3	Primitive Types	Noted
5.3.2.4.1.4	Parameter Definition	Title
5.3.2.4.1.4.1	Priority Indicator	Noted
5.3.2.4.1.4.2	Message Unit	Noted
5.3.2.4.2	Primitive Procedure	Title
5.3.2.4.2.1	General	Noted
(4 of 7)		



Compliance with TS014 Layer 2 - data link layer requirements (continued)

Clause	Title	Compliance
5.3.2.4.2.2	Layer 3 - Data Link Layer Interaction	Noted
5.3.2.5	Peer-to-peer Procedures of the Data Link Layer	Title
5.3.2.5.1	Elements	Noted
5.3.2.5.2	Procedure of the Use of the P/F Bit	Fully complies
5.3.2.5.3	Automatic Negotiation of Data Link Layer Parameter Values	Fully complies
5.3.2.5.4	Procedure for Establishment and Release of Multiple Frame Operation	Title
5.3.2.5.4.1	Establishment of Multiple Frame Operation	Title
5.3.2.5.4.1.1	General	Noted
5.3.2.5.4.1.2	Establishment Procedures	Fully complies
5.3.2.5.4.1.3	Procedure on Expiry of Timer T200	Fully complies
5.3.2.5.4.2	Information Transfer	Fully complies
5.3.2.5.4.3	Termination of Multiple Frame Operation	Title
5.3.2.5.4.3.1	General	Noted
5.3.2.5.4.3.2	Release Procedure	Fully complies
5.3.2.5.4.3.3	Procedure on Expiry of Timer T200	Fully complies
5.3.2.5.4.4	TEI-Assigned State	Fully complies
5.3.2.5.4.5	Collision of Unnumbered Commands and Responses	Title
5.3.2.5.4.5.1	Identical Transmitted and Received Commands	Fully complies
5.3.2.5.4.5.2	Different Transmitted and Received Commands	Fully complies
5.3.2.5.4.5.3	Unsolicited DM Response and SABME or DISC Command	Fully complies
5.3.2.5.5	Procedure for Information Transfer in Multiple Frame Operation	Title
5.3.2.5.5.1	Transmitting I Frame	Fully complies
5.3.2.5.5.2	Receiving I Frame	Fully complies
(5 of 7)		



Compliance with TS014 Layer 2 - data link layer requirements (continued)

Clause	Title	Compliance
5.3.2.5.5.2.1	P Bit set to 1	Fully complies
5.3.2.5.5.2.2	P Bit set to 0	Fully complies
5.3.2.5.5.3	Sending and Receiving Acknowledgments	Title
5.3.2.5.5.3.1	Sending Acknowledgments	Fully complies
5.3.2.5.5.3.2	Receiving Acknowledgment	Fully complies
5.3.2.5.5.4	Receiving REJ Frames	Fully complies
5.3.2.5.5.5	Receiving RNR Frames	Fully complies
5.3.2.5.5.6	Data Link Layer Own Receiver Busy Condition	Fully complies
5.3.2.5.5.7	Waiting Acknowledgment	Fully complies
5.3.2.5.6	Re-establishment of Multiple Frame Operation	Title
5.3.2.5.6.1	Criteria for Re-establishment	Fully complies
5.3.2.5.6.2	Procedures	Fully complies
5.3.2.5.7	Exception Condition Reporting and Recovery	Noted
5.3.2.5.7.1	N(S) Sequence Error	Fully complies
5.3.2.5.7.2	N(R) Sequence Error	Fully complies
5.3.2.5.7.3	Timer Recovery Condition	Fully complies
5.3.2.5.7.4	Invalid Frame Condition	Fully complies
5.3.2.5.7.5	Frame Rejection Condition	Fully complies
5.3.2.5.7.6	Receipt of an FRMR Response Frame	Fully complies
5.3.2.5.7.7	Unsolicited Response Frames	Fully complies
5.3.2.5.8	System Parameters	Title
5.3.2.5.8.1	General	Noted
5.3.2.5.8.2	Timer T200	Fully complies
5.3.2.5.8.3	Maximum Number of Retransmissions (N200)	Fully complies
5.3.2.5.8.4	Maximum Number of Octets in an Information Field (N201)	Fully complies
(6 of 7)		



Compliance with TS014 Layer 2 - data link layer requirements (continued)

Clause	Title	Compliance
5.3.2.5.8.5	Maximum Number of Outstanding I Frames (K)	Fully complies
5.3.2.5.8.6	Timer T203	Fully complies
5.3.2.5.9	Data Link Layer Monitor Function	Title
5.3.2.5.9.1	General	Noted
5.3.2.5.9.2	Data Link Layer Supervision in the Multiple-frame-establishment State	Fully complies
5.3.2.5.9.3	Connection Verification Procedures	Title
5.3.2.5.9.3.1	Start of Timer T203	Fully complies
5.3.2.5.9.3.2	Stop Timer T203	Fully complies
5.3.2.5.9.3.3	Expiry of Timer T203	Fully complies
5.3.3	An SDL Representation of the Point-to-point Procedure of the Data Link Layer Primary Rate Access	Title
5.3.3.1	General	Noted
5.3.3.2	An Overview of the State of the Point-to-point Data Link Layer Entity	Noted
5.3.3.3	The Use of the Queues	Fully complies
5.3.3.4	Occurrence of MDL-ERROR-INDICATION Within the Basic States	Fully complies
(7 of 7)		

Layer 3 requirements

The clauses related to the ISO network layer (Layer 3) specification are in TS014 chapter 5, paragraph 4.

Compliance with TS014 Layer 3 - general information and definitions

Clause	Title	Compliance
5.4	Layer 3	Title
5.4.1	General	Noted
5.4.2	Network Interface Specification Primary Rate Access	Title
5.4.2.1	General	Noted
(1 of 2)		



Compliance with TS014 Layer 3 - general information and definitions

Clause	Title	Compliance
5.4.2.2	Message Functions Definitions and Content	Title
5.4.2.2.1	Overview	Noted
5.4.2.2.2	Message Length	Fully complies
5.4.2.2.3	Messages	Noted
(2 of 2)		

Compliance with TS014 Layer 3 - message definitions

Clause	Title	Compliance
5.4.2.2.3.1	ALERTING • receipt of • generation of	• Not supported • Not supported The receipt of ALERTING can only occur in an outgoing call (which is not supported)
5.4.2.2.3.2	CALL PROCEEDING • receipt of • generation of	• Complies, with exception: message triggers a STATUS message • Fully complies The receipt of CALL PROCEEDING can only occur in an outgoing call (which is not supported)
5.4.2.2.3.3	CONNECT • receipt of • generation of	• Complies, with exception: message triggers a STATUS message • Fully complies The receipt of CONNECT can only occur in an outgoing call (which is not supported)
(1 of 2)		



Compliance with TS014 Layer 3 - message definitions (continued)

Clause	Title	Compliance
5.4.2.2.3.4	CONNECT ACKNOWLEDGE • receipt of • generation of	• Fully complies • Complies, with exception: message triggers a STATUS message The generation of CONNECT ACKNOWLEDGE can only occur in an outgoing call (which is not supported)
5.4.2.2.3.5	DISCONNECT	Fully complies
5.4.2.2.3.6	INFORMATION	Not supported Overlap signaling not supported
5.4.2.2.3.7	PROGRESS	Not supported No support for interworking event
5.4.2.2.3.8	RELEASE	Fully complies
5.4.2.2.3.9	RELEASE COMPLETE	Fully complies
5.4.2.2.3.10	RESTART	Fully complies
5.4.2.2.3.11	RESTART ACKNOWLEDGE	Fully complies
5.4.2.2.3.12	SETUP	Complies, with exception: Calling Party Number is mandatory
5.4.2.2.3.13	SETUP ACKNOWLEDGE	Not supported No support for overlap signaling
5.4.2.2.3.14	STATUS	Fully complies
5.4.2.2.3.15	STATUS ENQUIRY	Fully complies
(2 of 2)		



Compliance with TS014 Layer 3 - message format and Information Elements coding

Clause	Title	Compliance
5.4.2.3	Message Format and Information Elements Coding	Title
5.4.2.3.1	Overview	Noted
5.4.2.3.2	Protocol Discriminator	Fully complies
5.4.2.3.3	Call Reference	Fully complies
5.4.2.3.4	Message Type	Fully complies
5.4.2.3.5	Other Information Elements	Title
5.4.2.3.5.1	Coding Rules	Fully complies
5.4.2.3.5.2	Bearer Capability	Fully complies
5.4.2.3.5.3	Call State	Fully complies
5.4.2.3.5.4	Called Party Number	Fully complies
5.4.2.3.5.5	Called Party Subaddress	Complies, with exception: IE ignored (not used)
5.4.2.3.5.6	Calling Party Number	Fully complies
5.4.2.3.5.7	Calling Party Subaddress	Complies, with exception: IE ignored (not used)
5.4.2.3.5.8	Cause	Fully complies
5.4.2.3.5.9	Channel Identification	Fully complies
5.4.2.3.5.10	Display	Fully complies
5.4.2.3.5.11	High Layer Compatibility	Complies, with exception: IE ignored (not used)
5.4.2.3.5.12	Low Layer Compatibility	Fully complies
5.4.2.3.5.13	Progress Indicator	Complies, with exception: IE ignored (not used)
5.4.2.3.5.14	Restart Indicator	Fully complies

Compliance with TS014 Layer 3 - call control procedure

Clause	Title	Compliance
5.4.2.4	Circuit Switched Call Control Procedures	Title
5.4.2.4.1	General	Noted
(1 of 3)		



Compliance with TS014 Layer 3 - call control procedure (continued)

Clause	Title	Compliance
5.4.2.4.2	Call Establishment at the Originating Interface	Title
5.4.2.4.2.1	Call Request	Not supported Outgoing calls are not supported
5.4.2.4.2.2	B-channel Selection - Originating	Not supported Outgoing calls are not supported
5.4.2.4.2.3	Overlap Sending	Not supported Overlap signaling is not supported
5.4.2.4.2.4	Invalid Call Information	Noted
5.4.2.4.2.5	Call Proceeding	Title
5.4.2.4.2.5.1	En-bloc Sending	Not supported Outgoing calls are not supported
5.4.2.4.2.5.2	Overlap Sending	Not supported Overlap signaling is not supported
5.4.2.4.2.6	Notification of Interworking at the Originating Interface	Title
5.4.2.4.2.6.1	Calls Leaving the ISDN Environment	Noted
5.4.2.4.2.6.2	Call Entering the ISDN Environment	Noted
5.4.2.4.2.7	Call Confirmation Indication	Noted Outgoing calls are not supported
5.4.2.4.2.8	Call Connected	Noted Outgoing calls are not supported
5.4.2.4.2.9	Call Rejection	Noted Outgoing calls are not supported
(2 of 3)		



Compliance with TS014 Layer 3 - call control procedure (continued)

Clause	Title	Compliance
5.4.2.4.3	Call Establishment at the Destination Interface	Title
5.4.2.4.3.1	Incoming Call	Fully complies
5.4.2.4.3.2	Compatibility Checking	Fully complies
5.4.2.4.3.3	B-channel Selection	Fully complies
5.4.2.4.3.4	Overlap Receiving	Not supported Overlap signaling is not supported
5.4.2.4.3.5	Call Confirmation	Title
5.4.2.4.3.5.1	Response to En-bloc SETUP	Fully complies
5.4.2.4.3.5.2	Receipt of CALL PROCEEDING	Not supported. No support for outgoing calls
5.4.2.4.3.5.3	Call Failure Procedures	Not supported. No support for outgoing calls
5.4.2.4.3.6	Call Acceptance	Fully complies
5.4.2.4.3.7	Active Indication	Fully complies
5.4.2.4.4	Call Clearing	Title
5.4.2.4.4.1	Terminology	Noted
5.4.2.4.4.2	Exception Conditions	Fully complies
5.4.2.4.4.3	Clearing by the CE	Fully complies
5.4.2.4.4.4	Clearing by the Network	Fully complies
5.4.2.4.4.5	Clear Collision	Fully complies
5.4.2.4.5	In-band Tones and Announcements	Not supported N/A to data
5.4.2.4.6	Restart Procedure	Title
5.4.2.4.6.1	Overview	Noted
5.4.2.4.6.2	Sending RESTART	Fully complies
5.4.2.4.6.3	Receipt of a RESTART	Fully complies
5.4.2.4.6.4	RESTART Collision	Fully complies
5.4.2.4.7	Call Collision	Fully complies
(3 of 3)		



Compliance with TS014 Layer 3 - handling error conditions

Clause	Title	Compliance
5.4.2.4.8	Handling of Error Conditions	Title
5.4.2.4.8.1	Overview	Fully complies
5.4.2.4.8.2	Protocol Discriminator Error	Fully complies
5.4.2.4.8.3	Message Too Short	Fully complies
5.4.2.4.8.4	Call Reference Error	Title
5.4.2.4.8.4.1	Invalid Call Reference Format	Fully complies
5.4.2.4.8.4.2	Call Reference Procedural Error	Fully complies
5.4.2.4.8.5	Message Type or Message Sequence Errors	Title
5.4.2.4.8.5.1	Message Type Not Implemented	Fully complies
5.4.2.4.8.5.2	Message Incompatible with Call State	Fully complies
5.4.2.4.8.6	Information Element Errors	Title
5.4.2.4.8.6.1	Essential Information Element Missing	Fully complies
5.4.2.4.8.6.2	Essential Information Element Content Error	Fully complies
5.4.2.4.8.6.3	Information Element Not Implemented	Fully complies
5.4.2.4.8.6.4	Non-essential Information Element Content Error	Fully complies
5.4.2.4.8.6.5	Information Element Length Error	Fully complies
5.4.2.4.8.7	Data Link Reset	Fully complies
5.4.2.4.8.8	Data Link Failure	Complies, with exceptions: - overlap signaling condition is not applicable - timer T309 is not supported
5.4.2.4.8.9	STATUS ENQUIRY Procedure	Fully complies
5.4.2.4.8.10	Receiving a STATUS Message in Response to a STATUS ENQUIRY	Fully complies



Compliance with TS014 Layer 3 - system parameters

Clause	Title	Compliance
5.4.2.5	System Parameters	Title
5.4.2.5.1	Timers and Counters in the Network Side	Noted
5.4.2.5.2	Timers and Counters in the CE Side	Title
	T304	Not supported
	T305	Fully complies
	T308	Fully complies
	T309 (Optional)	Not supported
	T313	Fully complies
	T316	Fully complies
	T317	Fully complies
	N316	Fully complies

Compliance with TS014 Layer 3 - SDL diagrams

Clause	Title	Compliance
5.4.3	Primary Rate Access, ISDN CE-Network Interface Layer 3 SDL Diagrams	Title
5.4.3.1	General	Noted
5.4.3.2	Layer 3 Architecture	Noted
5.4.3.3	Process Interactions	Fully complies

Compatibility requirements

The clauses related to this group are in TS014 chapters 6 and 7.

Compliance with TS014 Layer 3 - compatibility requirements

Clause	Title	Compliance
6	Testing	Noted
7	Compliance with International Standards	Noted



Appendices requirements

The clauses related to this group are in TS014 appendices A to K.

Compliance with TS014 appendices A to K

Clause	Title	Compliance
A	CRC Multiframe	Noted
B	Cause Definition	Noted
C	Layer 2 - Test Matrix 1	Noted
D	Layer 2 - Test Matrix 2	Noted
E	Layer 2 - State Initialization (States 4-8)	Noted
F	Layer 2 - State Check Sequences (States 4-8)	Noted
G	Layer 3 - Primary Rate Test Matrix 1	Noted
H	Layer 3 - Primary Rate Test Matrix 2	Noted
I	Layer 3 - Primary Rate Test Matrix 3	Noted
J	Layer 3 - Primary Rate Test Matrix 4	Noted
K	Layer 3 - Initialization Sequences for All User States	Noted



European ISDN compliance statement

Use the following sections to learn more about Nortel Multiservice Switch frame relay ISDN switched access compliance.

Navigation

- [Terminology \(page 305\)](#)
- [European ISDN standards for PRI access \(page 306\)](#)
- [French Euro-numeris ISDN network standards \(page 333\)](#)
- [German ISDN network standards \(page 336\)](#)

Terminology

This appendix uses the following terms to define compliance:

- **Noted.** The standard provides background information or clarification.
- **Fully complies.** Frame relay ISDN switched access fully complies with this section.
- **Complies, with exception(s).** Frame relay ISDN switched access does not completely comply with the text. Exceptions are listed.
- **Not supported.** Frame relay ISDN switched access does not support this section of the standard.
- **N/A.** Not applicable to frame relay ISDN switched access.
- **Title.** This is a clause number and title of the standard.

There are some additional explanatory notes. When no text accompanies a particular section number, this means the original document has no text within this section. The number appears for clarity and completeness only.

The clauses of the ETS 300 125 are bundled into functional groupings as in SAR of ETS 300 156. Each grouping appears in a separate table. In this appendix, the following terms define the SAR:

- **GID** The section provides General Information and Definitions
- **M** Mandatory



- **O** Optional
- **N/A** Not applicable

European ISDN standards for PRI access

The following three standards define the European ISDN primary rate interface (PRI) access:

- Layer 1 specifications:
ETS 300 011, *Integrated Services Digital Network (ISDN); User-network interface data link layer specification Application of ITU-T Recommendation Q.920/I.441*, April 1992
- Layer 2 specifications:
ETS 300 125, *Integrated Services Digital Network (ISDN); User-network interface data link layer specification Application of ITU-T Recommendation Q.920/I.441*, September 1991
- Layer 3 specifications:
ETS 300 102, *Integrated Services Digital Network (ISDN); User-network interface layer 3 Specifications for basic call control*, December 1990, amended October 1993

In addition, the following documents define requirements for connecting terminal equipment to the European ISDN PRI access:

- ETS 300 156 and ETS 300 156/A1 (referred to as NET5), *Attachment requirements for terminal equipment to connect to an ISDN using ISDN primary rate access*, September 1992, amended March 1995
- i-CTR4, *Commission Decision of 18th November 1994 on a Common Technical Regulation for the pan-European Integrated Service Digital Network (ISDN) primary rate access (94/796/EC)*, published in the Official Journal of the European Communities No L-329, Vol 37, December 20th 1994

Document i-CTR4 also refers to:

- draft prTBR 4, *Integrated Services Digital Network (ISDN); Attachment requirements for terminal equipment to connect to an ISDN using ISDN primary rate access*, April 1994, published as TBR 4 on November 1995

Frame relay ISDN switched access compliance to the pan-European ISDN standard for PRI access addresses the i-CTR4 requirement. However, because i-CTR4 is basically a list of test cases along with a list of cross-references to the above-cited standards, the following compliance tables are based on NET-5 Static Attachment Requirements (SAR) with an indication of requirements that are included in i-CTR4.



Layer 1 requirements

See NN10600-170 *Nortel Multiservice Switch 7400 Hardware Description* for details on how the DS1C, E1C and 32-port MSA function processors comply to the applicable sections of the ISO physical layer (layer 1) specification.

Layer 2 requirements

This section describes frame relay ISDN switched access compliance to the European ISDN standard ETS 300 125, *Integrated Services Digital Network (ISDN): User-network interface data link layer specification Application of ITU-T Recommendation Q.920/I.441*, ETSI September 1991.

Layer 2 - service characteristics requirements

ETS 300 125 Part 1 (Q.920)	prETS 300 156 SAR	Included in i-CTR4	Title	Compliance
1	N/A		General	Noted
2	N/A		Concepts and Terminology	Noted
3	N/A		Overview description of LAPD functions and procedures	Noted
4			Service characteristics	
4.1	N/A		General	Noted
4.2			Services provided to Layer 3	
4.2.1	M		Unacknowledged information transfer service	Noted
4.2.2	M		Acknowledged information transfer service	Noted
4.3	M		Services provided to layer management	Noted
4.4	See [1]		Administrative services	Not Applicable
4.5	N/A		Model of the data link service	Noted
4.6	M		Service required from the physical layer	Noted
5	N/A		Data link layer - Management structure	Noted
[1] It is mandatory to provide administrative services for handling TEI assignment, and checking and removal procedures (BRI access). The provision of administrative services for data link parameter passing do not apply.				



Layer 2 - frame structure requirements

ETS 300 125 Part 2 (Q.921)	prETS 300 156 SAR	Included in i-CTR4	Title	Compliance
1	N/A		General	Noted
2			Frame structure for peer-to-peer communication.	Title
2.1	M		General	Fully complies
2.2	M		Flag sequence	Fully complies
2.3	M		Address field	Fully complies
2.4	M		Control field	Fully complies
2.5	M		Information I field	Fully complies
2.6	M		Transparency	Fully complies
2.7	M		FCS field	Fully complies
2.8			Format Convention	Title
2.8.1	N/A		Numbering convention	Noted
2.8.2	M		Order of bit transmission	Fully complies
2.8.3	M		Field mapping convention	Fully complies
2.9	M		Invalid frames	Fully complies
2.10	M		Frame abort	Fully complies

Layer 2 - procedural and formal requirements

ETS 300 125 Part 2 (Q.921)	prETS 300 156 SAR	Included in i-CTR4	Title	Compliance
3			Elements of procedures and formats of fields for data link layer peer-to-peer communication	
3.1	N/A		General	Noted
3.2	M		Address field format	Fully complies
3.3			Address field variables	
(1 of 3)				



Layer 2 - procedural and formal requirements (continued)

ETS 300 125 Part 2 (Q.921)	prETS 300 156 SAR	Included in i-CTR4	Title	Compliance
3.3.1	M		Address field extension bit (EA)	Fully complies
3.3.2	M		Command/Response field bit (C/R)	Fully complies
3.3.3	M		Service Access Point Identifier (SAPI)	Fully complies
3.3.4			Terminal Endpoint Identifier (TEI)	
3.3.4.1	M		TEI for broadcast data link connection	Not Applicable
3.3.4.2	M		TEI for point to point data link connection	Not Applicable
Annex A	M		Provision of point-to-point signalling connections	Fully complies
3.4	M		Control field formats	Fully complies
3.4.1	M		Information transfer (I) format	Fully complies
3.4.2	M		Supervisory (S) format	Fully complies
3.4.3	M		Unnumbered (U) format	Fully complies
3.5	M		Control field parameters and associated state variables	Noted.
3.5.1	M		Poll/Final bit	Fully complies
3.5.2			Multiple frame operation - variables and sequence numbers	Fully complies
3.5.2.1	M		Modulus	Fully complies
3.5.2.2	M		Send state variables V(S)	Fully complies
3.5.2.3	M		Acknowledge state variable V(A)	Fully complies
3.5.2.4	M		Send sequence number N(S)	Fully complies
3.5.2.5	M		Receive state variable V(R)	Fully complies
3.5.2.6	M		Receive sequence number N(R)	Fully complies
3.5.3	N/A		Unacknowledged operation - variables and parameters	Noted
3.6			Frame types	
3.6.1	M		Commands and responses	Fully complies
3.6.2	M		Information (I) command	Fully complies
(2 of 3)				



Layer 2 - procedural and formal requirements (continued)

ETS 300 125 Part 2 (Q.921)	prETS 300 156 SAR	Included in i-CTR4	Title	Compliance
3.6.3	M		Set Asynchronous Balance Mode Extended (SABME) command	Fully complies
3.6.4	M		DISConnect command (DISC)	Fully complies
3.6.5	M		Unnumbered information (UI) command	Fully complies
3.6.6	M		Receive Ready (RR) command/response	Fully complies
3.6.7	M		Reject (REJ) command/response	Fully complies
3.6.8	M		Receive Not Ready (RNR) command/response	Fully complies
3.6.9	M		Unnumbered Acknowledgment (UA) Response	Fully complies
3.6.10	M		Disconnect Mode (DM) response	Fully complies
3.6.11	N/A		Frame Reject (FRMR) response	Noted
3.6.12	N/A		eXchange IDentification (XID) command/response	Noted
(3 of 3)				

Layer 2 - layer-to-layer communication requirements

ETS 300 125 Part 2 (Q.921)	prETS 300 156 SAR	Included in i-CTR4	Title	Compliance
4	N/A		Elements for layer-to-layer communication	Noted



Layer 2 - parameter requirements

ETS 300 125 Part 2 (Q.921)	prETS 300 156 SAR	Included in i-CTR4	Title	Compliance
5.9	M		List of system parameters	Fully complies
5.9.1	M [1]		Timer T200	Fully complies
5.9.2	M [1]		Maximum number of retransmissions (N200)	Fully complies
5.9.3	M [1]		Maximum number of octets in an information field (N201)	Fully complies
5.9.4	M [1]		Maximum number of transmission of the TEI identity request message (N202)	Not Applicable
5.9.5	M [1]		Maximum number of outstanding I-frames (k)	Fully complies
5.9.6	N/A		Timer T201	Not supported
5.9.7	M		Timer T202	Fully complies
5.9.8	O		Timer T203	Fully complies
[1] The default values for the system parameters (appropriate to the SAP(s) implemented in the TE) shall be implemented in the TE.				

Layer 2 - unacknowledged information transfer requirements

ETS 300 125 Part 2 (Q.921)	prETS 300 156 SAR	Included in i-CTR4	Title	Compliance
5.2			Procedures for unacknowledged information transfer	
5.2.1	N/A		General	Noted
5.2.2	M	Yes	Transmission of unacknowledged information	Fully complies
5.2.3	M	Yes	Receipt of unacknowledged info	Fully complies
(1 of 2)				



Layer 2 - unacknowledged information transfer requirements (continued)

ETS 300 125 Part 2 (Q.921)	prETS 300 156 SAR	Included in i-CTR4	Title	Compliance
Annex C	[1]		An SDL representation of the broadcast procedures of the data link layer	Noted
[1] These annexes specify the protocols described in the main body of the standard using SDLs and state transition tables. As such, they place no requirements on the TE other than identified in the procedures as defined in ETS 300 125, PART2, Q.921, Clause 5.				
(2 of 2)				

Layer 2 - TEI management requirements

ETS 300 125 Part 2 (Q.921)	prETS 300 156 SAR	Included in i-CTR4	Title	Compliance
5.3			TEI management	
5.3.1	M		General	Noted
5.3.2	M		TEI assignment procedure	Not Applicable (TEI=0 is always used)
5.3.2.1	M	Yes	Expiry of timer T202	Fully complies
5.3.3			TEI check procedure	
5.3.3.1	N/A		Use of the TEI check procedure	Not Applicable (TEI=0 is always used)
5.3.3.2	M		Operation of the TEI check procedure	Not Applicable (TEI=0 is always used)
5.3.4	M		TEI removal procedure	Not Applicable (TEI=0 is always used)
5.3.4.1	M		Action taken by the data link layer entity receiving MDL-REMOVE-REQUEST primitive	Not Applicable (TEI=0 is always used)
5.3.4.2	M		Conditions for TEI removal	Not Applicable (TEI=0 is always used)
5.3.5	O		TEI identity verify procedure	Not supported
5.3.6			Formats and codes	
5.3.6.1	M		General	Fully complies
(1 of 2)				



Layer 2 - TEI management requirements (continued)

ETS 300 125 Part 2 (Q.921)	prETS 300 156 SAR	Included in i-CTR4	Title	Compliance
5.3.6.2	M	Yes	Layer management entity identifier	Fully complies
5.3.6.3	M		Reference number (Ri)	Fully complies
5.3.6.4	M		Message type	Fully complies
5.3.6.5	M		Action indicator (Ai)	Fully complies
(2 of 2)				

Layer 2 - negotiation of data link layer parameters requirements

ETS 300 125 Part 2 (Q.921)	prETS 300 156 SAR	Included in i-CTR4	Title	Compliance
5.4	N/A		Automatic negotiation of data link layer parameters	Not supported
Appendix IV	N/A		Automatic negotiation of data link layer parameters	Not supported

Layer 2 - establishment/release of MF operations requirements

ETS 300 125 Part 2 (Q.921)	prETS 300 156 SAR	Included in i-CTR4	Title	Compliance
5.5			Procedure for the establishment and release of multiple frame operations	Title
5.5.1	M		Establishment of multiple frame operation	Fully complies
5.5.1.1	M	Yes	General	Fully complies
5.5.1.2	M		Establishment procedures	Fully complies
5.5.1.3	M	Yes	Procedure on expiry of timer T200	Fully complies
5.5.2	M		Information transfer	Fully complies
(1 of 2)				



Layer 2 - establishment/release of MF operations requirements (continued)

ETS 300 125 Part 2 (Q.921)	prETS 300 156 SAR	Included in i-CTR4	Title	Compliance
5.5.3			Termination of multiple frame operation	Title
5.5.3.1	M		General	Fully complies
5.5.3.2	M		Release procedure	Fully complies
5.5.3.3	M		Procedure on expiry of timer T200	Fully complies
5.5.4	M		TEI-assigned state	Fully complies
5.5.5			Collision of unnumbered commands and responses	
5.5.5.1	M		Identical transmitted and received commands	Fully complies
5.5.5.2	M		Different transmitted and received commands	Fully complies
5.5.6	M		Unsolicited DM response and SABME or DISC command	Fully complies
(2 of 2)				

Layer 2 - frame structure requirements

ETS 300 125 Part 2 (Q.921)	prETS 300 156 SAR	Included in i-CTR4	Title	Compliance
5.6			Procedure for information transfer in multiple frame operation	
5.6.1	M	Yes	Transmitting I-frames	Fully complies
5.6.2	M	Yes	Receiving I-frames	Fully complies
5.6.2.1	M		P bit set to 1	Fully complies
5.6.2.2	M	Yes	P bit set to 0	Fully complies
5.6.3			Sending and receiving acknowledgments	
5.6.3.1	M		Sending acknowledgments	Fully complies
5.6.3.2	M	Yes	Receiving acknowledgments	Fully complies
(1 of 2)				



Layer 2 - frame structure requirements (continued)

ETS 300 125 Part 2 (Q.921)	prETS 300 156 SAR	Included in i-CTR4	Title	Compliance
5.6.4	M	Yes	Receiving REJ-frames	Fully complies
5.6.5	M	Yes	Receiving RNR-frames	Fully complies
5.6.6	M		Data link layer own receiver busy condition	Fully complies
5.6.7	M	Yes	Waiting acknowledgment	Fully complies
(2 of 2)				

Layer 2 - re-establishment of MF operation requirements

ETS 300 125 Part 2 (Q.921)	prETS 300 156 SAR	Included in i-CTR4	Title	Compliance
5.7			Re-establishment of multiple frame operation	
5.7.1	M		Criteria for re-establishment	Fully complies
5.7.2	M	Yes	Procedures	Fully complies

Layer 2 - handling of exception condition requirements

ETS 300 125 Part 2 (Q.921)	prETS 300 156 SAR	Included in i-CTR4	Title	Compliance
5.8	N/A		Exception condition reporting and recovery	Noted
5.8.1	M	Yes	N(S) sequence error	Fully complies
5.8.2	M		N(R) sequence error	Fully complies
5.8.3	M		Timer recovery condition	Fully complies
5.8.4	M	Yes	Invalid frame condition	Fully complies
5.8.5	M	Yes	Frame rejection condition	Fully complies
(1 of 2)				



Layer 2 - handling of exception condition requirements (continued)

ETS 300 125 Part 2 (Q.921)	prETS 300 156 SAR	Included in i-CTR4	Title	Compliance
5.8.6	N/A		Receipt of an FRMR response frame	Fully complies FRMR frame is never generated. Upon receipt, however, it is handled as required in ETS 300 125
5.8.7	M		Unsolicited response frames	Fully complies
5.8.8	M		Multiple-assignment of TEI value	Fully complies
App. I	N/A		Retransmission of REJ response frames	Not supported
App. II	N/A		Occurrence of MDL-ERROR-Indication within the basic states & actions to be taken by the management entity	Fully complies
(2 of 2)				

Layer 2 - monitor requirements

ETS 300 125 Part 2 (Q.921)	prETS 300 156 SAR	Included in i-CTR4	Title	Compliance
5.10	O		Data link monitor function	Fully complies

Layer 3 requirements

This section describes frame relay ISDN switched access compliance to the European ISDN standard ETS 300 102-1 and ETS 300 102-2.

Because ETS 300 156 SAR relies mostly on ETS 300 104 SAR, an asterisk (*) is used in references to ETS 300 156, to mark cases where ETS 300 156 differs from ETS 300 104.



Layer 3 - general Layer 2 - monitor requirements

ETS 300 102	prETS 300 156 SAR	Included in i-CTR4	Title	Compliance
1	GID		General	Noted

Layer 3 - overview of call control

ETS 300 102	prETS 300 156 SAR	Included in i-CTR4	Title	Compliance
2	GID		Overview of call control	Noted
2.1	GID		Circuit switched calls	Noted
2.1.1	GID		Call states at the user side of the interface	Noted
2.1.1.1	M		Null state (U0)	Fully complies
2.1.1.2	M		Call Initiated (U1)	Fully complies However, there is No support for outgoing calls (see [1])
2.1.1.3	O		Overlap Sending (U2)	Not supported No support for overlap signalling
2.1.1.4	Note1		Outgoing call proceeding (U3)	Fully complies However, there is No support for outgoing calls (see N[1])
2.1.1.5	Note1		Call delivered (U4)	Fully complies However, there is No support for outgoing calls (see [1])
2.1.1.6	M		Call present (U6)	Fully complies
2.1.1.7	O		Call received (U7)	Not supported
2.1.1.8	M		Connect request (U8)	Fully complies
2.1.1.9	O		Incoming call proceeding (U9)	Fully complies
(1 of 2)				



Layer 3 - overview of call control (continued)

ETS 300 102	prETS 300 156 SAR	Included in i-CTR4	Title	Compliance
2.1.1.10	M		Active (U10)	Fully complies
2.1.1.11	M		Disconnect request (U11)	Fully complies
2.1.1.12	M		Disconnect indication (U12)	Fully complies
2.1.1.13	N/A * [3]		Suspend request (U15)	Not Applicable
2.1.1.14	N/A * [3]		Resume request (U17)	Not Applicable
2.1.1.15	M		Release request (U19)	Fully complies
2.1.1.16	O		Overlap receiving (U25)	Not supported
2.1.2	N/A		Call states at the network side of the interface	N/A
2.2	N/A		Packet mode access connections	N/A
2.3	N/A		Temporary signalling connections	N/A
2.4	GID *		States associated with global call reference	Noted
2.4.1	GID *		Call states at the user side of the interface	Noted
2.4.1.1	M *		Null (Rest 0)	Fully complies
2.4.1.2	O [2]		Restart request (R1)	Fully complies
2.4.1.3	M *		Restart (Rest 2)	Fully complies
2.4.2	N/A		Call states at the network side of the interface	N/A
[1] The ETS 300 156 states that to fulfill procedures specified in ETS 300 102, this state must be implemented even if a given TE does not need this call state. [2] The ETS 300 156 states that the user shall implement the Restart Procedures (see subclause 5.5 of ETS 300 102-1). Whether the user actually sends a RESTART message is optional, though. So, the RESTART request (REST 1) state is also optional. [3] The use of the call rearrangement procedure is restricted to basic access (see ETSI Requirement, subclause 5.6, of ETS 300 102-1).				
(2 of 2)				



Layer 3 - message definition and content

ETS 300 102	prETS 300 156 SAR	Included in i-CTR4	Title	Compliance
3	GID		Message functional definitions and content	Noted
3.1	GID		Messages for circuit mode connections	Noted
3.1.1	M O		ALERTING - receipt of - generation of	- Complies with exception: message triggers a STATUS message - Not supported The receipt of ALERTING can only occur in dial-out, which is not supported.
3.1.2	M O		CALL PROCEEDING - receipt of - generation of	- Complies with exception: message triggers a STATUS message. - Fully complies The receipt of CALL PROCEEDING can only occur in an outgoing call, which is not supported.
3.1.3	N/A		CONGESTION CONTROL	Not supported
3.1.4	M		CONNECT	Fully complies
3.1.5	M O		CONNECT ACKNOWLEDGE - receipt of - generation of	- Fully complies - Not supported The generation of CONNECT ACKNOWLEDGE can only occur in an outgoing call, which is not supported.
(1 of 3)				



Layer 3 - message definition and content (continued)

ETS 300 102	prETS 300 156 SAR	Included in i-CTR4	Title	Compliance
3.1.6	M		DISCONNECT	Fully complies
3.1.7	N/A		FACILITY	Not supported
3.1.8			INFORMATION	Complies with exception: message triggers a STATUS message. Overlap Signaling not supported
3.1.9			NOTIFY	Complies with exception: message triggers a STATUS
3.1.10	M O		PROGRESS - receipt of - generation of	- Complies with exception: message triggers a STATUS - Not supported No support to interworking event
3.1.11	M		RELEASE	Fully complies
3.1.12	M		RELEASE COMPLETE	Fully complies
3.1.13	N/A*		RESUME	Not supported Message triggers a STATUS
3.1.14	N/A*		RESUME ACKNOWLEDGE	Not supported Message triggers a STATUS
3.1.15	N/A*		RESUME REJECT	Not supported Message triggers a STATUS
3.1.16	M		SETUP	Fully complies (see [1])
3.1.17	O		SETUP ACKNOWLEDGE	Not supported No support to overlap signaling
3.1.18	M		STATUS	Fully complies
(2 of 3)				



Layer 3 - message definition and content (continued)

ETS 300 102	prETS 300 156 SAR	Included in i-CTR4	Title	Compliance
3.1.19	O		STATUS ENQUIRY	Fully complies
3.1.20	N/A*		SUSPEND	Not supported Message triggers a STATUS
3.1.21	N/A*		SUSPEND ACKNOWLEDGE	Not supported Message triggers a STATUS
3.1.22	N/A*		SUSPEND REJECT	Not supported Message triggers a STATUS
3.1.23	N/A		USER INFORMATION	Not supported Message triggers a STATUS
3.2	N/A		Messages for packet mode connections	Not supported No support to packet mode
3.3	N/A		Messages for user to user signalling not associated with circuit switched calls	Not supported
3.4	GID*		Messages used with the global call reference	
3.4.1	O		RESTART	Fully complies
3.4.2	O		RESTART ACKNOWLEDGE	Fully complies
3.4.3	O		STATUS	Fully complies
[1] If the calling party number IE is not part of the SETUP message, the call can be accepted (depending on SETUP validation) if an empty ChannelGroup exists.				
(3 of 3)				



Layer 3 - message format and coding

ETS 300 102	prETS 300 156 SAR	Included in i-CTR4	Title	Compliance
4	[1]		General Message Format and Information Element Coding	
4.1	M		Overview	Fully complies
4.2	M		Protocol Discriminator	Fully complies
4.3	M		Call reference	Fully complies
4.4	M		Message Type	Fully complies
4.5			Other Information Elements	
4.5.1	M		Coding rules	Fully complies
4.5.1.1	M		Codeset 0	Fully complies
4.5.1.2	[6]		Codeset 5 - receipt of - generation of	- Fully complies - Not Applicable
4.5.2	M		Extensions of Codesets	Fully complies
4.5.3	[2] M O		Locking Shift Procedure - receipt of - generation of	- Fully complies - Not Applicable
4.5.4	[2] M O		Non-locking Shift Procedure - receipt of - generation of	- Fully complies - Not Applicable
4.5.5	M		Bearer Capability	Fully complies
4.5.6	N/A*		Call Identity	Not Applicable
4.5.7	M		Call State	Fully complies
4.5.8	M		Called Party Number	Fully complies
4.5.9	N/A		Called Party Subaddress	Not supported
4.5.10	N/A		Calling Party Number	Fully complies
4.5.11	N/A		Calling Party Subaddress	Not supported
4.5.12	M		Cause	Fully complies
4.5.13	M		Channel Identification	Fully complies
4.5.14	N/A		Congestion Level	Not supported
(1 of 3)				



Layer 3 - message format and coding (continued)

ETS 300 102	prETS 300 156 SAR	Included in i-CTR4	Title	Compliance
4.5.15	[7]		Display - receipt of - generation of	- Fully complies - Not Applicable
4.5.16	O		High Layer Compatibility	Complies with exceptions: IE ignored.
4.5.17	N/A		Keypad facility	Not supported
4.5.18	O		Low Layer Compatibility	Fully complies
4.5.19	N/A		More Data	Not supported
4.5.20	N/A		Network-specific facilities	Not supported
4.5.21	[2]		Notification Indicator - receipt of - generation of	- Complies with exceptions: IE ignored. - Not supported
4.5.22	[2]		Progress Indication - receipt of - generation of	- Complies with exceptions: IE ignored. - Not Applicable
4.5.23	N/A		Repeat Indicator	Not supported
4.5.24	M*		Restart Indicator	Fully complies
4.5.25	N/A		Segmented Message	Not supported
4.5.26	[4]		Sending Complete	Not supported
4.5.27	N/A		Signal	Not supported
4.5.28	O		Transit Network Selection	Not supported
4.5.29	N/A		User to user	Not supported
4.6.1	[5]		Date/Time - receipt of - generation of	- Complies with exceptions: IE ignored. - Not Applicable
(2 of 3)				



Layer 3 - message format and coding (continued)

ETS 300 102	prETS 300 156 SAR	Included in i-CTR4	Title	Compliance
4.6.2	N/A		Facility	Not supported
4.7	N/A		Information element for packet communications	Not supported
[1] Refer to ETS 300 102-1, subclause 4.2 to subclause 4.5.28 for the content of each information element. [2] It is optional whether a TE ever transmits this information element, but all TEs must be able to receive the information element and handle it correctly as defined in the procedures specified in ETS 300 102-1, clause 5 and all its subsections. [3] The use of call rearrangement procedure is restricted to basic access (see ETSI Requirement, subclause 5.6 of ETS 300 102-1). [4] While it is optional for a TE to generate this information element, TEs that do implement the overlap receiving procedure shall recognize the information element and handle it correctly as defined in the procedures specified in ETS 300 102-1, clause 5 and all its subsections. [5] It is mandatory that the TE recognizes the Date/time information element. This means that no error handling procedures shall initiate upon its receipt. ETSs/NETs that relate to a specific terminal type can require the use of the contents of the Date/time information element, but otherwise it can be ignored. [6] It is mandatory for the TE to handle correctly the shifting procedures between different codesets. It is optional whether the TE recognizes or generates any of the information elements in codeset 5. [7] It is mandatory that the TE recognizes the display information element meaning, and that no error handling procedures shall initiate upon its receipt, but display of the contents of the information element is optional.				
(3 of 3)				

Layer 3 - system parameters

ETS 300 102	prETS 300 156 SAR	Included in i-CTR4	Title	Compliance
9			List of System Parameters	
9.1	N/A		Timers in the network side	N/A
9.2			Timers in the user side	
	N/A		T301	Not Applicable No support to overlap signaling
(1 of 3)				



Layer 3 - system parameters (continued)

ETS 300 102	prETS 300 156 SAR	Included in i-CTR4	Title	Compliance
	[1]		T302	Not Applicable No support to overlap signaling
	O		T303	Not supported No support for outgoing calls
	O		T304	Not supported No support for outgoing calls
	M		T305	Fully complies
	M		T308	Fully complies
	O		T309	Not supported
	O		T310	Not supported No support for outgoing calls
	M		T313	Fully complies
	[1]		T314	Not Applicable No support to segmented messages
	M* [2]		T316	Fully complies
	M*		T317	Fully complies
	N/A* [3]		T318	Not Applicable
	N/A* [3]		T319	Not Applicable
	N/A		T321	Not Applicable
(2 of 3)				



Layer 3 - system parameters (continued)

ETS 300 102	prETS 300 156 SAR	Included in i-CTR4	Title	Compliance
	[1]		T322	Fully complies
[1] Mandatory if the corresponding procedure is implemented, otherwise not applicable. [2] If the TE is capable of sending a RESTART message, timer T316 is mandatory, otherwise it is not applicable. [3] The use of the call rearrangement procedure is restricted to basic access (see ETSI Requirement, subclause 5.6 of ETS 300 102-1).				
(3 of 3)				

Layer 3 - call establishment

ETS 300 102	prETS 300 156 SAR	Included in i-CTR4	Title	Compliance
5			Circuit Switched Call Control Procedures	Fully complies with mandatory elements
5.1	M		Call Establishment at Origination Interface	Not Applicable No support for outgoing calls
5.1.1	M		Call Request	Not Applicable No support for outgoing calls
5.1.2	M	Yes	B-channel selection originating	Not Applicable No support for outgoing calls
5.1.3	O	Yes	Overlap sending	Not Applicable No support for outgoing calls
5.1.4	N/A		Invalid Call Information	N/A
5.1.5			Call Proceeding	Title
5.1.5.1	M	Yes	Call Proceeding, en-bloc sending	Not Applicable No support for outgoing calls
(1 of 4)				



Layer 3 - call establishment (continued)

ETS 300 102	prETS 300 156 SAR	Included in i-CTR4	Title	Compliance
5.1.5.2	M	Yes	Call Proceeding, overlap sending	Not Applicable No support for outgoing calls
5.1.6	M O		Notification of interworking at the originating interface • receipt of • generation of	• Not Applicable • No support for outgoing calls N/A (for user side)
5.1.7	M	Yes	Call Confirmation Indication	Not Applicable No support for outgoing calls
5.1.8	M	Yes	Call Connected	Not Applicable No support for outgoing calls
5.1.9	N/A		Call Rejection	N/A (for user side)
5.1.10	N/A		Transit Network Selection	Not Applicable.
5.2	M	Yes	Call Establishment at Destination Interface	Fully complies with
5.2.1	M	Yes	Incoming Call	Fully complies with
5.2.2	M	Yes	Compatibility Checking	Fully complies with
5.2.3			B-channel selection - destination	
5.2.3.1	M		SETUP message delivered by point to point data link	Fully complies with
5.2.3.2	N/A*		SETUP message delivered by broadcast data link	Not Applicable (no support to broadcast data link)
5.2.4	O	Yes	Overlap receiving	Not supported (no support to overlap signalling)
5.2.5			Call confirmation	
(2 of 4)				



Layer 3 - call establishment (continued)

ETS 300 102	prETS 300 156 SAR	Included in i-CTR4	Title	Compliance
5.2.5.1	M	Yes	Response to en-bloc SETUP or Completion of Overlap Receiving	Fully complies with
5.2.5.2	N/A		Receipt of CALL PROCEEDING and ALERTING	N/A (no support to broadcast data link)
5.2.5.3	N/A		Called user clearing during call establishment	N/A
5.2.5.4	N/A		Call Failure procedures	N/A
5.2.6	M O		Notification of interworking at terminating interface • receipt of • generation of	• Noted • N/A
5.2.7	M		Call Accept	Fully complies with
5.2.8		Yes	Active Indication	Fully complies with
5.2.9	N/A		Non-selected User Clearing	Noted
5.7	M		Call collision	Fully complies with
Annex B	M	Yes	Compatibility checking	Fully complies with
Annex C	O		Transit Network Selection	Not supported
Annex D	N/A		Extension for Symmetric Call	Noted
Annex E	N/A		Network specific facility selection	Not Applicable
Annex F	N/A		D channel backup procedures	Not Applicable
Annex G	GID		Cause definitions	Noted
Annex H	GID		Examples of information element coding	Noted
Annex I	M		Use of progress indicator	Noted
Annex J	N/A		Examples of cause value and location for busy condition	Noted
Annex K	N/A		Message segmentation procedures	Not Applicable
Annex L	M		Low Layer Information coding principles	Noted
Annex M	O		Low Layer Compatibility Negotiation	Not supported
(3 of 4)				



Layer 3 - call establishment (continued)

ETS 300 102	prETS 300 156 SAR	Included in i-CTR4	Title	Compliance
Annex N	O		Procedures for establishment of bearer connection prior to call acceptance	Not supported
Annex O	N/A		Optional procedures for bearer service change	Not Applicable
Annex N	O		Procedure for Establishment of Bearer Connection prior to Call Acceptance	Not supported
Annex O	N/A		Optional Procedures for Bearer Service Change	Not Applicable
(4 of 4)				

Layer 3 - call clearing

ETS 300 102	prETS 300 156 SAR	Included in i-CTR4	Title	Compliance
5.3			Call Clearing	
5.3.1	GID		Terminology	Noted
5.3.2	M	Yes	Exception conditions	Fully complies
5.3.3	M	Yes	Clearing initiated by the User	Fully complies
5.3.4	M		Clearing by the Network	Fully complies
5.3.4.1	M		Clearing when Tones/Announcements provided	Not Applicable No in-band tones-announcements
5.3.4.2	M	Yes	Clearing when Tones/Announcements not provided	Fully complies
5.3.4.3	M	Yes	Completion of Clearing	Fully complies
5.3.5	M	Yes	Clear Collision	Fully complies
Appendix I	[1]		Usage of Cause Values	Fully complies
[1] Appendix I summarizes the use of Cause Values in ETS 300 102-1 and thus place no additional requirements on the TE other than those which have already been identified in this document.				



Layer 3 - tones and announcements

ETS 300 102	prETS 300 156 SAR	Included in i-CTR4	Title	Compliance
5.4	N/A		In-band Tones and Announcements	Not Applicable

Layer 3 - restart

ETS 300 102	prETS 300 156 SAR	Included in i-CTR4	Title	Compliance
5.5	M* [1]	Yes	Restart Procedure	Fully complies
5.5.1	O		Sending RESTART	Fully complies
5.5.2	M* [1]	Yes	Receiving RESTART	Fully complies
[1] The procedures of subclause 5.5.1 and 5.5.2 of ETS 300 102-1 shall be implemented (see subclause 5.5c of ETS 300 102-1). It is optional, however, whether sending a RESTART will ever be invoked.				

Layer 3 - call re-arrangement

ETS 300 102	prETS 300 156 SAR	Included in i-CTR4	Title	Compliance
5.6	GID* [1]		Call Re-arrangements	Not Applicable
5.6.1	N/A* [1]		Call Suspension	Not Applicable
5.6.2	N/A [1]		Call Suspended	Not Applicable
5.6.3	N/A* [1]		Call Suspend Error	Not Applicable
(1 of 2)				



Layer 3 - call re-arrangement (continued)

ETS 300 102	prETS 300 156 SAR	Included in i-CTR4	Title	Compliance
5.6.4	N/A* [1]		Call Re-establishment	Not Applicable
5.6.5	N/A* [1]		Call Resume Error	Not Applicable
5.6.6	N/A* [1]		Double Suspension	Not Applicable
5.6.7	O [1]		Call Re-arrangement Controlled by an NT2	Not Applicable
5.9	M O	Yes	User Notification Procedure • receipt of • generation of	• Fully complies • Not supported
[1] The use of the call rearrangement procedure is restricted to basic access, see ETS Requirements, subclause 5.6 of ETSI 300 102-1.				
(2 of 2)				

Layer 3 - user-to-user signaling

ETS 300 102	prETS 300 156 SAR	Included in i-CTR4	Title	Compliance
7	N/A		Procedures for user-to-user signaling	Not applicable

Layer 3 - handling of error conditions

ETS 300 102	prETS 300 156 SAR	Included in i-CTR4	Title	Compliance
5.8	M		Handling of Error Conditions	Fully complies
5.8.1	M		Protocol Discriminator Error	Fully complies
5.8.2	M		Message too Short	Fully complies
(1 of 2)				



Layer 3 - handling of error conditions (continued)

ETS 300 102	prETS 300 156 SAR	Included in i-CTR4	Title	Compliance
5.8.3			Call Reference Error	Title
5.8.3.1	M		Invalid Call Reference Format	Fully complies
5.8.3.2	M	Yes	Call Reference Procedural Error	Fully complies
5.8.4	M	Yes	Message Type or Message Sequence Errors	Fully complies
5.8.5			General Information Element Errors	
5.8.5.1	M		Information Element out of Sequence	Fully complies
5.8.5.2	M		Duplicated Information Element	Fully complies
5.8.6			Mandatory Information Element Errors	
5.8.6.1	M	Yes	Mandatory Information Element Missing	Fully complies
5.8.6.2	M		Mandatory Information Element Content Error	Fully complies
5.8.7			Non-mandatory Information Element Error	
5.8.7.1	M	Yes	Unrecognized Information Element	Fully complies
5.8.7.2	M	Yes	Non-mandatory Information Element Content Error	Fully complies
5.8.8	M		Data Link Reset	Fully complies
5.8.9	M		Data Link Failure	Fully complies
5.8.10	O	Yes	Status Enquiry Procedure	Fully complies
5.8.11	M	Yes	Receiving a STATUS message	Fully complies
(2 of 2)				

Layer 3 - packet communications

ETS 300 102	prETS 300 156 SAR	Included in i-CTR4	Title	Compliance
6	N/A		Procedures for Packet Communication	Not Applicable Packet mode not supported



Layer 3 - supplementary service

ETS 300 102	prETS 300 156 SAR	Included in i-CTR4	Title	Compliance
8	N/A		Application of Circuit-switched Supplementary Service to Terminals using Stimulus Procedures	Not supported

Layer 3 - user side and network side SDL diagrams

ETS 300 102	prETS 300 156 SAR	Included in i-CTR4	Title	Compliance
Annex A	[1]		User side and network side SDL diagrams	Noted
[1] Annex A of ETS 300 102-1 refers out to ETS 300 102-2, the contents of which are a specification using SDL of the procedures specified in clause 5 of ETS 300 102-1. This Annex, therefore, places no additional requirements of the TE other than those that have already been identified in this document.				

French Euro-numeris ISDN network standards

This section describes frame relay ISDN switched access compliance to the French Euro-numeris ISDN network standards. The following documents define these standards:

- CSE P 10-20 A, *List of additional tests applicable to terminals complying with NET3 and NET 5 for connection to the Euro-Numeris network*, June 1994
- CSE P 10-21 A, *List of additional requirements applicable to terminals complying with NET3 and NET 5 for connection to the Euro-Numeris network*, June 1994
- ETS 300 156 (also known as NET 5), *Attachment requirements for terminal equipment to connect to an ISDN using ISDN primary rate access*, September 1992, amended March 1995
- ETS 300 102, *Integrated Services Digital Network (ISDN); User-network interface layer 3 specifications for basic call control*, December 1990, amended October 1993



Compliance statements for French ISDN delta

CSE P 10-21 A	French, additional requirements to ETS 300 156 (NET5)	Title	Compliance
1	None	Scope	Noted
2	None	Applicable conditions	Noted
3	None	Complement to NET3 (and NET5) or selected options	Noted
3.1		Layer 1	Title
3.1.1	5.2.3	CRC4 procedure Abstract: This procedure is not allowed.	Fully complies
3.2		Layer 2	Title
3.2.1	None	Management of the point-to-point configuration Abstract: TEI assignment	N/A to PRI access TEI always equal 0 (point- to-point configuration)
3.3		Layer 3	Title
3.3.1	10.2.2.4 A.3.21.1.2.1 A.3.21.1.2.2 A.3.21.2	RESTART procedure Abstract: Any TE connected to Euro- Numeris is not allowed to send a RESTART message, nor to activate a RESTART procedure.	Not compliant [1] A RESTART procedure is used at the initialization and whenever a lock/ unlock is used.
3.3.2	10.2.1.4	Maximum length of the Called and Calling Party Subaddress Abstract: the maximum length must be less than 8 octets.	N/A The Called (Calling) Party Subaddress IE is not supported
3.3.3	10.2.2.7	STATUS ENQUIRY procedure Abstract: the STATUS ENQ procedure must never been activated unless it is associated to unsynchronized call.	Fully complies
3.3.4	10.2.1.4	Specific cause values Abstract: New cause values are used (not Q.931 defined)	Fully complies
3.3.5	10.2.1.4	Procedure for codeset extension or modification Abstract: the codeset escape procedure is not used.	N/A Only codeset 0 is supported
(1 of 3)			



Compliance statements for French ISDN delta (continued)

CSE P 10-21 A	French, additional requirements to ETS 300 156 (NET5)	Title	Compliance
3.3.6	None	RESUME procedure Abstract: A RELEASE COMPLETE is sent back instead of a RESUME REJECT.	N/A RESUME, RESUME ACK, and RESUME REJECT are not supported
3.3.7	10.2.1.4	Management of the extension bit Abstract: Management of the extension bit is not used.	Fully complies
3.3.8	10.2.1.4	Maximum length of High Layer Compatibility IE Abstract: The maximum length must be less than 4 octets.	N/A this IE is not used
3.3.9	10.2.1.5	T310 timer Abstract: when the TEI is different than 0 timer T310 is not used at the network side.	Fully complies TEI is always 0
3.3.10	10.2.1.4	Maximum length of Calling and Called Party Number IE Abstract: The maximum length must be less than 13 and 21 octets respectively.	Fully complies
3.3.11	10.2.1.4	Maximum length of Channel Identification IE Abstract: The maximum length must be less than 5 octets (primary access).	Fully complies
3.3.12	10.2.2.7	Identification of the originator of the call Abstract: A TE connected to Euro-Numeris must accept a SETUP with 2 Calling Party Number.	Not compliant [2]
3.3.13	10.2.2.1	Segmentation procedure Abstract: The segmentation procedure is not used.	N/A
3.3.14	10.2.2.1	Bearer Capability and High Layer Compatibility negotiation Abstract: the negotiation procedure is not used.	Fully complies
(2 of 3)			



Compliance statements for French ISDN delta (continued)

CSE P 10-21 A	French, additional requirements to ETS 300 156 (NET5)	Title	Compliance
3.3.15	10.2.1.4	Acceptation of messages related to Euro-Numeris supplementary Service Abstract: a TE may send a STATUS with cause 99 for any unrecognized message or IE (specific to Euro-Numeris supplementary Service)	Fully complies
3.3.16	10.2.1.4	Maximum length of Low Layer Compatibility Abstract: The maximum length must be less than 15 octets.	Fully complies
[1] The Euro-Numeris ISDN network does not allow any TE to activate the RESTART procedure. The TE must never send a RESTART. [2] The Euro-Numeris network can use two calling party number information elements in a SETUP message.			
(3 of 3)			

German ISDN network standards

This section describes frame relay ISDN switched access compliance to the German ISDN network standards. The following documents define these standards:

- BAPT 223 ZV 25, *Type Approval Specification for ISDN equipment for attachment to ISDN Basic Rate Accesses of Deutsch Bundespost Telekom in conjunction with the technical regulations according to the Commission Decision on a common technical regulation for the pan-European ISDN Basic access and Commission Decision on a common technical regulation for the pan-European ISDN Primary Rate Access and national supplementary requirements*, Edition October 1994.
- ETS 300 156 (also known as NET 5), *Attachment requirements for terminal equipment to connect to an ISDN using ISDN primary rate access*, September 1992, amended March 1995
- ETS 300 102, *Integrated Services Digital Network (ISDN); User-network interface layer 3 specification for basic call control*, December 1990, amended October 1993



Compliance statements for German ISDN delta

BAPT 223 ZV 25	Title	Compliance
1	Scope	Noted
2	Additional requirements for terminal equipment for attachment to basic access	Title
2.1	Receipt of a SETUP message via the point-to-point data link	N/A Only primary rate access is supported
2.2	Test case for the receipt of a SETUP message via the point-to-point data link	N/A Only primary rate access is supported
2.3	Restart procedure	N/A Only primary rate access is supported
2.4	Test case for the receipt of a RESTART message	N/A Only primary rate access is supported
2.5	Type approval requirements for the DSS1 supplementary services of a terminal equipment	N/A Supplementary service are not supported.
2.5.1	Requirements for the implementation of supplementary services	N/A
3	Additional requirements for terminal equipment for attachment at the primary rate access	Title
3.1	Requirements for the assignment of TEI values	Fully complies with
3.2	Testing of assignment of TEI values	Fully complies with
3.3	Type approval requirements for the DSS1 supplementary services of a terminal equipment	N/A Supplementary service are not supported.
3.3.1	Requirement for the implementation of supplementary services	N/A



Japanese ISDN compliance statement

Use the following sections to learn more about frame relay ISDN switched access Compliance with *INS-Net Interface and Services*, issued by Nippon Telegraph and Telephone (NTT) Corporation - ISDN Promotion Department, March, 1993.

Navigation

- [Terminology \(page 338\)](#)
- [Layer 1 requirements \(page 338\)](#)
- [Layer 2 requirements \(page 339\)](#)
- [Layer 3 requirements \(page 343\)](#)

Terminology

In this appendix, the following terms define Compliance with INS-Net:

- Noted. The standard provides background information or clarification.
- Fully complies. Nortel Multiservice Switch frame relay ISDN switched access fully complies with this section.
- Complies, with exception. Multiservice Switch frame relay ISDN switched access does not completely comply with the text. Exceptions are listed.
- Not supported. Multiservice Switch frame relay ISDN switched access does not support this section of the standard.
- N/A. Not applicable to Multiservice Switch frame relay ISDN switched access.
- Title. This is a clause number and title of the standard.

Layer 1 requirements

See NN10600-170 *Nortel Multiservice Switch 7400 Hardware Description* for details on how the DS1C, E1C and 32-port MSA function processors comply to the applicable sections of the ISO physical layer (layer 1) specification.



Layer 2 requirements

The clauses related to the ISO network layer (layer 2) specification are in *INS-Net Interface and Services, Part 4*.

Compliance with INS-Net Layer 2 - frame structure requirements

Clause	Title	Compliance
4.1	General	Noted
4.2	Frame structure for peer-to-peer communication	Title
4.2.1	Flag sequence	Fully complies
4.2.2	Address field	Fully complies
4.2.3	Control field	Fully complies
4.2.4	Transparency	Fully complies

Compliance with INS-Net Layer 2 - procedural and format requirements

Clause	Title	Compliance
4.3	Elements of procedures and formats of fields for data link layer peer-to-peer communication	Title
4.3.1	Address field variables	Title
4.3.1.1	Terminal Endpoint Identifier (TEI)	Title
4.3.1.1.1	TEI for point to point data link connection	Fully complies (only TEI 0 is used)
4.3.2	Control field formats	Fully complies
4.3.2.1	Unnumbered (U) format	Fully complies
4.3.3	Control field parameters and associated state variables	Title
4.3.3.1	Multiple frame operation - variables and sequence numbers	Fully complies
4.3.3.2	Unacknowledged operation - variables and parameters	Noted
4.3.4	Frame types	Title
4.3.4.1	Commands and responses	Fully complies
(1 of 2)		



Compliance with INS-Net Layer 2 - procedural and format requirements

Clause	Title	Compliance
4.3.4.2	Receive Not Ready (RNR) command/response	Fully complies
4.3.4.3	eXchange IDentification (XID) command/response	Noted
(2 of 2)		

Compliance with INS-Net Layer 2 - layer-to-layer communication requirements

Clause	Title	Compliance
4.4	Elements for layer-to-layer communication	Title

Compliance with INS-Net Layer 2 - TEI management requirements

Clause	Title	Compliance
4.5	Definition of the peer-to-peer procedures of the data link layer	Noted
4.5.1	TEI management procedures	Title
4.5.1.1	General	Noted
4.5.1.2	TEI assignment procedure	N/A (TEI=0 is always used)
4.5.1.3	TEI check procedure	Title
4.5.1.3.1	Use of the TEI check procedure	N/A (TEI=0 is always used)
4.5.1.4	TEI removal procedure	N/A (TEI=0 is always used)
4.5.1.4.1	Action taken by the data link layer entity receiving MDL-REMOVE-REQUEST primitive	N/A (TEI=0 is always used)
4.5.1.4.2	Conditions for TEI removal	N/A (TEI=0 is always used)
4.5.1.5	TEI identity verify procedure	Title
4.5.1.5.1	General	N/A (TEI=0 is always used)
4.5.1.5.2	Operation of the TEI identity verify procedure	N/A (TEI=0 is always used)
(1 of 2)		



Compliance with INS-Net Layer 2 - TEI management requirements (continued)

Clause	Title	Compliance
4.5.1.6	Format and codes	Title
4.5.1.6.1	Action indicator (Ai)	N/A (TEI=0 is always used)
(2 of 2)		

Compliance with INS-Net Layer 2 - negotiation of data link layer parameters requirements

Clause	Title	Compliance
4.5.2	Automatic negotiation of data link layer parameters	Not supported

Compliance with INS-Net Layer 2 - establishment/release of MF operation requirements

Clause	Title	Compliance
4.5.3	Procedure for the establishment and release of multiple frame operations	Title
4.5.3.1	Establishment of multiple frame operation	Title
4.5.3.1.1	General	Fully complies
4.5.3.1.2	Establishment procedures	Fully complies
4.5.3.2	Termination of multiple frame operation	Title
4.5.3.2.1	Procedure on expiry of timer T200	Fully complies
4.5.3.3	TEI-assigned state	Fully complies

Compliance with INS-Net Layer 2 - information transfer in MF operation requirements

Clause	Title	Compliance
4.5.4	Procedure for information transfer in multiple frame operation	Title
4.5.4.1	Receiving REJ-frames	Fully complies
4.5.4.2	Receiving RNR-frames	Fully complies
4.5.4.3	Waiting acknowledgment	Fully complies



Compliance with INS-Net Layer 2 - handling of exception conditions requirements

Clause	Title	Compliance
4.5.5	Exception condition reporting and recovery	Title
4.5.5.1	N(S) sequence error	Fully complies
4.5.5.2	Frame rejection condition	Fully complies

Compliance with INS-Net Layer 2 - parameter requirements

Clause	Title	Compliance
4.5.6	List of system parameters	Fully complies

Compliance with INS-Net Layer 2 - monitor requirements

Clause	Title	Compliance
4.5.7	Data link monitor function	Fully complies

Compliance with INS-Net Layer 2 - SDL for point-to-point procedures requirements

Clause	Title	Compliance
4.6	SDL for point-to-point procedures	Title
4.6.1	Cover notes	Noted
4.6.2	SDL representation	Fully complies with
4.7	State transition table of the point-to-point procedures of the data link layer	Title
4.7.1	Definition of basic states	Noted



Compliance with INS-Net Layer 2 - state transition of the point-to-point procedures requirements

Clause	Title	Compliance
4.8	Retransmission of REJ response frames	Noted
4.9	Optional basic access deactivation procedures	Not supported

Layer 3 requirements

The clauses related to the ISO network layer (layer 3) specification are in *INS-Net Interface and Services, Part 5*.

Compliance with INS-Net Layer 3- general information and definitions

Clause	Title	Compliance
5	ISDN User-Network Interface: Layer 3 Specification	Title
5.1	General	Noted
5.2	Overview of call control	Noted
5.2.1	Circuit Switched Call	Noted
5.2.1.1	Call States at the user side of the interface	Fully complies
5.2.1.2	Network Call States	N/A Only incoming calls are supported
5.2.2	Temporary signaling connections	N/A Only circuit switched calls are supported

Compliance with INS-Net Layer 3- message definition and content requirements

Clause	Title	Compliance
5.3	Message functional definitions and content	Noted
5.3.1	Messages for circuit mode connections	Noted
(1 of 4)		



Compliance with INS-Net Layer 3- message definition and content requirements (continued)

Clause	Title	Compliance
5.3.1.1	ALERTING • receipt of • generation of	• Complies with exception: message trigger is a STATUS message. • Not supported The receipt of ALERTING can only occur in a dial-out call, which is not supported.
5.3.1.2	CALL PROCEEDING • receipt of • generation of	• Complies with exception: message trigger is a STATUS message. • Fully complies The receipt of CALL PROCEEDING can only occur in an outgoing call, which is not supported.
5.3.1.3	CONGESTION CONTROL	Not supported
5.3.1.4	CONNECT	Fully complies
5.3.1.5	CONNECT ACKNOWLEDGE • receipt of • generation of	• Fully complies • Not supported The generation of CONNECT ACKNOWLEDGE can only occur in an outgoing call, which is not supported.
5.3.1.6	DISCONNECT	Fully complies
5.3.1.7	FACILITY	Not supported
(2 of 4)		



Compliance with INS-Net Layer 3- message definition and content requirements (continued)

Clause	Title	Compliance
5.3.1.8	INFORMATION	Complies with exception: message trigger is a STATUS message Overlap Signaling is not supported
5.3.1.9	NOTIFY	Complies with exception: message trigger is a STATUS message.
5.3.1.10	PROGRESS • receipt of • generation of	• Complies with exception: message trigger is a STATUS message. • Not supported No support to interworking event
5.3.1.11	RELEASE	Fully complies
5.3.1.12	RELEASE COMPLETE	Fully complies
5.3.1.13	RESUME	Not supported Message trigger is a STATUS message
5.3.1.14	RESUME ACKNOWLEDGE	Not supported Message trigger is a STATUS message
5.3.1.15	RESUME REJECT	Not supported Message trigger is a STATUS message
5.3.1.16	SETUP	Fully complies (see [1])
5.3.1.17	SETUP ACKNOWLEDGE	Not supported No support for overlap signaling
5.3.1.18	STATUS	Fully complies
5.3.1.19	STATUS ENQUIRY	Fully complies
5.3.1.20	SUSPEND	Not supported Message trigger is a STATUS message
5.3.1.21	SUSPEND ACKNOWLEDGE	Not supported Message trigger is a STATUS message
(3 of 4)		



Compliance with INS-Net Layer 3- message definition and content requirements (continued)

Clause	Title	Compliance
5.3.1.22	SUSPEND REJECT	Not supported Message trigger is a STATUS message
5.3.1.23	USER INFORMATION	Not supported Message trigger is a STATUS message
5.3.2	Messages for user to user signalling not associated with circuit switched calls	Not supported
5.3.3	Messages used with the global call reference	Title
5.3.3.1	RESTART	Fully complies
5.3.3.2	RESTART ACKNOWLEDGE	Fully complies
5.3.3.3	STATUS	Fully complies
[1] If the calling party number IE is not in the SETUP message, the network can accept the call if an empty ChannelGroup exists (depending on SETUP validation).		
(4 of 4)		

Compliance with INS-Net Layer 3- message format and coding requirements

Clause	Title	Compliance
5.4	General Message Format and Information Element Coding	Title
5.4.1	Call reference	Fully complies
5.4.2	Message Type	Fully complies
5.4.3	Other Information Elements	Title
5.4.3.1	Coding Rules	Fully complies
5.4.3.2	Extensions of Codesets	Fully complies
5.4.3.3	Locking Shift Procedure	N/A
5.4.3.4	Non-Locking Shift Procedure	Fully complies
5.4.3.5	Bearer Capability	Fully complies
5.4.3.6	Call State	Fully complies
5.4.3.7	Called Party Number	Complies with exception: unknown is supported and the value received is interpreted as ISDN/ telephony numbering plan - recommendation E.164.
(1 of 2)		



Compliance with INS-Net Layer 3- message format and coding requirements (continued)

Clause	Title	Compliance
5.4.3.8	Calling Party Number	Complies with exception: unknown is supported and the value received is interpreted as ISDN/ telephony numbering plan - recommendation E.164.
5.4.3.9	Cause	Fully complies
5.4.3.10	Channel Identification	Fully complies
5.4.3.11	Congestion Level	Not supported
5.4.3.12	More Data	Not supported
5.4.3.13	Network-specific Facilities	Not supported
5.4.3.14	Notification Indicator	Not supported
5.4.3.15	Progress Indication	N/A
5.4.3.16	Repeat Indicator	Not supported
5.4.3.17	Restart Indicator	Fully complies
5.4.3.18	Segmented Message	Not supported
5.4.3.19	Sending Complete	Fully complies (Only en-bloc signaling is supported)
5.4.3.20	Signal	Not supported
5.4.3.21	Transit Network Selection	Not supported
5.4.3.22	User to user	Not supported
5.4.4	Supplementary service information elements	N/A No support for supplementary services
5.4.5	NTT's INS-Net specific information element	N/A IEs are ignored
(2 of 2)		



Compliance with INS-Net Layer 3- call establishment requirements

Clause	Title	Compliance
5.5	Circuit Switched Call Control Procedures	Fully complies (mandatory elements)
5.5.1	Call Establishment at Origination Interface	N/A No support for outgoing calls
5.5.1.1	Call Request	N/A No support for outgoing calls
5.5.1.2	B-channel selection originating	N/A No support for outgoing calls
5.5.1.3	Overlap sending	N/A No support for overlap signaling
5.5.1.4	Call Proceeding	Title
5.5.1.4.1	Call Proceeding, overlap sending	N/A No support for overlap signaling
5.5.1.5	Transit Network Selection	N/A
5.5.2	Call Establishment at Destination Interface	Fully complies
5.5.2.1	Incoming Call	Fully complies
5.5.2.2	B-channel selection - destination	Title
5.5.2.2.1	SETUP message delivered by point to point data link	Fully complies
5.5.2.2.2	SETUP message delivered by broadcast data link	N/A No support to broadcast data link
5.5.2.3	Overlap receiving	Not supported No support for overlap signaling
5.5.2.4	Call confirmation	Title
5.5.2.4.1	Called user clearing during call establishment	Fully complies (for point to point data link)
5.5.2.5	Notification of interworking at terminating interface	N/A



Compliance with INS-Net Layer 3- call clearing requirements

Clause	Title	Compliance
5.5.3	Call Clearing	
5.5.3.1	Exception conditions	Fully complies
5.5.3.2	Clearing initiated by the User	Fully complies
5.5.3.3	Clearing by the Network	Fully complies
5.5.3.3.1	Clearing when Tones/Announcements provided	N/A No in-band tones-announcements supported
5.5.3.3.2	Clearing when Tones/Announcements not provided	Fully complies
5.5.3.3.3	Completion of Clearing	Fully complies
5.5.3.4	Clear Collision	Fully complies

Compliance with INS-Net Layer 3- tones and announcements requirements

Clause	Title	Compliance
5.5.4	In-band Tones and Announcements	N/A

Compliance with INS-Net Layer 3- restart requirements

Clause	Title	Compliance
5.5.5	Restart Procedure	Fully complies

Compliance with INS-Net Layer 3- call re-arrangement requirements

Clause	Title	Compliance
5.5.6	Call Re-arrangements	N/A
5.5.6.1	Call Suspension	N/A



Compliance with INS-Net Layer 3- handling of error conditions requirements

Clause	Title	Compliance
5.5.8	Handling of Error Conditions	Title
5.5.8.1	Message Type or Message Sequence Errors	Fully complies
5.5.8.2	General Information Element Errors	Title
5.5.8.2.1	Information Element out of Sequence	Fully complies
5.5.8.3	Mandatory Information Element Errors	Title
5.5.8.3.1	Mandatory Information Element Missing	Fully complies
5.5.8.4	Non-mandatory Information Element Error	Title
5.5.8.4.1	Unrecognized Information Element	Fully complies
5.5.8.4.2	Non-mandatory Information Element Content Error	Fully complies
5.5.8.5	Data Link Failure	Fully complies
5.5.8.11	Receiving a STATUS message	Fully complies

Compliance with INS-Net Layer 3- packet communications requirements

Clause	Title	Compliance
5.6	Packet Communication at the S/T reference point	N/A Packet mode not supported

Compliance with INS-Net Layer 3- user-to-user signaling and supplementary services requirements

Clause	Title	Compliance
5.7	User-to-user Signaling Procedures	N/A
5.8	Application of Circuit-switched supplementary services to terminals using stimulus procedures	N/A No support for supplementary services



Compliance with INS-Net Layer 3- system parameters requirements

Clause	Title	Compliance
5.9	List of System Parameters	Title
5.9.1	Timers in the network side	N/A Only incoming calls are supported
5.9.2	Timers in the user side	
	T301	N/A No support for overlap signaling
	T302	N/A No support for overlap signaling
	T303	Not supported No support for outgoing calls
	T304	Not supported No support for outgoing calls
	T305	Fully complies
	T308	Fully complies
	T309	Not supported
	T310	Not supported No support for outgoing calls
	T313	Fully complies
	T314	N/A No support for segmented messages.
	T316	Fully complies
	T317	Fully complies
	T318	N/A
	T319	N/A
	T321	N/A
	T322	Fully complies



Compliance with INS-Net Layer 3- annex requirements

Clause	Title	Compliance
5.10.1	Annex A: User side and Network side SDL	Fully complies
5.10.2	Annex B: Compatibility checking	Title
5.10.2.1	Called side compatibility checking	Title
5.10.2.1.1	Compatibility checking with addressing information	Fully complies
5.10.2.1.2	User-to-User Compatibility checking	N/A
5.10.2.1.3	User Action Tables	N/A No support to LLC negotiation.
5.10.2.2	Interworking with existing networks	Noted
5.10.3	Annex C: Transit network selection	N/A
5.10.4	Annex D: Extensions for symmetric call operations	Not supported
5.10.5	Annex E: Network specific facility selection	Not supported
5.10.6	Annex F: D channel backup procedures	N/A
5.10.7	Annex G: Cause definitions	Fully complies
5.10.8	Annex H: Examples of information elements coding	Noted
5.10.9	Annex I: Use of progress indicators	Not supported
5.10.10	Annex K: Message segmentation procedures	N/A
5.10.11	Annex M: Low Layer Compatibility Negotiation	Not supported
5.10.12	Annex N: Procedures for establishment of bearer connection prior to call acceptance	Not supported
5.10.13	Annex O: Optional procedures for bearer service change	N/A



North American NI-2 compliance statement

This appendix describes Nortel Multiservice Switch frame relay ISDN switched access compliance with Telcordia standards for North American primary rate access.

The following standards define ISDN PRI access:

- TR-TSY-00754 - ISDN Primary Rate Access Transport System Requirements (Layer 1), July 1990
- TR-TSY-000793 - ISDN D-channel Exchange Access Signaling and Switching Requirements (Layer 2) November 1994
- TR-NWT-001268 - ISDN Primary Rate Interface Call Control Switching and Signaling Generic Requirements for Class II Equipment (Layer 3) May 1997

Navigation

- [Terminology \(page 353\)](#)
- [Layer 1 requirements \(page 354\)](#)
- [Layer 2 requirements \(page 354\)](#)
- [Layer 3 requirements \(page 361\)](#)

Terminology

The following terms define NI-2 compliance:

- Noted. The standard provides background information or clarification.
- Fully complies. Nortel Multiservice Switch frame relay ISDN switched access fully complies with this section.
- Complies, with exception. Multiservice Switch frame relay ISDN switched access does not completely comply with the text. Exceptions are listed.
- Not supported. Multiservice Switch frame relay ISDN switched access does not support this section of the standard.



- N/A. Not applicable to Multiservice Switch frame relay ISDN switched access.
- Title. This is a clause number and title of the standard.

Layer 1 requirements

See NN10600-170 *Nortel Multiservice Switch 7400 Hardware Description* for details on how the DS1C, E1C and 32-port MSA function processors comply to the applicable sections of the ISO physical layer (layer 1) specification.

Layer 2 requirements

The clauses related to the ISO network layer (layer 2) specification are in sections one through eight of the Bellecore standard, "ISDN D-channel Exchange Access Signaling and Switching Requirements (Layer 2)" November 1994

- [NI-2 Layer 2 Introduction \(page 354\)](#)
- [NI-2 Layer 2 - Frame structure requirements \(page 355\)](#)
- [NI-2 Layer 2 Commands and responses requirements \(page 356\)](#)
- [NI-2 Layer 2 Exchange of information requirements \(page 358\)](#)
- [NI-2 Layer 2 DLCI management requirements \(page 359\)](#)
- [NI-2 Layer 2 Error detection and recovery requirements \(page 360\)](#)
- [NI-2 Layer 2 SPCS layer 2 system requirements \(page 361\)](#)
- [NI-2 Layer 2 Services requirements \(page 361\)](#)

NI-2 Layer 2 Introduction

Clause	Title	Compliance
1	Introduction	Noted
1.1	Purpose	Noted
1.2	Scope	Noted
1.3	The Layer 2 Model	Noted
1.4	ISDN D-channel Services	Noted
1.5	Reference Architectures	Noted
1.6	Related Documents	Noted
1.7	Organization Of This Document	Noted
1.8	Requirements Conventions	Noted



NI-2 Layer 2 - Frame structure requirements

Clause	Title	Compliance
2	Frame Structure	Noted
2.1	Numbering Conventions	Fully complies
2.2	Frame Format Types	Fully complies
2.2.1	Type I Frame Format	Fully complies
2.2.2	Type U Frame Format	Fully complies
2.2.3	Type S Frame Format	Fully complies
2.3	Frame Field Types	Noted
2.3.1	Flags	Fully complies
2.3.2	Address	Fully complies
2.3.2.1	Address Field Extension Bit	Fully complies
2.3.2.2	Command/ Response Field Bit	Fully complies
2.3.2.3	Service Access Point Identifier	Fully complies Only SAPI=0 is supported. Other values apply to BRI only.
2.3.2.4	Terminal Endpoint Identifier	Fully complies Only TEI=0 is supported. Other values apply to BRI only.
2.3.2.4.1	Dynamic TEIs	Noted
2.3.2.4.2	Broadcast TEI	Noted
2.3.2.4.3	Fixed TEIs	Noted
2.3.3	Frame Check Sequence	Fully complies
2.3.4	Control Field	Fully complies
2.3.4.1	N (S) Sub-field	Fully complies
2.3.4.2	N (R) Sub-field	Fully complies
2.3.4.3	Poll/Final Sub-field	Fully complies
2.3.4.4	Supervisory Function Sub-field	Fully complies
2.3.4.5	Reserved Sub-field	Fully complies
2.3.4.6	Modifier Sub-field	Fully complies
2.3.5	Information Field	Fully complies
2.4	State Variables	Fully complies
2.4.1	Send State Variable V (S)	Fully complies
(1 of 2)		



NI-2 Layer 2 - Frame structure requirements (continued)

Clause	Title	Compliance
2.4.2	Acknowledge State Variable V (A)	Fully complies
2.4.3	Receive State Variable V (R)	Fully complies
2.4.4	Recovery State Variable V (M)	Fully complies
2.5	Frame Coding and Decoding	Fully complies
2.5.1	Zero-bit Stuffing and Removal (Transparency)	Fully complies
2.5.2	FCS Processing	Fully complies
2.5.3	Frame Type Decoding and Octet Counts	Fully complies
2.5.4	Address Decoding	Fully complies
2.6	Invalid Frames	Fully complies
2.6.1	Frame Delimiting	Fully complies
2.6.1.1	Idle Codes	Fully complies
2.6.1.2	Flag Detection and Processing	Fully complies
2.6.1.3	Abort Code	Fully complies
2.6.2	Zero-bit Decoding	Fully complies
2.6.3	FCS Calculation	Fully complies
2.6.4	Valid Frame Length	Fully complies
2.6.5	Valid Frame Type Check	Fully complies
2.6.6	Address Field Check	Fully complies
2.6.7	DLCI Check	Fully complies
2.6.7.1	Currently Active	Fully complies
2.6.7.2	Change to Active	Fully complies
2.6.7.3	Deactivate	Fully complies
(2 of 2)		

NI-2 Layer 2 Commands and responses requirements

Clause	Title	Compliance
3	Commands and Responses	Fully complies
3.1	Type I Command	Fully complies
3.2	Type U Commands and Responses	Fully complies
(1 of 2)		



NI-2 Layer 2 Commands and responses requirements (continued)

Clause	Title	Compliance
3.2.1	Set Asynchronous Balanced Mode Extended Command	Fully complies
3.2.1.1	SABME Transmitted by SPCS	Fully complies
3.2.1.2	SPCS Receives SABME	Fully complies
3.2.2	Disconnect Command	Fully complies
3.2.2.1	SPCS Receives DISC	Fully complies
3.2.2.2	DISC Transmitted by SPCS	Fully complies
3.2.3	Unnumbered Information Command	Fully complies
3.2.3.1	Identity Request	N/A Only PRI is supported
3.2.3.2	Identity Assigned	N/A Only PRI is supported
3.2.3.3	Identity Denied	N/A - Only PRI is supported
3.2.3.4	Identity Check Request	N/A - Only PRI is supported
3.2.3.5	Identity Check Response	N/A - Only PRI is supported
3.2.3.6	Identity Remove	N/A - Only PRI is supported
3.2.3.7	Identity Verify	N/A - Only PRI is supported
3.2.4	Unnumbered Acknowledgment Response	Fully complies
3.2.5	Disconnected Mode Response	Fully complies
3.2.6	Frame Reject Response	Fully complies
3.2.6.1	SPCS Receives FRMR	Fully complies
3.2.6.2	SPCS Transmits FRMR	Fully complies
3.2.7	Exchange Identification Command/ Response	Fully complies
3.2.7.1	XID Format and Protocol	Noted
3.2.7.2	Automatic Notification	Noted
3.3	Type S Commands/ Responses	Fully complies
3.3.1	Receive Ready Command/ Response	Fully complies
3.3.2	Reject Command/ Response	Fully complies
3.3.3	Receive Not Ready Command/ Response	Fully complies
(2 of 2)		



NI-2 Layer 2 Exchange of information requirements

Clause	Title	Compliance
4	Exchange of Information	Noted
4.1	Unacknowledged Frame Transfer	Noted
4.2	Acknowledged Frame Transfer	Noted
4.2.1	Initiating Multiple Frame Operation	Noted
4.2.1.1	SPCS Requests Multiple Frame Operation	Fully complies
4.2.1.2	User Requests Multiple Frame Operations	Fully complies
4.2.1.3	Checking Subscription Limits	Fully complies
4.2.1.3.1	Active	Fully complies
4.2.1.3.2	Inactive	Fully complies
4.2.1.3.3	Disabled	Fully complies
4.2.1.3.4	Out-of-Service	Fully complies
4.2.1.4	Establishment Procedures	Noted
4.2.1.4.1	SPCS-originated Establishment	Fully complies
4.2.1.4.2	Timer T200Expiration	Fully complies
4.2.1.4.3	User-Originated Establishment Procedures	Fully complies
4.2.1.4.4	Information Transfer	Fully complies
4.2.1.4.5	Unsolicited DM Response	Fully complies
4.2.2	Information Transfer During Multiple Frame Operations	Noted
4.2.2.1	Transmitting I Frames	Fully complies
4.2.2.2	Receiving I Frames	Fully complies
4.2.2.3	Sending and Receiving Acknowledgments	Fully complies
4.2.2.4	Receiving REJ Frames	Noted
4.2.2.4.1	Multiple Frame Operations	Fully complies
4.2.2.4.2	Test N (R)	Fully complies
4.2.2.4.3	Valid N (R)	Fully complies
4.2.2.4.4	I - Frame Re-Transmission	Fully complies
4.2.2.4.5	Other than Multiple Frame State	Fully complies
(1 of 2)		



NI-2 Layer 2 Exchange of information requirements (continued)

Clause	Title	Compliance
4.2.2.5	Receiving RNR Frames	Noted
4.2.2.5.1	While in Multiple Frame Operations	Fully complies
4.2.2.5.2	While in Timer Recovery	Fully complies
4.2.2.5.3	Other States	Fully complies
4.2.2.6	Polls	Fully complies
4.2.2.7	Link Controller Own Receiver Busy Condition	Fully complies
4.2.2.8	T200 Expiration	Fully complies
4.2.2.9	Clearing Timer Recovery	Fully complies
4.2.3	Termination of Multiple Frame Operation	Fully complies
4.2.3.1	T200 Expiration	Fully complies
4.2.3.2	T203 Expiration	Fully complies
4.2.3.3	TEI Removal	Fully complies
4.2.3.4	Layer 1 Failure	Fully complies
4.2.3.5	User-Transmitted DISC	Fully complies
4.2.3.6	User-Transmitted FRMR	Fully complies
4.2.4	TEI Assigned State	Noted
4.2.5	Link Monitor Function	Fully complies
4.2.5.1	Start Timer T203	Fully complies
4.2.5.2	Stop Timer T203	Fully complies
4.2.5.3	Expiration of Timer 203	Fully complies
(2 of 2)		

NI-2 Layer 2 DLCI management requirements

Clause	Title	Compliance
5	Data Link Connection Identifier (DLCI) Management	Noted
5.1	Active DLCI Table	Fully complies. Only TEI = 0 and SAP = 0 is supported.
5.2	SPCS TEI Management	Noted
(1 of 2)		



NI-2 Layer 2 DLCI management requirements (continued)

Clause	Title	Compliance
5.2.1	TEI Assignment Procedures	N/A TEI = 0 is always used.
5.2.1.1	Identity Request	N/A TEI = 0 is always used.
5.2.1.2	Identity Assigned	N/A TEI = 0 is always used.
5.2.1.3	Identity Denied	N/A TEI = 0 is always used.
5.2.2	TEI Check Procedures	N/A TEI = 0 is always used.
5.2.2.1	Query for a Specific TEI Value	N/A TEI = 0 is always used.
5.2.2.2	Query for Redundant TEIs	N/A TEI = 0 is always used.
5.2.2.3	Query to Perform a TEI Audit	N/A TEI = 0 is always used.
5.2.3	TEI Removal	N/A TEI = 0 is always used.
5.2.4	TEI Identity Verify	N/A TEI = 0 is always used.
(2 of 2)		

NI-2 Layer 2 Error detection and recovery requirements

Clause	Title	Compliance
6	Error Detection and Recovery	Noted
6.1	N (S) Sequence Error	Fully complies
6.2	N (R) Sequence Error	Fully complies
6.3	Lost Transmitted Frames	Fully complies
6.4	Frame Rejection	Fully complies
6.5	Unsolicited Response Frames	Noted



NI-2 Layer 2 SPCS layer 2 system requirements

Clause	Title	Compliance
7	The SPCS Layer 2 System	Noted
7.1	Specification and Description Language	Noted
7.2	The Layer 3 Interface	Noted
7.3	The Layer 1 Interface	Noted
7.4	The Signaling Address Controller	Noted
7.5	The Link Controllers	Noted

NI-2 Layer 2 Services requirements

Clause	Title	Compliance
8	Services	Noted
8.1	ISDN Exchange Access Signaling Service	Noted
8.2	Timers	Noted
8.3	Counters	Noted

Layer 3 requirements

The clauses related to the ISO network layer (layer 3) specification are in sections one through five of the Bellecore standard, "ISDN Primary Rate Interface Call Control Switching and Signaling Generic Requirements for Class II Equipment", May 1997.

- [NI-2 Layer 3 introduction, general descriptions and assumptions \(page 362\)](#)
- [NI-2 Layer 3 call processing \(page 362\)](#)
- [NI-2 Layer 3 procedures and parameters \(page 364\)](#)
- [NI-2 Layer 3 messages and information elements \(page 366\)](#)
- [NI-2 Layer 3 appendix \(page 369\)](#)



NI-2 Layer 3 introduction, general descriptions and assumptions

Clause	Title	Compliance
1	Introduction	Noted
2	General Descriptions and Assumptions	Noted

NI-2 Layer 3 call processing

Clause	Title	Compliance
3	Call Processing	Noted
3.1	General Rules for Message Processing	Noted
3.2	Origination Treatment	N/A Outgoing calls are not supported.
3.3	Termination Treatment	Noted
3.3.1	Termination To An ISDN Primary Rate Interface	Noted
3.3.1.1	Subscription Parameters	Noted
3.3.1.2	Interface Busy/ Idle Status	Noted
3.3.1.3	Call Delivery	Noted
3.3.1.4	Expiration of Timer T303 or T-delay	N/A T303 and T-delay only apply to outgoing calls which are not supported.
3.3.1.5	Call Rejected by the Called Party	Noted
3.3.1.6	Channel Negotiation Procedures	Noted
3.3.1.7	Call Proceeding Indication	Noted
3.3.1.8	Called Party Alerting	Noted
3.3.1.9	Call Progress Indication	Noted
3.3.1.10	Called Party Answer	Noted
3.3.1.11	Delivery of an Incoming X.25 Call	N/A Packet mode is not supported.
3.3.1.12	Remote Network Interworking	N/A Interworking is not supported.
3.3.2	Termination To Non-ISDN Line	N/A
3.3.3	Termination To Non-ISDN Trunk	N/A
(1 of 3)		



NI-2 Layer 3 call processing (continued)

Clause	Title	Compliance
3.3.4	Termination of Calls That Cannot Be Completed	Noted
3.3.4.1	Abandoned Calls	Noted
3.3.4.2	Busy Line Treatment	Noted
3.3.4.3	Intercept Treatment	Noted
3.3.4.4	Unavailable Equipment/ Facilities, Internal or Other Failures	Noted
3.4	Call Clearing Requirements	Noted
3.4.1	Call Clearing of Answered Calls	Noted
3.4.1.1	ISDN PRI Class II-to-ISDN PRI Class II	Noted
3.4.1.2	ISDN PRI Class II-to-Non-ISDN Call	Noted
3.4.1.3	Non-ISDN-to-ISDN PRI Class II Call	Noted
3.4.2	General Call Clearing Procedures	Fully complies with
3.4.2.1	Call Clearing Request Received From an ISDN Class II Equipment	Fully complies with
3.4.2.2	SPCS Initiates Call Clearing to an ISDN PRI Class II Equipment	Fully complies with
3.4.2.3	Clear Collision	Fully complies with
3.4.3	Clearing Packet-Mode Calls	N/A Packet mode is not supported.
3.5	Error Treatment	Noted
3.5.1	Protocol Error Treatment	Noted
3.5.1.1	Error Treatments Associated With an Established Call Reference	Fully complies with
3.5.1.2	Errors During Call Establishment	Fully complies with
3.5.1.3	Errors During an Active Call	Fully complies with
3.5.1.4	Errors During Call Clearing	Fully complies with
3.5.1.5	Errors Not Associated With A Call	Fully complies with
3.5.1.6	System Status Failures	Fully complies with
3.5.1.7	Nonlocking Shift Procedures	Fully complies with
3.5.2	Layer 3 Actions Resulting From A Malfunctioning Data Link	Noted
(2 of 3)		



NI-2 Layer 3 call processing (continued)

Clause	Title	Compliance
3.5.2.1	Layer 3 Procedures	Fully complies with
3.5.2.2	Call Clearing By Remote User During Recovery Timing	Noted
(3 of 3)		

NI-2 Layer 3 procedures and parameters

Clause	Title	Compliance
4	Procedures and Parameters	Noted
4.1	Determination of Interface Characteristics	Noted
4.2	Office and Subscription Parameters	Noted
4.2.1	Parameters per Office for ISDN PRI to Class II Equipment	Noted
4.2.2	Subscription Parameters per Interface	Noted
4.2.2.1	General subscription Parameters per Interface	Noted
4.2.2.2	Packet-Mode Data Service Parameters	N/A Packet mode is not supported.
4.2.3	Packet-Mode Parameters Per DN	N/A Packet mode is not supported.
4.3	Call Reference Administration	Noted
4.3.1	Allocation and Release of Call References	Noted
4.3.2	Call References Used In Messages Sent by the SPCS	Noted
4.3.3	Call References In Messages Received by the SPCS	Noted
4.4	SPCS Channel Assignment	Noted
4.4.1	Originating B-channel Selection	N/A Outgoing calls are not supported.
4.4.2	D/B Channel Selection For Terminating Calls	Fully complies
4.4.2.1	Bearer Capability Indicates Channel Type	Noted
(1 of 3)		



NI-2 Layer 3 procedures and parameters (continued)

Clause	Title	Compliance
4.4.2.2	Throughput Class Indication of D-channel Subscription Parameter	N/A Packet mode is not supported.
4.4.2.3	Selection of a B-channel when the D-channel is Not Applicable	Noted
4.4.2.4	Selection of a B-channel when the D-channel May be Applicable	N/A Packet mode is not supported.
4.4.2.5	Selection of the D-channel	N/A Packet mode is not supported.
4.5	ISDN Code Interpretation	Noted
4.5.1	General	Noted
4.5.2	Enbloc Signaling Method	Noted
4.5.2.1	General Use	Noted
4.5.2.2	Called Party Number Information Element Interpretation	Noted
4.5.2.3	Transit Network Selection Information Element Interpretation	N/A Transit Network Selection is not supported.
4.5.2.4	Operator System Access Information Element Interpretation	N/A Operator System Access is not supported.
4.5.3	ISDN-to-ISDN Packet Connections	N/A Packet mode is not supported.
4.5.4	Interworking with Packet-Switched Public Data Networks	N/A Interworking is not supported.
4.5.5	Subaddressing	N/A Subaddressing is not supported.
4.6	ISDN Routing	N/A
4.7	Network Determination of Interface Busy	Noted
4.7.1	Busy for Circuit-Switched Calls	Noted
4.7.1.1	Busy Limit for Originating Circuit-Switched Calls	Noted
4.7.1.2	Busy Limit for Terminating Circuit-Switched Calls	Noted
4.7.2	Busy for Packet Connections	N/A Packet mode is not supported.
(2 of 3)		



NI-2 Layer 3 procedures and parameters (continued)

Clause	Title	Compliance
4.8	Timers	Fully complies exceptions to the following timers that are not supported: T301, no support for overlap signaling T303, T-delay, T310, no support for outgoing calls T306, T-PROG, no support for interworking T309, T408, T-rest, not supported T402, no support for on-hook T320, no support for packet switching
4.9	Tones and Announcements	N/A
4.10	Notification of D-channel Failure	Noted
(3 of 3)		

NI-2 Layer 3 messages and information elements

Clause	Title	Compliance
5	Layer 3 Messages and Information Elements	Noted
5.1	Message Functional Definitions	Title
5.1.1	Overview	Noted
5.1.2	Messages	Complies, with exception: messages SERVICE and SERVICE ACKnowledge are not supported
5.1.2.1	ALERTing - receipt of - generation of	Complies, with exception. The message triggers a STATUS message Not supported. ALERTING messages only apply to dial-outs which are not supported.
(1 of 4)		



NI-2 Layer 3 messages and information elements (continued)

Clause	Title	Compliance
5.1.2.2	CALL PROCeeding - receipt of - generation of	Not supported. CALL PROCeeding messages only apply to outgoing calls, which are not supported. Complies, with exception. The message triggers a STATUS message
5.1.2.3	CONNect - receipt of - generation of	Not supported. CONNECT messages only apply to outgoing calls which are not supported. Fully complies
5.1.2.4	CONNect ACKnowledge - receipt of - generation of	Fully complies Not supported CONNECT ACKNOWLEDGE messages only occur in outgoing calls which are not supported.
5.1.2.5	DISConnect	Fully complies
5.1.2.6	PROGress - receipt of - generation of	Complies, with exception. The message triggers a STATUS message Not supported. The interworking event is not supported.
5.1.2.7	RELease	Fully complies
5.1.2.8	RELease COMplete	Fully complies
5.1.2.9	REStart	Fully complies
5.1.2.10	REStart ACKnowledge	Fully complies
(2 of 4)		



NI-2 Layer 3 messages and information elements (continued)

Clause	Title	Compliance
5.1.2.11	SETUP • receipt of • generation of	Fully complies Not supported. SETUP messages apply to outgoing calls which are not supported. If an empty ChannelGroup exists, a call may be accepted, even if the Calling Party Number IE is not included in the SETUP message.
5.1.2.12	STATus	Fully complies
5.1.2.13	STATus ENquiry	Fully complies
5.2	Message Structure	Fully complies
5.2.1	Overview	Fully complies
5.2.2	Protocol Discriminator	Fully complies
5.2.3	Call Reference	Fully complies
5.2.4	Message Type	Fully complies
5.2.5	Other Information Elements	Fully complies
5.2.5.1	Shift Information Elements	Complies, with exception: generation of it is not supported.
5.2.5.2	Bearer Capability	Fully complies
5.2.5.3	Call State	Fully complies
5.2.5.4	Called Party Number	Fully complies
5.2.5.5	Called Party Subaddress	Not supported The IE is ignored.
5.2.5.6	Calling Party Number	Fully complies
5.2.5.7	Calling Party Subaddress	Not supported The IE is ignored.
5.2.5.8	Cause	Fully complies
5.2.5.9	Channel Identification	Fully complies
5.2.5.10	Closed User Group	N/A Packet mode is not supported.
(3 of 4)		



NI-2 Layer 3 messages and information elements (continued)

Clause	Title	Compliance
5.2.5.11	High-Layer Compatibility	Fully complies
5.2.5.12	Information Rate	N/A Packet mode is not supported.
5.2.5.13	Low Layer Compatibility	Fully complies
5.2.5.14	Packet-Layer Binary Parameters	N/A Packet mode is not supported.
5.2.5.15	Packet-Layer Window Size	N/A Packet mode is not supported.
5.2.5.16	Packet Size	N/A Packet mode is not supported.
5.2.5.17	Progress Indicator	N/A Interworking is not supported.
5.2.5.18	Restart Indicator	Fully complies
5.2.5.19	Reverse Charging Indication	N/A Packet mode is not supported.
5.2.5.20	Transit Delay Selection and Indication	N/A Packet mode is not supported.
5.2.5.21	Transit Network Selection	Not supported. The IE is ignored
5.2.5.22	End-to-End Transit Delay	N/A Packet mode is not supported.
5.2.6	National-Specific Information Elements (Codeset 5)	Title
5.2.6.1	Operator System Access	N/A Codeset 5 is not supported.
(4 of 4)		

NI-2 Layer 3 appendix

Clause	Title	Compliance
Appendix A	SDL Diagrams	Noted
Appendix B	D-channel Backup Procedure	Not supported
Appendix C	Data Dictionary	Noted



Standards

Nortel Multiservice Switch FR UNI and FR NNI conform to several standards. See the following sections for related documents.

Navigation

- [Frame relay technical references \(page 370\)](#)
- [ITU-T references \(page 371\)](#)
- [ANSI references \(page 371\)](#)

Frame relay technical references

Refer to the following frame relay technical references for more information.

- *Frame Relay User-to-Network Interface With Extensions, First Issue of Joint specification*, Rev 1.0, September 1990, document number 001-208966, Nortel Networks Inc., Digital Equipment Corporation, StrataCom Inc. and Cisco Systems.
- FRF.12, "Frame Relay PVC Fragmentation Implementation Agreement.", December 1997, Frame Relay Forum Technical Committee.
- FRF.16, "Multilink Frame Relay UNI/NNI Implementation Agreement.", August 1999, Frame Relay Forum Technical Committee.

Nortel Multiservice Switch frame relay service conforms to the following standards:

- American National Standard Institute for Telecommunications (ANSI, as recited in the Frame Relay Forum Implementation agreement)
- International Telecommunications Union (ITU, as recited in the Frame Relay Forum Implementation agreement)
- Vendor Specification (original four party specification)
- Frame Relay Forum Technical Committee



For information about the network element running hyperstream software, see the following documents:

- “Configuring Frame Relay Switching Services”, Software Version BNX 7.2, February 1998, Bay Networks, Part No. 115392-C Rev. A.
- “Event Messages”, February 1998, Bay Networks, Part No. 115412-C Rev. A.
- “MIBs”, February 1998, Bay Networks, Part No. 115416-C Rev. A.
- “Statistics”, February 1998, Bay Networks, Part No. 115410-B Rev. A.

ITU-T references

Refer to the following ITU-T references for more information.

- Recommendation I.122, *Framework For Providing Additional Packet Mode Bearer Services*, Blue Book, ITU, Geneva, 1988.
- Recommendation I.233.1, *ISDN Frame Mode Bearer Services (FMBS) - ISDN Frame Relaying Bearer Service*, ITU, Geneva, 1992.
- Recommendation I.370, *Congestion Management for Frame Relaying Bearer Service*.
- Recommendation Q.922, *ISDN Data Link Layer Specifications for Frame Mode Bearer Services*, ITU, Geneva, 1992.
- Recommendation Q.931, *Digital Subscriber Signalling System No. 1 - (DSS 1) User-Network Interface Layer 3 Specification for Basic Call Control*, ITU, Geneva, 1993.
- Recommendation Q.933, *Integrated Services Digital Network (ISDN) - Digital Subscriber Signaling System No. 1 (DSSI) - Signaling Specification for Frame Relay Bearer Service*, ITU, Geneva, 1992.
- Recommendation X.76, *Network-to-Network Interface Between Public Data Networking Providing the Frame Relay Data Transmission Service*, ITU, Geneva, 1995.
- Recommendation X.76 Amendment 1, *Network-to-Network Interface Between Public Data Networking Providing the Frame Relay Data Transmission Service Amendment 1: Switch virtual circuits*, ITU, Geneva, 1997.

ANSI references

Refer to the following ANSI references for more information.

- ANSI T1.606 *Integrated Services Digital Network (ISDN) - Architectural Framework and Service Description for Frame relaying Bearer Service*, ANSI 1990. Plus Addendum 1 on Congestion Management, ANSI 1992.
- ANSI T1.617 *Integrated Services Digital Network - Digital Subscriber Signaling System No. 1 - Signaling Specification for Frame Relay Bearer*



Service for Digital Subscriber Signaling System Number 1 (DSS1), June 18, 1991.

- *ANSI T1.618 Core Aspects of Frame Mode Bearer Service Protocol for use with Frame Relay Bearer Service, ANSI 1992.*

Nortel Multiservice Switch 7400/15000/20000

Frame Relay Technology Fundamentals

Copyright © 2005 Nortel.
All Rights Reserved.

Publication: NN10600-900
Document status: Standard
Document issue: 7.1S1
Document date: October 2005
Product release: PCR7.1 and up
Job function: Product Fundamentals
Type: NTP
Language type: U.S. English

NORTEL, the globemark design, and the NORTEL corporate logo are
trademarks of Nortel.

